



Published in final edited form as:

Neuroimage. 2010 February 15; 49(4): 3210. doi:10.1016/j.neuroimage.2009.11.065.

Competition between learned reward and error outcome predictions in anterior cingulate cortex

William H. Alexander and Joshua W. Brown

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

Abstract

The anterior cingulate cortex (ACC) is implicated in performance monitoring and cognitive control. Non-human primate studies of ACC show prominent reward signals, but these are elusive in human studies, which instead show mainly conflict and error effects. Here we demonstrate distinct appetitive and aversive activity in human ACC. The error likelihood hypothesis suggests that ACC activity increases in proportion to the likelihood of an error, and ACC is also sensitive to the consequence magnitude of the predicted error. Previous work further showed that error likelihood effects reach a ceiling as the potential consequences of an error increase, possibly due to reductions in the average reward. We explored this issue by independently manipulating reward magnitude of task responses and error likelihood while controlling for potential error consequences in an incentive change signal task. The fMRI results ruled out a modulatory effect of expected reward on error likelihood effects in favor of a competition effect between expected reward and error likelihood. Dynamic causal modeling showed that error likelihood and expected reward signals are intrinsic to the ACC rather than received from elsewhere. These findings agree with interpretations of ACC activity as signaling both perceptions of risk and predicted reward.

Introduction

Executive control theories require the ability to monitor behavioral consequences and, when necessary, exert goal-directed control over behavior. Anterior cingulate cortex (ACC) has been implicated in performance monitoring and cognitive control (Carter, Braver et al., 1998; Botvinick, Nystrom et al., 1999). Research has identified the ACC and surrounding medial prefrontal cortex (mPFC) as an area that responds to error commission and error feedback (Gemba, Sasaki et al., 1986; Hohnsbein, Falkenstein et al., 1989; Gehring, Coles et al., 1990) as well as response conflict (Botvinick, Nystrom et al., 1999; MacDonald, Cohen et al., 2000). Recently, combined fMRI and computational modeling studies showed that ACC activity is proportional to the likelihood of committing an error, even controlling for error and conflict effects (Brown & Braver, 2005). Additional studies motivated by *a priori* predictions of the error likelihood model have shown that ACC is, more generally, sensitive to expected risk (Brown & Braver, 2007), i.e. the combination of error likelihood and the potential severity of the error. ACC activity related to anticipation of risk seems to drive risk avoidance (Magno, Foxe et al., 2006; Brown & Braver, 2007). Despite the success of the error likelihood model, previous neuroimaging results showed an under-additive interaction of anticipated error

Address correspondence to: Joshua W. Brown, Dept. of Psychological & Brain Sciences, 1101 E Tenth St., Bloomington, IN 47405 USA, +1 812 855-9282 (office), +1 812 855-4691 (FAX), jwmbrown@indiana.edu, www.indiana.edu/~cclab.

Publisher's Disclaimer: This is a PDF file of an unedited manuscript that has been accepted for publication. As a service to our customers we are providing this early version of the manuscript. The manuscript will undergo copyediting, typesetting, and review of the resulting proof before it is published in its final citable form. Please note that during the production process errors may be discovered which could affect the content, and all legal disclaimers that apply to the journal pertain.

consequence magnitude and error likelihood (Brown & Braver, 2007), which was not predicted by the model. This suggests that an additional factor may be involved in driving ACC activity. Evidence from single-unit recording, fMRI, and ERP studies suggests that ACC is not only sensitive to error commission, but also to prediction and processing of rewarding events. Neurons in monkey ACC and nearby supplementary motor area become increasingly activated in proportion to the temporal proximity (Amador, Schlag-Rey et al., 2000; Shidara & Richmond, 2002; Ito, Stuphorn et al., 2003) and predicted level of reward (Amiez, Joseph et al., 2005).

While effects of both error likelihood and, in monkeys, the level of reward have been observed in ACC and related areas, it remains unclear how these effects combine in ACC. Here we test between two competing hypotheses of how reward prediction signals might combine with risk prediction (error likelihood and anticipated error consequence magnitude) signals in human ACC. One possibility, the *modulation* hypothesis, suggests that increased anticipated reward will increase the sensitivity to error likelihood and potential error consequence magnitude. Intuitively, a subject might not care about error likelihood if there is no reward to be gained. This would account for the previous finding that as the potential magnitude of error consequences increases, sensitivity to error likelihood decreases, because expected value decreases. An alternative hypothesis, the *competition* hypothesis, suggests that reward anticipation and error likelihood can each activate ACC, and activation by one reduces sensitivity to another. In support of this hypothesis, competition between reward-seeking and risk-avoidant behavior has been proposed to explain group differences in decision-making tasks (Yechiam, Busemeyer et al., 2005). This would also account for the underadditive interaction between error likelihood and anticipated consequence magnitude. Nonetheless, the two hypotheses make strong competing predictions: If anticipated error magnitude is held constant, increases in anticipated reward magnitude should *increase* error likelihood effects under the modulation hypothesis but *decrease* error likelihood effects under the competition hypothesis.

A question raised by the competing hypotheses outlined above is, in the event that ACC activity is influenced by anticipated reward magnitude, either through competition or modulation, then what is the source of information about reward magnitude to the ACC? One possibility is that ACC receives input from additional areas in the brain, and that these signals are integrated by the ACC along with information about error likelihood and risk prediction. The *integration* hypothesis suggests that brain regions whose activity reflects predictions of reward magnitude should have a causal influence on ACC activity. Alternately, the ACC itself may compute a prediction of reward magnitude independent of similar calculations which occur elsewhere in the brain. The *computation* hypothesis suggests that brain areas outside ACC which show effects of reward magnitude are causally independent of ACC activity, or that they themselves may be causally affected by ACC activity. These additional hypotheses may be differentiated by analyses designed to determine causation amongst brain regions (e.g., DCM; Friston, Harrison et al., 2003).

Methods

To examine the potential role of reward magnitude of a task on ACC activity related to error likelihood, we implement a modified version of the Incentive Change Signal Task (ICST; (Brown & Braver, 2005; Brown & Braver, 2007), as shown in Figure 1. The modified ICST here manipulates error likelihood and changes in reward magnitude while controlling for error consequence magnitude.

Participants

Participants (N=24, 13 female, ages 19 to 36, average age 23, right handed) were recruited from the campus of Indiana University, Bloomington and surrounding areas, using flyers posted in public spaces. Participants were paid \$25 per hour plus a performance bonus (see below) averaging approximately \$6.70. Recruiting and experimental procedures were approved by the Indiana University Institutional Review Board.

Behavioral Task

The Incentive Change Signal Task (ICST) (Brown & Braver, 2007) is a modified version of the change signal task (Brown & Braver, 2005) and was implemented in E-Prime (Psychology Software Tools, Pittsburgh, PA). The ICST consisted of four phases: color cue, target, response, and feedback (Fig. 1). At the beginning of each trial two horizontal dashes were displayed in the center of the screen. Dashes were one of four colors: white, brown, yellow, or light blue. Each color was paired with one of the four possible combinations of error likelihood (high and low) and average reward magnitude (high and low). These pairings were counterbalanced across all participants, and the pairings were constant across all trials for an individual participant. Trials were presented pseudo-randomly. After the dashes were displayed for 1000 msec, an angle brace appeared to the right or left of the dashes, forming an arrow pointing either left or right (48 pt font). The direction of the arrow indicated which response the participant was to make. On change signal trials (1/3 of all trials), an additional arrow (96 pt font) appeared above the first arrow and pointing in the opposite direction, indicating that the participant was to cancel the initial response and make a response according to the second arrow. The stimuli remained visible for 1000 msec after the appearance of the first arrow. The change signal delay (CSD) between onset of the initial arrow and the second arrow was adjusted by an asymmetric staircase algorithm to maintain target error rates, and the CSD was adjusted independently for each of the four colors. For the low error likelihood (EL) conditions, the CSD was adjusted to achieve an error rate of 5% on change signal trials, while an error rate of 50% was maintained for the high error likelihood conditions. On each change trial, the CSD was increased for a correct trial, while incorrect trials decreased the CSD. After presentation of the stimuli and expiration of the response deadline, the screen was blank for 500 msec, after which visual feedback was provided to the participants for 1000 msec. For correct trials, feedback consisted of the word 'Correct' and 4 digits indicating how many points the participant earned for the trial. For incorrect trials, participants saw the word 'Incorrect' in addition to the number of points earned. The number of points earned on each trial depended both on the outcome (correct or incorrect) of the trial as well as the average reward magnitude (RM) condition. For the high RM condition, subjects earned 2000 points for a correct trial and 1000 points for an incorrect trial, while in the low RM condition subjects earned 1000 points for a correct trial and 0 points for an incorrect trial. Participants were informed that their points were to be converted directly to a cash payment at the end of the session. Points were converted at the rate of 1000 points for each US \$0.01. Participants were not informed of the conversion rate of points to dollars prior to participation, nor were they given direct information regarding their accumulated point total. We found in pilot studies that subjects performed with greater motivation for large amounts of points, with conversion factors revealed after the session, than for the equivalent relatively small monetary payment. After feedback, the screen remained black for a minimum of 1500 msec until the start of the next trial. Intertrial intervals (ITI) were jittered by adding 0, 2000, 4000, or 6000 msec (3 TRs) to the ITI. Jitter delays were chosen by a weighted random selection of each of the possible durations; the weights for each of the jitter durations were 30, 12, 5 and 2, respectively, allowing for efficient estimation of the HRF (Burock et al., 1998).

Participants performed 6 blocks of 82 trials per block in the scanner. Participants were trained on the task prior to scanning in order to familiarize them with the task instructions, but not the

specific reward magnitude and error likelihood conditions. Training typically consisted of fewer than the 82 trials comprising a single block. Subjects learned the payoff amounts and probabilities associated with each color cue condition solely by experience while performing the task in the scanner, as in previous studies (Brown & Braver, 2005; Brown & Braver, 2007). Differences in BOLD signals due to effects of reward magnitude and error likelihood are therefore the result of experience with the task during scanning, and not previous training.

In the task design, reward magnitude was manipulated by adding 1000 points to the outcomes such that a correct response in the high reward magnitude condition was worth 1000 points more than a correct response in the low reward magnitude condition, and, similarly, an error was worth 1000 points more in the high reward magnitude condition than in the low reward magnitude condition. Average reward is commonly calculated as the sum of the probability of each potential outcome multiplied by the value of that outcome; in the current task, manipulation of error likelihood necessarily affects the actual expected value of each condition. In high error likelihood conditions, a participant is more likely to commit an error, leading to a lower average reward than in the low error likelihood condition. Critically, however, changes in average reward are the same across conditions: the difference between the average reward (high RM and low RM) in the low error likelihood conditions is the same as the difference in the high error likelihood condition (see figure 1). By manipulating the predicted reward magnitude, the task design allows us to distinguish between the modulation hypothesis and competition hypothesis, as follows. Greater predicted reward magnitude should lead to greater error likelihood effects under the modulation hypothesis but smaller error likelihood effects under the competition hypothesis.

Individual Differences

Previous studies have found that error likelihood-related activity in ACC may vary with individual differences related to risky behavior (Brown & Braver, 2007; Brown & Braver, 2008). Since risk-taking behavior may influence error likelihood effects, participants were given the Domain-Specific Risk Taking inventory (DOSPERT; (Weber, Blais et al., 2002) in order to assess individual propensities for risk. The DOSPERT measures risk-taking behavior within different domains (Social, Recreational, Gambling, Investment, Ethical, and Health/Safety). In the context of reinforcement learning, the propensity to engage in risky behaviors may also be related to impaired or altered function of neuromodulatory systems such as dopamine that underlie reinforcement learning (Riba, Kramer et al., 2008). For the ICST, aversion to risky financial behavior, and especially gambling, is most relevant to the incentive component of the task.

Functional Imaging

Functional images were acquired using a Siemens 3T Trio MRI scanner with images slices tilted 30° toward the coronal plane from the AC-PC line for whole brain coverage (EPI, 33 slices, 3mm slice thickness, TR=2000 msec, TE=25, flip angle = 70, FOV=220×220mm, 64×64 voxel in-plane resolution). T1-weighted structural images for each participant were also acquired (160 sagittal slices, 1mm slice thickness, TR=2300 msec, TE=3.93, flip angle = 12, pixel width in-plane = 0.5mm).

Event-related responses were estimated using a general linear model approach and analyses conducted using SPM5 and the Marsbar toolkit for ROI analyses (Brett, Anton et al., 2002). A GLM was estimated for each subject using a total of 17 regressors: a constant term, 6 regressors for movement, and 10 regressors for experimental conditions. 8 regressors were used to model correct trials for all combinations of levels of high vs. low reward magnitude, high vs. low error likelihood, and change vs. go trials (i.e., trials in which a change signal was either presented or not presented). Events were time-locked to the onset of each trial

(appearance of angle brace indicating which response the subject should make) and were modeled as having duration of 0 seconds (as is standard in SPM). Error trials were modeled by two regressors, one for errors made for change trials, and another for errors committed when no change signal was presented or when no response was made. Beta values for model regressors were estimated using the SPM canonical hemodynamic response function (HRF). Analyses for main effects, interactions, and pairwise comparisons were done at the 2nd-level (random effects), and performed only for correct go trials at the whole-brain level. Planned analyses included tests for error likelihood effects (correct/go/high EL – correct/go/low EL), tests for main effects of reward magnitude (correct/go/highRM – correct/go/low RM) as well as the interaction of reward magnitude and error likelihood for correct go trials. Except where noted, regions of interest for additional analyses were selected by the peak area of activation for clusters of activation that passed whole-brain (family-wise error) correction for planned analyses.

Dynamic Causal Modeling

We investigated potential causal relationships between regions showing a significant main effect of reward magnitude and ACC using dynamic causal modeling (DCM; Friston, Harrison et al., 2003). DCM treats interconnected brain regions as a nonlinear input-state-output system which is sensitive to experimental perturbations. Bayesian estimation is used to estimate parameters for the direct influence of exogenous inputs (e.g., experimental manipulations) on system states, the coupling of states, and parameters that modulate state coupling. The Bayesian framework is combined with a forward model of how neural activity is affected by input and produces measured BOLD responses.

For each extra-cingulate brain area showing significant main effects of reward magnitude we constructed four DCM models embodying four possible causal relationships (summarized in figure 2): (1) unidirectional causation originating from ACC, (2) unidirectional causation originating from outside the cingulate, and reciprocal causation originating either from (3) ACC or (4) extracingulate areas. DCM models were estimated individually for each subject, taking the average time course of activity of voxels within a sphere (5mm radius) centered on the peak coordinates found for group level contrast. For each model, specific conditions (levels of reward magnitude and error likelihood) were included as modulatory influences on connection strengths. For reciprocally connected models, modulatory parameters were estimated for both connections.

Model parameters for individual subjects were estimated using an Expectation Maximization algorithm (Friston, Harrison et al., 2003) under the SPM5 framework. In order to determine the optimal of two candidate models, the evidence for each model, approximated by either the Akaike Information Criterion (AIC) and Bayesian Information Criterion (BIC), is used to compute a Bayes Factor (Penny, Stephan et al., 2004). Generally, the BIC penalizes a model more than the AIC for added model complexity. In our analyses, we adopt the convention suggested by Penny et al., that evidence for one model over another exists if the AIC and BIC agree, and the minimum (if both AIC and BIC are greater than 1) or maximum (if less than 1) of the two is taken. If the AIC and BIC disagree about which model is optimal, the Bayes Factor for that model comparison is set to 1 (equal evidence for both models).

For each set of four models described above, a Bayes Factor was computed for each model against the three alternative models for each subject. An average Bayes Factor for each model comparison was computed across all subjects to determine the overall evidence for that model (Penny, Stephan et al., 2004; Smith, Stephan et al., 2006), computed as the *n*th-root of the product of the Bayes Factor for *N* individual subjects:

$$\text{Average } B_{ij} = \sqrt[n]{\prod_n B_{ij}}$$

where B is the Bayes Factor; n is the number of models, and i and j are models being compared. Group average Bayes Factors that exceeded a critical threshold (2.72, (Penny, Stephan et al., 2004)) for one model versus the other three candidate models were selected for further analysis, with the additional requirement that the model was selected as the optimal for a majority of subjects (13 or more). The computation of the average Bayes Factor is sensitive to outliers, and one subject was removed due to extreme values.

Results

Behavioral Results

A two-way analysis of variance showed no significant main effect of error likelihood, reward magnitude or interaction effect on reaction time (RT) for correct, go trials, ($f(3, 92) = 0.15$, $p > 0.05$) consistent with previous results (Brown & Braver, 2007), indicating that ACC activity related to error likelihood effects was not confounded with RT effect. Mean RTs were 730.06ms ($sd = 127.82$ ms) for High RM/Low EL trials, 737.61ms ($sd = 123.40$ ms) for High RM/High EL trials, 729.34ms ($sd = 126.62$ ms) for Low RM/High EL trials, and 725.73ms ($sd = 127.16$ ms) for Low RM/Low EL trials (for all go, correct trials). Additionally, observed error rates were 50.32% for the High EL conditions and 9.09% for Low EL conditions. These were significantly different from each other and consistent with target error rates.

The change signal delay (CSD) between cue onset and presentation of the delay signal (if any) was manipulated dynamically in order to maintain target error rates. A potential confound might exist if the change signal delay period were influenced by levels of reward magnitude or interactions of reward magnitude and error likelihood. A two-way analysis of variance was performed on each subject's final CSD in each RM/EL condition for go/correct trials. A main effect of error likelihood on CSD was observed as expected ($F(1, 92) = 176.48$, $p < 0.01$). However, there was neither a main effect of reward magnitude ($F(1, 92) = 0.04$, $p > 0.05$) nor was the interaction (reward magnitude \times error likelihood) significant ($F(1, 92) = 0.13$, $p > 0.05$).

The DOSPERT gambling subscale has a range of 4 (most risk averse) to 20 (most risk seeking). Subjects scored an average of 6.00 on the DOSPERT gambling subscale, with a standard deviation of 2.96. Sixteen subjects scored at or below the mean. Overall, the majority of subjects were strongly averse to gambling risk-taking, suggesting that a failure to find error likelihood effects would not be due to individual differences.

Error Likelihood

We found error likelihood effects consistent with previous findings (Brown & Braver, 2005; Brown & Braver, 2007) in dorsal ACC. Results for the main effect of error likelihood, (High EL – Low EL) showed significant activation at the cluster level (peak activation at MNI +4, -4, 52; 207 voxels; $t(23) = 4.52$, $p < 0.001$, corrected). Of note, this region was slightly more caudal than regions with similar effects found in previous studies (Brown & Braver, 2005; Brown & Braver, 2007). Tests for the main effect of reward magnitude (High RM – Low RM) failed to show significant effects in the dorsal ACC region. Pairwise comparisons of error likelihood for both levels of reward magnitude suggested a potential interaction of error likelihood and reward magnitude. The contrast for error likelihood in the low reward magnitude conditions (Low RM/High EL – Low RM/Low EL) yielded significant results at the cluster level ($t(23) = 5.59$, $p < 0.005$ corrected). For the low RM error likelihood contrast, a region of

interest was observed with a peak activation at +6, -4, 46 (MNI coordinates; figure 3A) and containing 186 voxels, while no other regions showed significant effect for this contrast. The full-brain error likelihood contrast for the high RM conditions (High RM/High EL – High RM/Low EL) yielded a qualitatively weaker result for this same region ($t(23)=1.505$, $p>0.05$ corrected). Since effects of error likelihood are more pronounced for the pairwise contrast (Low RM/High EL – Low RM/Low EL), we use the ROI with peak activation at +6, -4, 46 for our subsequent investigation of causal influences on ACC.

One potential concern with the task design is that ACC activity during the pre-response period is not easily discriminated from feedback signals due to correct or error responses. So might reward magnitude signals in ACC reflect greater *actual* reward instead of greater *anticipated* reward? To address this, we note that there was no main effect of reward magnitude in the identified ACC region, so it is not the case that ACC reflects differences in the value of the actual reward for that trial. Furthermore, we note that ACC activity for correct Go trials is highest for the Low RM/High EL condition, in which average reward is lowest, so again ACC activity does not correlate positively with the actual reward outcome of the trial.

Competition vs. modulation

As suggested by pairwise comparisons showing weaker error likelihood effects in the high reward magnitude condition, further analyses were conducted to investigate a potential interaction between reward magnitude and error likelihood in ACC. In order not to bias the results, the cluster in ACC identified for the main effect of Error Likelihood with peak of activity at +4, -4, +52 was used to conduct an ROI analysis. Within this region, we discovered a significant interaction ($t(23)=2.79$, $p<0.01$, uncorrected) for the interaction contrast (High RM/Low EL + Low RM/High EL) - (High RM/High EL + Low RM/Low EL), suggesting that error likelihood effects are smaller in the high RM condition than in the low RM condition. Figure 3C shows that increases in reward magnitude lead to an apparent saturation, such that increases in error likelihood cause a proportionally smaller increase in ACC activity.

Is ACC activity saturating? If so, this may suggest that the modulation hypothesis cannot be ruled out, since a similar pattern would be produced if there were such a ceiling effect. In order to rule out the possibility that ACC activity does indeed saturate, we tested whether activity in the same ROI (+4, -4, -52) for Change/Error Trials was greater than activity in the same region for Change/High RM/High EL/Correct Trials. There was a significant effect of Error within the region ($t(23)=4.67$, $p<0.001$, peak at MNI coordinates -2, +6, 44). This indicates that activity in ACC does not saturate since we would expect no significant difference in activation for error trials if activity were already at peak for Correct trials.

Furthermore, it may be seen from the pattern of activation that ACC activity is not merely proportional to average reward in the task (figure 3C). If ACC responded proportionally to increases in average reward, we would expect activity to be greatest in the High RM/Low EL condition, whereas if ACC activity is inversely proportional to average reward (i.e., low activity for conditions with high average reward), we would expect activity to be lowest in the High RM/Low EL condition. From figure 3C, we can see that the High RM/Low EL condition yields neither the greatest nor the least activation, indicating that ACC does not simply track reward magnitude or lack thereof.

Overall, these results are consistent with the competition hypothesis as discussed above in the introduction but cannot be accommodated by the modulation hypothesis, which incorrectly predicts that error likelihood effects in ACC should be greater in the High RM condition.

Origin of reward magnitude effects in ACC

Given that ACC activity appears to reflect competition between reward anticipation and error likelihood, the next question is where the anticipatory signals related to expected reward originate from. Previous studies (Seymour, O'Doherty et al., 2004; Knutson, Taylor et al., 2005), in addition to the present findings, suggest that ACC is part of a network of brain areas that process reward information. The presence of an interaction between reward magnitude and error likelihood in ACC suggests that regions which encode reward magnitude alone may be functionally connected to ACC, and that the reward magnitude of a task could be expected to contribute to cognitive control by influencing activity in ACC. Information about expected value appears to be encoded in a distributed fashion throughout the brain (Knutson, Taylor et al., 2005); one possibility is that ACC receives signals from one or more of these areas pertaining to reward magnitude. The alternative hypothesis is that ACC computes the predicted reward magnitude internally. These hypotheses are tested below. Our approach to this question is to first identify regions with effects of reward magnitude, then ascertain whether these regions exert a causal influence on ACC activity using dynamic causal modeling (DCM).

Regions showing main effects of reward magnitude

Tests for the main effect of reward magnitude (High RM-Low RM) showed no differences in the region of ACC which showed effects of error likelihood in the low RM condition (ROI +6, -4, +46, $t(23)=2.5095$, $p>0.05$ corrected). However, main effects of reward magnitude were observed in three regions (figure 4; locations of peak activation given): parahippocampal gyrus (PH; ROI -26, -18, -20, $t(23)=4.45$, $p=0.01$ corrected), ventral striatum (VStr; ROI 8, 2, -6, $t(23)=4.79$, $p<0.05$ corrected), and inferior frontal gyrus (IFG; ROI -54, 30, 6, $t(23)=3.75$, $p<0.05$ corrected). These areas showing main effects of reward magnitude are consistent with previous imaging studies using reinforcement learning tasks (Elliott, Friston et al., 2000; Knutson, Taylor et al., 2005; Rolls, McCabe et al., 2008).

DCM results

Table 1 shows sets of model comparisons testing for a causal relationship between ACC and the three areas showing effects of reward magnitude. Of the three sets of models tested, only two sets yielded positive evidence for a model in which the Bayes Factor for that model exceeded the threshold in comparison to all other three candidate models, and for which that model was preferred for a majority of the subjects. The results show a causal influence of ACC on PH and VStr, while there was insufficient evidence to determine a direction of causality between ACC and IFG. Average parameter estimates for the optimal model in each set are additionally given in table 1. Positive parameter estimates suggest excitatory connections from ACC to VStr but not *vice versa*. Similarly, a negative connection parameter estimate for the ACC→PH model suggests that ACC inhibits activity in PH. However, this may also indicate that higher levels of reward magnitude are associated with less inhibitory influence. The contrasts between levels for the reward magnitude parameter estimate, however, were not significant, so this remains an open question. The pattern of parameter estimates for modulatory influence of reward magnitude in this case suggests that for higher levels of reward magnitude, less inhibition occurs. However, a t-test of this relationship failed to yield significant results.

Since none of the areas showing main effects of reward magnitude were shown to have a causal influence on ACC, these results support the *computation* hypothesis, while the *integration* hypothesis cannot be accommodated by the present analyses. Nonetheless, we were concerned that reward magnitude signals might also originate from medial orbitofrontal cortex (MOFC), or that midbrain structures involved in reward processing might deliver reward magnitude information to ACC. We attempted to address this question with a second set of DCM analyses which included a region in medial orbitofrontal cortex (MOFC) which was observed for the contrast High RM-Low RM, but did not survive corrections for multiple comparisons. Previous

studies (Knutson, Taylor et al., 2005) have found expected value-related activity in a similar area. An ROI analysis in this area showed significant differences for the High RM-Low RM contrast ($t(1,23)=4.10$, $p<0.01$, uncorrected). Averaged activity time courses were extracted from individual subjects as described above, and four additional sets of DCM models were analyzed (MOFC to ACC, VStr, PH, and IFG). However, no evidence was found for a causal relationship between ACC and MOFC, indicating that ACC does not receive reward magnitude signals from MOFC. Similar analyses which attempt to localize midbrain structures underlying reward processing (e.g., substantia nigra) likewise yielded no evidence for a causal role on ACC activity. These results are discussed in more detail in the supplementary material. Given the above, we conclude that ACC computes the reward magnitude of an action internally.

DISCUSSION

The present findings of distinct anticipatory reward and error effects in human ACC provide a stronger bridge between the human and monkey findings on performance monitoring. On the one hand, earlier human studies mostly emphasized error and conflict detection (Hohnsbein, Falkenstein et al., 1989; Gehring, Coles et al., 1990; Carter, Braver et al., 1998; Botvinick, Nystrom et al., 1999; MacDonald, Cohen et al., 2000; Holroyd & Coles, 2002; Yeung, Cohen et al., 2004). More recent human studies have begun to emphasize anticipatory (Brown & Braver, 2005; Sohn, Albert et al., 2007; Aarts, Roelofs et al., 2008) and regulatory (Roelofs, van Turenout et al., 2006; Behrens, Woolrich et al., 2007) functions of ACC. On the other hand, monkey neurophysiology studies including our own (Ito, Stuphorn et al., 2003) have uniformly shown that ACC provides distinct signals related to anticipated and actual reward (Matsumoto, Suzuki et al., 2003) (Amador, Schlag-Rey et al., 2000; Procyk, Tanaka et al., 2000; Shidara & Richmond, 2002; Amiez, Joseph et al., 2005; Amiez, Joseph et al., 2006; Kennerley, Dahmubed et al., 2009). Our findings as a whole are consistent with a common function of both human and monkey ACC, namely the evaluation of the relative risks versus rewards of an anticipated action (Kennerley, Walton et al., 2006; Crosson, Walton et al., 2009; Kennerley, Dahmubed et al., 2009; Kouneiher, Charron et al., 2009).

Our analysis showed that in agreement with previous studies (Brown & Braver, 2005; Brown & Braver, 2007), a main effect of error likelihood was found in ACC. While this finding provides additional support for the error likelihood hypothesis, we note that the locus of activation is somewhat more dorsal and posterior to areas of ACC which have previously been observed to show effects of error likelihood, and extends into extracingle regions such as SMA. One reason for this may be the differences in experimental manipulations in the present study; previous studies manipulated error likelihood (Brown & Braver, 2005) as well as expected risk (Brown & Braver, 2007) without controlling for reward magnitude. A recent study (Kouneiher, Charron et al., 2009) has suggested that more anterior aspects of mPFC code the longer-term cost and value of behavior, but the more posterior mPFC codes the more immediate reward and motivational factors of an action. Our results are consistent with those findings. A more recent study (Fujiwara, Tobler et al., 2009) found activation in dorsal ACC (Brodmann area 32) which integrated both gains and losses, similar to the present task. In any case, the area identified is within the region identified by (Bush, Luu et al., 2000) as being part of the cognitive division of ACC, and extending into the posterior rostral cingulate zone (RCZp) (Picard & Strick, 1996; Fan, Hof et al., 2008), consistent with other ACC areas involved in cognitive and motor function (Beckmann, Johansen-Berg et al., 2009).

A previous study has reported a failure to replicate error likelihood effects (Nieuwenhuis, Schweizer et al., 2007). How can the present findings be reconciled with the apparent failure to replicate? In one of our follow-up studies, we measured individual differences in risk tolerance and found that error likelihood effects were strongly present in risk-averse individuals but virtually absent in risk-tolerant individuals (Brown & Braver, 2007). This suggests the

possibility that our sample may have been more risk avoidant, and this was confirmed by a gambling likelihood self-report (Weber, Blais et al., 2002), consistent with our prior findings (Brown & Braver, 2007). The same study that questioned the replicability of error likelihood effects also raised an important issue, which is whether error likelihood effects are predicted by the paired cue or whether they are confounded with the difficulty of performing the task itself at the time of response. This remains an important open question which the present study does not address: as shown in figure 1, the interval between trial onset and the limit for responses, 2000msec, is too brief to effectively differentiate between the two intervals. Other studies (e.g., Aarts et al., 2008) more directly address this question, and appear to show that error likelihood-type effects are more directly related to the performance of a task rather than to predictive cues.

The present study was motivated by the question of whether and how anticipated reward might interact with error likelihood effects in ACC. Previous findings showed underadditive effects of error likelihood and expected risk on ACC activity. Specifically, as error likelihood and consequence severity continue to increase, the ACC response appears to reach a plateau (Brown & Braver, 2007). This finding differed from the predictions of the Error Likelihood computational model, which did not predict a plateau but rather a linear relationship between expected risk and ACC activity (Brown & Braver, 2007). ACC is implicated in the processing of rewarding as well as aversive events (Ito, Stuphorn et al., 2003; Amiez, Joseph et al., 2006; Berns, Capra et al., 2008), and appears to participate in a network of brain areas underlying reinforcement learning and eliciting behaviors necessary for avoiding undesirable outcomes (Magno, Foxe et al., 2006; Brown & Braver, 2007).

One possible explanation for the discrepancy between computational model predictions and observed results is that, as a part of a distributed reinforcement learning network, ACC is activated by the likelihood and potential severity of an error, and the magnitude of this effect could have been modulated by the predicted reward magnitude associated with a condition. In other words, if an action is not likely to lead to significant reward, why should the ACC respond to error likelihood in that condition in the first place? This is the *modulation* hypothesis referred to above. Despite the apparent plausibility, the modulation hypothesis could not account for the results of the present study.

While the ROI analyses for the main effect of reward magnitude failed to achieve significance, tests for the interaction (reward magnitude \times error likelihood) yielded significant results. Furthermore, pairwise comparisons show that the effect of reward magnitude is significant only in the low error likelihood condition, while absent for high error likelihood conditions. This pattern of activity shows that reward magnitude has an under-additive effect on ACC activity and is consistent with the *competition* hypothesis referred to above. Thus, there appears to be a tradeoff between reward and punishment sensitivity, such that greater activity in response to anticipated reward allows less dynamic range for responses to anticipated punishment in the form of error likelihood and potential consequence magnitude (Croxson, Walton et al., 2009). In that case, individuals who are more sensitive to anticipated reward might show reduced error likelihood and error consequence magnitude sensitivity in ACC. This seems to be the case for substance-dependent individuals in particular (Yechiam, Busemeyer et al., 2005).

Further evidence for the competition hypothesis is provided by our analysis of causal relationships between regions showing effects of reward magnitude in the current study and ACC. The results of the DCM analysis are consistent with the hypothesis that ACC computes an internal estimate of the reward value of a given action, and the results provided no evidence that ACC integrates signals coding reward magnitude computed elsewhere in the brain. To the contrary, for two regions (PH and VStr), the optimal DCM suggests that ACC exerts a causal

influence on these regions, rather than *vice versa*. Although there appears to be a causal relationship in these two instances, this does not mean that ACC has direct anatomical connectivity to PH and VStr; rather, it merely implies that ACC is only *functionally* connected to these areas. Functional connectivity may imply either direct projections from one brain region to another, or that there are intermediate areas between two causally linked regions. The lack of input to ACC containing direct reward magnitude signals suggests that, in addition to learning representations of error likelihood, ACC may also learn representations of predicted reward magnitude, and that these representations may compete within ACC for limited neural representation.

It may be somewhat surprising that our DCM analyses suggest that ACC appears not to be the target of functional connections from regions encoding levels of reward magnitude, especially considering the large number of regions in the brain which are known to be connected to ACC (Beckmann, Johansen-Berg et al., 2009). One possible explanation for our findings is that our DCM analyses modeled activity in ACC from the onset of each trial; as a locus of performance monitoring, it may be that ACC proactively exerts control, or signals the need for increased control, to other brain areas prior to response generation and feedback. It may be the case that if we instead modeled events in the task based on response or outcome timing, we may find the opposite pattern of causal interactions. More work is needed to address the question of how ACC and extracingulate areas of the brain interact at various task periods.

ACC is thought to play a role in cognitive control and executive function by signaling the need for increased control, while other brain areas, especially dorsolateral prefrontal cortex (DLPFC), are responsible for implementing control (MacDonald, Cohen et al., 2000; Botvinick, Braver et al., 2001). In the present study, DLPFC was not considered in our analyses, although previous work suggests that ACC should show a strong causal relationship with it. Rather, we focused instead on potential causal relationships between ACC and regions which were observed to respond to information regarding reward, which did not include DLPFC.

A key goal of computational modeling is to account for observed empirical results and to generate additional, testable predictions. The present study was motivated in part by predictions of a computational model – the Error Likelihood model – that suggested an approximately linear relationship between ACC activity and expected risk. Expected risk effects observed in human participants showed an under-additive influence of anticipated error consequence magnitude and error likelihood in ACC, suggesting that an additional factor such as reward magnitude is involved in the ACC signal beyond that predicted by the error likelihood model. In this paper, we investigated two alternative hypotheses about the effect of reward magnitude on ACC activity, namely the modulation and competition hypotheses. Our results support the competition hypothesis, i.e. that predicted reward magnitude and error likelihood both activate ACC, and that increased activity in response to one decreases sensitivity to the other. While the competition hypothesis is supported by the current evidence, the mechanism by which such competition occurs is not yet clear. One possibility is that signals encoding reward magnitude from regions outside the ACC may drive ACC activity toward saturation. This seems unlikely since, for trials in which an error was committed, the percent signal change in the same ROI was greater than for non-error trials, indicating that the observed pattern of effects was not the result of saturation. Another possibility is areas projecting to ACC encode components of reward magnitude (e.g., reward feedback for correct vs. incorrect trials,) and that ACC uses these components to learn representations of reward magnitude which compete with similarly learned representations of error likelihood. The present study failed to find conclusive evidence of predicted reward magnitude signals which could influence ACC activity, lending support to the hypothesis that ACC computes its own representation of predicted reward magnitude, versus an alternative hypothesis that ACC integrates external signals. While the Error Likelihood model contained no mechanism by which varying levels of reward magnitude could

influence ACC activity), the present study suggests that future models of ACC should incorporate such a mechanism.

Supplementary Material

Refer to Web version on PubMed Central for supplementary material.

Acknowledgments

Supported in part by AFOSR FA9550-07-1-0454, A NARSAD Young Investigator Award, the Sidney R. Baer, Jr. Foundation, R03 DA023462, R01 DA026457, and the Indiana METACyt Initiative of Indiana University, funded in part through a major grant from the Lilly Endowment, Inc. The authors would like to thank Derek Nee for helpful comments in the preparation of this manuscript and E. Dinh for help with data collection.

References

- Aarts E, Roelofs A, et al. Anticipatory activity in anterior cingulate cortex can be independent of conflict and error likelihood. *J Neurosci* 2008;28(18):4671–8. [PubMed: 18448644]
- Amador N, Schlag-Rey M, et al. Reward-predicting and reward-detecting neuronal activity in the primate supplementary eye field. *J Neurophysiol* 2000;84(4):2166–70. [PubMed: 11024104]
- Amiez C, Joseph JP, et al. Anterior cingulate error-related activity is modulated by predicted reward. *Eur J Neurosci* 2005;21(12):3447–52. [PubMed: 16026482]
- Amiez C, Joseph JP, et al. Reward encoding in the monkey anterior cingulate cortex. *Cereb Cortex* 2006;16(7):1040–55. [PubMed: 16207931]
- Beckmann M, Johansen-Berg H, et al. Connectivity-based parcellation of human cingulate cortex and its relation to functional specialization. *J Neurosci* 2009;29(4):1175–90. [PubMed: 19176826]
- Behrens TE, Woolrich MW, et al. Learning the value of information in an uncertain world. *Nat Neurosci* 2007;10(9):1214–21. [PubMed: 17676057]
- Berns GS, Capra CM, et al. Nonlinear neurobiological probability weighting functions for aversive outcomes. *Neuroimage* 2008;39(4):2047–57. [PubMed: 18060809]
- Botvinick MM, Braver TS, et al. Conflict monitoring and cognitive control. *Psychological Review* 2001;108:624–652. [PubMed: 11488380]
- Botvinick MM, Nystrom L, et al. Conflict monitoring versus selection-for-action in anterior cingulate cortex. *Nature* 1999;402(6758):179–181. [PubMed: 10647008]
- Brett, M.; Anton, J-L., et al. Region of interest analysis using an SPM toolbox 8th International Conference on Functional Mapping of the Human Brain; Sendai, Japan. 2002.
- Brown J, Braver TS. Risk prediction and aversion by anterior cingulate cortex. *Cog Aff Behav Neurosci* 2007;7(4):266–277.
- Brown JW, Braver TS. Learned Predictions of Error Likelihood in the Anterior Cingulate Cortex. *Science* 2005;307(5712):1118–1121. [PubMed: 15718473]
- Brown JW, Braver TS. A computational model of risk, conflict, and individual difference effects in the anterior cingulate cortex. *Brain Res* 2008;1202:99–108. [PubMed: 17707352]
- Bush G, Luu P, et al. Cognitive and emotional influences in anterior cingulate cortex. *Trends Cogn Sci* 2000;4(6):215–222. [PubMed: 10827444]
- Carter CS, Braver TS, et al. Anterior cingulate cortex, error detection, and the online monitoring of performance. *Science* 1998;280:747–749. [PubMed: 9563953]
- Crosson PL, Walton ME, et al. Effort-based cost-benefit valuation and the human brain. *J Neurosci* 2009;29(14):4531–41. [PubMed: 19357278]
- Elliott R, Friston KJ, et al. Dissociable neural responses in human reward systems. *Journal of Neuroscience* 2000;20:6159–6165. [PubMed: 10934265]
- Fan J, Hof PR, et al. The functional integration of the anterior cingulate cortex during conflict processing. *Cereb Cortex* 2008;18(4):796–805. [PubMed: 17652463]
- Friston KJ, Harrison L, et al. Dynamic causal modelling. *Neuroimage* 2003;19(4):1273–302. [PubMed: 12948688]

- Fujiwara J, Tobler PN, et al. Segregated and integrated coding of reward and punishment in the cingulate cortex. *J Neurophysiol* 2009;101(6):3284–93. [PubMed: 19339460]
- Gehring WJ, Coles MGH, et al. The error-related negativity: An event-related potential accompanying errors. *Psychophysiology* 1990;27:S34.
- Gemba H, Sasaki K, et al. 'Error' potentials in limbic cortex (anterior cingulate area 24) of monkeys during motor learning. *Neurosci Lett* 1986;70(2):223–7. [PubMed: 3774229]
- Hohnsbein J, Falkenstein M, et al. Error processing in visual and auditory choice reaction tasks. *Journal of Psychophysiology* 1989;3:32.
- Holroyd CB, Coles MG. The neural basis of human error processing: Reinforcement learning, dopamine, and the error-related negativity. *Psych Rev* 2002;109(4):679–709.
- Ito S, Stuphorn V, et al. Performance Monitoring by Anterior Cingulate Cortex During Saccade Countermanding. *Science* 2003;302:120–122. [PubMed: 14526085]
- Kennerley SW, Dahmubed AF, et al. Neurons in the frontal lobe encode the value of multiple decision variables. *J Cogn Neurosci* 2009;21(6):1162–78. [PubMed: 18752411]
- Kennerley SW, Walton ME, et al. Optimal decision making and the anterior cingulate cortex. *Nat Neurosci* 2006;9(7):940–7. [PubMed: 16783368]
- Knutson B, Taylor J, et al. Distributed neural representation of expected value. *J Neurosci* 2005;25(19):4806–12. [PubMed: 15888656]
- Kouneiher F, Charron S, et al. Motivation and cognitive control in the human prefrontal cortex. *Nat Neurosci* 2009;12(7):939–45. [PubMed: 19503087]
- MacDonald AW, Cohen JD, et al. Dissociating the role of the dorsolateral prefrontal cortex and anterior cingulate cortex in cognitive control. *Science* 2000;288:1835–1838. [PubMed: 10846167]
- Magno E, Foxe JJ, et al. The anterior cingulate and error avoidance. *J Neurosci* 2006;26(18):4769–73. [PubMed: 16672649]
- Matsumoto K, Suzuki W, et al. Neuronal correlates of goal-based motor selection in the prefrontal cortex. *Science* 2003;301(5630):229–32. [PubMed: 12855813]
- Nieuwenhuis S, Schweizer T, et al. Error-likelihood prediction in the medial frontal cortex: A critical evaluation. *Cereb Cortex* 2007;17:1570–81. [PubMed: 16956979]
- Penny WD, Stephan KE, et al. Comparing dynamic causal models. *Neuroimage* 2004;22(3):1157–72. [PubMed: 15219588]
- Picard N, Strick PL. Motor areas of the medial wall: a review of their location and functional activation. *Cereb Cortex* 1996;6(3):342–53. [PubMed: 8670662]
- Procyk E, Tanaka YL, et al. Anterior ingulate activity during routine and non-routine sequential behaviors in macaques. *Nature Neuroscience* 2000;3(5):502–508.
- Riba J, Kramer UM, et al. Dopamine agonist increases risk taking but blunts reward-related brain activity. *PLoS ONE* 2008;3(6):e2479. [PubMed: 18575579]
- Roelofs A, van Turenout M, et al. Anterior cingulate cortex activity can be independent of response conflict in Stroop-like tasks. *Proc Natl Acad Sci U S A* 2006;103(37):13884–9. [PubMed: 16954195]
- Rolls ET, McCabe C, et al. Expected value, reward outcome, and temporal difference error representations in a probabilistic decision task. *Cereb Cortex* 2008;18(3):652–63. [PubMed: 17586603]
- Seymour B, O'Doherty J, et al. Temporal difference models describe higher-order learning in humans. *Nature* 2004;429:664–667. [PubMed: 15190354]
- Shidara M, Richmond BJ. Anterior cingulate: single neuronal signals related to degree of reward expectancy. *Science* 2002;296(5573):1709–11. [PubMed: 12040201]
- Smith AP, Stephan KE, et al. Task and content modulate amygdala-hippocampal connectivity in emotional retrieval. *Neuron* 2006;49(4):631–8. [PubMed: 16476670]
- Sohn MH, Albert MV, et al. Anticipation of conflict monitoring in the anterior cingulate cortex and the prefrontal cortex. *Proc Natl Acad Sci U S A* 2007;104(25):10330–4. [PubMed: 17563353]
- Weber E, Blais A, et al. A Domain-specific Risk-attitude Scale: Measuring Risk Perceptions and Risk Behaviors. *J Behav Decision Making* 2002;15:263–290.
- Yechiam E, Busemeyer JR, et al. Using cognitive models to map relations between neuropsychological disorders and human decision-making deficits. *Psychol Sci* 2005;16(12):973–8. [PubMed: 16313662]

Yeung N, Cohen JD, et al. The neural basis of error detection: conflict monitoring and the error-related negativity. *Psychol Rev* 2004;111(4):931–59. [PubMed: 15482068]

INCENTIVE CHANGE-SIGNAL TASK

LOW Reward Magnitude
(0000 pts error/1000 pts correct)

HIGH Reward Magnitude
(1000 pts error/2000 pts correct)

HIGH Error Likelihood
(50% Errors)



\$0.00 error/\$0.01 correct
\$0.005 Avg. Rew.



\$0.01 error/\$0.02 correct
\$0.015 Avg. Rew.

LOW Error Likelihood
(5% Errors)



\$0.00 error/\$0.01 correct
\$0.0095 Avg. Rew.



\$0.01 error/\$0.02 correct
\$0.0195 Avg. Rew.

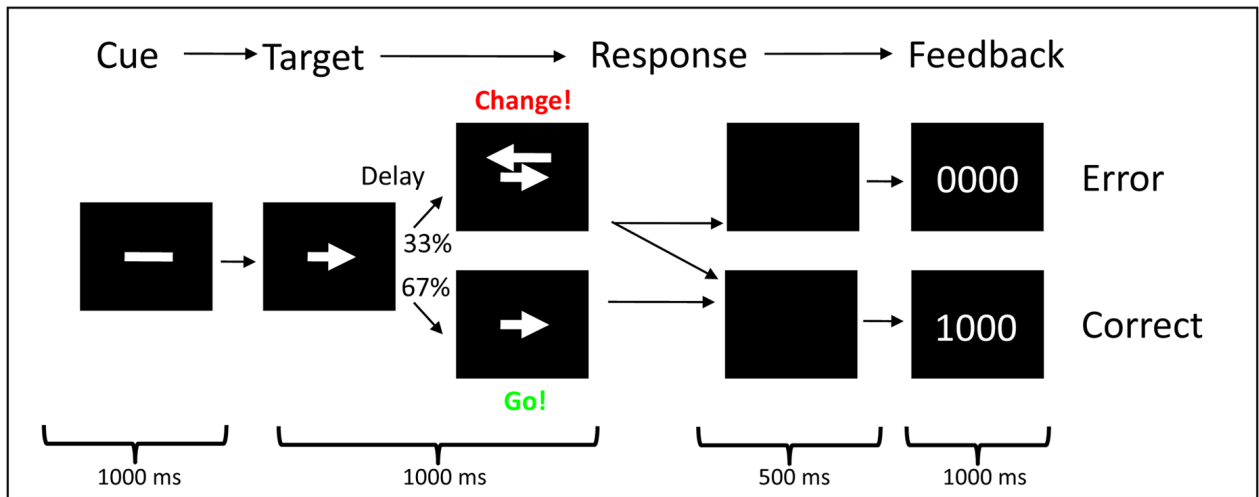


Figure 1. Incentive Change Signal Task

A modified version of the Change Signal Task presented in Brown & Braver, 2007. Subjects earn \$0.01 for correct trial and \$0.00 for incorrect trials in the low reward magnitude condition, and \$0.02 for correct trials and \$0.01 for incorrect trials in the high reward magnitude condition. Error rates were controlled at 50% for the high error likelihood condition and 5% for the low error likelihood condition.




Figure 2. Dynamic Causal Modeling

DCM models were created to examine the causal structure (if any) between ACC and extracingulate regions. Information between two areas may flow in a single direction (Models 1 and 3), or information may originally be available to one region, and the subsequent activity of both regions is influenced by reciprocal connectivity (Models 2 & 4). Furthermore, effective connectivity between the two regions might be modulated by one or more task variables.

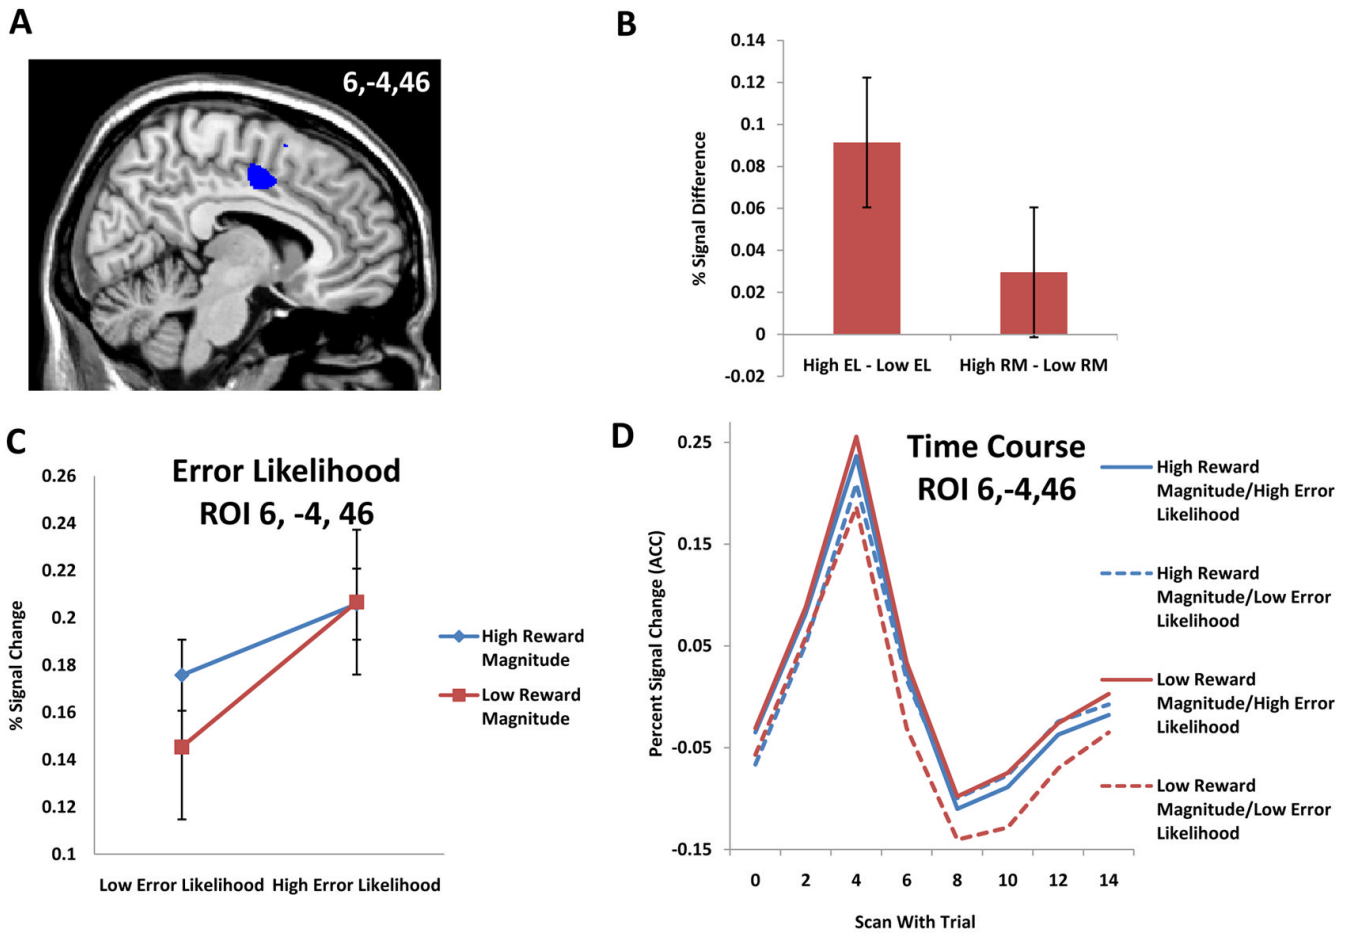


Figure 3. Reward Magnitude and Error Likelihood

A) Error likelihood effects were observed in dorsal ACC. No other areas showed significant activation for error likelihood. **B)** Tests for main effects of error likelihood and reward magnitude within this region confirm error likelihood effects, but show no significant difference in activity for reward magnitude. Error bars reflect standard error. **C)** ACC responds to both reward magnitude and error likelihood. However, these effects appear to saturate for high levels of RM and EL. Error bars reflect standard error. All analyses were for correct, go trials only.




Figure 4. Main Effect of Reward Magnitude

Three clusters (left panels) showed significant activation for reward magnitude (High RM - Low RM): A) ventral striatum, B) parahippocampal gyrus, C) inferior frontal gyrus.

Table 1

Dynamic Causal Modeling

Of the seven sets of models initially tested, four yielded positive evidence of a causal relationship between two regions. Group average Bayes Factors for the comparison between the optimal model in each set and the three alternative models are given in the left four columns. Parameter estimates for the optimal model, averaged across all subjects are given in the middle, and contrasts between levels of modulation are shown on the right. An asterisk indicates significance at the .05 level. Complete results are included in supplementary material.

Selected Model	Alternative Models				Input Wt.	Conn. Wt.	Modulators				Contrast of Modulators	
	→ACC→VStr	ACC↔VStr	ACC→VStr←	ACC↔VStr←			High RM	Low RM	High EL	Low EL	HRM-LRM t(1,23)	HEL-LEL t(1,23)
→ACC→VStr	69.92	60.19	727.94	ACC↔VStr←	0.114**	0.249**	-0.0833	-0.0835	-0.038	-0.129	-0.006	2.08*
Avg. BF												
→ACC→PH	78.23	9.27	390.47	ACC→PH←	0.11**	-0.051	-0.0725	-0.0434	-0.008	-0.108	-0.789	2.06†
Avg. BF												
→ACC→IFG	67.56	2.35	64.27	ACC↔IFG←	0.119**	-0.197*	-0.011	-0.006	0.039	-0.056	-0.102	1.36
Avg. BF												

** p<0.01

* p<0.05

† p<0.1