

The Neural Basis of Human Error Processing: Reinforcement Learning, Dopamine, and the Error-Related Negativity

Clay B. Holroyd

University of Illinois at Urbana–Champaign

Michael G. H. Coles

University of Illinois at Urbana–Champaign
and F. C. Donders Centre for Cognitive Neuroimaging

The authors present a unified account of 2 neural systems concerned with the development and expression of adaptive behaviors: a mesencephalic dopamine system for reinforcement learning and a “generic” error-processing system associated with the anterior cingulate cortex. The existence of the error-processing system has been inferred from the error-related negativity (ERN), a component of the event-related brain potential elicited when human participants commit errors in reaction-time tasks. The authors propose that the ERN is generated when a negative reinforcement learning signal is conveyed to the anterior cingulate cortex via the mesencephalic dopamine system and that this signal is used by the anterior cingulate cortex to modify performance on the task at hand. They provide support for this proposal using both computational modeling and psychophysiological experimentation.

Human beings learn from the consequences of their actions. Thorndike (1911/1970) originally described this phenomenon with his law of effect, which made explicit the commonsense notion that actions that are followed by feelings of satisfaction are more likely to be generated again in the future, whereas actions that are followed by negative outcomes are less likely to reoccur. This fundamental *reinforcement learning* principle has been developed by the artificial intelligence community into a body of algorithms used to train autonomous systems to operate independently in complex and uncertain environments (Barto & Sutton, 1997; Sutton & Barto, 1998). Research has also evaluated the neural mechanisms underlying reinforcement learning in biological systems, but these mechanisms are still poorly understood.

In this article, we provide a framework for understanding the neural basis of reinforcement learning in humans. Our proposal links together two areas of research that have, until now, been considered separately. On the one hand, we have previously inferred the existence of a generic, high-level error-processing system in humans from the error-related negativity (ERN), a negative

deflection in the ongoing electroencephalogram (EEG) seen when human participants commit errors in a wide variety of psychological tasks. The ERN appears to be generated in the anterior cingulate cortex. On the other hand, other researchers have argued that the mesencephalic dopamine system conveys reinforcement learning signals to the basal ganglia and frontal cortex, where they are used to facilitate the development of adaptive motor programs. Although the reinforcement learning function attributed to the mesencephalic dopamine system and the error-processing function associated with the ERN appear to be concerned with the same problem—namely, evaluating the appropriateness of ongoing events, and using that information to facilitate the development and expression of adaptive behaviors—a possible relationship between these two systems remains to be explored.

In this article, we propose a hypothesis that unifies the two accounts by explicitly linking the generation of the ERN to the activity of the mesencephalic dopamine system. Specifically, we suggest that when human participants commit errors in reaction-time tasks, the mesencephalic dopamine system conveys a negative reinforcement learning signal to the frontal cortex, where it generates the ERN by disinhibiting the apical dendrites of motor neurons in the anterior cingulate cortex. Furthermore, we suggest that the error signals are used to train the anterior cingulate cortex, ensuring that control over the motor system will be released to a motor controller that is best suited for the task at hand.

To address this issue, we follow a parallel approach involving both computational modeling and empirical investigation. On the one hand, we develop a computational model that we use to simulate the ERN and human behavior in two psychological tasks. On the other hand, we describe two event-related brain potential (ERP) experiments involving the same tasks, and we compare the predictions derived from the model with the experimental results. Although the model is compatible with the neurophysiology and neuroanatomy of the mesencephalic dopamine system, it is not a detailed neural model. Instead, it depicts in an abstract fashion the

Clay B. Holroyd, Neuroscience Program, University of Illinois at Urbana–Champaign; Michael G. H. Coles, Department of Psychology, University of Illinois at Urbana–Champaign and F. C. Donders Centre for Cognitive Neuroimaging, Nijmegen, the Netherlands.

Clay B. Holroyd especially thanks Jesse Reichler for several years of fruitful, and enjoyable, discussion of reinforcement learning and the error-related negativity. Both authors are also grateful to Gary Dell for reading and commenting on parts of a previous version of the article. This research was supported in part by National Institutes of Mental Health (NIMH) predoctoral fellowship MH11530 and NIMH Grant MH41445. Preliminary versions of this study have been presented in poster format (Holroyd, Reichler, & Coles, 1998, 1999).

Correspondence concerning this article should be addressed to Clay B. Holroyd, who is now at the Department of Psychology, Princeton University, Green Hall, Princeton, New Jersey 08544. E-mail: cholroyd@princeton.edu

function we ascribe to the error-processing system underlying the ERN. Specifically, the model adopts a reinforcement learning algorithm, the *method of temporal differences*, that has previously been used by others to describe the activity of the mesencephalic dopamine system. Within this framework, we propose that the ERN is elicited when a neural system first detects that the consequences of an action are worse than expected, and that the associated error signal is used to train the motor system in a way that is consistent with neural network implementations of reinforcement learning principles.

We begin by reviewing several experiments that have provided insight into the system that gives rise to the ERN. We then present a brief overview of the role of the mesencephalic dopamine system in reinforcement learning, and of recent computational models of this system. Next, we review the anatomy and function of the anterior cingulate cortex. Motivated by previous computational models, we then present our own simulations of the ERN and human behavior in two tasks, and we compare the results of the simulations with empirical observations from two separate experiments. We find that the model accounts for much of the experimental data. We conclude that, when it is taken together with neurophysiological and neuroanatomical considerations, the evidence is consistent with our proposal that the ERN is elicited by a dopaminergic system for reinforcement learning. These results suggest promising avenues for future research.

Background

The ERN

When human beings make errors in reaction-time tasks, a negative deflection appears in the ongoing EEG at the time of error commission. This phenomenon is best seen by averaging together several epochs of the EEG, each of which is associated with an erroneous response, to create an average ERP (for an introduction to ERPs, see Coles & Rugg, 1995). The component of the ERP revealed by this procedure is called the ERN, or Ne (Falkenstein, Hohnsbein, Hoormann, & Blanke, 1990; Gehring, 1992; Gehring, Goss, Coles, Meyer, & Donchin, 1993). The onset of the ERN coincides with response initiation, as determined by the onset of the electromyogram (EMG) associated with the responding hand, and peaks roughly 80 ms thereafter. Its spatial distribution lies over frontal-central regions of the scalp, reaching maximum amplitude in a region over the supplementary motor area.

Over the past 10 years or so, we and others have elucidated the nature of the system that generates the ERN, both in terms of structure and function (for reviews, see Coles, Scheffers, & Holroyd, 1998; Falkenstein, Hoormann, Christ, & Hohnsbein, 2000; Gehring, Coles, Meyer, & Donchin, 1995). We believe that the ERN is generated by a high-level, generic, error-processing system. By *error-processing*, we mean that the system is involved in detecting the fact that an error has occurred in a given task and in using that error information to improve performance at the task. By *generic*, we mean that the system is highly flexible, capable of processing errors in a wide variety of contexts. And by *high-level*, we mean that the system is associated with the executive control processes mediated by frontal areas of the brain.

Much of our research has been devoted to elucidating the nature of the error-processing system underlying the ERN. Gehring

(1992) showed that the amplitude of the ERN increased as participants were motivated, via a payoff function, to strive for accuracy over speed in a choice reaction-time task (see also Gehring et al., 1993). His work has suggested that the system that produces the ERN is sensitive to the importance of error commission to the participant. In a different study, participants performed a four-choice reaction-time task by pressing buttons using either of two fingers on either of their left and right hands. Errors could be committed with the wrong hand, the wrong finger, or both the wrong hand and the wrong finger. It was found that the magnitude of the ERN increased with the degree of error, being largest when the incorrect response was committed with both the wrong hand and the wrong finger (Bernstein, Scheffers, & Coles, 1995). Thus, the error-processing system appears to be sensitive to the degree of error.

In another experiment, participants were required to press a button when they estimated that 1 s had elapsed following presentation of a warning stimulus. At the end of the trial, a feedback stimulus was presented indicating whether or not their estimate on that trial was within a criterion. The stimulus that indicated to the participants that their response was not within the criterion (*negative feedback*) elicited an ERN (Miltner, Braun, & Coles, 1997). It is important to note that, because the feedback was delivered some time after the response occurred, the ERN elicited by the feedback was dissociated from the response generation process. This result demonstrates that the ERN is not elicited by the process that causes the error in the first place, by the absence of a process needed to effect the correct response, nor by the execution of a remedial action made in conjunction with the error. Rather, the experiment reveals that the system appears to be concerned with an aspect of error processing that is not directly tied to error commission. This function might include detection of the error itself or the use of the error information to prevent future error repetition.

ERNs are elicited by incorrect responses in a wide variety of tasks (for reviews, see Coles et al., 1998; Gehring et al., 1995). Hence, the error processing system that gives rise to the ERN must be flexible enough to be programmable with arbitrary goals. Moreover, the system can be made sensitive to various sources of error information. As we have mentioned, ERNs can be elicited both by negative feedback and by error commission itself, but ERNs can also be elicited by errors committed with the feet (Holroyd, Dien, & Coles, 1998) or eyes (Nieuwenhuis, Ridderinkhof, Blom, Band, & Kok, 2001; Van 't Ent & Apkarian, 1999) as well as with the hands, indicating that the system is equally sensitive to different output modalities of error commission. Similarly, ERNs can be elicited by negative feedback stimuli presented in the auditory, visual, and somatosensory modalities, so the system must be indifferent to the input modality of the error information (Miltner, Braun, & Coles, 1997). When speed is emphasized, ERNs are also elicited by late responses (Johnson, Otten, Boeck, & Coles, 1997; Luu, Flaisch, & Tucker, 2000). Taken together, the evidence paints a picture of a highly flexible error-processing system.

Converging evidence indicates that the ERN may be generated in the anterior cingulate cortex. Studies using the brain electric source analysis (BESA) technique for source localization (Scherg, 1990; Scherg & Picton, 1991) have suggested that the ERN is generated within that region (Dehaene, Posner, & Tucker, 1994; see also Gehring, Himle, & Nisenson, 2000). The BESA technique yields an anterior cingulate location even when the ERN is elicited

by errors committed with the feet as well as with the hands (Holroyd, Dien, & Coles, 1998) and by negative feedback presented in the auditory, somatosensory, and visual modalities (Miltner, Braun, & Coles, 1997). When a source localization technique was applied to the magnetic equivalent of the ERN, an anterior cingulate location was again suggested (Miltner, Lemke, et al., 1997).

Neurophysiological evidence is also consistent with the position that the ERN is generated in the anterior cingulate cortex. Niki and Watanabe (1979, Figure 7) found error-recognition units in the monkey anterior cingulate sulcus that were activated when the animals received negative feedback in the form of an absence of an expected reward. Similarly, Gamba, Sasaki, and Brooks (1986) found that when monkeys made errors in a simple response task, error-related potentials were generated in the anterior cingulate sulcus. This discovery led Brooks (1986, 1989) to suggest that the anterior cingulate cortex functions as a comparator, comparing the outcome of an action against its intent. Recently, several functional neuroimaging studies found that the anterior cingulate cortex of human participants engaged in reaction-time tasks was more activated on error trials than on correct trials (Braver, Barch, Gray, Molfese, & Snyder, 2001; Carter et al., 1998; Kiehl, Liddle, & Hopfinger, 2000; Menon, Adelman, White, Glover, & Reiss, 2001; Ullsperger & von Cramon, 2001). When taken together, the evidence argues in favor of an anterior cingulate location for ERN generation.

Frontal systems of the brain, including the prefrontal cortex (Damasio, 1994; Luria, 1973; Stuss & Benson, 1986), the anterior cingulate cortex (Devinsky, Morrell, & Vogt, 1995; Posner & DiGirolamo, 1998), and the basal ganglia (L. L. Brown, Schneider, & Lidsky, 1997; Cummings, 1993), are believed to contribute to executive control. Systems concerned with executive function are thought to regulate the most global aspects of human behavior, such as planning and decision making, and are said to come into play when a task is novel or difficult (Baddeley, 1986; Logan, 1985; Norman & Shallice, 1986). One aspect of executive control concerns *response monitoring*, or ensuring that the consequences of an action are consistent with the intent. We believe that the highly flexible nature of the error-processing system associated with the ERN is consistent with what would be expected of an executive control system concerned with response monitoring. The frontal-central scalp distribution of the ERN and its putative generation within the anterior cingulate cortex also point to an association between the ERN process and executive functions implemented in frontal regions. Furthermore, the ERN is abnormal both in individuals with obsessive-compulsive disorder (Gehring et al., 2000) and in individuals with dorsolateral prefrontal damage (Gehring & Knight, 2000), conditions that are associated with executive dysfunction. When considered in this context, the system underlying the ERN appears to contribute to the regulation of the most global aspects of human behavior.

The observations outlined in this section support two critical assumptions of our hypothesis, namely, that the ERN is generated by the anterior cingulate cortex and that the ERN is elicited by a high-level error processing system.

The Mesencephalic Dopamine System

The mesencephalic dopamine system is composed of a small collection of nuclei that project diffusely to the basal ganglia and

cortex (Berger, Gaspar, & Verney, 1991; Gaspar, Berger, Febvret, Vigny, & Henry, 1989; Haber & Fudge, 1997; S. M. Williams & Goldman-Rakic, 1993). These nuclei include the substantia nigra pars compacta and the ventral tegmental area (VTA). In the rat, the cortical projections tend to synapse on the cortex in the frontal midline, including the anterior cingulate cortex; in primates, the projections are spread more widely but still reach their highest density over the medial regions of the frontal cortex (Berger et al., 1991; Gaspar et al., 1989; S. M. Williams & Goldman-Rakic, 1993). The terminals of the dopamine neurons in the basal ganglia (Smith & Bolam, 1990) and prefrontal cortex (Goldman-Rakic, Leranth, Williams, Mons, & Geffard, 1989; Smiley, Williams, Szigeti, & Goldman-Rakic, 1992) end in "synaptic triads," in which the axon of each dopamine neuron makes contact with a local synapse. In the basal ganglia, the synapses are composed of the spine of a local spiny cell abutting the axon of a pyramidal cell; in the frontal cortex, they consist of a pyramidal cell axon synapsing onto the spine of another pyramidal cell.

In the early 1950s, Olds and Milner (1954) discovered that rats will engage in behaviors, like pressing a lever, that result in the delivery of current via stimulating electrodes into certain regions of their brains. The consequence of the stimulation was seen to be reinforcing, in that the rats gradually increased the frequency with which they engaged in the behavior that led to the stimulation. This observation precipitated a search for the neural substrates of reinforcement learning, the results of which implicated the mesencephalic dopamine system (for reviews, see Berridge & Robinson, 1998; Phillips & Fibiger, 1989; Robbins & Everitt, 1996; R. A. Wise & Rompre, 1989).

Accumulating evidence suggests that the mesencephalic dopamine system can facilitate both long-term potentiation and long-term depression at its terminal synapses (Calabresi, Pisani, Mercuri, & Bernardi, 1996; Wickens & Kotter, 1995) and that this learning follows stringent timing requirements, in which the reinforcing signal must coincide with the activity of the pre- and postsynaptic processes (Wickens, Begg, & Arbuthnott, 1996). Dopamine also appears important for the basal ganglia to express learned conditioned responses that might be used for developing new motor programs (Graybiel, Aosaki, Flaherty, & Kimura, 1994; Graybiel & Kimura, 1995; Kawagoe, Takikawa, & Hikosaka, 1998; Kimura & Matsumoto, 1997).

Originally, the mesencephalic dopamine system was widely believed to contribute to reinforcement learning by mediating the feelings of satisfaction experienced by an animal on receiving a reward. As R. A. Wise and colleagues put it,

In introspective language we would say that neuroleptics [that disrupt the midbrain dopamine system] appear to take the pleasure out of normally rewarding brain stimulation, take the euphoria out of normally rewarding amphetamine, and take the "goodness" out of normally rewarding food. (R. A. Wise, Spindler, DeWit, & Gerber, 1978, p. 263)

However, our understanding of the role that mesencephalic dopamine plays in reinforcement learning has evolved significantly over the past 10 or 15 years, and although the "hedonia" hypothesis is still often repeated, it has since been rejected (see Wickelgren, 1997).

The work of Schultz and colleagues has been particularly influential in bringing about this change in view (for reviews, see

Schultz, 1997, 1998; Schultz et al., 1995). Schultz and colleagues recorded spike activity from mesencephalic dopamine cells in monkeys as the monkeys learned to perform various delayed response tasks. In a typical task, a monkey may be required to press a lever every time a green light appears and is rewarded when it does so. Before the monkey has learned how to perform the task, presentation of the reward to the monkey elicits a phasic response in the dopamine neurons. This observation is consistent with the hypothesis that the activity of the mesencephalic dopamine system codes for the hedonic aspects of reward. However, once the monkey has learned to perform the task correctly, then presentation of the reward no longer elicits the phasic dopaminergic response (Ljungberg, Apicella, & Schultz, 1992; Schultz, Apicella, & Ljungberg, 1993). Rather, the conditioned stimulus elicits the phasic activity. With learning, the dopaminergic signal "propagates back in time" from the time the reward is delivered to the onset of the trial, when the conditioned stimulus is presented. Thus the mesencephalic dopamine system can become active in anticipation of a forthcoming reward rather than on delivery of the reward itself. Because the hedonic aspect of reinforcement is presumably associated only with reward consumption, the presence of dopaminergic activity prior to reward presentation is inconsistent with the hypothesis that the mesencephalic dopamine system actually codes for the pleasure elicited by the reward. Similar conclusions have been reached by other investigators (e.g., Phillips, Pfaus, & Blaha, 1991; Richardson & Gratton, 1996).

After the monkey has had some practice at the task in the same experiments, then when a reward is not given, the mesencephalic dopamine neurons decrease their firing rate at the time the reward would normally have been delivered (Hollerman & Schultz, 1998; Ljungberg, Apicella, & Schultz, 1991; Schultz et al., 1993). Dopaminergic activity also falls below baseline when the monkey is presented with a stimulus that predicts punishment (Mirenowicz & Schultz, 1996; but see Guarraci & Kapp, 1999).

Schultz and colleagues have proposed that the dopamine neurons are sensitive to changes in the prediction of the "goodness" of ongoing events: A positive dopamine signal is elicited when an event is better than predicted, and a negative dopamine signal is elicited when an event is worse than predicted. Because modern learning theories specify that learning occurs when an event is unpredicted (e.g., Rescorla & Wagner, 1972), these observations provide insight into how a reinforcement learning algorithm might be implemented in natural systems. Accordingly, Schultz and colleagues have suggested that the phasic responses seen in the dopamine neurons might serve as error signals, used for adjusting the associative strength of stimuli and responses in neural areas that receive input from the mesencephalic dopamine system (Schultz et al., 1995).

Several groups of investigators (Barto, 1995; Friston, Tononi, Reeke, Sporns, & Edelman, 1994; Houk, Adams, & Barto, 1995; Montague, Dayan, & Sejnowski, 1996; Schultz, Dayan, & Montague, 1997; Suri & Schultz, 1998; cf. J. Brown, Bullock, & Grossberg, 1999) have noted similarities between the phasic activity of the mesencephalic dopamine system and a particular error signal, called a *temporal difference error* (TD error), associated with a reinforcement learning algorithm called the *method of temporal differences* (Sutton, 1988; for reviews, see Barto, 1995; Sutton & Barto, 1998). This algorithm has theoretical roots in both artificial intelligence (Sutton & Barto, 1998) and animal learning

theory (Sutton & Barto, 1990) and is a generalization of the Rescorla–Wagner learning rule (Rescorla & Wagner, 1972) to the continuous time domain. In neural network models, TD errors are computed by an "adaptive critic," which associates a value with the ongoing events and outputs a TD error when it changes its own prediction: Positive TD errors indicate that ongoing events are "better" than expected, and negative TD errors indicate that ongoing events are "worse" than expected. Typically, a response selection module called an *actor* or a *motor controller* uses the error signal to reinforce behaviors that elicit reward. The method of temporal differences accounts for a wide range of behavioral phenomena, including all of the observations predicted by the Rescorla–Wagner model and several observations that are inconsistent with it (Sutton & Barto, 1990).

It is important to note that, like the phasic dopaminergic activity, TD errors propagate back in time from the reward to the conditioned stimulus with learning. Also, like the phasic dopaminergic activity, a negative TD error is elicited by the absence of an expected reward. These observations led naturally to the hypothesis that the mesencephalic dopamine system carries a TD error. As a TD error, the signal would be sensitive to changes in the value of ongoing events, when things are suddenly better or worse than expected. This hypothesis has been formalized in several neural network models (Berns, McClure, Pagnoni, & Montague, 2001; Braver & Cohen, 2000; Egelman, Person, & Montague, 1998; Montague et al., 1996; O'Reilly, Noelle, Braver & Cohen, 2002; Schultz et al., 1997; Suri, Bargas, & Arbib, 2001; Suri & Schultz, 1998, 1999, 2001; Waelti, Dickinson, & Schultz, 2001; cf. Dehaene & Changeux, 1991; Dehaene, Kerszberg, & Changeux, 1998).

Although not universally accepted (e.g., Berridge & Robinson, 1998; Pennartz, 1996; Redgrave, Prescott, & Gurney, 1999), the temporal-difference hypothesis of mesencephalic dopaminergic function has been highly influential. The hypothesis has had a positive response in part because the idea dovetails with previous conceptions about mesencephalic dopamine and reinforcement learning and because the neural network implementation of the method of temporal differences contains other formal properties that are consistent with the structure and function of the midbrain dopamine system (see Holroyd, 2001).

The observations outlined in this section support two critical assumptions of our hypothesis, namely, that the mesencephalic dopamine system carries predictive error signals, and that these predictive error signals are used by other parts of the brain for reinforcement learning.

The Anterior Cingulate Cortex

Impressed by the diversity of cortical and subcortical input into the anterior cingulate cortex, researchers have long regarded this area as a neural locus where motor intentions are transformed into action (Bentivoglio, Kultas-Illinsky, & Illinsky, 1993; Morecraft & Van Hoesen, 1998; Morecraft, Van Hoesen, & Maynard, 1989; Van Hoesen, Morecraft, & Vogt, 1993). In primates, the anterior cingulate cortex contains somatotopically represented motor areas buried in the depths of the cingulate sulcus (Dum & Strick, 1993; Picard & Strick, 1996). The human cingulate sulcus is a prominent feature of the medial wall, composing about half of its surface (Vogt, Nimchinsky, Vogt, & Hof, 1995; see also Paus et al., 1996),

and the anterior cingulate sulcus is distinguished by large Layer V pyramidal cells with extensive dendritic arborizations (Vogt, 1993). These neurons project to other motor areas, including the basal ganglia, supplementary motor area, primary motor area, and spinal cord (Dum & Strick, 1993; Van Hoesen et al., 1993). Single-cell, intracortical stimulation, and functional neuroimaging studies have revealed that the cingulate motor areas contribute to movement generation and execution (reviewed in Dum & Strick, 1993; Picard & Strick, 1996; see also Diehl et al., 2000; Fink, Frackowiak, Pietrzyk, & Passingham, 1997).

The anterior cingulate cortex receives input from several areas concerned with directing motor behavior. For example, the anterior cingulate motor areas receive widespread projections from the limbic lobe, including the orbitofrontal cortex, from other limbic areas such as the amygdala, and from nociceptive sources; for this reason, the anterior cingulate cortex has been said to provide a critical pathway for emotional and motivational factors to influence motor activity (Morecraft & Van Hoesen, 1998; Pandya, Van Hoesen, & Mesulam, 1981; Van Hoesen et al., 1993; Vogt, Sikes, & Vogt, 1993). Likewise, the ventral bank of the cingulate sulcus is richly interconnected with the dorsolateral prefrontal cortex (Bates & Goldman-Rakic, 1993; Morecraft & Van Hoesen, 1991; Van Hoesen et al., 1993). The dorsolateral prefrontal cortex is thought to be involved in the generation of contextually appropriate behaviors in the absence of external stimulation and in the production of new or novel responses (Deiber et al., 1991; Eslinger & Grattan, 1993; Frith, Friston, Liddle, & Frackowiak, 1991a, 1991b; Goldman-Rakic, 1989, 1995; Jahanshahi et al., 1995; Jenkins, Brooks, Nixon, Frackowiak, & Passingham, 1994; Milner & Petrides, 1984; Petersen, Fox, Posner, Mintun, & Raichle, 1989; see also Channon & Crawford, 1999), and the anterior cingulate cortex tends to be concurrently activated with the dorsolateral prefrontal cortex in functional neuroimaging experiments (e.g. Corbetta, Miezin, Dobmeyer, Shulman, & Petersen, 1991; Deiber et al., 1991; D'Esposito et al., 1995; Frith et al., 1991a, 1991b; Jenkins et al., 1994; Koski & Paus, 2000; Paus, Petrides, Evans, & Meyer, 1993; Petersen et al., 1989; Raichle et al., 1994). As a result, some researchers have suggested that the anterior cingulate motor areas provide an important route for the dorsolateral prefrontal cortex to influence motor output (Bates & Goldman-Rakic, 1993; Morecraft & Van Hoesen, 1991; Paus et al., 1993; Van Hoesen et al., 1993). The anterior cingulate cortex is also highly activated when two or more incompatible responses are simultaneously activated (for a review, see Botvinick, Braver, Barch, Carter, & Cohen, 2001), suggesting that the anterior cingulate cortex is involved in detecting (Botvinick et al., 2001) or resolving (Paus et al., 1993) response conflict.

Involvement of the anterior cingulate cortex in high-level motor control is supported by neuropsychological research. For example, frontal medial lesions that include the anterior cingulate cortex can produce a condition called *akinetie mutism*, in which the afflicted person appears to lack the will or motivation to generate behavior, even though he or she is physically capable of doing so (Damasio & Van Hoesen, 1983; Degos, da Fonseca, Gray, & Cesaro, 1993; Devinsky et al., 1995). Conversely, anterior cingulate dysfunction is associated with obsessive-compulsive disorder, a condition characterized by an excessive preoccupation with motor output (Devinsky et al., 1995; Machlin et al., 1991; Rapoport & Wise, 1988; Swedo et al., 1989; S.P. Wise & Rapoport, 1991). In one

person with a focal lesion of the anterior cingulate motor area for manual behavior, performance on two executive control tasks was impaired when the tasks demanded manual, but not verbal, responses (Turken & Swick, 1999). Taken together, the evidence indicates that anterior cingulate motor areas are in a pivotal position for using their diverse neural inputs to contribute to the planning, generation, and execution of behavior (Devinsky et al., 1995; Dum & Strick, 1993).

Learning in the Anterior Cingulate Cortex

The involvement of the anterior cingulate cortex in the acquisition of new behaviors appears to be limited to the earliest stages of learning, when the task is most novel and difficult. In studies by Gabriel and his colleagues (e.g., see Gabriel, 1993), rabbits are required to learn to differentiate between one of two cues (typically tones), only one of which (CS+) is consistently followed by an aversive stimulus (e.g., a foot shock). Gabriel and colleagues have shown that neurons in the anterior cingulate cortex discriminate between the CS+ and the other cue during the early stages of this task, and that the discriminative activity "moves" from anterior to posterior areas of the brain with increased learning (Gabriel, 1993; see also Bussey, Muir, Everitt, & Robbins, 1996).

Functional neuroimaging studies in humans have revealed a comparable sequence of activations. For example, in a trial-and-error learning experiment, the anterior cingulate cortex (and dorsolateral prefrontal cortex) was activated during the earliest stages of the task, but its contribution diminished as the task became overlearned whereas other areas of the brain became more involved as the task progressed (Jenkins et al., 1994; cf. Raichle et al., 1994; Toni & Passingham, 1999). In another study, anterior cingulate activity associated with response conflict in the Stroop task was seen to diminish with practice (Bush et al., 1998).

During the early stages of learning, anterior cingulate motor areas appear specifically to use reward and error information to identify and select appropriate behaviors. In monkeys, the error potentials discovered by Brooks and colleagues in the anterior cingulate sulcus occurred only during the earliest stages of the task as the monkeys learned what was required of them (Brooks, 1986, 1989; Gemba et al., 1986). Brooks (1986) argued that the anterior cingulate cortex contributes to learning only "when certainty is poised against uncertainty" during "the process of distinguishing between what is appropriate and inappropriate in the task context" (p. 31). This position is supported by the observation that, as monkeys learn a sequence of movements by trial and error, one group of neurons in the anterior cingulate sulcus is most active during an early "search" phase, before the animal has identified the sequence of responses that predict the correct solution and reward, whereas another group of neurons is most activated during a later "repetition" phase, once the monkey has just learned the response contingencies (Procyk, Tanaka, & Joseph, 2000). Furthermore, following a reduction of an expected reward in a different trial-and-error learning experiment, but not following control conditions, cells in the monkey anterior cingulate sulcus exhibit a change in activity when the monkey prepares to switch from the unsuccessful response to a different response; when this area is inactivated (by topical application of muscimol, a γ -aminobutyric acid agonist), the monkey is unable to switch to the more adaptive response, evidently because it is unable to use the information

relating to the reduced reward (Shima & Tanji, 1998). Thus, the anterior cingulate motor areas appear to be involved in selecting a new response when the outcome of a trial is not satisfactory.

In humans, functional neuroimaging experiments have also suggested that the anterior cingulate cortex uses reward and error information to identify and select appropriate behaviors. The anterior cingulate cortex is activated by both monetary reward and monetary punishment, with the right anterior cingulate cortex relatively more sensitive to monetary punishment than to monetary reward (Knutson, Westdorp, Kaiser, & Hommer, 2000). In another study (Elliott & Dolan, 1998), subjects used intertrial feedback (the words *correct* and *incorrect*) to develop hypotheses about how best to respond in the task; it was found that a region corresponding to the ventral bank of the anterior cingulate cortex (see Elliott & Dolan, 1998, p. 24, Figure 3b) was most activated when subjects made both a hypothesis and a choice, compared with when they made a hypothesis only and when they made a choice only. The authors suggested that this "ventral anterior cingulate activation in part reflects evaluative processing related to the emotional consequences of making a choice" (p. 25).

In the rat, neurons in the medial prefrontal cortex predict reward acquisition in a radial-arm maze, with some neurons differentiating between high- and low-rewarded arms (Pratt & Mizumori, 2001). The anterior cingulate cortex in the rat appears to be organized topographically, with rostral and ventral parts representing stimulus attributes that predict reward or no reward, and caudal and dorsal parts related to the execution of learned instrumental behaviors (Takenouchi et al., 1999). Importantly, dopaminergic input to the medial prefrontal cortex changes whenever presentation of a response-contingent food reward deviates from what the animal has come to expect (Richardson & Gratton, 1998). It has also been shown that the first pairing of an unconditioned stimulus with a foot shock elicits a large increase in dopamine in the rat medial prefrontal cortex and that this level of dopamine decreases significantly over the next two pairings. This result has been interpreted to mean that the dopamine system may facilitate the medial prefrontal cortex's coping with novel situations, when new learning may be required (Wilkinson et al., 1998).

Given that mesencephalic dopamine neurons synapse on motor neurons in the anterior cingulate cortex, one might expect the anterior cingulate cortex to contribute to reinforcement learning. In fact, the medial prefrontal cortex has long been thought to be involved in reinforcement learning because it is one of the few cortical areas that supports intracranial self-stimulation in the rat (Olds & Milner, 1954; Routtenberg & Sloan, 1972) and because increased dopamine levels have been found in the region following self-stimulation (Mora & Myers, 1977). When conditioned stimuli of very brief duration (200 ms) are used in discriminative avoidance learning, moreover, rabbits that are neonatally exposed to cocaine exhibit impaired avoidance response acquisition as well as abnormal anterior cingulate processing of the conditioned stimulus (Gabriel & Taylor, 1998; C. Taylor, Freeman, Holt, & Gabriel, 1997). These authors have suggested that mesencephalic dopamine may facilitate the development of stimulus associations that capture the participant's attention. Similarly, in a monkey operant conditioning task, anterior cingulate neurons were shown to respond differentially not only to the delivery of positive and negative rewards (e.g., to food vs. an aversive stimulus), but also to the presentation of conditioned stimuli that predicted those rewards

(e.g., to a blue-colored object that predicts the reward vs. a red-colored object that predicts the aversive stimulus; Nishijo et al., 1997). In principle, the learning of such associations could be driven by the mesencephalic dopamine system, perhaps fostering the development of new behaviors that are appropriate to the conditioned stimuli.

Using a metabolic mapping technique (Sokoloff et al., 1977), Porrino and colleagues examined brain activation patterns following electrical stimulation of the VTA. They found that when rats injected current into the VTA by pressing a lever, and when the experimenter delivered current to that same area in a way that was not contingent on the animal's behavior, several areas of the brain became activated (Porrino et al., 1984); however, the difference between the two conditions was largest in anterior portions of the limbic cortex (Porrino, 1993). Porrino (1993) concluded:

Because the distribution of changes in metabolic activity in self-stimulating animals is specific to the reinforced behavior of these animals, and not merely the result of the electrical stimulation to the VTA, the changes in functional activity observed in cortex specifically reflect reinforcement processing and confirm the role of limbic cortex in this processing. (pp. 453–454)

The observations outlined in this section support three critical assumptions of our hypothesis, namely, that anterior cingulate motor areas comprise a neural locus where high-level motor intentions are mapped into actions; that the anterior cingulate cortex is involved in learning these mappings; and that this learning is driven in part by reward-related information carried to the anterior cingulate cortex by the mesencephalic dopamine system.

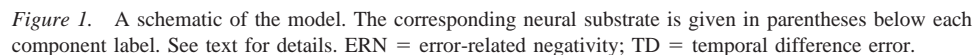
Summary

We have described a neurophysiological phenomenon, the ERN, that appears to be elicited by activation of a system concerned with error detection and compensation and that is generated in the anterior cingulate motor cortex. We have also described the mesencephalic dopamine system, its putative role in reinforcement learning, and a recent hypothesis that holds that a phasic increase in dopaminergic activity occurs when the system detects that ongoing events are better than expected and that a phasic decrease occurs when the system detects that ongoing events are worse than expected. We have further discussed the nature of the anterior cingulate motor areas, including their strategic position for channeling motor commands issued from various command structures to the motor system for execution and their potential contribution to dopamine-dependent reinforcement learning.

Although the reinforcement learning functions associated with the mesencephalic dopamine system and the error-processing functions attributed to the anterior cingulate cortex appear to be concerned with the same problem—namely, evaluating the appropriateness of ongoing events, and using that information to facilitate the development of adaptive behaviors—a possible relationship between these two systems remains to be explored. We now propose a hypothesis that explicitly links the generation of the ERN to the activity of the mesencephalic dopamine system.

Proposal

We assume that the human nervous system is composed, in part, of multiple motor controllers (Figure 1). We envision each of these



We propose that the anterior cingulate cortex, at the confluence of all this information, decides which motor commands are actu-

In keeping with previous simulations, we assume that the reinforcement learning signal conveyed by the mesencephalic dopamine system is specifically a TD error. In one account of the neural

basis of temporal difference learning, the adaptive critic and the actor were assumed to be implemented within “striosome” and “matrisome” modules of the basal ganglia, respectively (Barto, 1995; Houk et al., 1995; cf. J. Brown et al., 1999). Although we assume that several actors (motor controllers) are implemented throughout the brain, our model conforms to the previous account by attributing the role of adaptive critic to the basal ganglia, which computes the value and the change in value of ongoing events.

We propose that errors induce phasic decreases in mesencephalic dopaminergic activity when the system first determines that ongoing events are worse than expected. Conversely, correct responses induce phasic increases in dopaminergic activity. The ERN is generated on error trials, but not on correct trials, when the reduction of dopaminergic input disinhibits neurons in the anterior cingulate cortex. In both cases, the anterior cingulate cortex uses these predictive error signals to select and reinforce the motor controller that is most successful at carrying out the task at hand.

The fundamental points of our hypothesis are (a) that the ERN reflects the transmission of a reinforcement learning signal to the anterior cingulate cortex; (b) that this error signal is carried by the mesencephalic dopamine system; and (c) that it is used to train the anterior cingulate motor cortex to optimize performance on the task at hand.

Computational Simulations and Experimental Data

In this section, we formalize the assumptions described above in a computational model. We use the model to simulate both human behavior and the amplitude of the ERN in two experimental tasks. We also present the results of two experiments that involve human participants engaged in the same tasks, and the predictions produced by the simulations are compared with the experimental results.

The ERN in a Probabilistic Learning Task

The first of the two studies examined the behavior of the ERN in a simple probabilistic learning task. Previous experiments have demonstrated that both error commission (e.g., Falkenstein et al., 1990; Gehring, 1992; Gehring et al., 1993) and negative feedback stimuli (Miltner, Braun, & Coles, 1997) elicit ERNs, but none have investigated both types of error potentials in the same experiment. This study afforded us the opportunity to examine the relationship between the two varieties of ERN as learning progressed throughout the course of a block of trials. Moreover, the task lent itself to simulation using the method of temporal differences, which we used to make specific predictions of ERN amplitude and human behavior.¹

On each trial in this task, an imperative stimulus appeared in front of the participant on a computer screen. The participant was then required to make a two-choice decision by pressing one of two buttons. At the end of the trial, a feedback stimulus indicated to the participant that he or she was either rewarded or penalized one cent of bonus money on that trial. The participants were not informed of the appropriate stimulus–response mappings and had to infer the optimal response strategy by trial and error. The experiment was divided into 10 blocks of 300 trials each. Each block was characterized by a new set of six imperative stimuli, each of which were presented 50 times in random order. In every

block, one of the six stimuli was mapped to the left button, so that participants were rewarded if they pressed the left button and penalized if they pressed the right button. Conversely, another of the six stimuli was mapped to the right button, so that participants were rewarded if they pressed the right button and penalized if they pressed the left. We call both of these two mappings *100% mappings*. For two other stimuli in each block, feedback was delivered at random to the participant, independently of how the participant actually responded. In this case, the participant was rewarded on 50% of the trials and penalized on the other 50% of the trials (*50% mapping*). When the fifth stimulus appeared, the participant was always rewarded, regardless of how he or she responded, and when the sixth stimulus appeared, the participant was always penalized, also independently of how the person actually responded (*always correct/always incorrect mapping*).

We predicted that the negative feedback stimuli would elicit ERNs for all three mapping conditions (100%, 50%, and always correct/always incorrect) at the start of each block. However, we predicted that, as the blocks progressed, newly learned associations in each of the three conditions would differentially affect the amplitude of the ERN. In the 50% mapping condition, the system must wait for the feedback to determine the outcome of the trial. Therefore, we predicted that negative feedback stimuli in this condition would continue to elicit the ERN throughout the course of each block. In contrast, in the 100% mapping condition, the response itself determines the outcome of the trial. Therefore, we predicted that as the system learned the associations between response and feedback, the ERN associated with the response would increase, whereas the ERN associated with the feedback would decrease. Finally, in the always correct/always incorrect mapping condition, the imperative stimulus determines the outcome of the trial. Therefore, we predicted that as the system learned the associations between the imperative and feedback stimuli, neither the response nor the feedback would elicit the ERN.

Experiment. Fifteen participants were paid \$5 an hour to participate, plus a bonus for good performance. The experiment consisted of one 3.5-hr session. Participants sat in front of a CRT in a dimly lit room and performed the probabilistic learning task described above. Each trial consisted of the presentation of an imperative stimulus and the execution of a button-press response by the participant. Imperative stimuli were visual images generated from a Neuroscan (Neurosoft Inc., Sterling, VA) library, including pictures of buildings, animals, vegetables, clothing, and parts of the body. Participants sat 1m from the display such that each image subtended about 3° of visual angle. Each block was associated with images from multiple categories. Stimulus onset asynchrony was 2 s. Participants were required to press one of two buttons (Buttons 2 and 3 on a Neuroscan button box) as soon as possible after stimulus presentation. One second after the stimulus was displayed, a feedback stimulus was presented to the participant. The feedback stimuli were an image of a carrot and an image of a head of lettuce, indicating to the participant that they either were rewarded or were penalized one cent on that trial; the mappings

¹ Note that Kopp and Wolff (2000) have recently examined the effects of error-driven learning on other ERP components, for example, P300.

between the reward and feedback stimuli were counterbalanced across participants.

Rewards and penalties were determined on each trial according to the schedule described above. Participants were not informed of the stimulus–response mappings, and on each block they had to determine the optimal response strategy by trial and error. Participants were told that some mappings might be difficult to figure out, but to “do the best you can.” Additionally, if the participant failed to respond within 600 ms following stimulus presentation, then on trials with a 100% or 50% mapping, a cherry feedback stimulus was presented in lieu of the carrot or lettuce feedback. The cherry feedback stimulus communicated to participants that they were penalized two cents on that trial, providing motivation for them to respond faster. The speed requirement ensured that the participants made some impulsive errors to stimuli with 100% mappings even after the mappings had been learned. Trials in which no response was generated within the 600-ms window were not included in the analysis.

Participants began the experiment with a bonus of \$1. At the end of each block, participants were provided with information indicating the amount of money they earned or lost on that block and the total amount of bonus money they earned throughout the experiment. Bonus money was paid to the participants on their completion of the experiment.

For each participant and for each condition, response-locked and stimulus-locked ERP averages were created by averaging the data recorded at channel Cz with respect to the onset of the response and to the onset of the feedback stimulus. Difference waveforms were generated by subtracting the activity elicited on trials with positive feedback from the activity elicited on trials with negative feedback. The amplitude of the response-locked ERN was determined by identifying the peak negativity of the difference waveform between 0 ms and 100 ms following the response, and the amplitude of each feedback-locked ERN was determined by identifying the peak negativity of the difference waveform between 200 ms and 300 ms following feedback onset. In addition, running averages of difference-wave amplitude were created for each participant and each condition by averaging the data into 10-trial-wide bins, that is, Bin 1 contained data from Trials 1–10, Bin 2 contained data from Trials 11–20, Bin 3 contained data from Trials 21–30, and so on. Changes in ERN amplitude during the block were tested as follows: (a) For each participant, the average amplitude of the response-locked difference-wave over the first 10 trials was subtracted from the average amplitude of the feedback-locked difference-wave over the same period; (b) for each participant, the average amplitude of the response-locked difference-wave over the last 10 trials was subtracted from the average amplitude of the feedback-locked difference-wave over the same period; and (c) a paired *t* test was conducted on these two measures. (See Appendix A for a complete description of recording procedures and data analysis).

Simulation. A schematic of the network is depicted in Figure 1. An input layer composed of six units, one for each of the six stimuli appearing in a block of trials, was fully interconnected with each of five motor control modules. At the start of each trial, the identity of the stimulus was noted by the input module and communicated to all of the motor controllers. Each of the motor controllers was composed of two units, one for each possible response (left- vs. right-button presses). Each controller selected

one of the two response options probabilistically, with the unit receiving the highest net activation most likely to be selected. To simulate the differing ability of the controllers to generate appropriate behavior in the task, a noise term associated with response selection in each controller was manipulated, so that each of the controllers was capable of performing the task with varying degrees of success (see Appendix B). A control filter probabilistically selected one of the five controllers and communicated the response option selected by that controller to the output layer for execution (see Appendix B). An adaptive critic received the stimulus information processed by the input layer, information about the response executed by the output layer, and feedback information presented to the network at the end of the trial. The adaptive critic computed the value of ongoing events and outputted a TD error when it detected a change in value. The TD error was distributed to three parts of the network: to the adaptive critic itself, where it was used to refine the ongoing predictions; to the motor controllers, where it was used to adjust the stimulus–controller mappings; and to the control filter, where it was used to train the filter to select the motor controller with the best performance. Together, the adaptive critic, the motor controllers, and the control filter learned how best to perform the probabilistic learning task.

In the model, the motor-control filter corresponded to the anterior cingulate cortex, and the motor controllers corresponded to each of the neural command areas projecting to the anterior cingulate cortex. Thus, the job of the anterior cingulate was to identify which of its inputs were best suited for carrying out the task. In keeping with previous simulations (e.g., Barto, 1995; J. Brown et al., 1999; Friston et al., 1994; Houk et al., 1995; Montague et al., 1996; Schultz et al., 1997; Suri et al., 2001; Suri & Schultz, 1998, 1999, 2001), the adaptive critic was implemented within the basal ganglia, and the TD error was carried by the mesencephalic dopamine system. An ERN was elicited when the anterior cingulate cortex received a negative TD error from the mesencephalic dopamine system, during or after the response-generation processes, when an eligibility trace associated with the winning motor controller was nonzero.

For statistical purposes, the simulation was run 15 times. On each run, the model was trained anew on 10 blocks of 300 trials. Thus, each simulation corresponded to the data of an individual “participant.” The ERN was the amplitude of the TD error multiplied by the magnitude of the eligibility trace (see Appendix B). However, for comparison with the empirical results (see above), “difference waves” were determined by subtracting the simulated ERN on trials with positive feedback from the simulated ERN on trials with negative feedback. The ERN was then taken as the magnitude of this difference. In addition, running averages of difference-wave amplitude were created for each participant and each condition by averaging the data into 10-trial-wide bins, that is, Bin 1 contained data from Trials 1–10, Bin 2 contained data from Trials 11–20, Bin 3 contained data from Trials 21–30, and so on. Changes in ERN amplitude during the block were tested as follows: (a) For each participant, the average amplitude of the response-locked difference-wave over the first 10 trials was subtracted from the average amplitude of the feedback-locked difference-wave over the same period; (b) for each participant, the average amplitude of the response-locked difference-wave over the last 10 trials was subtracted from the average amplitude of the

feedback-locked difference-wave over the same period; and (c) a paired t test was conducted on these two measures.

Results. Of the three stimulus types (with 100%, 50%, and the always correct/always incorrect mappings), performance could improve only on those trials in which the imperative stimuli were mapped with 100% probability to one of the two response options. Figure 2 shows the accuracy values associated with the 100% mapping condition, averaged across the two stimuli, across blocks, and across participants, for the first and second halves of each block separately. As can be seen, participants' accuracy in this condition improved as the blocks progressed, from 68% during the first half of the blocks to 79% during the second half, $t(14) = -8.75$, $p < .01$. The model's performance also improved as the blocks progressed, from 64% to 83%, $t(14) = -24.4$, $p < .01$.

Figure 3 illustrates typical ERNs elicited in the experiment. Shown are the ERPs associated with positive and negative feedback, averaged across trials, across stimuli, and across blocks. Figure 3A shows the ERN elicited by the feedback in the 50% mapping condition. Confidence intervals (.95) confirmed that the amplitude of the associated difference wave was less than zero ($M = -4.1 \mu V$, $SD = 3.0 \mu V$, interval = $-5.2 \mu V$, $-2.2 \mu V$). The ERN reached peak amplitude about 250 ms following presentation of the negative feedback stimulus (cf. Miltner, Braun, & Coles, 1997). Figure 3B illustrates the ERN associated with the response in the 100% mapping condition. Confidence intervals (.95) also confirmed that the amplitude of the associated difference wave was less than zero ($M = -6.4 \mu V$, $SD = 4.4 \mu V$, interval = $-8.6 \mu V$, $-4.2 \mu V$). The ERN reached maximum amplitude about 80 ms following the button press.

Figure 4 illustrates the amplitudes of the ERN associated with the response and with the feedback for each mapping condition. Shown in each figure are the magnitudes of the difference waves, as determined by averaging the experimental and the simulated data across trials, blocks, stimuli, and participants. Figure 4A illustrates the amplitude of the ERN on trials with 50% mapping probabilities. In this condition, the experimental ERN was elicited primarily by the feedback ($M = -4.1 \mu V$) and not by the response ($M = -0.4 \mu V$), $t(14) = 4.57$, $p < .01$. Conversely, in the 100% mapping condition (Figure 4B), the experimental ERN was elicited




Figure 2. Simulated and empirical accuracy data for the probabilistic learning task. For both the model and human participants, accuracy in the 100% mapping condition improved in the second half of each block relative to the first half. Exp = Experiment.

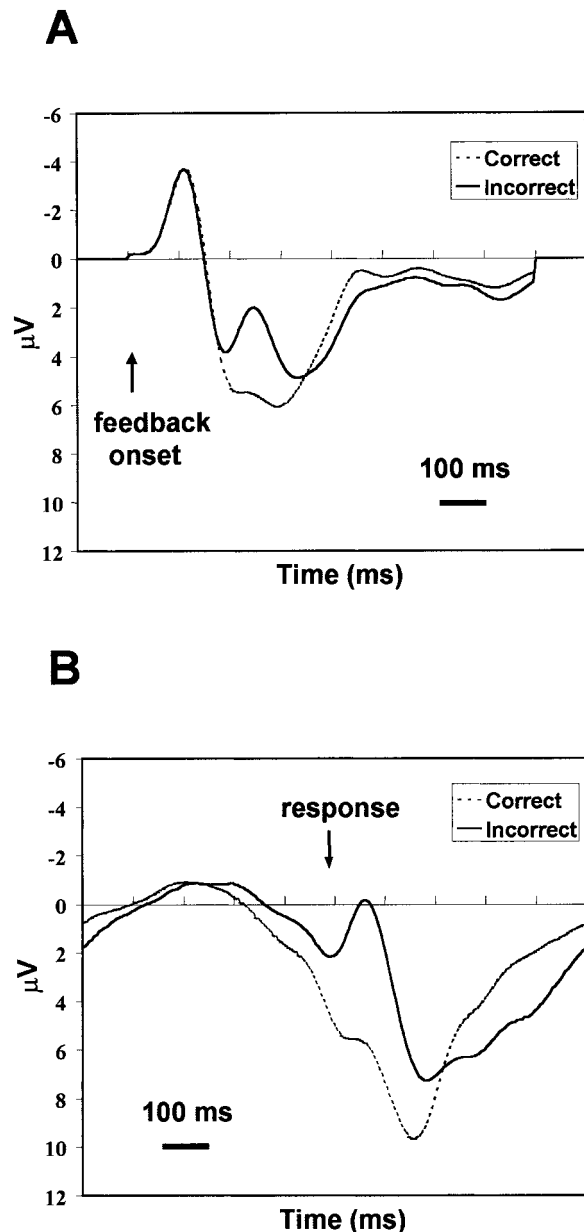


Figure 3. The error-related negativity (ERN) in the probabilistic learning task elicited by feedback stimuli in the 50% mapping condition (A) and the response in the 100% mapping condition (B). The waveforms are more negative on error trials than on correct trials. The ERN peaks about 250 ms after the onset of the feedback stimulus and about 80 ms after the onset of the incorrect response. Waveforms were recorded from channel Cz.

primarily by the response ($M = -6.4 \mu V$), and not by the feedback ($M = -1.5 \mu V$), $t(14) = -3.46$, $p < .01$. Lastly, the experimental ERN in the always correct/always incorrect condition (Figure 4C) was relatively small for both the response ($M = -1.7 \mu V$) and the feedback ($M = -2.6 \mu V$), and there was no significant difference between the two conditions, $t(14) = 1.35$, $p > .05$. Confidence intervals indicated that the small ERNs in this condition associated with the response ($M = -1.7 \mu V$, $SD = 1.5$




Figure 4. Simulated (Model) and empirical (Exp) error-related negativity (ERN) amplitudes in the probabilistic learning task. Amplitudes are determined from the difference in magnitudes on trials with positive feedback

μV , interval = $-2.3 \mu\text{V}$, $-0.9 \mu\text{V}$), and with the feedback ($M = -2.7 \mu\text{V}$, $SD = 2.3 \mu\text{V}$, interval = $-3.8 \mu\text{V}$, $-1.4 \mu\text{V}$), were statistically reliable.

The simulated data follow a similar pattern. In the 50% mapping condition (Figure 4A), the simulated ERN was elicited primarily by the feedback ($M = -0.84$), and not by the response ($M = -0.003$), $t(14) = 81.2$, $p < .01$. Conversely, in the 100% mapping condition (Figure 4B), the simulated ERN was elicited primarily by the response ($M = -1.35$), and not by the feedback ($M = -0.20$), $t(14) = -61.5$, $p < .01$. The simulated ERN in the always correct/always incorrect condition (Figure 4C) was also relatively small for both the response ($M = -0.21$) and the feedback ($M = -0.16$), but the ERN to the response was the larger of the two, $t(14) = -5.3$, $p = .01$. Confidence intervals confirmed that both simulated ERNs were statistically reliable: response ($M = -0.209$, $SD = .030$, interval = -0.225 , -0.194), feedback ($M = -0.164$, $SD = .021$, interval = -0.175 , -0.154).

Figure 5 shows running averages of ERN amplitude, for both the simulated and empirical data, associated with the response and with the feedback for the three mapping conditions. The running averages illustrate the evolution of the error potentials as learning progressed throughout the course of each block. In the 50% mapping condition (Figure 5A and 5B), the magnitude of the ERN elicited by the feedback was greater than the magnitude of ERN elicited by the response during the entire blocks, for both the simulated (difference during first 10 trials, $M = -0.81$, difference during last 10 trials, $M = -0.82$), $t(14) = 0.14$, $p > .05$, and empirical data (difference during first 10 trials, $M = -2.31 \mu\text{V}$, difference during last 10 trials, $M = -3.23 \mu\text{V}$), $t(14) = -1.91$, $p > .05$, data. In contrast, in the 100% mapping condition (Figure 5C and 5D) the magnitude of the ERN elicited by the response increased relative to the magnitude of the ERN elicited by the feedback as the blocks progressed, for both the simulated (difference during first 10 trials, $M = 0.16$, difference during last 10 trials, $M = 1.83$), $t(14) = -35.2$, $p < .01$, and empirical (difference during first 10 trials, $M = 2.08 \mu\text{V}$, difference during last 10 trials, $M = 7.76 \mu\text{V}$), $t(14) = -4.89$, $p < .01$, data. For the always correct/always incorrect condition (Figure 5E and 5F), the simulated results deviated from the empirical results. By the end of the blocks, for the simulated data, the magnitude of the ERN associated with both the response (-0.016) and the feedback (-0.003) approached zero, whereas for the empirical data, the magnitude of the ERN associated with the response ($-3.11 \mu\text{V}$) and with the feedback ($-1.54 \mu\text{V}$) remained relatively far from zero. However, as the blocks progressed, the magnitude of the ERN elicited by the response increased relative to the magnitude of the ERN elicited by the feedback, for both the simulated (difference during first 10 trials, $M = -0.10$, difference during last 10 trials, $M = 0.01$),

and trials with negative feedback. A: 50% mapping condition; B: 100% mapping condition; C: always correct/always incorrect (AC/AI) mapping conditions. In the 50% mapping condition, the ERN is elicited primarily by the feedback, and not by the response. In the 100% mapping condition, the ERN is elicited primarily by the response, and not by the feedback. In the AC/AI mapping condition, neither the response nor the feedback elicit a large ERN. Simulated results are given in multiples of reward (r). Empirical results are given in microvolts. Exp = Experiment.




Figure 5. Running averages of error-related negativity (ERN) amplitude for simulated (Model column) and empirical (Experiment column) data. Running average bin size = 10 trials (see text). Solid line: response-locked ERN amplitude. Dashed line: feedback-locked ERN amplitude. A and B: 50% mapping condition; C and D: 100% mapping condition; E and F: always correct/always incorrect (AC/AI) mapping condition. See text for details.

$t(14) = -3.7, p < .01$, and empirical (difference during first 10 trials, $M = -0.64 \mu\text{V}$, difference during last 10 trials, $M = 1.36 \mu\text{V}$), $t(14) = -2.61, p < .05$, data. In other words, for both the simulated and empirical data, the feedback-related ERN was larger than the response-related ERN at the beginnings of the blocks, but was smaller than the response-related ERN at the ends of the blocks.

We chose a running average bin size of 10 to collect enough trials per condition per bin to generate empirical ERNs with good signal-to-noise ratios. However, the simulated ERNs were not bound by the same constraint. For this reason, the simulated data in Figure 5 were replotted in Figure 6 (bin size = 1) to show the trial-by-trial changes in amplitude. Replotting the data in this fashion revealed some transient phenomena at the beginning of each block that were obscured when the data were plotted with the larger bin size. In particular, the response-locked ERN in the 100%

mapping condition (Figure 6B) started at about zero ($M = 0.01$) on the first trial and rose thereafter, whereas the feedback-locked ERN was large at the beginning of the blocks ($M = -0.84$), $t(14) = 26.1, p < .01$, and decreased as the blocks progressed. In the always correct/always incorrect condition (Figure 6C), the feedback-ERN again was larger at the beginning of the blocks ($M = -0.81$) than the response-locked ERN ($M = 0.01$), $t(14) = 13.8, p < .01$, and decreased as the blocks progressed, but the response-locked ERN increased in magnitude until about Trial 8, whereupon it started to decrease. Finally, in the 50%




Figure 6. Error-related negativity (ERN) amplitude for simulated data plotted by trial within block. Solid line: response-locked ERN amplitude. Dashed line: feedback-locked ERN amplitude. A: 50% mapping condition; B: 100% mapping condition; C: always correct/always incorrect (AC/AI) mapping condition. See text for details.

condition (Figure 6A) the feedback-locked ERN started the blocks at maximum amplitude ($M = -0.82$), the response-locked ERN began significantly smaller at zero ($M = 0.00$), $t(14) = 20.2$, $p < .01$, and both measures remained constant throughout the blocks.

If the amplitude of the ERN was correlated with the learning process, then we might expect to see the effects of such learning on individual trials. In particular, the amplitude of the ERN should be largest when a feedback stimulus disconfirms a prediction induced by a previous feedback stimulus, when both stimuli are encountered under the same conditions. Figure 7 illustrates the effects of different sequences of feedback stimuli on the amplitude of the ERN. For both the simulation and the human participant data, we determined the amplitudes of the ERN elicited by feedback on trials with 50% mapping probabilities according to the type of feedback seen by the system on the immediately previous encounter with the same stimulus, when both encounters were associated with an identical response. Thus, the data were averaged according to four conditions: error trials preceded by error trials, error trials preceded by correct trials, correct trials preceded by correct trials, and correct trials preceded by error trials. We then determined the magnitude of the difference between the waveforms on error trials preceded by error trials and the waveforms on correct trials preceded by correct trials (*same condition*) and between the waveforms on error trials preceded by correct trials and the waveforms on correct trials preceded by error trials (*opposite condition*).




Figure 7. Sequence effects on error-related negativity (ERN) amplitude induced by feedback presentation in the 50% mapping condition of the probabilistic learning task. Data are averaged according to the valence of the feedback stimulus on the most recent trial associated with the same imperative stimulus and response. When successive encounters with the same imperative stimulus and response are associated with feedback stimuli of opposite valence, the ERN elicited by the second feedback stimulus is larger than when the two feedback stimuli have the same valence. Simulated results are given in multiples of reward (r). Empirical results are given in microvolts. Same = ERN amplitude determined by subtracting data associated with trials with positive feedback when preceded by positive feedback from data associated with trials with negative feedback when preceded with negative feedback; Opposite = ERN amplitude determined by subtracting data associated with trials with positive feedback when preceded by negative feedback from data associated with trials with negative feedback when preceded by positive feedback. Exp = Experiment.

Figure 7 shows the magnitude of the resulting difference associated with the same and opposite conditions, for both the model and human participants. The amplitude of the observed ERN tended to be smaller when two successive encounters with the same stimulus (and response) were followed by feedback with the same valence ($M = -3.0 \mu V$), compared with when the two successive encounters were followed by feedback with opposite valence ($M = -4.3 \mu V$), $t(14) = 2.52$, $p < .05$. This difference in amplitude was reflected in the simulated data as well (same, $M = -0.68$, opposite, $M = -1.01$), $t(14) = 35.1$, $p < .01$. Note that the consecutive encounters were often separated by trials associated with a different stimulus or response type.

In the simulated data, these sequential learning effects arise from weight changes associated with the response. Figure 8 illustrates the change in weight of the value unit associated with the response in the 50% mapping condition, for both positive feedback trials and negative feedback trials. As can be seen, the value associated with the given response increases following reward but decreases following punishment. Thus, on subsequent encounters with the same imperative stimulus, the response that was previously rewarded generates a relatively positive prediction, whereas the response that was previously punished generates a relatively negative prediction. If the ERN indeed reflects a prediction error, feedback stimuli that disconfirm these predictions should elicit relatively large ERNs, whereas feedback stimuli that confirm these predictions should elicit relatively small ERNs. This was the pattern of results we obtained.

Discussion. In general, there was considerable correspondence between the simulated and experimental data. For both humans and the model, on trials with 50% mapping probabilities, the ERN was elicited primarily by feedback presentation (Figures 4A, 5A, and 5B). In this condition, the system waited until the feedback was presented to determine the outcome of each trial. In contrast, on trials with 100% mapping probabilities, the accuracy of the system increased across the blocks (Figure 2), and the amplitude of the ERN was larger to the response than it was to the feedback (Figures 4B, 5C, and 5D). Evidently, as the system learned the correct response mappings associated with these stimuli, it tended to rely less on the feedback and more on its own representation of what the response should be to determine the outcome of each trial. The simulated data in Figure 6B suggested that most of this learning occurred during the first 10 or so encounters with the imperative stimulus, a plausible result given the simplicity of the task. In the always correct and always incorrect mappings, the amplitude of the ERN was small to both the response and the feedback for both the experimental and simulated data (Figure 4C). However, in this case, the simulated ERNs decreased to zero as the blocks progressed (Figure 5E), whereas the empirical ERNs appeared constant throughout the blocks (Figure 5F). In the model, the system ceased to rely on the response and the feedback as it learned that the imperative stimulus predicted the outcome of the trial. The fact that neither the response nor feedback ERNs decreased by much in the empirical data suggests that an additional cognitive process must have been at work in the human system, continuing the search for an appropriate response strategy even after the simpler system had given up. On the other hand, the ERN in this analysis was computed across two different stimulus conditions (always correct and always incorrect), and may reflect




Figure 8. Weight changes to the value units associated with each response derived from the simulation. Data are for the 50% mapping condition. Positive = change in weight induced by positive feedback; Negative = change in weight induced by negative feedback.

methodological artifacts in the empirical data that did not occur in the simulated data.

Trial-to-trial variation in the learning process can be seen in Figure 7, which illustrates that the magnitude of the ERN elicited by feedback on trials with 50% mapping probabilities tended to be larger when the previous encounter with the same imperative stimulus and response was associated with feedback with opposite valence, compared with when it was associated with feedback with the same valence. These variations reflect the development of transient (and incorrect) predictions in the random feedback condition (Figure 8). The ERN tended to be larger when the feedback stimulus disconfirmed, rather than confirmed, a prediction induced by a previous feedback stimulus.

The ERN in a Modified Version of the Eriksen Flankers Task

In the typical version of the Eriksen Flankers Task (Eriksen & Eriksen, 1974), participants are exposed to a series of four different stimulus arrays, presented randomly with equal probabilities: five *H*s in a row (HHHHH), five *S*s in a row (SSSSS), an *H* with two *S*s on both sides (SSHSS), and an *S* with two *H*s on both sides (HSHHH). Participants are directed to focus their attention onto the letter in the center of each stimulus and to respond with one hand (typically, by pressing a button or squeezing a dynamometer) whenever they see a stimulus with an *H* in the center (HHHHH, SSHSS), and with the contralateral hand whenever they see a stimulus with an *S* in the center (SSSSS, HSHHH). In the present experiment, we modified the probabilities so that half the participants saw each of the stimuli with a central *H* on 10% of the trials and each of the stimuli with a central *S* on 40% of the trials; the remaining half of the participants saw the same stimuli with the converse set of probabilities. Stimuli in which the central and flanker letters are the same are referred to as *compatible*, and stimuli in which the central letter is different from the flankers are called *incompatible*. Likewise, we call the highly probable stimuli *frequent* and the less probable stimuli *infrequent*. Thus, there are

four stimulus conditions, which we refer to with the following shorthand: *infrequent compatible*, *infrequent incompatible*, *frequent incompatible*, and *frequent compatible*.

When this experiment was originally conducted, we predicted that the amplitude of the ERN would be smaller on error trials in the incompatible conditions than in the compatible conditions, because we thought the error detection process might be compromised when the flanking letters differed from the central target letter. In fact, we were surprised to find that the empirical results did not conform to our expectations (see *Discussion* section below). The task was then modeled using the method of temporal differences to see if the simulation could account for, and provide insight into, the observed results.

Experiment. The experiment consisted of two 3.5-hr sessions. Fifteen participants were paid \$5 an hour for their participation, plus a \$5 bonus for completing both sessions. Participants sat in front of a CRT in a dimly lit room and performed the modified version of the Eriksen Flankers Task described above. Each session consisted of a practice block followed by 12 blocks of 200 trials each, with 5–10-min breaks between blocks. Hence, for each participant, 4,800 trials of data were collected. Of these, 480 trials were associated with each of the two infrequent stimulus arrays (10% probability of appearance) and 1,920 with each of the two frequent stimulus arrays (40% probability of appearance). Stimulus onset asynchrony was 1.5 s. Each stimulus array appeared on the screen for 50 ms, and participants sat 1 m from the display such that each letter in the array subtended about .5° of visual angle. Participants responded by squeezing two zero-displacement dynamometers (Linear Velocity Force Transducers, Model 152A, Daytronic Corp., Dayton, OH) connected to an amplifier system (Conditioner Amplifiers, Model 830A, Daytronic Corp., Dayton, OH). During the experiment, overt responses were registered when the participant's squeeze force exceeded 25% of his or her maximum squeeze force, which was determined for each participant at the start of the session.

Following each block, feedback informing the participants of their accuracy (percent correct) and average speed (in milliseconds) was presented on the CRT. For the purposes of the feedback, responses generated during the first 50 ms following stimulus onset were considered errors. Participants were asked to respond as quickly as possible while maintaining an accuracy of about 85%. Participants were told that if their accuracy fell to 80%, then on the following block they should respond slower to improve their performance. Conversely, they were told that if their accuracy were to rise to 90%, then on the following block they should improve their speed. Verbal feedback was also provided, for example, as encouragement to the participants to break their personal records in speed and accuracy.

The amplitude of the ERN was determined from the response-locked ERP averages for each participant with an algorithm that identified the peak negativity recorded at channel Cz between 0 ms and 200 ms following the response. This algorithm also identified the onset of the component and computed its base-to-peak magnitude. (Recording procedures and data analysis are described in Appendix A; a more complete description will be presented in a separate article, Holroyd & Coles, 2002).

Simulation. Fifteen simulations were run, one for each participant. The model was trained using the data collected in the experiment from each individual participant. For each trial for a

particular participant, the adaptive critic received the stimulus encountered by the participant on that trial (HHHHH, SSHSS, HHSHH, or SSSSS), the response produced (left vs. right hand), and the associated outcome (correct vs. incorrect). As in the previous simulation, the adaptive critic computed the value of ongoing events and outputted a TD error when it detected a change in value. The TD error was used by the adaptive critic itself to refine its ongoing predictions. ERN amplitude was taken as the magnitude of the temporal difference error associated with response generation (see Appendix B). To simulate the effect of task instructions, which specified the appropriate stimulus–response mapping, the value layer weights associated with the “correct” conjunctions of stimuli and responses were initialized at +1, and the value layer weights associated with the “incorrect” conjunctions of stimuli and responses were initialized at –1. Because the adaptive critic was trained using the participant’s own responses, we did not model the response-generation process itself (i.e., the motor controllers, control filter, and output modules; see Figure 1). However, under these circumstances, the adaptive critic should learn to generate prediction errors associated with the outcomes of each of the participant’s responses that mirror the response-locked ERNs generated by each participant in the experiment (see Appendix B).

Results. In the experimental data, the ERN was larger on incorrect ($M = -4.0 \mu V$) than on correct ($M = -1.3 \mu V$) trials, $t(14) = 4.8$, $p < .01$ (cf. Falkenstein et al., 1990; Gehring, 1992; Gehring et al., 1993). Figure 9 illustrates the amplitude of the ERN on infrequent compatible, infrequent incompatible, and frequent incompatible error trials, as generated by both the simulation and the human participants. (Very few errors were generated on frequent compatible trials, so an analysis of the ERN in this condition was omitted.) A one-way analysis of variance (ANOVA) on ERN amplitude as a function of stimulus type revealed that not all the mean amplitudes were equal, $F(2, 28) = 16.8$, $p < .01$, $\epsilon = .61$;




Figure 9. Error-related negativity (ERN) amplitude in the modified Eriksen Flankers Task. For both the model and the experimental data, ERN amplitude is larger on frequent incompatible error trials than on infrequent compatible and infrequent incompatible error trials. Simulated results are given in multiples of reward (r). Empirical results are given in microvolts. III = infrequent compatible condition; FIF = infrequent incompatible condition; IFI = frequent incompatible condition; Exp = Experiment.

Duncan’s Multiple Range Test indicated that ERNs elicited on infrequent compatible ($M = -4.5 \mu V$) and infrequent incompatible ($M = -3.3 \mu V$) error trials were significantly smaller than ERNs elicited on frequent incompatible ($M = -9.2 \mu V$) error trials, but not significantly different from each other. To ensure that the observed difference in ERN amplitude between the frequent incompatible and infrequent incompatible conditions did not depend on differences in reaction time between the two conditions, we also matched each of the errors in the frequent incompatible and infrequent incompatible conditions by reaction time (see Coles, Scheffers & Holroyd, 2001). When so matched, the ERN elicited by errors on frequent incompatible ($M = -9.1 \mu V$) trials was still larger than when it was elicited by errors on infrequent incompatible ($M = -5.0 \mu V$) trials, $t(14) = 3.8$, $p < .01$.

As can be seen, the output of the model corresponded qualitatively to the amplitude of the ERN observed in the experiment, with the largest simulated ERN occurring on frequent incompatible error trials ($M = -1.83$), the smallest on infrequent incompatible error trials ($M = -0.60$), and an intermediate magnitude on infrequent compatible error trials ($M = -1.16$). A one-way ANOVA on ERN amplitude as a function of stimulus type revealed that not all the mean amplitudes are equal, $F(2, 28) = 291$, $p < .01$, $\epsilon = .93$, and Duncan’s Multiple Range Test indicated that all three means were significantly different from one another.

Further analysis revealed that the amplitude of the error signal depended on the probability of generating the correct response in each stimulus condition. Table 1 presents accuracy and reaction times (of the empirical data) for each condition; a one-way ANOVA on participant accuracy as a function of stimulus type revealed that not all the mean accuracies were equal, $F(2, 28) = 305$, $p < .01$, $\epsilon = .87$, and Duncan’s Multiple Range Test indicated that all three means were significantly different from one another. (A detailed analysis of the behavioral data will be presented in a separate article, Holroyd & Coles, 2002; see also Coles et al., 2001). Figure 10 illustrates the magnitude of the real and simulated ERNs plotted against participant accuracy within the infrequent compatible, infrequent incompatible, and frequent incompatible stimulus conditions. The magnitude of the error signal increases with the probability of making a correct response.

Figure 11 illustrates the weights of the value units associated with the imperative stimuli in the simulated data. As can be seen, the system learned to attribute a positive value to the frequent incompatible stimulus, a negative value to the infrequent incompatible stimulus, and an intermediate value to the infrequent compatible stimulus. Hence, the amplitudes of the ERN in Figures 9 and 10 reflect the change in prediction precipitated by error commission in each of these conditions.

Discussion. Originally, we had predicted that the amplitude of the ERN should be smaller on error trials in the incompatible conditions than in the compatible conditions because we thought the error detection process might be compromised when the flanking letters differed from the central target letter. Thus, we were surprised to find that the amplitude of the ERN was larger on frequent incompatible error trials than it was on infrequent compatible and infrequent incompatible error trials. In fact, the ERNs generated by the computational model mirrored those produced by the human participants (Figure 9). The variation in ERN amplitude appeared to be associated with the accuracy of the participant in each stimulus condition (Figure 10). Evidently, because most

Table 1
Accuracy (Percent Correct) and Reaction Times
(in Milliseconds)

Measure	FFF	IFI	FIF	III
Accuracy	99.1	93.9	31.4	60.8
Correct RT	241	284	401	353
Error RT	— ^a	319	219	198

Note. FFF = frequent compatible; IFI = frequent incompatible; FIF = infrequent incompatible; III = infrequent compatible; RT = reaction time.

^a Too few errors were generated in this condition to determine error RT.

frequent incompatible trials were correct, the model predicted that correct responses would be generated on those trials (Figure 11). However, when an incorrect response was generated instead, the prediction was revised and a large ERN was produced. In contrast, on infrequent compatible and infrequent incompatible trials, the system was less certain that the outcome would be favorable; hence, errors in these conditions generated much smaller ERNs.

General Discussion

We have presented an account of the ERN that depends on the reinforcement learning properties of the mesencephalic dopamine system and the anterior cingulate cortex. We began by reviewing the literature on the ERN, and we briefly described the role the mesencephalic dopamine system may play in reinforcement learning. We then discussed the nature of the anterior cingulate cortex, including the neuroanatomy of its motor areas, its sensitivity to response conflict, and its involvement in behavioral adaptation. We then proposed that the ERN is generated when a negative TD error is carried by the mesencephalic dopamine system to the




Figure 10. Error-related negativity (ERN) amplitude as a function of within-condition accuracy. For both the model and the experimental (Exp) data, ERN amplitude increases as the within-condition accuracy increases. Simulated results are given in multiples of reward (r). Empirical results are given in microvolts. IFI = frequent incompatible condition; III = infrequent compatible condition; FIF = infrequent incompatible condition.




Figure 11. Value unit weights associated with the imperative stimuli derived from the simulation. Shown are the initial weights at the start of the experiment and the final weights at the end of the experiment. III = infrequent compatible condition; FIF = infrequent incompatible condition; IFI = frequent incompatible condition.

anterior cingulate motor areas, during or after response generation. According to this view, the ERN is produced when the system first detects that the consequences of an action are worse than expected. We also suggested that the anterior cingulate cortex uses the error signals to improve the system's performance on the task at hand by shifting responsibility for the task to the motor controller that is most capable of performing it. We then presented two simulations that formalized these propositions, and we compared the output of the simulations with the results of two ERP experiments. The simulated and empirical results are largely consistent with each other and support our hypothesis.

We have provided a theoretical framework for understanding the error-processing system that produces the ERN. First, we have suggested a specific algorithm for describing the behavior of the ERN in various psychological tasks. The consistency of the model's data with those of human participants suggests that the algorithm is appropriate for modeling the ERN, one that can be used to elucidate the nature of the learning process in a wide variety of tasks (e.g., Herrnstein, 1997). Just as important, the model makes explicit the functional nature of the systems underlying ERN generation. Thus, our vocabulary is no longer limited to the general statement that the system that gives rise to the ERN is somehow involved in error processing. Rather, we can say explicitly that the ERN represents the first indication that the consequences of an action are worse than expected and that this signal is used to train the response production system in a manner that is consistent with neural network implementations of reinforcement learning. Second, we have presented a general account of how these systems are implemented in the brain. Although ours is not a detailed neurophysiological model, it is nonetheless consistent with current views on the structure and function of the specified neural systems. It thus provides an avenue for exploring the dependence of the ERN on the functional integrity of these systems. In this way, it may provide insight, for example, into why people with Parkinson's disease (in whom the midbrain dopamine

nuclei are damaged) are impaired at a reinforcement learning task (Knowlton, Mangels, & Squire, 1996). At the same time, our model also raises a number of important questions that need to be addressed in future research. In the following pages we discuss several of these unresolved issues.

Comparison With Other Theories

Our model should not be construed as a new theory of the anterior cingulate cortex. Rather, the model holds simply that the anterior cingulate cortex is trained to improve its performance using TD error signals carried to it by the mesencephalic dopamine system and that transfer of negative TD errors in this fashion elicits the ERN. In this respect, our theory is firmly grounded in previous computational work associating phasic dopaminergic activity with the method of temporal differences (Barto, 1995; Berns et al., 2001; Braver & Cohen, 2000; Friston et al., 1994; Houk et al., 1995; Montague et al., 1996; O'Reilly et al., 2002; Schultz et al., 1997; Suri et al., 2001; Suri & Schultz, 1998, 1999, 2001; Waelti et al., 2001; cf. J. Brown et al., 1999) and, more generally, with reinforcement learning signals (e.g., Dehaene & Changeux, 1991). Although we have tentatively suggested a particular function for the anterior cingulate cortex—namely, as a motor control filter—our theory is potentially compatible with several other accounts. Further investigation will be needed before any one theory of anterior cingulate function can be chosen over another.

Nevertheless, the relationship between our model and two prominent theories of the anterior cingulate cortex deserves consideration. The first of these two theories holds that the anterior cingulate cortex is involved in “selection for action,” which Allport (1987) defined as a cognitive mechanism that “can selectively designate a specified subset of the available, and potentially relevant, sensory information to have control of a given effector system, and can selectively decouple the remainder from such control” (p. 397). The motivation for this position comes from functional neuroimaging studies. In a positron-emission tomography (PET) study, blood flow in the anterior cingulate cortex increased with the number of targets in a cognitive task, suggesting to investigators that the anterior cingulate cortex was concerned with selecting among the targets (Posner, Petersen, Fox, & Raichle, 1988). Because blood flow to this area also increased relative to control conditions in a verb generation task (Petersen et al., 1989), it was also proposed that this area was involved in “selecting among several competing choices of an appropriate verb for each noun” (Posner & Raichle, 1994, p. 123). Thus, the anterior cingulate cortex was said to be concerned with “selection among competing, complex contingencies” (Posner & Raichle, 1994, p. 123), whether those contingencies were stimulus or response related. Though never stated explicitly, in this view, the anterior cingulate cortex is responsible for applying the mappings between stimuli and responses, becoming more activated as the nature of the stimuli or the response options make the mappings more difficult to implement.

Because the control filter in our model comprises a hidden layer placed between stimulus input and response output, our model is compatible with this very broad position on anterior cingulate function. More importantly, our theory of the ERN is consistent with the view advanced by Paus et al. (1993), who took the selection-for-action theory one step further. They suggested that

motor areas in the anterior cingulate cortex “funnel” executive commands originating from the dorsolateral prefrontal cortex to the effector system, and they speculated that influence of the mesencephalic dopamine system on the anterior cingulate cortex would result in “facilitation of the execution of the appropriate responses and/or suppression of the inappropriate ones” (p. 467). By using the mesencephalic dopamine system to train the anterior cingulate cortex to support appropriate behaviors while suppressing inappropriate ones, our model achieves much the same result.

Conflict monitoring (Barch, Braver, Sabb, & Noll, 2000; Botvinick et al., 2001; Botvinick, Nystrom, Fissell, Carter, & Cohen, 1999), the second theory of anterior cingulate function, shares several important characteristics with the selection for action theory but is computationally more specific. The theory holds that the anterior cingulate cortex is responsible for detecting the simultaneous activation of incompatible processing channels. In contrast to the selection for action theory, the anterior cingulate cortex is not given the role of applying the mapping between stimulus and response. Instead, the anterior cingulate cortex detects when crosstalk between stimulus channels or between response channels impedes smooth application of the stimulus–response mapping. This information is then sent to the relevant channels where it is used to reduce the conflict. In this view, the anterior cingulate cortex performs a monitoring function, initiating intervention in other systems to adjust their behavior as needed. Thus, the conflict theory is distinguished from the action selection theory, and from the hypothesis advanced in this article, by attributing a monitoring role rather than a response-selection role to the anterior cingulate cortex.

The conflict model makes specific predictions about the ERN, and so we compare it here with our own model in detail. With respect to the ERN, the two models appear to be in fundamental disagreement. Whereas our theory holds that the ERN is elicited by transmission of training signals to the anterior cingulate cortex, the Botvinick et al. (2001) model maintains that the ERN is generated by response conflict—that is, the more conflict, the larger the ERN. These positions lead to several incompatible predictions. First, the conflict model predicts that ERNs should be elicited on some correct trials, when conflict is high. This prediction was supported by the observation of a so-called “correct ERN,” which—as the name indicates—is an apparent ERN elicited in conjunction with a correct response (e.g., Vidal, Hasbroucq, Grapperon, & Bonnet, 2000). However, we have recently shown that ERNs can occur on some correct trials when (a) stimulus-locked activity creates an ERN-like artifact in the response-locked ERP or (b) the participant’s representation of the correct or actual response is compromised (Coles et al., 2001). Moreover, both the conflict and reinforcement learning accounts predict that ERNs can occur on correct trials when the correct response is produced simultaneously with an incorrect response: in the former case, because of conflict associated with the simultaneous activation of incompatible response channels, and in the latter case, because of error processing associated with the generation of the incorrect response. To test these predictions, Scheffers (1999) obtained a measure of conflict by multiplying together the EMG activity associated with both response channels in a two-choice reaction-time task. He showed that, when matched for conflict, the ERN was larger on incorrect trials (when activation of the incorrect response was larger than that of the correct response) than on

correct trials (when activation of the correct response was larger than that of the incorrect response). Thus, the process that generates the ERN appears to be more sensitive to error commission than to response conflict. The same conclusion was reached by Falkenstein and colleagues in a different set of experiments (Falkenstein et al., 2000). Finally, our model accounts nicely for the observation that negative feedback elicits an ERN (Miltner, Braun, & Coles, 1997), and, as shown in this article, that this ERN propagates with learning to the erroneous response. It is difficult to see how the conflict-monitoring theory could account for ERNs elicited by negative feedback, much less demonstrate how the behavior of the ERN evolves with learning.

Nevertheless, the conflict-monitoring hypothesis accounts for a wealth of functional magnetic resonance imaging (fMRI) and PET data that the reinforcement learning–ERN model does not address, especially data that suggest activation of the anterior cingulate cortex in conditions of high response conflict. In general, our position allows for activation of the anterior cingulate cortex in the absence of ERN generation. Thus, our model is not incompatible with a conflict model of anterior cingulate function—only with the position that the ERN is a manifestation of such conflict. Note that we could have implemented our model as a simple adaptive critic–actor combination, with the anterior cingulate cortex serving as the actor. The anterior cingulate cortex in this design might be thought to mediate selection for action because it would function only to translate stimulus input into response output. We decided against adopting this simpler architecture for two reasons. First, the single motor-controller model does not provide a role for the convergence of cognitive and limbic input passing through the anterior cingulate cortex on the way to the motor system. Second, the single-controller model does not specify how conflict could arise from the simultaneous activation of incompatible processing channels. We suggest that the design we implemented, with multiple motor controllers, is more consistent with these observations. On the one hand, the motor controllers in our model provide an avenue for cognitive and limbic input to direct motor output via the anterior cingulate cortex; on the other hand, the anterior cingulate cortex decides between any conflicting directives issued by the competing controllers. A measure of such motor controller conflict could be developed in a future model.

Recently, Braver and Cohen (2000) independently developed a computational model of dopaminergic modulation of the prefrontal cortex that is similar in some important respects to our model of the ERN. In their model, as in ours, the mesencephalic dopamine system carries a TD error to frontal areas of the brain. This phasic dopamine activity both gates information into and out of working memory in the prefrontal cortex and subserves a reinforcement learning function that enables the system to discover what information is relevant in the task context. Braver and Cohen (1999) also suggested that a reduction of dopaminergic activity on error trials could “reset” information held in working memory, facilitating the search for a more adaptive response (see also O’Reilly et al., 2002). Braver and Cohen’s model appears to be compatible with our model, and current work involves unifying them, together with a third model of the basal ganglia, prefrontal cortex, and working memory (Frank, Loughry, & O’Reilly, 2001), into a common framework.

Structure of the Model

As we noted, our conception of the anterior cingulate cortex as a control filter was motivated by several observations, including the convergence of cognitive and limbic input passing through it on the way to the motor system and its contribution to reinforcement learning. Our interpretation is consistent with prevailing notions of anterior cingulate function. It does, however, require future elaboration. For example, response selection by the anterior cingulate cortex need not occur in a “winner-take-all” fashion, in which a single motor controller is given responsibility over the motor system. Instead, more sophisticated algorithms exist for blending the output of several modules (e.g., Whitehead et al., 1993; Wolpert & Kawato, 1998). A future simulation could adopt such an algorithm, with the anterior cingulate cortex combining the output of several motor controllers to guide the response selection process. Moreover,² the “temperature” used by the control filter for selecting between controllers (τ in Equation 3, Appendix B) could be gradually reduced over the course of learning, thereby improving the system’s performance once an appropriate motor controller has been identified.

Our treatment of the controllers themselves is even more abstract. We envision only that control over the motor system occurs in a modular fashion, that certain controllers are capable of performing a given task better than others, and that some of the output from all the modules passes through the anterior cingulate cortex. Substantial neuropsychological, neurophysiological, and functional neuroimaging evidence indicates that motor control does proceed in such a modular fashion. For example, areas in the orbitofrontal cortex appear to be involved in guessing (Elliott, Rees, & Dolan, 1999) and in decision making under uncertainty (Bechara, Damasio, Tranel, & Anderson, 1998; Bechara, Damasio, Tranel, & Damasio, 1997; Damasio, 1994). These phenomena require sensitivity to the value associated with different possible response options (Damasio, 1994). Moreover, neurons in the orbitofrontal cortex code for the relative preference between different rewards (Tremblay & Schultz, 1999). The orbitofrontal cortex appears to apply such information for deciding between large, unlikely rewards and small, likely rewards (Rogers et al., 1999) and for abandoning obsolete response patterns in favor of newly rewarded behaviors (Rolls, Hornak, Wade, & McGrath, 1994). In contrast, the dorsolateral prefrontal cortex appears to be more concerned with generating responses when the external context is impoverished, both by brainstorming novel response alternatives when necessary (Eslinger & Grattan, 1993; Milner & Petrides, 1984; see also Channon & Crawford, 1999; Deiber et al., 1991; Frith et al., 1991a, 1991b; Jahanshahi et al., 1995; Jenkins et al., 1994; Petersen et al., 1989) and by maintaining the external context in working memory to generate a response in the absence of external stimulation (Goldman-Rakic, 1989, 1995; Owen, 1997; see also Bechara et al., 1998). Our model asserts only that such different controllers exist, without detailing the specific nature of every one.

Although all the motor controllers in the model receive a copy of the TD error, we do not intend to suggest that this is the case for all the motor controllers in the human brain. Rather, we expect that

² We thank one of the reviewers for making this suggestion.

some neural controllers can execute their normal function without depending on reinforcement learning signals. Nevertheless, given that the mesencephalic dopamine system projects throughout the basal ganglia and frontal cortex, we do suggest that many of the controllers in these regions use the reinforcement signals to carry out their normal function. Disruption of the mesencephalic dopamine system severely upsets the working memory and executive control functions of the prefrontal cortex (Brozoski, Brown, Rosvold, & Goldman, 1979; Dolan et al., 1995; Sawaguchi, 1995; Sawaguchi & Goldman-Rakic, 1991; Schneider & Kovelowski, 1990; Simon, Scatton, & Le Moal, 1980; J. R. Taylor, Elsworth, Roth, Sladek, & Redmond, 1990; G. V. Williams & Goldman-Rakic, 1995; Zahrt, Taylor, Mathew, & Arnsten, 1997; for review, see Arnsten, 1997). Thus, although the impact of the dopamine signal may not be limited to reinforcement learning per se, it may nevertheless facilitate error-processing functions mediated by these areas.

Clearly, many of these regions do depend on error-processing functions that could be facilitated by reinforcement learning. For example, even though individuals with prefrontal damage make repeated errors when performing certain psychological tasks, they can still verbally report when their responses are correct and incorrect (Konow & Pribram, 1970; Milner, 1963; Picton, Stuss, & Marshall, 1986; but see Slachevsky et al., 2001). This observation has motivated the proposition that the prefrontal cortex, although not involved in error detection, is concerned with applying error information toward behavioral modification (Konow & Pribram, 1970; see also Diamond & Goldman-Rakic, 1989; E. Goldberg & Barr, 1991; T. E. Goldberg, Weinberger, Berman, Pliskin, & Podd, 1987; Nauta, 1971; Stuss & Benson, 1986). Error-related activity has also been identified in several neural areas throughout the frontal cortex. For example, two functional neuroimaging studies (Carter et al., 1998; Kiehl et al., 2000) found error-related activity in the human dorsolateral prefrontal cortex, as well as in the anterior cingulate cortex, when human participants committed errors in reaction-time tasks. Single-unit studies in monkeys have also identified error-sensitive units in the dorsolateral prefrontal cortex (Kubota & Komatsu, 1985; Watanabe, 1989) and the supplementary eye field (Stuphorn, Taylor, & Schall, 2000). These results suggest that the dorsolateral prefrontal cortex and the supplementary eye field, in addition to the anterior cingulate cortex, contribute to improving performance at some tasks. Similarly, other studies have found error-related units in the orbitofrontal cortex (Mora, Avrith, Phillips, & Rolls, 1979; Niki, 1982; Rosenkilde, Bauer, & Fuster, 1981; Thorpe, Rolls, & Maddison, 1983; see also Zald & Kim, 1996). These neural signals may reflect contributions by the mesencephalic dopamine system to orbitofrontal function, perhaps by developing a value system needed to guide decision making under uncertainty (Damasio, 1994; see also Elliott et al., 1999).

Our model does not include many of the connections known to exist between the neural areas that are represented in it. Of particular importance is the fact that the model does not include well-known striatocortical loops. These appear to be segregated into two more or less independent systems (Haber & Fudge, 1997). On the one hand, a set of segregated loops project in parallel through "matrisome" modules in the basal ganglia (Graybiel et al., 1994; Graybiel & Kimura, 1995) via the substantia nigra pars reticulata and internal segment of the globus pallidus to the thal-

amus. From there they project in parallel to the frontal cortex and then return to the basal ganglia (Alexander & Crutcher, 1990). These loops are thought by many to form a neural substrate for working memory (e.g. Houk, 1995). On the other hand, "striosome" modules in the basal ganglia compose the primary input into the mesencephalic dopamine system—which, as we have seen, conveys a scalar error signal to the basal ganglia and to the frontal cortex (Graybiel et al., 1994; Graybiel & Kimura, 1995). The striosome modules primarily convey limbic information to the mesencephalic dopamine system, and it has been suggested that they comprise the neural area where TD errors are computed (Barto, 1995; Houk et al., 1995; see also J. Brown et al., 1999). Importantly, the only cortical regions seen to project strongly to the striosomes derive from the caudal orbitofrontal–anterior insular cortex and the mediofrontal prelimbic–anterior cingulate cortex (Eblen & Graybiel, 1995). Thus, neural activity in the anterior cingulate cortex could in principle drive the activity of the mesencephalic dopamine system, which would then be reflected back to the anterior cingulate cortex and other areas. Indeed, stimulation of the medial prefrontal cortex and the anterior cingulate cortex has been shown to induce burst firing in midbrain dopamine neurons (Gariano & Groves, 1988; Tong, Overton, & Clark, 1996; see also Murase, Grenhoff, Chouvet, Gonon, & Svensson, 1993; Svensson & Tung, 1989). This feedback complicates the picture of where exactly error detection occurs. Further research will be needed to elucidate how the existence of this loop bears on the generation of the ERN.

Neurophysiological Considerations

We suggest that dopaminergic disinhibition of the ventral bank of the anterior cingulate sulcus can generate the ERN by enabling large portions of the apical dendrites of Layer V neurons to become depolarized. The EEG is generated by the synchronous activity of postsynaptic potentials acting on the dendrites or soma of cortical pyramidal cells (Goff, Allison, & Vaughan, 1978; Lopes da Silva & Van Rotterdam, 1982; Martin, 1991; Speckmann & Elger, 1982; Wood, 1987). Because pyramidal cells are aligned in parallel, postsynaptic potentials acting across several of them tend to summate by the principle of superposition (Griffiths, 1989) to generate measurable scalp potentials. Their elongated structure also ensures that negative potentials associated with current sinks at one end of the neuron suffer minimal cancellation from positive potentials associated with current sources at the opposite end of the cell. This geometry fixes the dipole in parallel with the direction of pyramidal cell orientation, which is invariably perpendicular to the cortical surface. We suggest that the dopaminergic disinhibition of the apical dendrites of anterior cingulate motor neurons is consistent with these requirements. Indeed, the mechanism we offer here for consideration would seem to be a textbook example of ERP generation.

As we have seen, the mesencephalic dopamine projection terminates diffusely in the human frontal cortex, but its output is focused over the frontal midline (i.e., anterior cingulate and supplementary motor area; Berger et al., 1991; Gaspar et al., 1989; S. M. Williams & Goldman-Rakic, 1993). Dopamine neurons innervate Layer I of the cortical mantle densely, preferentially synapsing there on the apical dendrites of pyramidal cells (Berger et al., 1991; Crino, Morrison, & Hof, 1993; Gaspar et al., 1989;

Goldman-Rakic et al., 1989; Smiley et al., 1992; Verney, Alvarez, Geffard, & Berger, 1990) including Layer V motor neurons (Berger et al., 1991; S. M. Williams & Goldman-Rakic, 1993), which express high levels of dopamine D1 receptors (Berger, Febvret, Greengard, & Goldman-Rakic, 1990). Although the issue has not been resolved, dopamine's influence on deep cells in rat medial prefrontal areas appears to be inhibitory (Bunney & Aghajanian, 1979; Reader, Ferron, Descarries, & Jasper, 1979; but see also Aou, Oomura, Nishino, Inokuchi, & Mizuno, 1983; Cepeda, Radisavljevic, Peacock, Levine, & Buchwald, 1992; Otani, Blond, Desce, & Crepel, 1998; Sawaguchi & Matsumura, 1985; Sawaguchi, Matsumura, & Kubota, 1986; Servan-Schreiber, Printz, & Cohen, 1990). Dopamine reduces the peak amplitude of excitatory postsynaptic potentials evoked by stimulating the Layer I inputs to these neurons (Law-Tho, Hirsch, & Crepel, 1994; cf. Bernardi, Cherubini, Marciani, Mercuri, & Stanzioni, 1982; Otani et al., 1998; Pirot, Glowinski, & Thierry, 1996) and reduces their rate of action potential generation (Bernardi et al., 1982; Gullledge & Jaffe, 1998). Furthermore, stimulation of the VTA inhibits deep cells in the rat medial prefrontal cortex (Ferron, Thierry, Le Douarin, & Glowinski, 1984; Finch, 1993; Mantz, Milla, Glowinski, & Thierry, 1988; Pirot et al., 1992, 1996). Conversely, a reduction of dopaminergic input from the VTA likely disinhibits deep cingulate cortical neurons, resulting in their increased activity (Crino et al., 1993; see also G. V. Williams & Goldman-Rakic, 1995). Hence, in trial-and-error learning tasks, it is possible that the reduction in dopaminergic activity seen on error trials when an expected reward is not delivered (Hollerman & Schultz, 1998; Ljungberg et al., 1991; Schultz et al., 1993) could give rise to the increased activity of anterior cingulate neurons seen in the same circumstances (Niki & Watanabe, 1979).

Layer V neurons in the anterior cingulate motor area are especially capable of contributing to the EEG. These neurons are the defining feature of the cingulate sulcus, each having a single apical dendrite that stretches from the cell body for more than 1 mm to Layer I, where they bifurcate profusely (Vogt, 1993; Vogt et al., 1995). In area 24c'g, these neurons are called "gigantopyramidal" cells in deference to their enormous girth and extensive dendritic arborizations (Braak, 1976; Vogt et al., 1995). Excitation of Layer V neurons generates large amplitude (20.5 ± 1.0 mV) excitatory postsynaptic potentials lasting longer than 50 ms (Vogt & Gorman, 1982; see also Cauller & Connors, 1992). Thus, their elongated structure and strong excitability, together with massive synaptic input onto their apical dendrites (Vogt, 1993; see also Shepherd, 1988), make activity in these neurons plausible candidates for ERP generation. Furthermore, a reduction of dopaminergic input to the apical dendrites of Layer V neurons in the anterior cingulate cortex, in conjunction with strong cortical excitation, activates voltage-dependent Ca^{2+} channels, precipitating a cascade of high-threshold spikes that depolarizes extended portions of their dendritic arbor (Yang & Seamans, 1996). When taken together, the morphology and physiology of anterior cingulate motor neurons in the presence of mesencephalic dopamine would seem to be compatible with the physical principles that underlie ERP generation.

We have described a scenario in which the apical dendrites of the pyramidal cells are depolarized following disinhibition induced by a reduction from baseline of dopaminergic input. In this case, the negative side of the potential field associated with the equivalent dipole lies above the apical dendrites of the pyramidal cells.

Because the ERN is a negativity with a frontal central distribution, the negative pole of the equivalent dipole must point in this direction, where the scalp activity reaches a maximum. In practice, because the equivalent dipole is always oriented in parallel to the pyramidal cells, this constraint implies that the cortical layer that generates the ERN must run parallel to the area of the scalp where the ERN is maximal. Within the medial frontal cortex, the only cortical areas that are oriented in parallel to that area are the sulci. Thus, if the ERN is generated by a single focal dipole layer in the frontal midline, then it must be generated within a sulcus there. Within the frontal midline, only the cingulate sulcus, followed by the paracingulate sulcus, are consistent across the human population (Paus et al., 1996; Vogt et al., 1995). Moreover, because of the orientation of the dipole layer, the ERN must be generated within the ventral bank of the sulcus, where the apical dendrites of the pyramidal cells point toward the scalp. Thus, we suggest that the cingulate motor areas, which belong to the ventral bank of the cingulate sulcus, which itself composes about a quarter of the surface of the medial wall (Vogt et al., 1995), generate the ERN (Coles et al., 1998).

Computational Considerations

Because the method of temporal differences was developed by researchers in the artificial intelligence community, the algorithm might seem to learning theorists to be an odd choice for modeling human behavior. However, the algorithm is appropriate, for several reasons. First, the roots of the algorithm do in fact stem from animal learning theory. As we have noted, the method of temporal differences is a generalization of the Rescorla-Wagner learning rule, an extension of the model from trial-by-trial learning to the continuous time domain (Sutton & Barto, 1990). The algorithm accounts for all of the behavioral observations predicted by the Rescorla-Wagner rule, and more (Sutton & Barto, 1990). Therefore, any perceived dissociation between the algorithm and animal learning theory may be more apparent than real. Second, in addition to accounting for overt behavior, the method of temporal differences simultaneously describes the dynamics of the neural system that gives rise to that behavior (Houk et al., 1995; Montague et al., 1996; Schultz et al., 1997; Waelti et al., 2001). The algorithm also belongs to a class of biologically plausible learning rules (e.g., Mazzoni, Andersen, & Jordan, 1991). Thus the algorithm provides a compact description of learning phenomena at both the behavioral and neurophysiological levels. Third, the algorithm is computationally robust. Several studies have shown that the algorithm converges rapidly under many conditions (e.g., Dayan & Sejnowski, 1994; Jaakkola, Jordan, & Singh, 1994; Sutton, 1988). Fourth, the computational power of the algorithm has been amply demonstrated. For example, the method of temporal differences has been used to teach a computer program to play world-class backgammon (Tesauro, 1994). With 30 game pieces, 26 possible locations for each of those pieces, and on average 20 different ways of playing each dice roll, the number of possible states of the game is huge (Sutton & Barto, 1998). Therefore the algorithm should perform well in real-world situations, where the number of possible world states that predict any particular outcome are presumably at least as great. Lastly, as a reinforcement learning algorithm, the method of temporal differences was designed to solve the kinds of executive control prob-

lems associated with the frontal system, that is, guiding the organism in complex and uncertain environments. Thus the algorithm is appropriate for simulating the function of the neural system in our model.

Our model makes use of eligibility traces. Eligibility traces function as a kind of memory, activating when an associated state is encountered and decaying thereafter (Sutton & Barto, 1998). Each of the motor controllers in the model is associated with an eligibility trace, which all have zero activation at the start of every trial. When the anterior cingulate cortex selects a motor controller during the response generation process, the eligibility trace associated with the winning controller reaches maximum value; when the response is concluded, the trace decreases to a fraction of its maximum value. Thus, the representation of the selected controller is strongest during the selection process, when the winning controller is actually used to generate a response, but is less strong when maintained as a memory. Eligibility traces are typically thought to represent a synaptic mechanism whereby a given synapse remains eligible for learning after the pre- and postsynaptic elements have been activated (e.g., Wickens, 1990). In principle, however, eligibility traces can also be sustained in working memory, at once maintaining an abstract representation of the encountered state and keeping activated the neural units that describe that state. Given the well-known involvement of the prefrontal cortex and the basal ganglia in working memory (L. L. Brown et al., 1997; Goldman-Rakic, 1989, 1995; Houk, 1995; Owen, 1997; S. P. Wise, Murray, & Gerfen, 1996), these systems could in fact contribute to maintaining states eligible for learning.

In our model, an ERN is generated only when (a) the anterior cingulate receives a TD error from the mesencephalic dopamine system and (b) an eligibility trace associated with any of the motor controllers is nonzero. We suggest that the apical dendrites (in Layer I) of Layer V motor neurons in the anterior cingulate cortex receive excitatory input from the motor controllers during the response generation process, and continue to receive input from the winning motor controller after the response generation process has been completed, when the representation of the selected controller is maintained in working memory. We further suggest that the ERN is generated when this excitatory input is disinhibited by a phasic reduction from baseline activity of the mesencephalic dopamine system. In practice, this means that although the TD error propagates back to the imperative stimulus to the degree the stimulus predicts the outcome of the trial, the ERN is never elicited before a response is generated. Thus, the ERN, unlike a negative TD error, specifically indicates when the consequences of a response are worse than expected. This position is consistent with a fundamental principle of operant conditioning, which holds that learning should only occur when the reward or punishment is contingent on the animal's behavior. It also is supported by the work of Porrino (1993), who found that reward-induced activity in the rat anterior limbic cortex is greatest when the reward is contingent on the rat's behavior, compared with when it is delivered independently of what the animal is doing.

Another issue concerns the timing of the ERN with respect to the error. The onset of the ERN coincides with the onset of EMG activity that gives rise to the erroneous response (Gehring, 1992; Gehring et al., 1993). If error detection depended on feedback information generated by the overt response, then ERN onset might be expected to be delayed by about 100 ms following the

onset of error commission (cf. Jeannerod, 1991). For this reason, it has been argued that the process that gives rise to the ERN depends on *effference copy*—a copy of the motor command issued to the motor system, processed by some other part of the brain. Such a command signal could be evaluated even before an erroneous movement has begun (e.g., Angel, 1976). For didactic reasons, however, the system that detects the error in our model depends on feedback from the overt response (Figure 1). This fact might suggest that the ERN should be delayed following error commission. To account for this apparent discrepancy, the model could be modified such that a copy of the control filter output (and/or the motor controller output) could be sent directly to the adaptive critic. In this way, the adaptive critic could monitor internal *effference copy* to detect task-inappropriate neural activity before it has developed into an overt behavior.

The timing of the ERN with respect to the response may have important functional consequences. Unlike reinforcement learning algorithms that wait until an episode has terminated for learning to occur, temporal difference learning can modify behavior on the fly. Thus, the neural signal reflected by the ERN could in principle be used to modulate the activity of the motor system while the behavior is still in progress without having to wait until the end of the trial to observe its outcome. For example, the error signal could initiate a remedial action even before an erroneous response has reached completion. Although this possibility is not realized in our current version of the model, it could be in future simulations. For both computational (Servan-Schreiber et al., 1990) and physiological (Yang & Seamans, 1996) reasons, it has been suggested that the level of dopaminergic input into the prefrontal cortex could modulate the gain of synaptic transmission there. Indeed, Yang and Seamans (1996) have suggested that a reduction of dopaminergic input to the apical dendrites of Layer V pyramidal cells—in the way we have argued generates the ERN—will in fact decrease the gain of the pyramidal cell, enabling a wider range of external input to stimulate the neuron. Lowering the gain, together with noisy input, would increase the variability in the response selection process (Servan-Schreiber et al., 1990). We speculate that a reduction of gain on error trials would enable the anterior cingulate cortex to search for an optimal motor controller over a wider variety of controller inputs. In principle, several controllers might be compared and evaluated before the response has reached the overt stage of execution (cf. Egelman et al., 1998; Suri et al., 2001).³

Other Considerations

Several studies have noted that aversive stimulation can increase the level of extracellular dopamine in animals (e.g., Abercrombie,

³ It is interesting to note that the apical dendrites of Layer V cells have been hypothesized to form a "distal integrative region" in Layer I (Cauller & Connors, 1992; Vogt, 1993). These investigators have suggested that the Layer I interface is computationally significant because it is electrically isolated from the soma, enabling the dendrites there to perform simple logic operations on incoming information (Cauller & Connors, 1992; Shepherd, Woolf, & Carnevale, 1989). They have further argued that synaptic triads on the apical dendrites of Layer V cells form logic gates strategically situated to inhibit cortical-cortical excitation (Shepherd et al., 1989). If this were indeed the case, then the ERN might reflect massive disinhibition of this computational interface.

Keefe, DiFrischia, & Zigmond, 1989; Doherty & Gratton, 1992; Louilot, Le Moal, & Simon, 1986; Young, Joseph, & Gray, 1993). If negative feedback stimuli are understood to be punishment signals, these observations may seem in disagreement with our position that negative feedback should reduce dopaminergic activity. However, we suggest that this potential criticism is not valid, for two reasons. First, our argument is based on the observation that phasic dopaminergic activity decreases from baseline when a predicted reward is not delivered (Hollerman & Schultz, 1998; Ljungberg et al., 1991; Schultz et al., 1993); to our knowledge, this important discovery has not yet been challenged. From a functional standpoint, we suggest that the absence of a predicted reward conveys the same information to a monkey as a negative feedback stimulus does to a human participant, namely, that an error has just been committed. This error information need not be associated with physical pain or discomfort. Thus, although we might loosely call the penalties in our feedback learning paradigm "punishments," these stimuli are not expected to induce the increases in extracellular dopamine that result from physical distress. Second, although the absence of a predicted reward elicits dopaminergic activity on a short time scale (~ 100 ms), the increase in extracellular dopamine induced by punishment occurs on a much longer time scale (> 1 s). This dissociation between fast and slow dopaminergic activity suggests that different neural mechanisms underlie the two processes. Because our argument is agnostic about how such slow activity should affect the ERP, it makes no predictions about whether or not such increases would disrupt the ERN. Specifically, it is unclear whether the slow increases in extracellular dopamine involve D1 receptors located on the apical dendrites of Layer V pyramidal cells and, even if they did, how those increases would affect local signal transduction.

A final note concerns the predictions we used to evaluate the model. Although the simulated ERNs behaved quite similarly to what we see in humans, the variance of those data was relatively small, despite variability in the model's starting weights and a noise term added to the TD error (see Appendix B). It might be argued that the simulation should capture the variance of the observed ERNs in addition to their central tendencies. We disagree with this criticism for the following reasons. First, the ERN (like any ERP) is subject to noise, the unwanted effects of which are typically reduced, but never eliminated, by averaging. Thus the observations obtained with our measurement technique contain variance not associated with the underlying functional process, which we assume conforms to the output of the simulation. Second, as with most biological phenomena, individual variation between participants can be quite large; for example, considerable variation exists between participants in the structure and orientation of sulci on the medial wall (Paus et al., 1996; Vogt et al., 1995). Although such factors can clearly influence the amplitude and distribution of the ERN, it is less clear if they affect the associated neural function. It might also be argued that the simulation should produce ERNs that are given in microvolts, like real ERNs, instead of in arbitrary units. However, because ours is a functional model, not a biophysical model, we are interested mainly in the ordinal relationship between the predictions and not their absolute magnitude. Of course, we could introduce a simple multiplicative factor to bring the predictions in line with the observations, but we feel such an addition would be gratuitous and obscure the core properties of the model.

Conclusion

Demonstrating the existence of biological error signals is fundamentally important to computational neuroscience (Crick, 1989; Thorpe & Imbert, 1989). Our work on the ERN has revealed one such signal in normal human participants, has highlighted the important role it appears to play in executive control, and has suggested a neural area where it might be generated (Coles et al., 1998). This article contributes to that account by grounding the observations in a computational framework that (a) formalizes a general set of statements about the error signal into a coherent model with increased explanatory power and (b) makes predictive statements about the error processing system that can be tested in future experiments. Furthermore, the model makes specific claims about the neural systems that implement the computations and thus provides an avenue for testing the integrity of the theory using patient populations. Although the model requires future elaboration, our initial simulations and experiments have yielded positive results that provide motivation for further research. If valid, this account of the system that gives rise to the ERN may provide insight into how large-scale neural networks for motor control are trained and modulated in biological systems.

References

- Abercrombie, E. D., Keefe, K. A., DiFrischia, D. S., & Zigmond, M. J. (1989). Differential effect of stress on in vivo dopamine release in striatum, nucleus accumbens, and medial frontal cortex. *Journal of Neurochemistry*, 52, 1655–1658.
- Alexander, G. E., & Crutcher, M. D. (1990). Functional architecture of basal ganglia circuits: Neural substrates of parallel processing. *Trends in Neurosciences*, 13, 266–271.
- Allport, A. (1987). Selection for action: Some behavioral and neurophysiological considerations of attention and action. In H. Heuer & A. F. Sanders (Eds.), *Perspectives on perception and action* (pp. 395–419). Hillsdale, NJ: Erlbaum.
- Angel, R. W. (1976). Efference copy in the control of movement. *Neurology*, 26, 1164–1168.
- Aou, S., Oomura, Y., Nishino, H., Inokuchi, A., & Mizuno, Y. (1983). Influence of catecholamines on reward-related neuronal activity in monkey orbitofrontal cortex. *Brain Research*, 267, 165–170.
- Arnsten, A. F. T. (1997). Catecholamine regulation of the prefrontal cortex. *Journal of Psychopharmacology*, 11, 151–162.
- Baddeley, A. D. (1986). The central executive and its malfunctions. In A. D. Baddeley (Ed.), *Working memory* (pp. 224–253). New York: Oxford University Press.
- Barch, D. M., Braver, T. S., Sabb, F. W., & Noll, D. C. (2000). Anterior cingulate and the monitoring of response conflict: Evidence from an fMRI study of overt verb generation. *Journal of Cognitive Neuroscience*, 12, 298–309.
- Barto, A. G. (1995). Adaptive critics and the basal ganglia. In J. Houk, J. Davis, & D. Beiser (Eds.), *Models of information processing in the basal ganglia* (pp. 215–232). Cambridge, MA: MIT Press.
- Barto, A. G., & Sutton, R. S. (1997). Reinforcement learning in artificial intelligence. *Advances in Psychology*, 121, 358–386.
- Bates, J. F., & Goldman-Rakic, P. S. (1993). Prefrontal connections of medial motor areas in the rhesus monkey. *Journal of Comparative Neurology*, 336, 211–228.
- Bechara, A., Damasio, H., Tranel, D., & Anderson, S. W. (1998). Dissociation of working memory from decision making within the human prefrontal cortex. *Journal of Neuroscience*, 18, 428–437.
- Bechara, A., Damasio, H., Tranel, D., & Damasio, A. R. (1997, February

- 28). Deciding advantageously before knowing the advantageous strategy. *Science*, 275, 1293–1295.
- Bentivoglio, M., Kultas-Ilinsky, K., & Ilinsky, I. (1993). Limbic thalamus: Structure, intrinsic organization, and connections. In B. A. Vogt & M. Gabriel (Eds.), *Neurobiology of cingulate cortex and limbic thalamus: A comprehensive handbook* (pp. 71–122). Boston: Birkhauser.
- Berger, B., Febvret, A., Greengard, P., & Goldman-Rakic, P. S. (1990). DARPP-32, a phosphoprotein enriched in dopaminergic neurons bearing dopamine D1 receptors: Distribution in the cerebral cortex of the newborn and adult rhesus monkey. *Journal of Comparative Neurology*, 299, 327–348.
- Berger, B., Gaspar, P., & Verney, C. (1991). Dopaminergic innervation of the cerebral cortex: Unexpected differences between rodents and primates. *Trends in Neurosciences*, 14, 21–27.
- Bernardi, G., Cherubini, E., Marciani, M. G., Mercuri, N., & Stanzione, P. (1982). Responses of intracellularly recorded cortical neurons to the iontophoretic application of dopamine. *Brain Research*, 245, 267–274.
- Berns, G. S., McClure, S. M., Pagnoni, G., & Montague, P. R. (2001). Predictability modulates human brain response to reward. *Journal of Neuroscience*, 21, 2793–2798.
- Bernstein, P. S., Scheffers, M. K., & Coles, M. G. H. (1995). “Where did I go wrong?” A psychophysiological analysis of error detection. *Journal of Experimental Psychology: Human Perception and Performance*, 21, 1312–1322.
- Berridge, K. C., & Robinson, T. E. (1998). What is the role of dopamine in reward: Hedonic impact, reward learning, or incentive salience? *Brain Research Reviews*, 28, 309–369.
- Botvinick, M., Braver, T. S., Barch, D. M., Carter, C. S., & Cohen, J. D. (2001). Conflict monitoring and cognitive control. *Psychological Review*, 108, 624–652.
- Botvinick, M., Nystrom, L. E., Fissell, K., Carter, C. S., & Cohen, J. D. (1999, November 11). Conflict monitoring versus selection-for-action in anterior cingulate cortex. *Nature*, 402, 179–181.
- Braak, H. (1976). A primitive gigantopyramidal field buried in the depth of the cingulate sulcus of the human brain. *Brain Research*, 109, 219–233.
- Braver, T. S., Barch, D. M., Gray, J. R., Molfese, D. L., & Snyder, A. Z. (2001). Anterior cingulate and response conflict: Effects of frequency, inhibition, and errors. *Cerebral Cortex*, 11, 825–836.
- Braver, T. S., & Cohen, J. D. (1999). Dopamine, cognitive control, and the gating model. In J. A. Reggia, E. Rupp, & D. Glanzman (Eds.), *Progress in brain research* (pp. 327–349). Amsterdam: North-Holland/Elsevier Science.
- Braver, T. S. & Cohen, J. D. (2000). On the control of control: The role of dopamine in regulating prefrontal function. In S. Monsell & J. Driver (Eds.), *Attention and performance* (Vol. 18, pp. 713–737). Cambridge, MA: MIT Press.
- Bridle, J. S. (1990). Training stochastic model recognition algorithms as networks can lead to maximum mutual information estimation of parameters. In D. S. Touretzky (Ed.), *Advances in neural information processing systems* (Vol. 2, pp. 211–217). San Mateo, CA: Kaufmann.
- Brooks, V. B. (1986). How does the limbic system assist motor learning? A limbic comparator hypothesis. *Brain, Behavior and Evolution*, 29, 29–53.
- Brooks, V. B. (1989). Limbic assistance in task-related use of motor skill. In J. C. Eccles (Ed.), *Principles of design and operation of the brain* (pp. 343–368). Berlin, Germany: Springer-Verlag.
- Brown, J., Bullock, D., & Grossberg, S. (1999). How the basal ganglia use parallel excitatory and inhibitory learning pathways to selectively respond to unexpected rewarding cues. *Journal of Neuroscience*, 19, 10502–10511.
- Brown, L. L., Schneider, J. S., & Lidsky, T. I. (1997). Sensory and cognitive functions of the basal ganglia. *Current Opinion in Neurobiology*, 7, 157–163.
- Brozoski, T. J., Brown, R. M., Rosvold, H. E., & Goldman, P. S. (1979, August 31). Cognitive deficit caused by regional depletion of dopamine in prefrontal cortex of rhesus monkey. *Science*, 205, 929–932.
- Bunney, B. S., & Aghajanian, G. K. (1976). Dopamine and norepinephrine innervated cells in the rat prefrontal cortex: Pharmacological differentiation using microiontophoretic techniques. *Life Sciences*, 19, 1783–1792.
- Bush, G., Whalen, P. J., Rosen, B. R., Jenike, M. A., McInerney, S. C., & Rauch, S. L. (1998). The counting Stroop: An interference task specialized for functional neuroimaging—validation study with functional MRI. *Human Brain Mapping*, 6, 270–282.
- Bussey, T. J., Muir, J. L., Everitt, B. J., & Robbins, T. W. (1996). Dissociable effects of anterior and posterior cingulate cortex lesions on the acquisition of a conditional visual discrimination: Facilitation of early learning vs. impairment of late learning. *Behavioral Brain Research*, 82, 45–56.
- Calabresi, P., Pisani, A., Mercuri, N. B., & Bernardi, G. (1996). The corticostriatal projection: From synaptic plasticity to dysfunctions of the basal ganglia. *Trends in Neurosciences*, 19, 19–24.
- Carter, C. S., Braver, T. S., Barch, D. M., Botvinick, M. M., Noll, D., & Cohen, J. D. (1998, May 1). Anterior cingulate cortex, error detection, and the online monitoring of performance. *Science*, 280, 747–749.
- Cauler, L. J., & Connors, B. W. (1992). Functions of very distal dendrites: Experimental and computational studies of Layer I synapses on neocortical pyramidal cells. In T. McKenna, J. Davis, & S. F. Zornetzer (Eds.), *Single neuron computation* (pp. 199–229). New York: Academic Press.
- Cepeda, C., Radisavljevic, Z., Peacock, W., Levine, M. S., & Buchwald, N. A. (1992). Differential modulation by dopamine of responses evoked by excitatory amino acids in human cortex. *Synapse*, 11, 330–341.
- Channon, S., & Crawford, S. (1999). Problem-solving in real-life-type situations: The effects of anterior and posterior lesions on performance. *Neuropsychologia*, 37, 757–770.
- Coles, M. G. H., & Rugg, M. D. (1995). Event-related brain potentials: An introduction. In M. D. Rugg & M. G. H. Coles (Eds.), *Electrophysiology of mind* (pp. 1–26). Oxford, England: Oxford University Press.
- Coles, M. G. H., Scheffers, M. K., & Holroyd, C. (1998). Berger’s dream? The error-related negativity and modern cognitive psychophysiology. In H. Witted, U. Zweiner, B. Schack, & A. Döring (Eds.), *Quantitative and topological EEG and MEG analysis* (pp. 96–102). Jena-Erlangen, Germany: Druckhaus Mayer Verlag.
- Coles, M. G. H., Scheffers, M. K., & Holroyd, C. B. (2001). Why is there an ERN/Ne on correct trials? Response representations, stimulus-related components, and the theory of error-processing. *Biological Psychology*, 56, 173–189.
- Corbetta, M., Miezin, F. M., Dobmeyer, S., Shulman, G. L., & Petersen, S. E. (1991). Selective and divided attention during visual discriminations of shape, color, and speed: Functional anatomy by positron emission tomography. *Journal of Neuroscience*, 11, 2383–2402.
- Crick, F. (1989, January 12). The recent excitement about neural networks. *Nature*, 337, 129–132.
- Crino, P. B., Morrison, J. H., & Hof, P. R. (1993). Monoaminergic innervation of cingulate cortex. In B. A. Vogt & M. Gabriel (Eds.), *Neurobiology of cingulate cortex and limbic thalamus: A comprehensive handbook* (pp. 285–310). Boston: Birkhauser.
- Cummings, J. L. (1993). Frontal-subcortical circuits and human behavior. *Archives of Neurology*, 50, 873–880.
- Damasio, A. R. (1994). *Descartes’ error: Emotion, reason, and the human brain*. New York: Avon Books.
- Damasio, A. R., & Van Hoesen, G. W. (1983). Emotional disturbances associated with focal lesions of the limbic frontal lobe. In K. M. Heilman & P. Satz (Eds.), *Neuropsychology of human emotion* (pp. 85–110). New York: Guilford Press.
- Dayan, P., & Sejnowski, T. J. (1994). TD(λ) converges with probability 1. *Machine Learning*, 14, 295–301.
- Degos, J.-D., da Fonseca, N., Gray, F., & Cesaro, P. (1993). Severe frontal

- syndrome associated with infarcts of the left anterior cingulate gyrus and the head of the right caudate nucleus. *Brain*, 116, 1541–1548.
- Dehaene, S., & Changeux, J.-P. (1991). The Wisconsin Card Sort Task: Theoretical analysis and modeling in a neuronal network. *Cerebral Cortex*, 1, 62–79.
- Dehaene, S., Kerszberg, M., & Changeux, J.-P. (1998). A neuronal model of a global workspace in effortful cognitive tasks. *Proceedings of the National Academy of Sciences*, 95, 14529–14534.
- Dehaene, S., Posner, M. I., & Tucker, D. M. (1994). Localization of a neural system for error detection and compensation. *Psychological Science*, 5, 303–305.
- Deiber, M.-P., Passingham, R. E., Colebatch, J. G., Friston, K. J., Nixon, P. D., & Frackowiak, R. S. J. (1991). Cortical areas and the selection of movement: A study with positron emission tomography. *Experimental Brain Research*, 84, 393–402.
- D'Esposito, M., Detre, J. A., Alsop, D. C., Shin, R. K., Atlas, S., & Grossman, M. (1995, November 16). The neural basis of the central executive of working memory. *Nature*, 378, 279–281.
- Devinsky, O., Morrell, M. J., & Vogt, B. A. (1995). Contributions of anterior cingulate cortex to behaviour. *Brain*, 118, 279–306.
- Diamond, A., & Goldman-Rakic, P. S. (1989). Comparison of human infants and rhesus monkeys on Piaget's AB task: Evidence for dependence on dorsolateral prefrontal cortex. *Experimental Brain Research*, 74, 24–40.
- Diehl, B., Dinner, D. S., Mohamed, A., Najm, I., Klem, G., LaPresto, E., et al. (2000). Evidence of cingulate motor representation in humans. *Neurology*, 55, 725–728.
- Doherty, M. D., & Gratton, A. (1992). High-speed chronoamperometric measurements of mesolimbic and nigrostriatal dopamine release associated with repeated daily stress. *Brain Research*, 586, 295–302.
- Dolan, R. J., Fletcher, P., Frith, C. D., Friston, K. J., Frackowiak, R. S. J., & Grasby, P. M. (1995, November 9). Dopaminergic modulation of impaired cognitive activation in the anterior cingulate cortex in schizophrenia. *Nature*, 378, 180–182.
- Dum, R. P., & Strick, P. L. (1993). Cingulate motor areas. In B. A. Vogt & M. Gabriel (Eds.), *Neurobiology of cingulate cortex and limbic thalamus: A comprehensive handbook* (pp. 415–441). Boston: Birkhauser.
- Eblen, F., & Graybiel, A. M. (1995). Highly restricted origin of prefrontal cortical inputs to striosomes in the macaque monkey. *Journal of Neuroscience*, 15, 5999–6013.
- Egelman, D. M., Person, C., & Montague, P. R. (1998). A computational role for dopamine delivery in human decision-making. *Journal of Cognitive Neuroscience*, 10, 623–630.
- Elliott, R., & Dolan, R. J. (1998). Activation of different anterior cingulate foci in association with hypothesis testing and response selection. *NeuroImage*, 8, 17–29.
- Elliott, R., Rees, G., & Dolan, R. J. (1999). Ventromedial prefrontal cortex mediates guessing. *Neuropsychologia*, 37, 403–411.
- Eriksen, B. A., & Eriksen, C. W. (1974). Effects of noise letters upon the identification of a target letter in a nonsearch task. *Perception & Psychophysics*, 16, 143–149.
- Eslinger, P. J., & Grattan, L. M. (1993). Frontal lobe and frontal-striatal substrates for different forms of human cognitive flexibility. *Neuropsychologia*, 31, 17–28.
- Falkenstein, M., Hohnsbein, J., Hoormann, J., & Blanke, L. (1990). Effects of errors in choice reaction tasks on the ERP under focused and divided attention. In C. Brunia, A. Gaillard, & A. Kok (Eds.), *Psychophysiological brain research* (pp. 192–195). Tilburg, the Netherlands: Tilburg University Press.
- Falkenstein, M., Hoormann, J., Christ, S., & Hohnsbein, J. (2000). ERP components on reaction errors and their functional significance: A tutorial. *Biological Psychology*, 51, 87–107.
- Ferron, A., Thierry, A. M., Le Douarin, C., & Glowinski, J. (1984). Inhibitory influence of the mesocortical dopaminergic system on spontaneous activity or excitatory response induced from the thalamic mediodorsal nucleus in the rat medial prefrontal cortex. *Brain Research*, 302, 257–265.
- Finch, D. M. (1993). Hippocampal, subicular, and entorhinal afferents and synaptic integration in rodent cingulate cortex. In B. A. Vogt & M. Gabriel (Eds.), *Neurobiology of cingulate cortex and limbic thalamus: A comprehensive handbook* (pp. 224–248). Boston: Birkhauser.
- Fink, G. R., Frackowiak, R. S. J., Pietrzyk, U., & Passingham, R. E. (1997). Multiple nonprimary motor areas in the human cortex. *Journal of Neurophysiology*, 77, 2164–2174.
- Frank, M. J., Loughry, B., & O'Reilly, R. C. (2001). Interactions between frontal cortex and basal ganglia in working memory: A computational model. *Cognitive, Affective, and Behavioral Neuroscience*, 1, 137–160.
- Friston, K. J., Tononi, G., Reeke, G. N., Jr., Sporns, O., & Edelman, G. M. (1994). Value-dependent selection in the brain: Simulation in a synthetic neural model. *Neuroscience*, 59, 229–243.
- Frith, C. D., Friston, K. J., Liddle, P. F., & Frackowiak, R. S. J. (1991a). A PET study of word finding. *Neuropsychologia*, 29, 1137–1148.
- Frith, C. D., Friston, K. J., Liddle, P. F., & Frackowiak, R. S. J. (1991b). Willed action and the prefrontal cortex in man: A study with PET. *Proceedings of the Royal Society of London*, 244, 241–246.
- Gabriel, M. (1993). Discriminative avoidance learning: A model system. In B. A. Vogt & M. Gabriel (Eds.), *Neurobiology of cingulate cortex and limbic thalamus: A comprehensive handbook* (pp. 478–523). Boston: Birkhauser.
- Gabriel, M., & Taylor, C. (1998). Prefrontal exposure to cocaine impairs neuronal coding of attention and discriminative learning. In J. A. Harvey & B. E. Kosofsky (Eds.), *Annals of the New York Academy of Sciences: Vol. 846. Cocaine: Effects on the developing brain* (pp. 194–212). New York: Academy of Sciences.
- Gariano, R. F., & Groves, P. M. (1988). Burst firing induced in midbrain dopamine neurons by stimulation of the medial prefrontal and anterior cingulate cortices. *Brain Research*, 462, 194–198.
- Gaspar, P., Berger, B., Febvre, A., Vigny, A., & Henry, J. P. (1989). Catecholamine innervation of the human cerebral cortex as revealed by comparative immunohistochemistry of tyrosine hydroxylase and dopamine-beta-hydroxylase. *Journal of Comparative Neurology*, 279, 249–271.
- Gehring, W. J. (1992). *The error-related negativity: Evidence for a neural mechanism for error-related processing*. Unpublished doctoral dissertation, University of Illinois at Urbana-Champaign.
- Gehring, W. J., Coles, M. G. H., Meyer, D. E., & Donchin, E. (1995). A brain potential manifestation of error-related processing. In G. Karmos, M. Molnar, V. Csepe, I. Czizler, & J. E. Desmedt (Eds.), *Perspectives of event-related potentials research (EEG Suppl. 44)* (pp. 261–272). Amsterdam: North-Holland/Elsevier Science.
- Gehring, W. J., Goss, B., Coles, M. G. H., Meyer, D. E., & Donchin, E. (1993). A neural system for error detection and compensation. *Psychological Science*, 4, 385–390.
- Gehring, W. J., Himle, J., & Nisenson, L. G. (2000). Action-monitoring deficits in obsessive-compulsive disorder. *Psychological Science*, 11, 1–6.
- Gehring, W. J., & Knight, R. T. (2000). Prefrontal-cingulate interactions in action monitoring. *Nature Neuroscience*, 3, 516–520.
- Gemba, H., Sasaki, K., & Brooks, V. B. (1986). "Error" potentials in limbic cortex (anterior cingulate area 24) of monkeys during motor learning. *Neuroscience Letters*, 70, 223–227.
- Goff, W. R., Allison, T., & Vaughan, H. G. (1978). The functional neuroanatomy of event-related potentials. In E. Callaway, P. Tueting, & S. Koslow (Eds.), *Event-related brain potentials in man* (pp. 1–91). New York: Academic Press.
- Goldberg, E., & Barr, W. B. (1991). Three possible mechanisms of unawareness of deficit. In G. P. Prigatano & D. Schacter (Eds.), *Un-*

- awareness of deficit after brain injury (pp. 152–175). New York: Oxford University Press.
- Goldberg, T. E., Weinberger, D. R., Berman, K. F., Pliskin, N. H., & Podd, M. H. (1987). Further evidence for dementia of the prefrontal type in schizophrenia? *Archives of General Psychiatry*, 44, 1008–1014.
- Goldman-Rakic, P. S. (1989). The prefrontal contribution to working memory and conscious experience. In J. C. Eccles (Ed.), *Principles of design and operation of the brain* (pp. 389–410). Berlin, Germany: Springer-Verlag.
- Goldman-Rakic, P. S. (1995). Toward a circuit model of working memory and the guidance of voluntary motor action. In J. Houk, J. Davis, & D. Beiser (Eds.), *Models of information processing in the basal ganglia* (pp. 131–148). Cambridge, MA: MIT Press.
- Goldman-Rakic, P. S., Leranth, C., Williams, S. M., Mons, N., & Geffard, M. (1989). Dopamine synaptic complex with pyramidal neurons in primate cerebral cortex. *Proceedings of the National Academy of Sciences, USA*, 86, 9015–9019.
- Gratton, G., Coles, M. G. H., & Donchin, E. (1983). A new method for off-line removal of ocular artifact. *Electroencephalography and Clinical Neurophysiology*, 55, 468–484.
- Graybiel, A. M., Aosaki, T., Flaherty, A. W., & Kimura, M. (1994, September 23). The basal ganglia and adaptive motor control. *Science*, 265, 1826–1831.
- Graybiel, A. M., & Kimura, M. (1995). Adaptive neural networks in the basal ganglia. In J. Houk, J. Davis, & D. Beiser (Eds.), *Models of information processing in the basal ganglia* (pp. 103–116). Cambridge, MA: MIT Press.
- Griffiths, D. J. (1989). *Introduction to electrodynamics*. Englewood Cliffs, NJ: Prentice-Hall.
- Guarraci, F. A., & Kapp, B. S. (1999). An electrophysiological characterization of ventral tegmental area dopaminergic neurons during differential Pavlovian fear conditioning in the awake rabbit. *Behavioral Brain Research*, 99, 169–179.
- Gulledge, A. T., & Jaffe, D. B. (1998). Dopamine decreases the excitability of Layer V pyramidal cells in the rat prefrontal cortex. *Journal of Neuroscience*, 18, 9139–9151.
- Haber, S. N., & Fudge, J. L. (1997). The primate substantia nigra and VTA: Integrative circuitry and function. *Critical Reviews in Neurobiology*, 11, 323–342.
- Herrnstein, R. J. (1997). *The matching law*. New York: Russell Sage Foundation.
- Hollerman, J. R., & Schultz, W. (1998). Dopamine neurons report an error in the temporal prediction of reward during learning. *Nature Neuroscience*, 1, 304–309.
- Holroyd, C. B. (2001). *Reinforcement learning and the error-related negativity: A computational and neurophysiological investigation*. Doctoral dissertation, University of Illinois at Urbana-Champaign.
- Holroyd, C. B., & Coles, M. G. H. (2002). *Response selection in a modified version of the Eriksen Flankers Task*. Manuscript in preparation.
- Holroyd, C. B., Dien, J., & Coles, M. G. H. (1998). Error-related scalp potentials elicited by hand and foot movements: Evidence for an output-independent error-processing system in humans. *Neuroscience Letters*, 242, 65–68.
- Holroyd, C., Reichler, J., & Coles, M. G. H. (1998). Generation of error-related scalp potentials by a mesencephalic dopamine system for error processing: Hypothesis and model. In *Proceedings of the Second International Conference on Cognitive and Neural Systems* (p. 114). Boston: Department of Cognitive and Neural Systems, Boston University.
- Holroyd, C., Reichler, J., & Coles, M. G. H. (1999, April). Is the error-related negativity generated by a dopaminergic error signal for reinforcement learning? Hypothesis and model. In abstract program of the Cognitive Neuroscience Society 1999 annual meeting, Washington DC (p. 45).
- Houk, J. C. (1995). Information processing in modular circuits linking basal ganglia and cerebral cortex. In J. Houk, J. Davis, & D. Beiser (Eds.), *Models of information processing in the basal ganglia* (pp. 3–9). Cambridge, MA: MIT Press.
- Houk, J. C., Adams, J. L., & Barto, A. G. (1995). A model of how the basal ganglia generate and use neural signals that predict reinforcement. In J. Houk, J. Davis, & D. Beiser (Eds.), *Models of information processing in the basal ganglia* (pp. 249–270). Cambridge, MA: MIT Press.
- Jaakkola, T., Jordan, M. I., & Singh, S. P. (1994). On the convergence of stochastic iterative dynamic programming algorithms. *Neural Computation*, 6, 1185–1201.
- Jacobs, R. A., Jordan, M. I., & Barto, A. G. (1991). Task decomposition through competition in a modular connectionist architecture: The what and where vision tasks. *Cognitive Science*, 15, 219–250.
- Jahanshahi, M., Jenkis, I. H., Brown, R. G., Marsden, C. D., Passingham, R. E., & Brooks, D. J. (1995). Self-initiated versus externally triggered movements. *Brain*, 118, 913–933.
- Jasper, H. H. (1958). The ten twenty electrode system of the international federation. *Electroencephalography and Clinical Neurophysiology*, 10, 371–375.
- Jeannerod, M. (1991). The interaction of visual and proprioceptive cues in controlling reaching movements. In D. R. Humphrey, & H.-J. Freund (Eds.), *Motor control: Concepts and issues* (pp. 277–291). New York: Wiley.
- Jenkins, I. H., Brooks, D. J., Nixon, P. D., Frackowiak, R. S. J., & Passingham, R. E. (1994). Motor sequence learning: A study with positron emission tomography. *Journal of Neuroscience*, 14, 3775–3790.
- Johnson, T. M., Otten, L. J., Boeck, K., & Coles, M. G. H. (1997). Am I too late? The neural consequences of missing a deadline. *Psychophysiology*, 34, S48.
- Kalmar, Z., Szepesvari, C., & Lorincz, A. (1998). Module-based reinforcement learning: Experiments with a real robot. *Machine Learning*, 31, 55–85.
- Kawagoe, R., Takikawa, Y., & Hikosaka, O. (1998). Expectation of reward modulates cognitive signals in the basal ganglia. *Nature Neuroscience*, 1, 411–416.
- Kiehl, K. A., Liddle, P. F., & Hopfinger, J. B. (2000). Error processing and the rostral anterior cingulate: An event-related fMRI study. *Psychophysiology*, 37, 216–223.
- Kimura, M., & Matsumoto, N. (1997). Nigrostriatal dopamine system may contribute to behavioral learning through providing reinforcement signals to the striatum. *European Neurology*, 38, 11–17.
- Knowlton, B. J., Mangels, J. A., & Squire, L. R. (1996, September 6). A neostriatal habit learning system in humans. *Science*, 273, 1399–1402.
- Knutson, B., Westdorp, A., Kaiser, E., & Hommer, D. (2000). fMRI visualization of brain activity during a monetary incentive delay task. *NeuroImage*, 12, 20–27.
- Konow, A., & Pribram, K. H. (1970). Error recognition and utilization produced by injury to the frontal cortex in man. *Neuropsychologia*, 8, 489–491.
- Kopp, B., & Wolff, M. (2000). Brain mechanisms of selective learning: Event-related potentials provide evidence for error-driven learning in humans. *Biological Psychology*, 51, 223–246.
- Koski, L., & Paus, T. (2000). Functional connectivity of the anterior cingulate cortex within the human frontal lobe: A brain mapping analysis. *Experimental Brain Research*, 133, 55–65.
- Kubota, K., & Komatsu, H. (1985). Neuron activities of monkey prefrontal cortex during the learning of visual discrimination tasks with go/no-go performances. *Neuroscience Research*, 3, 106–129.
- Law-Tho, D., Hirsch, J. C., & Crepel, F. (1994). Dopamine modulation of synaptic transmission in rat prefrontal cortex: An in vitro electrophysiological study. *Neuroscience Research*, 21, 151–160.
- Ljungberg, T., Apicella, P., & Schultz, W. (1991). Responses of midbrain

- dopamine neurons during delayed alternation performance. *Brain Research*, 567, 337–341.
- Ljungberg, T., Apicella, P., & Schultz, W. (1992). Responses of monkey dopamine neurons during learning of behavioral reactions. *Journal of Neurophysiology*, 67, 145–163.
- Logan, G. D. (1985). Executive control of thought and action. *Acta Psychologica*, 60, 193–210.
- Lopes da Silva, F., & Van Rotterdam, A. (1982). Biophysical aspects of EEG and MEG generation. In E. Niedemeyer & F. Lopez da Silva (Eds.), *Electroencephalography: Basic principles, clinical applications, and related fields* (pp. 1–13). Baltimore: Urban & Schwarzenberg.
- Louilot, A., Le Moal, M., & Simon, H. (1986). Differential reactivity of dopaminergic neurons in the nucleus accumbens in response to different behavioral situations: An in vivo voltammetric study in free-moving rats. *Brain Research*, 397, 395–400.
- Luria, A. R. (1973). *The working brain: An introduction to neuropsychology* (B. Haigh, Trans.). New York: Basic Books.
- Luu, P., Flaisch, T., & Tucker, D. M. (2000). Medial frontal cortex in action monitoring. *Journal of Neuroscience*, 20, 464–469.
- Machlin, S. R., Harris, G. J., Pearson, G. D., Hoehn-Saric, R., Jeffrey, P., & Camargo, E. E. (1991). Elevated medial-frontal cerebral blood flow in obsessive-compulsive patients: A SPECT study. *American Journal of Psychiatry*, 148, 1240–1242.
- Mantz, J., Milla, C., Glowinski, J., & Thierry, A.M. (1988). Differential effects of ascending neurons containing dopamine and noradrenaline in the control of spontaneous activity and of evoked responses in the rat prefrontal cortex. *Neuroscience*, 27, 517–526.
- Martin, J. H. (1991). The collective electrical behavior of cortical neurons: The electroencephalogram and the mechanisms of epilepsy. In E. R. Kandel, J. H. Schwartz, & T. M. Jessell (Eds.), *Principles of neural science* (pp. 777–791). Norwalk, CT: Appleton & Lange.
- Mazzoni, P., Andersen, R. A., & Jordan, M. I. (1991). A more biologically plausible learning rule for neural networks. *Proceedings of the National Academy of Sciences, USA*, 88, 4433–4437.
- Menon, V., Adelman, N. E., White, C. D., Glover, G. H., & Reiss, A. L. (2001). Error-related brain activation during a go/no-go response inhibition task. *Human Brain Mapping*, 12, 131–143.
- Milner, B. (1963). Effects of different brain lesions on card sorting. *Archives of Neurology*, 9, 90–100.
- Milner, B., & Petrides, M. (1984). Behavioral effects of frontal-lobe lesions in man. *Trends in Neurosciences*, 7, 403–407.
- Miltner, W. H. R., Braun, C. H., & Coles, M. G. H. (1997). Event-related brain potentials following incorrect feedback in a time-estimation task: Evidence for a “generic” neural system for error detection. *Journal of Cognitive Neuroscience*, 9, 788–798.
- Miltner, W. H. R., Lemke, U., Weiss, T., Holroyd, C., Scheffers, M. K., & Coles, M. G. H. (1997). The source of the magnetic equivalent of the error-related negativity. *Psychophysiology*, 34, S65.
- Mirenowicz, J., & Schultz, W. (1996, February 1). Preferential activation of midbrain dopamine neurons by appetitive rather than aversive stimuli. *Nature*, 379, 449–451.
- Montague, P. R., Dayan, P., & Sejnowski, T. J. (1996). A framework for mesencephalic dopamine systems based on predictive Hebbian learning. *Journal of Neuroscience*, 16, 1936–1947.
- Mora, F., Avrieth, D. B., Phillips, A. G., & Rolls, E. T. (1979). Effects of satiety on self-stimulation of the orbitofrontal cortex in the rhesus monkey. *Neuroscience Letters*, 13, 141–145.
- Mora, F., & Myers, R. D. (1977, September 30). Brain self-stimulation: Direct evidence for the involvement of dopamine in the prefrontal cortex. *Science*, 197, 1387–1389.
- Morecraft, R. J., & Van Hoesen, G. W. (1991). A comparison of frontal lobe inputs to the primary, supplementary, and cingulate motor areas in the monkey. *Society for Neuroscience Abstracts*, 17, 1019.
- Morecraft, R. J., & Van Hoesen, G. W. (1998). Convergence of limbic input to the cingulate motor cortex in the rhesus monkey. *Brain Research Bulletin*, 45, 209–232.
- Morecraft, R. J., Van Hoesen, G. W., & Maynard, J. A. (1989). Cortical afferents to caudal area 24c (the cingulate motor area) and rostral area 23c. *Society for Neuroscience Abstracts*, 15, 73.
- Murase, S., Grenhoff, J., Chouvet, G., Gonon, F. G., & Svensson, T. H. (1993). Prefrontal cortex regulates burst firing and transmitter release in rat mesolimbic dopamine neurons studied in vivo. *Neuroscience Letters*, 157, 53–56.
- Nauta, W. J. H. (1971). The problem of the frontal lobe: A reinterpretation. *Journal of Psychiatric Research*, 8, 167–187.
- Nieuwenhuis, S., Ridderinkhof, K. R., Blom, J., Band, G. P. H., & Kok, A. (2001). Error-related potentials are differentially related to awareness of response errors: Evidence from an antisaccade task. *Psychophysiology*, 38, 752–760.
- Niki, H. (1982). Reward-related and error-related neurons in the primate frontal cortex. In S. Saito & T. Yanagita (Eds.), *Learning and memory drugs as reinforcer* (pp. 22–34). Amsterdam: Excerpta Medica.
- Niki, H., & Watanabe, M. (1979). Prefrontal and cingulate unit activity during timing behavior in the monkey. *Brain Research*, 171, 213–224.
- Nishijo, H., Yamamoto, Y., Ono, T., Uwano, T., Yamashita, J., & Yamashima, T. (1997). Single neuron responses in monkey anterior cingulate cortex during visual discrimination. *Neuroscience Letters*, 227, 79–82.
- Norman, D. A., & Shallice, T. (1986). Attention to action: Willed and automatic control of behavior. In R. J. Davidson, G. E. Schwartz, & D. Shapiro (Eds.), *Consciousness and self-regulation* (Vol. 4, pp. 1–19). New York: Plenum Press.
- Olds, J., & Milner, P. (1954). Positive reinforcement produced by electrical stimulation of septal area and other regions of rat brain. *Journal of Comparative and Physiological Psychology*, 47, 419–427.
- O'Reilly, R. C., Noelle, D. C., Braver, T. S., & Cohen, J. D. (2002). Prefrontal cortex and dynamic categorization tasks: Representational organization and neuromodulatory control. *Cerebral Cortex*, 12, 246–257.
- Otani, S., Blond, O., Desce, J.-M., & Crepel, F. (1998). Dopamine facilitates long-term depression of glutamatergic transmission in rat prefrontal cortex. *Neuroscience*, 85, 669–676.
- Owen, A. M. (1997). The functional organization of working memory processes within human lateral frontal cortex: The contribution of functional neuroimaging. *European Journal of Neuroscience*, 9, 1329–1339.
- Pandya, D. N., Van Hoesen, G. W., & Mesulam, M. M. (1981). Efferent connections of the cingulate gyrus in the monkey. *Experimental Brain Research*, 42, 319–330.
- Paus, T., Petrides, M., Evans, A. C., & Meyer, E. (1993). Role of the human anterior cingulate cortex in the control of oculomotor, manual, and speech responses: A positron emission tomography study. *Journal of Neurophysiology*, 70, 453–462.
- Paus, T., Tomaiuolo, F., Otaky, N., MacDonald, D., Petrides, M., Atlas, J., et al. (1996). Human cingulate and paracingulate sulci: Pattern, variability, asymmetry, and a probabilistic map. *Cerebral Cortex*, 6, 207–214.
- Pennartz, C. M. A. (1996). The ascending neuromodulatory systems in learning by reinforcement: Comparing computational conjectures with experimental findings. *Brain Research Reviews*, 21, 219–245.
- Petersen, S. E., Fox, P. T., Posner, M. I., Mintun, M., & Raichle, M. E. (1989). Positron emission tomographic studies of the processing of single words. *Journal of Cognitive Neuroscience*, 1, 153–170.
- Phillips, A. G., & Fibiger, H. C. (1989). Neuroanatomical bases of intracranial self-stimulation: Untangling the Gordian knot. In J. M. Leibman & J. J. Cooper (Eds.), *The neuropharmacological basis of reward* (pp. 66–105). New York: Oxford University Press.
- Phillips, A. G., Pfaus, J. G., & Blaha, C. D. (1991). Dopamine and motivated behavior: Insights provided by in vivo analyses. In P. Willner

- & J. Scheel-Kruger (Eds.), *The mesolimbic dopamine system: From motivation to action* (pp. 199–224). New York: Wiley.
- Picard, N., & Strick, P. L. (1996). Motor areas of the medial wall: A review of their location and functional activation. *Cerebral Cortex*, 6, 342–353.
- Picton, T. W., Stuss, D. T., & Marshall, K. C. (1986). Attention and the brain. In S. L. Friedman, K. A. Klivingson, & R. W. Peterson (Eds.), *The brain, cognition, and education* (pp. 19–79). Orlando, FL: Academic Press.
- Pirot, S., Glowinski, J., & Thierry, A.-M. (1996). Mediodorsal thalamic evoked responses in the rat prefrontal cortex: Influence of the mesocortical DA system. *Neuroreport*, 7, 1437–1441.
- Pirot, S., Godbout, R., Mantz, J., Tassin, J.-P., Glowinski, J., & Thierry, A.-M. (1992). Inhibitory effects of ventral tegmental area stimulation on the activity of prefrontal cortical neurons: Evidence for the involvement of both dopaminergic and gabaergic components. *Neuroscience*, 49, 857–865.
- Porrino, L. J. (1993). Cortical mechanisms of reinforcement. In B. A. Vogt & M. Gabriel (Eds.), *Neurobiology of cingulate cortex and limbic thalamus: A comprehensive handbook* (pp. 445–460). Boston: Birkhauser.
- Porrino, L. J., Esposito, R. U., Seeger, T. F., Crane, A. M., Pert, A., & Sokoloff, L. (1984, April 20). Metabolic mapping of the brain during rewarding self-stimulation. *Science*, 224, 306–309.
- Posner, M. I., & DiGirolamo, G. J. (1998). Executive attention: Conflict, target detection, and cognitive control. In R. Parasuraman (Ed.), *The attentive brain* (pp. 401–423). Cambridge, MA: MIT Press.
- Posner, M. I., Petersen, S. E., Fox, P. T., & Raichle, M. E. (1988, June 17). Localization of cognitive operations in the human brain. *Science*, 240, 1627–1631.
- Posner, M. I., & Raichle, M. E. (1994). *Images of mind*. New York: Scientific American Library.
- Pratt, W. E., & Mizumori, S. J. Y. (2001). Neurons in rat medial prefrontal cortex show anticipatory rate changes to predictable differential rewards in a spatial memory task. *Behavioral Brain Research*, 123, 165–183.
- Procyk, E., Tanaka, Y. L., & Joseph, J. P. (2000). Anterior cingulate activity during routine and non-routine sequential behaviors in macaques. *Nature Neuroscience*, 3, 502–508.
- Raichle, M. E., Fiez, J. A., Videen, T. O., MacLeod, A.-M. K., Pardo, J. V., Fox, P. T., & Petersen, S. E. (1994). Practice-related changes in human brain functional anatomy during non-motor learning. *Cerebral Cortex*, 4, 8–26.
- Rapoport, J. L., & Wise, S. P. (1988). Obsessive-compulsive disorder: Evidence for basal ganglia dysfunction. *Psychopharmacology Bulletin*, 24, 380–384.
- Reader, T. A., Ferron, A., Descarries, L., & Jasper, H. H. (1979). Modulatory role for biogenic amines in the cerebral cortex: Microiontophoretic studies. *Brain Research*, 160, 217–229.
- Redgrave, P., Prescott, T. J., & Gurney, K. (1999). Is the short-latency dopamine response too short to signal reward error? *Trends in Neurosciences*, 22, 146–151.
- Rescorla, R. A., & Wagner, A. R. (1972). A theory of Pavlovian conditioning: Variations in the effectiveness of reinforcement and nonreinforcement. In A. H. Black & W. F. Prokasy (Eds.), *Classical conditioning II: Current research and theory* (pp. 64–99). New York: Appleton-Century-Crofts.
- Richardson, N. R., & Gratton, A. (1996). Behavior-relevant changes in nucleus accumbens dopamine transmission elicited by food reinforcement: An electrochemical study in the rat. *Journal of Neuroscience*, 16, 8160–8169.
- Richardson, N. R., & Gratton, A. (1998). Changes in medial prefrontal cortical dopamine levels associated with response-contingent food reward: An electrochemical study in rat. *Journal of Neuroscience*, 18, 9130–9138.
- Robbins, T. W., & Everitt, B. J. (1996). Neurobehavioral mechanisms of reward and motivation. *Current Opinion in Neurobiology*, 6, 228–236.
- Rogers, R. O., Owen, A. M., Middleton, H. C., Williams, E. J., Pickard, J. D., Sahakian, B. J., & Robbins, T. W. (1999). Choosing between small, likely rewards and large, unlikely rewards activates inferior and orbital prefrontal cortex. *Journal of Neuroscience*, 19, 9029–9038.
- Rolls, E. T., Hornak, J., Wade, D., & McGrath, J. (1994). Emotion-related learning in patients with social and emotional changes associated with frontal lobe damage. *Journal of Neurology, Neurosurgery, and Psychiatry*, 57, 1518–1524.
- Rosenkilde, C. E., Bauer, R. H., & Fuster, J. M. (1981). Single cell activity in ventral prefrontal cortex of behaving monkeys. *Brain Research*, 209, 375–394.
- Routtenberg, A., & Sloan, M. (1972). Self-stimulation in the frontal cortex of *rattus norvegicus*. *Behavioral Biology*, 7, 567–572.
- Sawaguchi, T. (1995). The role of dopamine in frontal motor cortical functions of monkeys. In M. Kimura & A. M. Graybiel (Eds.), *Functions of the cortical-basal ganglia loop* (pp. 166–188). Tokyo: Springer-Verlag.
- Sawaguchi, T., & Goldman-Rakic, P. S. (1991, February 22). D1 dopamine receptors in prefrontal cortex: Involvement in working memory. *Science*, 251, 947–950.
- Sawaguchi, T., & Matsumura, M. (1985). Laminar distributions of neurons sensitive to acetylcholine, noradrenaline, and dopamine in the dorsolateral prefrontal cortex of the monkey. *Neuroscience Research*, 2, 255–273.
- Sawaguchi, T., Matsumura, M., & Kubota, K. (1986). Catecholamine sensitivities of motor cortical neurons of the monkey. *Neuroscience Letters*, 66, 135–140.
- Scheffers, M. K. (1999). *Performance monitoring: Error detection and the error-related negativity in choice-reaction time tasks*. Doctoral dissertation, University of Illinois at Urbana-Champaign.
- Scherg, M. (1990). Fundamentals of dipole source potential analysis. In F. Grandori, M. Hoke, & G. L. Romani (Eds.), *Auditory evoked magnetic fields and electric potentials* (pp. 40–69). Basel, Switzerland: Karger.
- Scherg, M., & Picton, T. W. (1991). Separation and identification of event-related potential components by brain electric source analysis. In C. H. M. Brunia, G. Mulder, & M. N. Verbaten (Eds.), *Event-related brain research (EEG suppl. 42)* (pp. 24–37). Amsterdam: North-Holland/Elsevier.
- Schneider, J. S., & Kovelowski, C. J., II. (1990). Chronic exposure to low doses of MPTP: I. Cognitive deficits in motor asymptomatic monkeys. *Brain Research*, 519, 122–128.
- Schultz, W. (1997). Dopamine neurons and their role in reward mechanisms. *Current Opinion in Neurobiology*, 7, 191–197.
- Schultz, W. (1998). Predictive reward signal of dopamine neurons. *Journal of Neurophysiology*, 80, 1–27.
- Schultz, W., Apicella, P., & Ljungberg, T. (1993). Responses of monkey dopamine neurons to reward and conditioned stimuli during successive steps of learning a delayed response task. *Journal of Neuroscience*, 13, 900–913.
- Schultz, W., Dayan, P., & Montague, P. R. (1997, March 14). A neural substrate of prediction and reward. *Science*, 275, 1593–1599.
- Schultz, W., Romo, R., Ljungberg, T., Mirenowicz, J., Hollerman, J. R., & Dickinson, A. (1995). Reward-related signals carried by dopamine neurons. In J. Houk, J. Davis, & D. Beiser (Eds.), *Models of information processing in the basal ganglia* (pp. 233–248). Cambridge, MA: MIT Press.
- Servan-Schreiber, D., Printz, H., & Cohen, J. D. (1990, August 24). A network model of catecholamine effects: Gain, signal-to-noise ratio, and behavior. *Science*, 249, 892–895.
- Shepherd, G. M. (1988). A basic circuit of cortical organization. In M. Gazzaniga (Ed.), *Perspectives in memory research* (pp. 93–134). Cambridge, MA: MIT Press.

- Shepherd, G. M., Woolf, T. B., & Carnevale, N. T. (1989). Comparisons between active properties of distal dendritic branches and spines: Implications for neuronal computations. *Journal of Cognitive Neuroscience*, 1, 273–286.
- Shima, K., & Tanji, J. (1998, November 13). Role for cingulate motor area cells in voluntary movement selection based on reward. *Science*, 282, 1335–1338.
- Simon, H., Scatton, B., & Le Moal, M. (1980, July 10). Dopaminergic A10 neurones are involved in cognitive functions. *Nature*, 286, 150–151.
- Slachevsky, A., Pillon, B., Fournier, P., Pradat-Diehl, P., Jeannerod, M., & Dubois, B. (2001). Preserved adjustment but impaired awareness in a sensory motor conflict following prefrontal lesions. *Journal of Cognitive Neuroscience*, 13, 332–340.
- Smiley, J. F., Williams, S. M., Szigeti, K., & Goldman-Rakic, P. S. (1992). Light and electron microscopic characterization of dopamine-immunoreactive axons in human cerebral cortex. *Journal of Comparative Neurology*, 321, 325–335.
- Smith, A. D., & Bolam, J. P. (1990). The neural network of the basal ganglia as revealed by the study of synaptic connections of identified neurones. *Trends in Neurosciences*, 13, 259–265.
- Sokoloff, L., Reivich, M., Kennedy, C., Des Rosiers, M. H., Patlak, C. S., Pettigrew, K., et al. (1977). The [^{14}C] deoxyglucose method for the measurement of local cerebral glucose utilization: Theory, procedure, and normal values in the conscious and anesthetized albino rat. *Journal of Neurochemistry*, 28, 897–916.
- Speckmann, E.-J., & Elger, C. E. (1982). Neurophysiological basis of the EEG and of DC potentials. In E. Niedermeyer & F. Lopez da Silva (Eds.), *Electroencephalography: Basic principles, clinical applications, and related fields* (pp. 1–13). Baltimore: Urban & Schwarzenberg.
- Steinmetz, H., Furst, G., & Meyer, B.-U. (1989). Craniocerebral topography within the international 10–20 system. *Electroencephalography and Clinical Neurophysiology*, 72, 499–506.
- Stuphorn, V., Taylor, T. L., & Schall, J. D. (2000, December 14). Performance monitoring by the supplementary eye field. *Nature*, 408, 857–860.
- Stuss, D. T., & Benson, D. F. (1986). *The frontal lobes*. New York: Raven Press.
- Suri, R. E., Bargas, J., & Arbib, M. A. (2001). Modeling functions of striatal dopamine modulation in learning and planning. *Neuroscience*, 103, 65–85.
- Suri, R. E., & Schultz, W. (1998). Learning of sequential movements by neural network model with dopamine-like reinforcement. *Experimental Brain Research*, 121, 350–354.
- Suri, R. E., & Schultz, W. (1999). A neural network model with dopamine-like reinforcement signal that learns a spatial delayed response task. *Neuroscience*, 91, 871–890.
- Suri, R. E., & Schultz, W. (2001). Temporal difference model reproduces anticipatory neural activity. *Neural Computation*, 13, 841–862.
- Sutton, R. S. (1988). Learning to predict by the methods of temporal differences. *Machine Learning*, 3, 9–44.
- Sutton, R. S., & Barto, A. G. (1990). Time-derivative models of Pavlovian reinforcement. In M. Gabriel & J. Moore (Eds.), *Learning and computational neuroscience: Foundations of adaptive networks* (pp. 497–537). Cambridge, MA: MIT Press.
- Sutton, R. S., & Barto, A. G. (1998). *Reinforcement learning: An introduction*. Cambridge, MA: MIT Press.
- Svensson, T. H., & Tung, C.-S. (1989). Local cooling of pre-frontal cortex induces pacemaker-like firing of dopamine neurons in rat tegmental area in vivo. *Acta Physiologica Scandinavica*, 136, 135–136.
- Swedo, S. E., Schapiro, M. B., Grady, C. L., Cheslow, D. L., Leonard, H. L., Kumar, A., et al. (1989). Cerebral glucose metabolism in childhood-onset obsessive-compulsive disorder. *Archives of General Psychiatry*, 46, 518–523.
- Takenouchi, K., Nishijo, H., Uwano, T., Tamura, R., Takigawa, M., & Ono, T. (1999). Emotional and behavioral correlates of the anterior cingulate cortex during associative learning in rats. *Neuroscience*, 93, 1271–1287.
- Taylor, C., Freeman, J. H., Jr., Holt, W., & Gabriel, M. (1997). Prenatal exposure to cocaine impairs anterior cingulate cortical neuronal activity and learning in rabbits. *Society for Neuroscience Abstracts*, 23, 262.
- Taylor, J. R., Elsworth, J. D., Roth, R. H., Sladek, J. R., & Redmond, D. E., Jr. (1990). Cognitive and motor deficits in the acquisition of an object retrieval/detour task in MPTP-treated monkeys. *Brain*, 113, 617–637.
- Tesauro, G. (1994). TD-gammon, a self-teaching backgammon program, achieves master-level play. *Neural Computation*, 6, 215–219.
- Thorndike, E. L. (1970). Laws and hypotheses for behavior. In E. L. Thorndike, *Animal intelligence* (pp. 241–281). Darien, CT: Hafner Publishing Co. (Original work published 1911)
- Thorpe, S. J., & Imbert, M. (1989). Biological constraints on connectionist modelling. In R. Pfeifer, Z. Schreier, F. Fogelman-Soulie, & L. Steels (Eds.), *Connectionism in perspective* (pp. 63–92). Amsterdam: North-Holland/Elsevier Science.
- Thorpe, S. J., Rolls, E. T., & Maddison, S. (1983). The orbitofrontal cortex: Neuronal activity in the behaving monkey. *Experimental Brain Research*, 49, 93–115.
- Tong, Z.-Y., Overton, P. G., & Clark, D. (1996). Stimulation of the prefrontal cortex in the rat induces patterns of activity in midbrain dopaminergic neurons which resemble natural burst events. *Synapse*, 22, 195–208.
- Toni, I., & Passingham, R. E. (1999). Prefrontal–basal ganglia pathways are involved in the learning of arbitrary visuomotor associations: A PET study. *Experimental Brain Research*, 127, 19–32.
- Tremblay, L., & Schultz, W. (1999, April 22). Relative reward preference in primate orbitofrontal cortex. *Nature*, 398, 704–708.
- Turken, A. U., & Swick, D. (1999). Response selection in the human anterior cingulate cortex. *Nature Neuroscience*, 2, 920–924.
- Ullsperger, M., & von Cramon, D. Y. (2001). Subprocesses of performance monitoring: A dissociation of error processing and response competition revealed by event-related fMRI and ERPs. *NeuroImage*, 14, 1387–1401.
- Van Hoesen, G. W., Morecraft, R. J., & Vogt, B. A. (1993). Connections of the monkey cingulate cortex. In B. A. Vogt & M. Gabriel (Eds.), *Neurobiology of cingulate cortex and limbic thalamus: A comprehensive handbook* (pp. 249–284). Boston: Birkhauser.
- Van 't Ent, D., & Apkarian, P. (1999). Motoric response inhibition in finger movement and saccadic eye movement: A comparative study. *Clinical Neurophysiology*, 110, 1058–1072.
- Verney, C., Alvarez, C., Geffard, M., & Berger, B. (1990). Ultrastructural double-labelling study of dopamine terminals and GABA-containing neurons in rat anteromedial cerebral cortex. *European Journal of Neuroscience*, 2, 960–972.
- Vidal, F., Hasbroucq, T., Grapperon, J., & Bonnet, M. (2000). Is the “error negativity” specific to errors? *Biological Psychology*, 51, 109–128.
- Vogt, B. A. (1993). Structural organization of cingulate cortex: Areas, neurons, and somatodendritic transmitter receptors. In B. A. Vogt & M. Gabriel (Eds.), *Neurobiology of cingulate cortex and limbic thalamus: A comprehensive handbook* (pp. 19–70). Boston: Birkhauser.
- Vogt, B. A., & Gorman, A. L. F. (1982). Responses of cortical neurons to stimulation of corpus callosum in vitro. *Journal of Neurophysiology*, 48, 1257–1273.
- Vogt, B. A., Nimchinsky, E. A., Vogt, L. J., & Hof, P. R. (1995). Human cingulate cortex: Surface features, flat maps, and cytoarchitecture. *Journal of Comparative Neurology*, 359, 490–506.
- Vogt, B. A., Sikes, R. W., & Vogt, L. J. (1993). Anterior cingulate cortex and the medial pain system. In B. A. Vogt & M. Gabriel (Eds.), *Neurobiology of cingulate cortex and limbic thalamus: A comprehensive handbook* (pp. 313–344). Boston: Birkhauser.
- Waelti, P., Dickinson, A., & Schultz, W. (2001, July 5). Dopamine re-

- sponses comply with basic assumptions of formal learning theory. *Nature*, 412, 43–48.
- Watanabe, M. (1989). The appropriateness of behavioral responses coded in post-trial activity of primate prefrontal units. *Neuroscience Letters*, 101, 113–117.
- Whitehead, S., Karlsson, J., & Tenenber, J. (1993). Learning multiple goal behavior via task decomposition and dynamic policy merging. In J. H. Connell & S. Mahadevan (Eds.), *Robot learning* (pp. 45–78). Boston: Kluwer Academic Publishers.
- Wickelgren, I. (1997, October 3). Getting the brain's attention. *Science*, 278, 35–37.
- Wickens, J. (1990). Striatal dopamine in motor activation and reward-mediated learning: Steps toward a unifying model. *Journal of Neural Transmission*, 80, 9–31.
- Wickens, J. R., Begg, A. J., & Arbuthnott, G. W. (1996). Dopamine reverses the depression of rat corticostriatal synapses which normally follows high-frequency stimulation of cortex in vitro. *Neuroscience*, 70, 1–5.
- Wickens, J., & Kotter, R. (1995). Cellular models of reinforcement. In J. Houk, J. Davis, & D. Beiser (Eds.), *Models of information processing in the basal ganglia* (pp. 187–214). Cambridge, MA: MIT Press.
- Wilkinson, L. S., Humby, T., Killcross, A. S., Torres, E. M., Everitt, B. J., & Robbins, T. W. (1998). Dissociations in dopamine release in medial prefrontal cortex and ventral striatum during the acquisition and extinction of classical aversive conditioning in the rat. *European Journal of Neuroscience*, 10, 1019–1026.
- Williams, G. V., & Goldman-Rakic, P. S. (1995, August 17). Modulation of memory fields by dopamine D1 receptors in prefrontal cortex. *Nature*, 376, 572–575.
- Williams, S. M., & Goldman-Rakic, P. S. (1993). Characterization of the dopaminergic innervation of the primate frontal cortex using a dopamine-specific antibody. *Cerebral Cortex*, 3, 199–222.
- Wise, R. A., & Rompre, P.-P. (1989). Brain dopamine and reward. *Annual Review of Psychology*, 40, 191–225.
- Wise, R. A., Spindler, J., DeWit, H., & Gerber, G. J. (1978, July 21). Neuroleptic-induced "anhedonia" in rats: Pimozide blocks reward quality of food. *Science*, 201, 262–264.
- Wise, S. P., Murray, E. A., & Gerfen, C. R. (1996). The frontal cortex–basal ganglia system in primates. *Critical Reviews in Neurobiology*, 10, 317–356.
- Wise, S. P., & Rapoport, J. L. (1991). Obsessive–compulsive disorder: Is it basal ganglia dysfunction? In J. L. Rapoport (Ed.), *Obsessive–compulsive disorder in children and adolescents* (pp. 327–344). Washington, DC: American Psychiatric Press.
- Wolpert, D. M., & Kawato, M. (1998). Multiple paired forward and inverse models for motor control. *Neural Networks*, 11, 1317–1329.
- Wood, C. (1987). Generators of event-related potentials. In A. M. Halliday, S. R. Butler, & R. Paul (Eds.), *Textbook of clinical neurophysiology* (pp. 535–567). New York: Wiley.
- Yang, C. R., & Seamans, J. K. (1996). Dopamine D1 receptor actions in layers V–VI rat prefrontal cortex neurons in vitro: Modulation of dendritic–somatic signal integration. *Journal of Neuroscience*, 16, 1922–1935.
- Young, A. M. J., Joseph, M. H., & Gray, J. A. (1993). Latent inhibition of conditioned dopamine release in rat nucleus accumbens. *Neuroscience*, 54, 5–9.
- Zahrt, J., Taylor, J. R., Mathew, R. G., & Arnsten, A. F. T. (1997). Supranormal stimulation of D1 dopamine receptors in the rodent prefrontal cortex impairs spatial working memory performance. *Journal of Neuroscience*, 17, 8528–8535.
- Zald, D. H., & Kim, S. W. (1996). Anatomy and function of the orbitofrontal cortex: II. Function and relevance to obsessive–compulsive disorder. *Journal of Neuropsychiatry*, 8, 249–261.

Appendix A

Recording and Data Analysis Procedures

Recording

The electroencephalogram (EEG) was recorded using disposable Ag/AgCl electrodes. Scalp electrodes were placed according to the 10–20 system (Jasper, 1958) in locations Fz, Cz, Pz, and Oz. Additionally, electrodes were placed at C3' and C4'; these are situated 4 cm to either side of the midline at Cz, over areas corresponding to the hand representations of the central sulcus (Steinmetz, Furst, & Meyer, 1989). An electrode was placed on the right mastoid and a ground electrode on the forehead. The scalp and right mastoid electrodes were referenced to an electrode placed over the left mastoid, and their impedances were less than 5 k Ω . The vertical and horizontal electrooculograms (EOGs) were collected with reciprocally referenced electrodes placed above and below the right eye and on the outer canthus of the left and right eye, respectively; EOG impedances were less than 10 k Ω . The EEG and EOG signals were sampled at 200 Hz. The Neuroscan (Neuroscan, El Paso, TX) and PEARL (University of Illinois at Urbana–Champaign) data acquisition and stimulus display programs were used for control of the probabilistic learning (Task 1) and modified Eriksen Flankers Task (Task 2) experiments, respectively.

Data Analysis

The EEG was rereferenced to linked mastoid electrodes off line by subtracting half of the activity recorded by the right mastoid from that collected by the scalp electrodes. Ocular artifact was removed from the EEG with the eye movement correction procedure (EMCP) described in Gratton, Coles, and Donchin (1983). For both tasks, the EEG was baseline corrected on each trial by subtracting the average activity of each channel during the 200 ms preceding the onset of the imperative stimulus from the remaining samples in the response period; for Task 1, the average activity of each channel during the 200 ms preceding feedback onset was subtracted from the remaining samples following feedback onset. The single-trial data were filtered with a bandpass of 1–10 Hz using the Interactive Data Language (Research Systems, Inc., Boulder, CO) digital filter algorithm. Stimulus-locked and response-locked event-related brain potential averages were derived for each participant and condition by averaging the data recorded at channel Cz across trials according to stimulus onset and response onset, respectively.

(Appendixes continue)

Appendix B

Simulation Details

Task 1: Probabilistic Learning

Ten blocks of 300 trials each were simulated for each of 15 “participants” (see text). Each trial was divided into 60 equal, discrete time steps, with stimulus presentation occurring at Time Step 10, response generation at Time Step 30, and feedback presentation at Time Step 60. A schematic of the model is depicted in Figure 1. An input layer consisted of six units, one for each stimulus in a given block. At Time Step 10, one of the six input units was selected at random, and its activation was set equal to 1. This unit remained activated until the end of the trial. Five motor control modules were each composed of two units, one for each response option (left vs. right button presses), and each of the input units were connected through feed-forward weights to all 10 motor controller units. Response selection proceeded as follows. At Time Step 30, the net activation of each of the motor control units was computed as the weighted sum of its inputs:

$$m_t^{j,k} = \sum_{i=1}^6 u_t^{k,j,i} x_i^j \quad (\text{B1})$$

Here, t was the current time step, x_i^j was the activation of input unit i , $m_t^{j,k}$ was the activation of unit j of motor controller k , and $u_t^{k,j,i}$ were the weights connecting the input units to each unit of each motor controller. Each motor controller k selected a possible response $j = j'$ according to the softmax probability function (Bridle, 1990):

$$P(\text{action } j, k \text{ selected}) = \frac{e^{m_t^{j,k}/\tau}}{\sum_{j=1}^2 e^{m_t^{j,k}/\tau}} \quad (\text{B2})$$

This function probabilistically selected one of each controller’s two responses according to their relative activations. In Equation B2, τ was a parameter called the *temperature*; large values of τ ensured that both response options were equally probable, whereas small values of τ made the unit with the higher activation most likely to be selected. We set different values of τ for each controller ($\tau = 0.1, 1.0, 5.0, 10.0, 100.0$), with the expectation that the controller with the smallest temperature should learn to perform the task best. For each controller k , the activation $m_t^{j',k}$ of the winning response option $j = j'$ was set to 1 and the activation of the losing response option $j \neq j'$ was set to 0.

Also at Time Step 30, a control filter (Figure 1) selected a motor controller $k = k'$, again according to a softmax probability function:

$$P(\text{controller } k \text{ selected}) = \frac{e^{w^k/\tau}}{\sum_{k=1}^5 e^{w^k/\tau}} \quad (\text{B3})$$

Here w^k was a weight associated with each of the five controllers, and $\tau = 2$. An output layer was composed of 2 units, one for each possible response, and the control filter activated (by setting its activation level to 1) the output unit associated with the winning response option j' of the winning controller k' . This output unit remained activated until the end of the trial. An eligibility trace ε associated with each of the k motor controllers was determined as follows:

$$\varepsilon(k \neq k', t) = 0 \quad \text{for all } t$$

$$\varepsilon(k = k', t) = \begin{cases} 0 & t < 30 \\ 1 & t = 30, \\ \omega & t > 30 \end{cases} \quad (\text{B4})$$

where ω was a fraction between zero and one (.42). Finally, a feedback layer consisted of two units, one for reward and one for punishment. At Time Step 60, one of two feedback units were activated (by setting its activation level to 1) according to the schedule described in the text.

The input layer, output layer, and feedback layer were connected with feed-forward weights to a “value layer” in the adaptive critic, which consisted of 22 units. These weights were constructed such that the value units were sensitive to the following input. Eight of the value layer units detected when any of the six stimuli, or either of the two responses, occurred in isolation. Another 12 units each computed a particular conjunction between one of the external stimuli and one of the possible response options. Finally, the last two units activated when either a positive or negative reward was presented, respectively. No more than one value unit was activated at any one time, and its activation was set equal to 1. Each of the value units projected with weight z to a summation unit \hat{V} ; the estimated value \hat{V}_t at the current time step t was taken as the linear sum of the activations v of each of the value units times their weight z :

$$\hat{V}_t = \sum_{l=1}^{22} z_l^l v_l^l \quad (\text{B5})$$

Thus the value of the summation unit \hat{V} was just the weight z^l associated with the activated unit l .

At each time step, the adaptive critic computed the temporal difference error (TD error).

$$\delta_t = r_t + \gamma \hat{V}_t - \hat{V}_{t-1} + \text{err} \quad (\text{B6})$$

Here r_t was the feedback signal presented at Time Step 60, taking on values of +1 (reward) and −1 (penalty); the “discount factor” γ was a constant (1.0); and “err” was a random value sampled uniformly at each time step from the range −.25 to +.25. The err term was included to simulate the observed variability in ERN amplitude. The TD error was sent to the value layer, where it was used to learn better predictions; to the motor controllers, where it was used to learn appropriate actions; and to the control filter, where it was used to learn to select the appropriate controller. At each time step t , the weights of the value layer were updated according to

$$z_{t+1}^l = z_t^l + \eta \delta_{t+1} v_t^l \quad (\text{B7})$$

the input layer-controller weights were updated as

$$u_{t+1}^{k,j,i} = u_t^{k,j,i} + \xi \delta_{t+1} x_i^j m_t^{k,j} \varepsilon(k, t), \quad (\text{B8})$$

and new controller-filter weights were given by

$$w_{t+1}^k = w_t^k + \zeta \delta_{t+1} \varepsilon(k, t). \quad (\text{B9})$$

Here η (.2), ξ (.5), and ζ (.2) were learning rate parameters. The value layer weights were initialized with random values sampled uniformly from the range −.25 to +.25, and other plastic weights were initialized with random values sampled uniformly from the range .25–.75.

The amplitude of the ERN was given as

$$\text{amp} = \delta_t \sum_{k=1}^5 \varepsilon(k, t). \quad (\text{B10})$$

Task 2: Modified Eriksen Flankers Task

This simulation adopted the model described in Task 1, but it only simulated the activity of the adaptive critic (see text). The simulation was

run 15 times, once for each participant in the experiment, using the actual experimental data. The input layer consisted of four units (one for each stimulus type), the output layer was composed of two units (for left-hand and right-hand responses), and the feedback layer was made up of two units (for correct and incorrect responses). On each trial, the input unit corresponding to the stimulus seen by the participant on that trial was activated at Time Step 10 and it remained activated until the end of the trial. At Time Step 30, the output unit corresponding to the response the participant made on that trial was activated, and it also remained activated until the end of the trial. At Time Step 60, the feedback unit associated with the outcome of the trial (correct or incorrect) was activated. Stimulus input, response output, and feedback information went to the value layer of the adaptive critic, which consisted of 16 "Value" units (v). Six of the units detected when any of the four stimuli, or either of the two responses, occurred in isolation. Another eight units each computed a particular conjunction between one of the external stimuli and one of the possible response options. Finally, the last two units were activated by the type of feedback (correct or incorrect). Activation of the value layer was winner-take-all.

Each of the value units projected with weight z to a summation unit; the estimated value \hat{V}_t at the current time step t was taken as the linear sum of the activations v of each of the value units times their weight z . At each time step, the adaptive critic outputted the TD error (Equation B6), where r_t was the feedback type, taking on values of $+1$ (correct) and -1 (incorrect); and the "discount factor" γ was a constant (1.0). The TD error was sent to the value layer, where it was used to learn better predictions (Equation B7; $\eta = .2$). Value weights associated with the "correct" conjunctions of stimuli and responses were initially set to $+1$, whereas value weights associated with the "incorrect" conjunctions of stimuli and responses were initially set to -1 ; all other value weights were initialized with random values sampled uniformly from the range $-.25$ to $+.25$. The amplitude of the ERN was given by Equation B10, taken at Time Step 30, when the response was generated.

Received August 15, 2000

Revision received September 10, 2001

Accepted October 10, 2001 ■

Call for Nominations

The Publications and Communications (P&C) Board has opened nominations for the editorships of *Contemporary Psychology: APA Review of Books*, *Developmental Psychology*, and *Psychological Review* for the years 2005–2010. Robert J. Sternberg, PhD, James L. Dannemiller, PhD, and Walter Mischel, PhD, respectively, are the incumbent editors.

Candidates should be members of APA and should be available to start receiving manuscripts in early 2004 to prepare for issues published in 2005. Please note that the P&C Board encourages participation by members of underrepresented groups in the publication process and would particularly welcome such nominees. Self-nominations are also encouraged.

Search chairs have been appointed as follows:

- *Contemporary Psychology: APA Review of Books*: Susan H. McDaniel, PhD, and Mike Pressley, PhD
- *Developmental Psychology*: Joseph J. Campos, PhD
- *Psychological Review*: Mark I. Appelbaum, PhD

To nominate candidates, prepare a statement of one page or less in support of each candidate. Address all nominations to the appropriate search committee at the following address:

Karen Sellman, P&C Board Search Liaison
Room 2004
American Psychological Association
750 First Street, NE
Washington, DC 20002-4242

The first review of nominations will begin November 15, 2002. The deadline for accepting nominations is November 25, 2002.