

Risk prediction and aversion by anterior cingulate cortex

JOSHUA W. BROWN

Indiana University, Bloomington, Indiana

AND

TODD S. BRAVER

Washington University, St. Louis, Missouri

The recently proposed *error-likelihood* hypothesis suggests that anterior cingulate cortex (ACC) and surrounding areas will become active in proportion to the perceived likelihood of an error. The hypothesis was originally derived from a computational model prediction. The same computational model now makes a further prediction that ACC will be sensitive not only to predicted error likelihood, but also to the predicted magnitude of the consequences, should an error occur. The product of error likelihood and predicted error consequence magnitude collectively defines the general “expected risk” of a given behavior in a manner analogous but orthogonal to subjective expected utility theory. New fMRI results from an incentive change signal task now replicate the error-likelihood effect, validate the further predictions of the computational model, and suggest why some segments of the population may fail to show an error-likelihood effect. In particular, error-likelihood effects and expected risk effects in general indicate greater sensitivity to earlier predictors of errors and are seen in risk-averse but not risk-tolerant individuals. Taken together, the results are consistent with an expected risk model of ACC and suggest that ACC may generally contribute to cognitive control by recruiting brain activity to avoid risk.

The anterior cingulate cortex (ACC) is critically involved in performance monitoring and cognitive control (Blakemore, Rees, & Frith, 1998; Botvinick, Nystrom, Fissel, Carter, & Cohen, 1999; Braver, Barch, Gray, Molfese, & Snyder, 2001; Carter et al., 1998; Carter, MacDonald, Ross, & Stenger, 2001; Gehring & Knight, 2000; Kerns et al., 2004; Liddle et al., 1992; MacDonald, Cohen, Stenger, & Carter, 2000; Menon, Adelman, White, Glover, & Reiss, 2001; Nordahl et al., 2001; Scheffers & Coles, 2000; Ullsperger & von Cramon, 2001; van Veen, Cohen, Botvinick, Stenger, & Carter, 2001). Performance monitoring is essential to theories of executive control in which a central executive or supervisory attentional system takes control when it detects that automated processes (or schema) may lead to undesirable outcomes (Norman & Shallice, 1986). The ACC was first highlighted as an area that responds to errors, stemming from electrophysiological studies in both monkeys (Gemba, Sasaki, & Brooks, 1986) and, later, humans (Gehring, Coles, Meyer, & Donchin, 1990; Hohnsbein, Falkenstein, & Hoorman, 1989). As of the late 1990s, one influential model of performance monitoring proposes that ACC detects response conflict. In this account, when two mutually incompatible response processes are active, the ACC detects the state of conflict and drives control processes to resolve the internal conflict and facilitate appropriate behavior (Botvinick et al., 1999; Carter et al., 1998; MacDonald et al., 2000). Doing

so allows individuals to suppress prepotent, automatic responses and instead generate more appropriate responses to achieve current goals. Subsequent monkey studies of ACC and the surrounding medial frontal cortex (MFC) have revealed error, reward, and conflict effects (Amiez, Joseph, & Procyk, 2005, 2006; Ito, Stuphorn, Brown, & Schall, 2003; Olson & Gettner, 2002; Shidara & Richmond, 2002; Stuphorn, Taylor, & Schall, 2000).

The ACC and other MFC regions have recently been found to also be important in decision making, as highlighted by a number of neuroimaging and lesion studies. A prominent recent study of framing effects (de Martino, Kumaran, Seymour, & Dolan, 2006) has shown greater ACC activity when participants make decisions that are framed as being more likely to result in loss. Another recent study of certainty equivalent choices (Paulus & Frank, 2006) has shown that greater ACC activity uniquely predicts more normative decision-making behavior and less risk seeking. Similarly, others have found that ACC is particularly active when individuals avoid errors (Frank, Woroch, & Curran, 2005; Hewig et al., 2007; Johansen & Fields, 2004; Kim, Shimojo, & O’Doherty, 2006; Magno, Foxe, Molholm, Robertson, & Garavan, 2006; Shima & Tanji, 1998). When the ACC is lesioned, animals have difficulty using preceding trial history to guide decision making in current trials (Kennerley, Walton, Behrens, Buckley, & Rushworth, 2006) and are unwilling to invest

as much effort to gain reward (Rudebeck, Walton, Smyth, Bannerman, & Rushworth, 2006).

Despite its successes, researchers have increasingly realized that the response conflict model also has limitations. First, the model does not provide a clear account of how the ACC contributes to decision making (Kennerley et al., 2006; Rudebeck et al., 2006; Walton, Devlin, & Rushworth, 2004). Second, the response conflict model does not provide a mechanism for context-sensitive learning within the ACC, which is critical for adaptive behavior. For example, pressing the left index finger and the right index finger simultaneously might constitute an error in one task context, such as two-alternative forced choice, but not in another (such as pulling both brake levers simultaneously to stop a bicycle). How does the ACC learn to contextualize the response patterns that signify conflict, so that control signals are implemented in an appropriate manner to avoid impending errors?

We set about to answer this question by proposing an alternative to the conflict hypothesis, namely the *error-likelihood* hypothesis (Brown & Braver, 2005). According to this proposal, ACC might not detect conflict or errors per se, but more generally represent a prediction of error likelihood. Put simply, we suggested that the ACC response to a given task condition will be proportional to the perceived likelihood of an error in that condition (Brown & Braver, 2005). Since errors leading to adverse outcomes are committed more frequently in certain circumstances, the ACC uses reinforcement learning signals to learn to respond more strongly to the conditions in which errors are more likely. Because conflict conditions generally lead to a higher probability of an error, conflict effects may be reinterpreted as a special case of error-likelihood prediction. We tested the error-likelihood hypothesis with a new task we developed, the *change signal* task (Brown & Braver, 2005), which is a variant of the stop signal paradigm (Husain, Parton, Hodgson, Mort, & Rees, 2003; Logan, Cowan, & Davis, 1984).

Initially, we implemented two competing computational models of performance monitoring: the conflict model and the error-likelihood model. We found that the computational models made strong and divergent predictions: Both accounted for conflict effects, but only the error-likelihood model predicted that ACC activity would be proportional to the likelihood of an error in the absence of conflict or actual error commission. We then tested the task in healthy human volunteers with fMRI. The results were striking: ACC showed a significant effect of anticipated error likelihood, even when controlling for both response conflict and error commission (Brown & Braver, 2005). This finding suggested a reinterpretation of the literature on response conflict effects in terms of learned error-likelihood prediction.

Nonetheless, our previous study (Brown & Braver, 2005) was limited in that there were no significant consequences for the participants when they made mistakes. In the real world, mistakes generally have consequences that incur an actual cost to the behaving agent. In this case, the risk associated with a behavior can be cast in terms of not only the likelihood of an error or mistake, but also the po-

tential consequences, should an error occur. According to *expected utility* theory (Bernoulli, 1954; von Neumann & Morgenstern, 1944), the utility of a choice is a function of both the likelihood and value of each potential outcome. *Normative decision* theory specifies that the utility of a potential outcome should be evaluated as the multiplicative product of the likelihood and value of the potential outcome. Nonetheless, human decision makers generally deviate from the predictions of normative theory by showing risk aversion, in that they are less willing to choose an option as the uncertainty of the expected outcome increases (Bernoulli, 1954; Kahneman & Tversky, 1979).

In this article, we explore the hypothesis that ACC activity mediates risk perception and aversion, using a combination of computational modeling and fMRI in humans. As a starting point, we define the quantity *expected risk* as the product of the error likelihood and the error consequence magnitude. Traditionally, risk is associated with potential loss. Risk increases with increasing likelihood and size of the potential loss, in comparison with a preferred outcome. Studies of the Iowa gambling task show that healthy humans generally avoid risky alternatives that lead on average to loss, whether losses are due to low probability, large penalty events or higher probability, smaller penalty events (Bechara, Damasio, Damasio, & Anderson, 1994). Some more recent studies of individual differences in risk-taking behavior have explored manipulations of both the probability of loss and the magnitude of the potential loss together (Hunt, Hopko, Bare, Lejuez, & Robinson, 2005; Lejuez, Aklin, Zvolensky, & Pedulla, 2003; Lejuez et al., 2002), but these methods do not control for the effects of likelihood and magnitude of loss separately as we do in the present article. For the present purposes, we explicitly define expected risk as the probability of a mistake multiplied by the consequences of the mistake. This definition is motivated by our computational model results, described below. For example, if mistakes are committed 50% of the time in a particular condition, and if a mistake entails a loss of \$10 in comparison with the correct response, then the expected risk is \$5. Of note, the expected risk is orthogonal to the expected value. For example, consider a condition in which correct responses are made 100% of the time, and the reward for correct responses is \$10. Then, the expected value is \$10, and the expected risk is \$0. In contrast, consider a condition in which correct responses and mistakes are each 50% likely, and correct responses yield \$15, but mistakes yield only \$5. In this case, the expected value is the same \$10 as that in the first scenario ($\$15 \times 0.5 + \5×0.5), but the expected risk is $(\$15 - \$5) \times 0.5 = \$5$. Thus, expected value and expected risk are conceptually similar, but can be manipulated independently of each other. If ACC indeed signals a prediction of the risk associated with a given behavior, then the concept of expected risk may provide a clearer conceptual and theoretical framework for understanding the role of ACC in decision making. This aspect of ACC function is only just beginning to be understood (Paulus & Frank, 2006; Walton et al., 2004).

Individuals differ in a variety of dispositional characteristics, including risk aversion. If ACC activity reflects risk

perception and aversion, then ACC activity may also be sensitive to individual differences in risk aversion. Those individuals who are most sensitive to expected risk may show stronger ACC effects. Specifically, we predicted that those who are more dispositionally risk averse will show stronger ACC effects of error likelihood, error magnitude, and response conflict. However, it is also possible that response conflict effects may show a different relationship to risk aversion. Specifically, strong effects of response conflict in ACC appear to represent a suboptimal state of proactive cognitive control (Botvinick et al., 1999; Carter et al., 2000). If error-likelihood and -magnitude effects serve to facilitate proactive cognitive control (i.e., even under conditions with no response conflict), then individuals that are high in risk aversion may actually show reduced response conflict effects. We explicitly tested these competing hypotheses.

Several instruments have been developed to measure personality traits that are related to sensation seeking, behavioral inhibition, and risk perception and aversion. The BIS/BAS scales (Carver & White, 1994) measure behavioral inhibition versus activation as a measure of sensitivity to impending punishment or reward, respectively. Zuckerman's sensation-seeking scale provides a measure of individual preferences for reward (Zuckerman, 1994). Previous findings have shown that ACC activity is sensitive to individual differences in personality traits, such as BIS/BAS (Gray & Braver, 2002), and in behavioral factors that are related to risk (Paulus & Frank, 2006).

Risk aversion is perhaps most easily studied in terms of financial decision making, where alternatives, risks, and payouts can be quantified in units of currency. In this article, we add an incentive component to our change signal task (Brown & Braver, 2005) so that mistakes lead to a reduced monetary reward for responding in the trial. Thus, for this study, we are most interested in aversion to financial loss. Existing survey instruments generally measure a range of behavioral preferences (Carver & White, 1994; Gray & Braver, 2002; Zuckerman, 1994), not just aversion to financial loss. A more recent instrument, the domain-specific risk taking inventory (DOSPERT) (Weber, Blais, & Betz, 2002) focuses more directly on distinct risk-taking behavior and attitudes within different content domains. In particular, the DOSPERT subdivides risky behavior into six categories: social, recreational, gambling, investment, ethical, and health/safety. In our human studies, we look specifically at individual differences in gambling aversion as a measure of risk aversion most relevant to our task. If ACC shows greater activity in individuals who are less likely to engage in risk-taking behavior, then ACC may provide not only an index of risk perception, but also of risk aversion.

A second motivation for examining the role of individual differences in ACC activity during the computation of expected risk is that there may be a high degree of variability in ACC activity across different samples. If so, it is important to assess this variability in order to determine whether it can be understood in terms of meaningful individual difference variables. This is particularly relevant, because our original report of error-likelihood effects (Brown & Braver, 2005)

has generated significant controversy. In particular, one study has called into question whether the error-likelihood effects can be replicated (Nieuwenhuis, Schweizer, Mars, Botvinick, & Hajcak, 2007). One potential explanation of the failure to replicate error-likelihood effects is the presence of significant individual difference effects that have an impact on the ACC.

In what follows, we first derive rigorous predictions of risk perception effects in ACC from our existing computational model, then, we test these predictions with fMRI and measures of individual differences in human participants.

Computational Simulations:

Predictions of Risk Perception Effects in ACC

In our original article reporting error-likelihood effects in ACC (Brown & Braver, 2005), we developed a computational model that predicted error-likelihood effects *a priori*. The predictions were subsequently confirmed with human fMRI. However, we did not examine whether increasing the severity of error consequences would have an impact on the development of representation in ACC. To address this question, we began with the model as published and asked whether systematic manipulations of the magnitude of the error signal as well as the likelihood of an error would affect ACC activity. Because there were no parameter or architectural changes, the architecture of the model was identical to that implemented in our previous work (Brown & Braver, 2005). As such, we direct the reader to the supplementary material of Brown and Braver (2005) for a detailed description of computational methods. Briefly, in the error-likelihood model ACC, error signals train a random sampling of model ACC cells to respond most strongly to those inputs that are active when errors occur. Thus, as errors occur more frequently in conjunction with particular inputs, the model ACC learns to respond more strongly to those inputs regardless of whether or not an error subsequently occurs.

Computational Model Results

For the present simulations, the published model was simulated with only the following manipulations. First, the original magnitude of the error signal as published was 1.0 in arbitrary units. This signal represented the pause in dopamine cell firing that occurs when expected reward is not received (Ljungberg, Apicella, & Schultz, 1992). Research has recently shown that the magnitude of phasic dopamine activity increases proportionally as the magnitude by which the actual reward exceeds the expected reward increases (Bayer & Glimcher, 2005). When expected reward is omitted, dopamine cell firing seems to be reduced in proportion to the magnitude of the deviation, except that the dopamine cell firing rate quickly saturates to a lower bound of zero spikes/second so that further reductions in reward below the expected value have no effect (Bayer & Glimcher, 2005). For the purposes of the computational model, we assume that reward omission causes a phasic depression of dopamine firing that is not strong enough to cause a saturating nonlinearity at the lower bound of zero spikes/sec.

In the present simulations (Figure 1), we compared various combinations of error likelihood and magnitude.

We used values of 0.5 (arbitrary units) for error signal magnitude (V_1) and 0.5 (arbitrary units) for error likelihood (E_1) to represent a baseline control condition. We then simulated larger error-consequence magnitudes (i.e., greater loss due to an error) as resulting in stronger error signals in the model. Additionally, we parametrically manipulated the experimental error-likelihood condition in separate runs to examine all combinations of error likelihood (E_2) and error magnitude (V_2), where error likelihood was 0.1, 0.5, and 0.75, and error magnitude was 0, 0.25, 0.5, 0.75, and 1. We then computed the product of each combination of error likelihood and error signal magnitude and plotted ACC activity in each of these conditions, using only correct trials and computing activation as a difference measure in comparison with the control condition (0.5, 0.5)

The result is a strong correlation (Figure 1). Essentially, the model predicts an approximately linear relationship between correct, go-trial ACC activity and the product of error likelihood and anticipated error consequence magnitude. This is a strong prediction. Essentially, the model predicts that ACC computes a measure of expected risk and, more specifically, that risk is computed as the product of error likelihood and the magnitude of the potential consequences. The product of likelihood and consequences suggests an analogy with expected utility theory, which predicts that expected utility is computed as the product of the likelihood of an outcome and the utility of the outcome. Note again, however, that expected utility and expected risk are orthogonal, just as mean is gener-

ally orthogonal to variance, as described previously. These further predictions of the error-likelihood computational model (Brown & Braver, 2005) stand in contrast with the conflict model, which was previously simulated to perform the change signal task for comparison with the error-likelihood model predictions (Brown & Braver, 2005). In these earlier simulations, the conflict model could not account for ACC error-likelihood effects, because the model incorporated no learning mechanism. Similarly, the conflict model would not be expected to show an effect of predicted error consequence magnitude in correct go trials for the same reason that the conflict model includes no learning mechanism by which it could associate stimulus cue information with a prediction of error magnitude.

NEUROIMAGING STUDY Expected Risk Effects in Human ACC

To test this striking model prediction, we developed a variant of the change signal task of Brown and Braver (2005), which we call the incentive change signal task (ICST, Figure 2). As the name suggests, in this variant, participants perform the task with the potential for monetary incentives for successful performance. The task was performed in four intermixed conditions that orthogonally varied the likelihood of an error and the magnitude of the consequence for the error. Thus, the task involved a 2×2 manipulation of high versus low error likelihood crossed with high versus low error consequences. For convenience, we refer to high and low error likelihood with the acronyms "HE" and "LE," respectively. Similarly, we refer to high and low error consequence magnitude with the acronyms "HM" and "LM," respectively. For example, the condition "HMLE" refers to the high error consequence magnitude, low-error-likelihood condition. Moreover, care was taken to ensure that error likelihood was not confounded with entropy or uncertainty of the outcome (Fiorillo, Tobler, & Schultz, 2003; Paulus, Hozack, Frank, & Brown, 2002). This was done by using error rates of 70% in the high-error-likelihood condition and 30% in the low-error-likelihood condition, since these two conditions are associated with equal uncertainty.

Method

Participants

Participants ($N = 21$, 9 female, ages 19–28) were recruited for the study after being screened for a history of neurological disorders, mental illness, substance abuse, and MRI contraindications. The recruitment base was the Washington University community and surrounding area of St. Louis, with recruitment occurring through newspaper advertisements and posted fliers. Participants were paid \$25/hr plus a performance bonus of \$0.01/1,000 points earned during task performance in the fMRI scanner. Doing so led to an additional bonus of approximately \$9 per participant. All procedures were approved by the Human Research Protection Office of Washington University in St. Louis.

Behavioral Task

The ICST was presented with PsyScope (Cohen, MacWhinney, Flatt, & Provost, 1993). Each trial of the ICST had four phases: color cue, target, response, and feedback (see Figure 2). Trials began with two horizontal dashes displayed in the center of the screen. The




Figure 1. Computational model predictions of expected risk. The published model of Brown and Braver (2005) was further manipulated to examine effects of the magnitude of the error signal as well as the likelihood of an error. The low-error-likelihood model input was fixed as the control condition, with an error signal magnitude (V_1) of 0.5 (arbitrary units), and an error likelihood (E_1) of 0.5. The high-error-likelihood condition was parametrically manipulated in separate runs to examine various combinations of error likelihood (E_2) and error magnitude (V_2). The plot shows the product of each combination of error likelihood and error magnitude (in comparison with the baseline control condition) versus the ACC activity for correct go trials (in comparison with the baseline). The result is a strong positive correlation. The model predicts an approximately linear relationship between correct go trial ACC activity and the product of error likelihood and error consequence magnitude.




Figure 2. Incentive change signal task. The task is a modified version of the change signal task of Brown and Braver (2005). Subjects earned \$0.02 for correct trials, \$0.00 for incorrect trials in the high-error-magnitude (HM) condition, and only \$0.01 for incorrect trials in the low-error-magnitude (LM) condition. Error rates were controlled at 70% for the high-error-likelihood condition (HE) and 30% for the low-error-likelihood condition (LE).

dashes were one of four colors: white, yellow, light blue, and brown, paired with the four possible combinations of high and low error likelihood and consequence magnitude. The pairings were counterbalanced across participants. All subsequent stimuli of the trial were the same color as the initial two dashes. After 1,000 msec, an angle brace appeared on the left or right of the two dashes, forming an arrow pointing either left or right (48-point font). On change signal trials, a second arrow appeared in 96-point font above the first, pointing in the opposite direction. Its appearance instructed the participants to withhold the response to the first arrow, if possible, and instead substitute a response to the second arrow. The arrow or arrows remained visible until a response deadline of 1,000 msec after the onset of the first arrow. The change signal delays (CSD) between the onset of the first and second arrow were adjusted by an asymmetric staircase algorithm to maintain target error rates. The CSDs were adjusted independently for each of the four color conditions. If the participant was correct on a change signal trial, then the CSD was increased. If instead the participant made an error on the change signal trial, then the CSD was decreased. In the low-error-likelihood conditions, the CSD was adjusted to achieve an error rate of approximately 30%. In the high-error-likelihood conditions, the CSD was adjusted to achieve an error rate of approximately 70%. After the response deadline, the screen was blank for 500 msec; then, participants were provided with visual feedback for 1,000 msec. The feedback consisted of four digits indicating how many points the subject earned for the trial: 0000, 1,000, or 2,000 points. The screen then remained blank for a minimum of 1,500 msec until the start of the next trial. The intertrial intervals (ITIs) were jittered by adding 0, 2,500, 5,000, or 7,500 msec (3 TRs), where the ITI followed an exponential distribution (40% probability of a subsequent blank screen for 2,500 msec) in order to allow efficient estimation of the HRF (Burock, Buckner, Woldorff, Rosen, & Dale, 1998).

Participants earned points for each trial according to their task performances and the specific task condition. They were informed that their points would be converted directly to a cash payment at the end of the session, with the conversion factor revealed after the session ended. The actual conversion was US \$0.01 paid for each 1,000 points earned. In all conditions, participants received 2,000

points for each correct response. In the high-error-magnitude condition, when errors were made, 0 points were obtained. In contrast, in the low-error-magnitude condition, participants still received 1,000 points, even if they committed an error. Thus, the consequences of an error were more severe in the high magnitude conditions, even though the reward for correct responses did not differ.

Participants performed 8 blocks of 66 trials in the scanner. Participants were trained on the task beforehand outside the fMRI scanner so that they learned the task instructions, but not the nature of the error-likelihood and error-magnitude manipulations. Training typically consisted of less than the 66 trials that comprise a block.

Individual Differences

To measure individual differences in risk aversion, we gave participants a reduced form of the DOSPERS survey (Weber et al., 2002). Although we were primarily interested in the gambling subscale, we also included questions from four of the six DOSPERS content domains: gambling, investment, recreational, and social risk taking. As described previously, the ICST involves behaviors with a significant likelihood of financial loss in comparison with the payout for correct responses, and the gambling subscore among all DOSPERS subscores most directly measures aversion to behavior that leads to financial loss. We also considered using the DOSPERS financial investment subscore, but we found it less relevant, since financial investment on average is less likely to lead to financial loss than gambling. Indeed, certain kinds of financial investing are intended to reduce rather than increase the risk of financial loss. Gambling also correlates more strongly with ethical and health/safety risk taking (Weber et al., 2002). Participants were not asked questions from the health/safety and ethical domains in order to minimize ethical and confidentiality concerns (i.e., these questions inquire into issues that are more personal and sensitive, such as criminal behavior and substance abuse). However, risk taking in the gambling domain correlates well with risk taking in health/safety ($R = .44$) and ethical domains ($R = .56$) (Weber et al., 2002), and is therefore of more interest clinically. All survey instruments were administered outside the scanner. Following the format of the DOSPERS questionnaire, participants were asked to self-report their likelihood of engaging

in certain risky behaviors, as well as to rate the perceived benefits and perceived risks of each behavior. Self-report scores of gambling likelihood were summed across questions, and the possible range of scores in the gambling likelihood subscale were from 4 (*not likely to gamble*) to 16 (*very likely to gamble*).

Given the results of the computational modeling, we predicted that participants would show effects of both error likelihood and error consequence magnitude in ACC, and that this would be most pronounced in individuals who were more risk averse. Because the study used monetary incentive for correct answers, our primary focus was on the group of participants ("low-gambling participants") who were the most averse to financial gambling (DOSPERT gambling subscore < 6 , $n = 8$). Although the use of a threshold to dichotomize groups can be controversial (Farrington & Loeber, 2000; MacCallum, Zhang, Preacher, & Rucker, 2002), we chose a score of 6 as the cutoff in order to restrict the analysis to only the most risk-averse participants while maintaining adequate sample size for analysis. This was important, since individuals with high gambling likelihoods may include clinical or subclinical populations (Lesieur & Blume, 1987), and we were concerned to avoid such potential confounds. We return to the question of this grouping validity later on.

Functional Imaging

Functional images were collected with a Siemens Allegra 3T head-only MRI scanner and image slices parallel to the AC-PC plane (32 contiguous slices, 4 mm slice thickness, TR = 2,500 msec, TE = 25, flip angle = 90). All functional images were corrected for movement using a rigid-body rotation and translation correction, and they were then registered to the participant's anatomical images. The data were then scaled to achieve a whole-brain mode value (used in place of mean because of its reduced sensitivity to variation in brain margin definition) of 1,000 for each scanning run (to reduce the effect of scanner drift or instability), resampled into 3-mm isotropic voxels, and spatially smoothed with a 9-mm FWHM Gaussian kernel. Next, participants' structural images were transformed into standardized atlas space, using a 12-parameter affine transformation (Woods, Cherry, & Mazziotta, 1992; Woods, Grafton, Holmes, Cherry, & Mazziotta, 1998). The functional images were then registered to the reference brain (Talairach & Tournoux, 1988) using the alignment parameters derived for the structural scans.

A general-linear model approach (Friston et al., 1995) was used to estimate event-related responses. These responses were analyzed by estimating values for the various time points within the hemodynamic response epoch. The duration of this epoch was taken to be 17.5 sec (7 scan frames). Statistical t tests and ANOVAs were computed from cross correlations of a gamma function (Boynton, Engel, Glover, & Heeger, 1996) with the estimated hemodynamic response function time course (Ollinger, Corbetta, & Shulman, 2001). This cross correlation provides a measure of the BOLD response magnitude. The analyses focused on contrasts that were related to the error-likelihood and error-magnitude effects, using only correct go trials. Trials with conflict (change signals) and error trials were modeled separately in the GLM. By doing so, the analysis controlled for and excluded potentially confounding effects of conflict and errors (Brown & Braver, 2005). For error-likelihood effects, we looked at the LMHE–LMLE contrast (acronyms described above). By restricting the analysis so that the irrelevant factor of error consequence magnitude was minimal, we hoped to avoid a saturation effect in which increasing both error likelihood and error magnitude may not lead to additional increases in ACC sensitivity. For example, the HMHE condition entails complete loss of reward in 70% of change signal trials. Similarly, we looked at the HMLE–LMLE contrast to evaluate error consequence magnitude effects, minimizing the irrelevant effects of error likelihood to minimize potential saturation.

The analysis focused on regions showing sensitivity to expected risk. The statistical procedure required voxels to pass multiple contrast tests for different effects. Specifically, we identified voxels

showing both error likelihood ($p < .05$, two-tailed, uncorrected) and error consequence magnitude ($p < .05$, two-tailed, uncorrected) effects. This multiple contrast approach provides increased specificity for complex predicted effects while maintaining high sensitivity because of the low threshold on statistical significance (Brown & Braver, 2005). Additionally, to improve false-positive protection, regions were only considered significant if they consisted of at least 21 contiguous 3-mm³ voxels. Because we were interested specifically in the dorsal ACC areas reported previously (Brown & Braver, 2005), we did not perform an additional correction for multiple comparisons. However, additional whole-brain analyses were conducted for exploratory purposes. The error-likelihood effect was measured as the percent MR signal difference between the low consequence magnitude, high-error-likelihood condition (LMHE) and the low consequence magnitude, low-error-likelihood condition (LMLE), for correct go trials only. Similarly, the error consequence magnitude effect was measured as the difference between the high error consequence magnitude, low error-likelihood condition (HMLE) and the low consequence magnitude, low error-likelihood condition (LMLE), for correct go trials only.

Results

Behavioral Performance

Overall, there was no response time (RT) effect of error likelihood—that is, difference between LMHE and LMLE correct go trials [$F(1,20) = 0.014$, $p = .90$]. There was likewise no RT effect of predicted error magnitude—that is, difference between HMLE and LMLE correct go trials [$F(1,20) = 0.44$, $p = .52$]. The error-likelihood RT difference between LMHE and LMLE correct go trials furthermore did not correlate significantly with DOSPERT gambling-likelihood scores [$R = .28$; $F(1,19) = 1.22$, $p = .28$]. Likewise, the error-magnitude RT difference between HMLE and LMLE correct go trials did not correlate significantly with DOSPERT gambling-likelihood scores [$R = .08$; $F(1,19) = 0.12$, $p = .73$]. Thus, error-likelihood and error-magnitude effects on ACC activity were not confounded with RT differences. There was also no effect of gambling aversion on CSD. A between-groups ANOVA with factors of error likelihood and predicted error magnitude yielded only a main effect of error likelihood. No other effects or interactions were significant ($F < 1$).

Neuroimaging Data

Expected risk effects. We found that in the low-gambling participants, a region of interest with strong effects in the dorsal ACC was observed with 32 voxels and center of mass at +9, 27, 29 (Talairach coordinates; see Figure 3A), in a region that overlapped with the region where we previously found error-likelihood effects (Brown & Braver, 2005). The results for the region as a whole replicated the error-likelihood effects that were previously reported [$t(7) = 3.97$, $p < .005$, one-tailed, uncorrected] and also showed strong effects of error-consequence magnitude [$t(7) = 3.69$, $p < .005$, one-tailed, uncorrected], as predicted by the computational model (Figure 3B). These effects were more consistent across low-gambling participants than across conflict [stop correct – go correct, $t(7) = 1.63$, $p = .07$] and error effects [stop error – stop correct, $t(7) = 1.65$, $p = .07$], which each just missed significance in the region. Thus, the results were consistent with the model predictions of




Figure 3. Expected risk effects in ACC. (A) Exploratory analysis of the dorsal ACC identified a region of interest in right ACC (Talairach 9, 27, 29) that showed significant effects of both error likelihood and anticipated consequence magnitude in the low-gambling (risk averse) individuals. (B) A confirmatory analysis showed that in contrast with the low-gambling group, both error-likelihood and anticipated error-magnitude effects were virtually absent in the more risk-tolerant high-gambling group. HM, high error magnitude; HE, high error likelihood; LM, low error magnitude; LE, low error likelihood. Except for conflict and error effects, all conditions reflect activation in correct go trials only.

error likelihood and anticipated consequence-magnitude effects in ACC. Furthermore, although we were specifically interested in the ACC, no other regions in the medial wall ($-18 < x < +18$) showed significant effects, as can be seen in Figure 3A. Two other brain regions showed expected risk effects (see Table 1), including one region in right dorsolateral PFC that also overlapped with that observed in our previous study (Brown & Braver, 2005). The overlap may suggest a common substrate of cognitive control driven by ACC activity.

We also performed a whole brain analysis in the 8 low-gambling participants to look separately for effects of error likelihood (LMHE–LMLE) or error magnitude (HMLE–LMLE) in correct go trials. Because this is not a conjunction analysis, we used a higher threshold of $p < .001$ uncorrected, with a minimum cluster size of 8 voxels. No regions were found in the 8 low-gambling participants that showed effects of error likelihood with this test. One region with 9 voxels was found that showed predicted error-magnitude effects, with a peak p value of $p < .0001$,

at Talairach coordinates 30, 39, 36 in the middle frontal gyrus (Brodmann's area 9).

The computational model predicted an approximately linear relationship between ACC activity and the product of error likelihood and error magnitude as predicted by the color cue (Figure 1). Although increased error likelihood and error magnitude both led to greater ACC activity, the two factors did not seem to be additive at high levels of error likelihood and magnitude. Whether over-additive or under-additive, the interaction between error likelihood and magnitude in the low-gambling group did not reach significance [$F(1,7) = 5.50$, $p = .052$]. The HMHE–LMLE (correct go trial) comparison (Figure 3B) did not reach significance [$t(7) = 1.65$, $p < .07$, one-tailed], as might have been expected from the model predictions, and it appeared to yield no greater ACC activity than separate manipulations of error likelihood or error magnitude alone. One possible account for the apparent discrepancy between the human and model results is that the human ACC activity may also be modulated by the average reward per trial. In the HMHE condition, fully 70% of trials result in no reward at all. Therefore, the average per trial reward is the lowest of any error-likelihood and -magnitude conditions. The model in its present form does not take average reward per trial into account as a multiplier of ACC responses, but this could be simulated in principle.

Participants were not trained on the task extensively before scanning, so we examined whether the error-likelihood and predicted error-magnitude effects increased significantly with learning throughout the course of a session. We looked for a learning effect by separately estimating the hemodynamic responses in four different time periods consisting of successive groups of two runs in each of the 8 low-gambling participants. In the region identified at Talairach +9, 27, 29, there was no significant correlation between error likelihood and time [$r = -.11$, $F(1,30) = 0.38$, $p = .54$]. There was also no significant correlation between predicted error magnitude and time [$r = -.10$, $F(1,30) = 0.33$, $p = .57$]. Given these results, it seems that the sample size was sufficient to yield a main effect of error likelihood and magnitude, but not to reveal the time-course of learning. In the interest of maximizing power, we therefore did not discard the earliest trials when estimating error-likelihood and magnitude effects.

Individual differences effects. For comparison with the low-gambling participants, we analyzed the remaining participants with DOSPRT gambling subscores ≥ 6 ; these

Table 1
Regions in Low-Gambling Subjects With Both Error-Likelihood and Error-Magnitude Effects With a Minimum of 21 Contiguous Voxels

Center x, y, z (Talairach)	Area	BA	No. Voxels
9, 27, 29	R anterior cingulate	32	32
31, -61, 26	R middle temporal gyrus	39	21
30, 37, 32	R middle/superior frontal gyrus	9	34

Note—Error likelihood = LMHE–LMLE correct go trials ($p < .05$, two-tailed, uncorrected). Error magnitude = HMLE–LMLE correct go trials ($p < .05$, two-tailed, uncorrected).

were identified as *high-gambling participants* ($n = 13$). In the same region, the high-gambling participants showed a remarkable absence of both error-likelihood [$t(12) = 0.10, p = .46$] and error-magnitude effects [$t(12) = 0.21, p = .42$]. Nonetheless, in the same region, conflict effects [$t(12) = 3.75, p < .002$, uncorrected] and error effects [$t(12) = 3.33, p < .004$, uncorrected] were strongly present (Figure 3B). We tested for a significant difference in error-likelihood and error-magnitude effects with group as a between-subjects factor. A MANOVA revealed a significant effect of group (high vs. low gambling) on the joint effects of error likelihood and error magnitude [$F(2,18) = 4.64$, Wilk's lambda = 0.66, $p < .03$]. However, when measured separately, the group differences in error likelihood [$F(1,19) = 2.76, p = .11$] and error magnitude [$F(1,19) = 2.69, p = .12$] failed to reach significance.

We further explored the validity of the high and low gambling-likelihood group assignments to determine whether the results were sensitive to the choice of gambling-likelihood cutoff score. A median split would seem the obvious method, but this was not possible strictly speaking, since 4 out of 21 participants scored at the median gambling-likelihood score of 6. Thus, a split has to have either 8 or 12 participants placed into the low gambling-likelihood group. If the participants are split on the other side of the median (i.e., 12 participants in the low-gambling group) does it change the results? We explored this question by reanalyzing the data with 12 participants in the low-gambling group. Paired t tests were run for the error likelihood (LMHE–LMLE) and error magnitude (HMLE–LMLE) contrasts. Again, an overlapping region at 12, 30, 30 was identified that showed significant effects of both error likelihood [$t(11) = 2.81, p = .017$, two-tailed] and error magnitude [$t(11) = 3.06, p = .011$, two-tailed]. As with the 8-subject split, no other regions in the medial PFC showed this effect in $-18 < x < +18$. Thus, the use of a different cutoff score on the other side of the median did not qualitatively change the results. The effects of error likelihood and error magnitude were quantitatively more significant in the 8-subject than in the 12-subject group, despite the smaller sample size. This finding is consistent with our hypotheses, because the 8-participant group has the lower mean gambling-likelihood score, and we predicted that those with the lowest gambling likelihoods would show the strongest error likelihood and predicted error magnitude effects.

We also analyzed the group differences by regressing the DOSPERT gambling likelihood score directly on the error-likelihood and error-magnitude effects across participants. The result is that although the 9, 27, 29 ROI shows a significant negative correlation for error magnitude [$r = -.4393, t(19) = 2.1315, p = .0463$] as expected, the correlation with error likelihood did not reach significance [$r = -.29, t = 1.342, p = .195$], apparently due to the high error-likelihood effect variations in the higher gambling-likelihood participants.

A further analysis examined the effects of individual differences by examining effects of error-likelihood (LMHE–LMLE) and error-magnitude (HMLE–LMLE) contrasts with the data pooled across high- and low-gambling par-

ticipants. With this method, the region identified in low-gambling participants at 9, 27, 29 did not show significant effects in the population as a whole. However, there was a nearby region of ACC centered on 15, 26, 29 (10 voxels) that showed a weak effect of error likelihood (LMHE–LMLE) [$t(20) = 2.02, p < .03$, one-tailed], consistent with our earlier findings (Brown & Braver, 2005). However, the error-magnitude effect (HMLE–LMLE) in this same region was not significant [$t(20) = 1.06, p = .15$, one-tailed].

Finally, to test whether any additional regions showed expected risk effects that might have been missed in the previous analyses, we conducted a whole-brain analysis in the pooled sample looking for areas showing a conjunction of significant effects of both error likelihood and error magnitude (21 voxels, $p < .05$, two-tailed). This analysis revealed two additional regions, one in right orbitofrontal cortex (OFC) area 11 (centered at 21, 48, -10), and one in the posterior cingulate cortex (PCC) area 31 (centered at 21, $-61, 25$). Previous studies have found effects of reward probability and reward magnitude in these areas (Kim et al., 2006; Knutson, Taylor, Kaufman, Peterson, & Glover, 2005). Of note, the effects of the right OFC also correlated positively across participants with gambling likelihood. Essentially, participants who were more likely to gamble showed greater OFC sensitivity to error likelihood [$r = .63, t(19) = 3.58, p < .002$] and error magnitude [$r = .53, t(19) = 2.69, p < .02$]. The PCC region effects did not correlate with gambling likelihood ($ps > .26$).

Conflict effects. We also examined whether conflict effects as well as error-likelihood effects in ACC varied with individual differences in risk aversion. To answer this question, we looked for regions in the ACC that showed a conjunction of two effects. First, we looked for regions in which gambling likelihood across participants correlated with the magnitude of the ACC error-likelihood effect (Figure 4A). Second, we looked for regions in which gambling likelihood across participants correlated with the magnitude of the ACC conflict effect (Figure 4B). Voxels passed the test if the correlations were both significant (each $p < .05$, two-tailed, uncorrected). We found a single region in left ACC (BA 24, center at TAL $-2, 4, 37$), shown in Figure 4. No other regions in the whole brain passed this test. Surprisingly, conflict activity was greater in this region for the high-gambling than for the low-gambling participants [$t(19) = 2.22, p < .04$], as shown in Figure 4C. The group $\times$ effect interaction (high vs. low gambling $\times$ conflict vs. error likelihood) was significant [$F(1,19) = 5.62, p < .03$, uncorrected]. At first, this negative correlation seems surprising from the perspective that response conflict effects are merely a reflection of error likelihood. However, in the Discussion section, we provide an explanation that appears to provide a satisfactory reconciliation.

Discussion

The present study was originally motivated by quantitative predictions of our previously published computational model. The computational model predictions of risk detection by ACC are notable for two reasons. First, they involved no changes from the previously published




Figure 4. Interactions of ACC effects and risk aversion. A single region of ACC showed significant correlations of error-likelihood effects and conflict effects with risk aversion. (A) Although error-likelihood effects correlated negatively with gambling likelihood as expected, the conflict effects (B) showed a surprising positive correlation with gambling likelihood. (C) The group $\times$ effect interaction was significant. (D) The region was slightly posterior to the area that showed the strongest expected risk effects. In high gambling-likelihood participants, a reduced sensitivity to expected risk may lead to a failure to implement early control over decision making to avoid risk. In that case, subsequent incongruent stimuli may elicit stronger immediate activity increases unchecked by prior control signals.

version of the model and so can be considered predictive in the strongest sense. Second, the computational-model results suggest a parallel between expected utility and expected risk, in that both involve the product of a likelihood and an outcome value. For expected utility, the value is the possible payoff. In contrast, for expected risk—as presently defined—the likelihood is multiplied not by the possible payoff, but instead by the possible deviation below the anticipated payoff if a mistake is made. Thus, expected utility and expected risk are computed similarly, but are nonetheless orthogonal, just as the variance of a statistical distribution is generally orthogonal to its mean. Some of the earliest results from the literature on judgment and decision making show that humans are generally risk averse (Bernoulli, 1954; Tversky & Kahneman, 1981), which leads to deviations from normative behavior under rational decision theory (Tversky & Kahneman, 1981), but the neural mechanisms by which risk is detected and avoided are only just beginning to be understood (Kuhnen

& Knutson, 2005; Paulus & Frank, 2006). The computational model provides a concrete, quantitative account of how risk is computed, which may provide a basis for understanding risk aversion in decision making.

The fMRI results are largely consistent with the prediction that ACC measures expected risk as presently defined. Notably, this effect was found prominently in more risk-averse participants, but was virtually absent in risk-tolerant participants. This result is consistent with the hypothesis that ACC not only detects general risk, but also drives risk aversion in decision making. The hypothesis that ACC activity reflects the uncertainty of a choice outcome (Paulus et al., 2002) cannot accommodate the present results, because outcome uncertainty was equivalent in the high-error-rate (70%) and low-error-rate (30%) conditions.

Although our results suggest that ACC encodes a measure of expected risk, previous work has suggested that related sets of regions encode the complementary expected value. Specifically, the orbitofrontal cortex

seems to represent the subjective desirability of an object (Padoa-Schioppa & Assad, 2006; Tremblay & Schultz, 1999). Other results suggest that computations of expected value are distributed so that the nucleus accumbens signals the value of a prospective gain, whereas the medial prefrontal cortex (which includes ACC) signals the likelihood of a gain (Knutson et al., 2005). Our results supported these interpretations in that we found OFC to be sensitive to expected risk, but with individual differences effects opposite to those observed in the ACC (greater activation associated with less risk aversion). Since expected value tended to negatively covary with expected risk in the present study (since reward for correct responses was constant across conditions), the results are consistent with the OFC coding information related to expected reward value. Specifically, one interpretation of the OFC effects is that individuals find engagement in the ICST task to be intrinsically rewarding, especially when punishments are successfully avoided in higher error-likelihood and -magnitude conditions (Kim et al., 2006). In that case, OFC might be understood as driving approach behavior such as engagement with the task, whereas ACC may drive error avoidance behavior (Magno et al., 2006), given the prior decision to engage a particular task. This account would also be consistent with effort-based accounts of ACC function (Walton, Bannerman, & Rushworth, 2002) in which the ACC has been shown to drive greater effort directed toward achieving a better outcome.

The present replication of Brown and Braver (2005) stands in contrast to a recent critique of the error-likelihood hypothesis (Nieuwenhuis et al., 2007), which failed to replicate the error-likelihood effect in ACC. The present results suggest one possible account of why the error-likelihood effect may not have been found, namely that the effect varies significantly with individual differences, a possibility that was in fact suggested by Nieuwenhuis et al. (2007). Cultural factors may also contribute to population differences in risk aversion between studies. For example, the Netherlands (where the Nieuwenhuis et al. study was performed) has on average a more tolerant view of certain kinds of potentially risky behavior than does the United States (where both the present study and our earlier study [Brown & Braver, 2005] were performed). In the same vein, it is also possible that our present participant sample may have excluded the most risk-tolerant participants. Our a priori selection criteria excluded participants with a prior history of substance abuse (see Participants section), which correlates positively with gambling likelihood (Weber et al., 2002). In any case, our results are consistent with the hypothesis that the participants of the Nieuwenhuis et al. study may have been more risk tolerant on average, although this is speculative, given that there were no individual differences measures reported for participants in that study (Nieuwenhuis et al., 2007).

A critical finding in the present study that was consistent with the Nieuwenhuis et al. (2007) results is that in a subset of participants (the high-gambling subgroup), error-likelihood effects were weak, although strong conflict effects could be detected. Moreover, we found that within a dorsal ACC region, error-likelihood and conflict

effects appeared to be negatively correlated across the full sample of participants. If pure error-likelihood effects are not found in some risk-tolerant individuals despite the presence of intact conflict effects, then can conflict effects still be understood as a special case of error-likelihood prediction? One potential way of reconciling these results is to consider them from within the framework of conflict-control feedback loops. In particular, a number of studies of conflict monitoring (Botvinick et al., 1999; Carter et al., 2000; Kerns et al., 2004) have shown that conflict effects are reduced when control is higher, such as when following the experience of a preceding incongruent trial. The conflict increases control in the next trial so that conflict is reduced in the subsequent trial, should an incongruent stimulus again appear. Of note, the increased control also entails reduced error rates—even for incongruent trials—in the subsequent trial (Botvinick, Braver, Barch, Carter, & Cohen, 2001; Botvinick et al., 1999).

In the same way, we suggest that risk-averse participants may be so because they are better able to recognize earlier predictors of error likelihood or general expected risk. In that case, they may be better able to increase control earlier in the course of a trial so that perceived error likelihoods are reduced and therefore less ACC activity is generated by a subsequent incongruent stimulus. The ACC of risk-tolerant individuals may still compute expected risk, but be less sensitive to more subtle cues. If risk-tolerant participants are less able to detect the increased error likelihood and consequent need for increased control given implicitly paired predictive color cues, then they may still show greater error-likelihood effects when the actual incongruent stimulus (in this case, the change signal) is presented, because it is a stronger predictor of impending error likelihood and potential consequences. In terms of the present task, if risk-averse participants detect the greater likelihood of an error on the basis of the color cue and increase control in response, then the appearance of the second arrow does not increase the perceived error likelihood as much. This hypothesis is speculative, because we did not find any variations in behavior (RT or CSD) with gambling likelihood and have no other behavioral measures of increased control within a trial. Nonetheless, the hypothesis is consistent with the negative correlation between gambling-likelihood and error-likelihood effects, and the positive correlation between gambling-likelihood and conflict effects in the dorsal ACC region identified previously. Moreover, in recent work with our computational model, we have found that simulations of individual difference effects in ACC error-driven learning rate can produce just this form of negative relationship between error likelihood/magnitude and conflict effects (Brown & Braver, *in press*). This result suggests that there is a functional relationship between the two types of effects, but that the nature of the relationship is more complex than we initially thought, and that it is dependent on both the dynamics of experience-based learning and on individual differences.

Recent studies have implicated ACC in decision making (Kennerley et al., 2006; Paulus & Frank, 2006), but it has not been clear how to reconcile these results with performance-monitoring theories of ACC (Botvinick

et al., 2001; Carter et al., 1998). Our results suggest a reconciliation of these seemingly disparate results, as follows. First, the expected risk hypothesis of ACC developed in the present work is a generalization of the error-likelihood hypothesis (Brown & Braver, 2005), which itself is a generalization of the conflict and error-detection hypotheses. Second, if ACC activity drives risk avoidance, then greater ACC activity should bias decision making against options that are more likely to lead to reduced reward, even if those options also entail less effort. Hence, ACC activity should bias decisions toward higher valued options, even when those options entail greater effort or cost (Kennerley et al., 2006; Walton et al., 2002).

As a whole, the present results are consistent with computational-model predictions that ACC serves a general function of detecting and avoiding risk. The ACC of more risk-averse participants may be more sensitive to earlier predictors of adverse outcomes and may therefore implement control processes earlier and more effectively to avoid potential errors and consequent losses. More generally, our results are also consistent with both the control loop theory of ACC and prefrontal cortex (Botvinick et al., 2001; Botvinick et al., 1999; Jones, Cho, Nystrom, Cohen, & Braver, 2002; Kerns et al., 2004), and with the hypothesis that ACC activity reflects in part an online evaluation of the expected risk of the current behavior.

AUTHOR NOTE

The present research was supported by ONR N00014-99-1-0715. J.W.B. was supported by a NARSAD Young Investigator award and the Sidney R. Baer, Jr. Foundation. T.S.B. was supported by P50 MH64445 and RO1 MH66088. Address correspondence to J. W. Brown, Department of Psychological and Brain Sciences, Indiana University, 1101 E. Tenth St., Bloomington, IN 47405 (e-mail: jwmbrown@indiana.edu).

REFERENCES

- AMIEZ, C., JOSEPH, J. P., & PROCYK, E. (2005). Anterior cingulate error-related activity is modulated by predicted reward. *European Journal of Neuroscience*, **21**, 3447-3452.
- AMIEZ, C., JOSEPH, J. P., & PROCYK, E. (2006). Reward encoding in the monkey anterior cingulate cortex. *Cerebral Cortex*, **16**, 1040-1055.
- BAYER, H. M., & GLIMCHER, P. W. (2005). Midbrain dopamine neurons encode a quantitative reward prediction error signal. *Neuron*, **47**, 129-141.
- BECHARA, A., DAMASIO, A. R., DAMASIO, H., & ANDERSON, H. W. (1994). Insensitivity to future consequences following damage to human prefrontal cortex. *Cognition*, **50**, 7-15.
- BERNOULLI, D. (1954). Exposition of a new theory on the measurement of risk. *Econometrica*, **22**, 23-36.
- BLAKEMORE, S. J., REES, G., & FRITH, C. D. (1998). How do we predict the consequences of our actions? A functional imaging study. *Neuropsychologia*, **36**, 521-529.
- BOTVINICK, M. M., BRAVER, T. S., BARCH, D. M., CARTER, C. S., & COHEN, J. C. (2001). Conflict monitoring and cognitive control. *Psychological Review*, **108**, 624-652.
- BOTVINICK, M. M., NYSTROM, L., FISSEL, K., CARTER, C. S., & COHEN, J. D. (1999). Conflict monitoring versus selection-for-action in anterior cingulate cortex. *Nature*, **402**, 179-181.
- BOYNTON, G. M., ENGEL, S. A., GLOVER, G. H., & HEEGER, D. J. (1996). Linear systems analysis of functional magnetic resonance imaging in human V1. *The Journal of Neuroscience*, **16**, 4207-4221.
- BRAVER, T. S., BARCH, D. M., GRAY, J. R., MOLFESE, D. L., & SNYDER, A. (2001). Anterior cingulate cortex and response conflict: Effects of frequency, inhibition, and errors. *Cerebral Cortex*, **11**, 825-836.
- BROWN, J. W., & BRAVER, T. S. (2005). Learned predictions of error likelihood in the anterior cingulate cortex. *Science*, **307**, 1118-1121.
- BROWN, J. W., & BRAVER, T. S. (in press). A computational model of risk, conflict, and individual difference effects in the anterior cingulate cortex. *Brain Research*.
- BUROCK, M. A., BUCKNER, R. L., WOLDORFF, M. G., ROSEN, B. R., & DALE, A. M. (1998). Randomized event-related experimental designs allow for extremely rapid presentation rates using functional MRI. *NeuroReport*, **9**, 3735-3739.
- CARTER, C. S., BRAVER, T. S., BARCH, D. M., BOTVINICK, M. M., NOLL, D. C., & COHEN, J. D. (1998). Anterior cingulate cortex, error detection, and the online monitoring of performance. *Science*, **280**, 747-749.
- CARTER, C. S., MACDONALD, A. M., BOTVINICK, M., ROSS, L. L., STENGER, A., NOLL, D., & COHEN, J. D. (2000). Parsing executive processes: Strategic versus evaluative functions of the anterior cingulate cortex. *Proceedings of the National Academy of Sciences*, **97**, 1944-1948.
- CARTER, C. S., MACDONALD, A. W., III, ROSS, L. L., & STENGER, V. A. (2001). Anterior cingulate cortex activity and impaired self-monitoring of performance in patients with schizophrenia: An event-related fMRI study. *American Journal of Psychiatry*, **158**, 1423-1428.
- CARVER, C. S., & WHITE, T. (1994). Behavioral inhibition, behavioral activation, and affective responses to impending reward and punishment: The BIS/BAS scales. *Journal of Personality & Social Psychology*, **67**, 319-333.
- COHEN, J., MACWHINNEY, B., FLATT, M., & PROVOST, J. (1993). PsychoScope: An interactive graphic system for designing and controlling experiments in the psychology laboratory using Macintosh computers. *Behavior Research Methods, Instruments, & Computers*, **25**, 257-271.
- DE MARTINO, B., KUMARAN, D., SEYMOUR, B., & DOLAN, R. (2006). Frames, biases, and rational decision-making in the human brain. *Science*, **313**, 684-687.
- FARRINGTON, D., & LOEBER, R. (2000). Some benefits of dichotomization in psychiatric and criminological research. *Criminal Behaviour & Mental Health*, **10**, 100-122.
- FIORILLO, C. D., TOBLER, P. N., & SCHULTZ, W. (2003). Discrete coding of reward probability and uncertainty by dopamine neurons. *Science*, **299**, 1898-1902.
- FRANK, M. J., WOROCH, B. S., & CURRAN, T. (2005). Error-related negativity predicts reinforcement learning and conflict biases. *Neuron*, **47**, 495-501.
- FRISTON, K. J., HOLMES, A. P., WORSLEY, K. J., POLINE, J.-P., FRITH, C. D., & FRACKOWIAK, R. S. J. (1995). Statistical parametric mapping in functional imaging: A general linear approach. *Human Brain Mapping*, **2**, 189-210.
- GEHRING, W. J., COLES, M. G. H., MEYER, D. E., & DONCHIN, E. (1990). The error-related negativity: An event-related potential accompanying errors. *Psychophysiology*, **27**, S34.
- 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.
- GRAY, J. R., & BRAVER, T. S. (2002). Personality predicts working-memory-related activation in the caudal anterior cingulate cortex. *Cognitive, Affective, & Behavioral Neuroscience*, **2**, 64-75.
- HEWIG, J., TRIPPE, R., HECHT, H., COLES, M. G., HOLROYD, C. B., & MILTNER, W. H. (2007). Decision-making in blackjack: An electrophysiological analysis. *Cerebral Cortex*, **17**, 865-877.
- HOHNSBEIN, J., FALKENSTEIN, M., & HOORMAN, J. (1989). Error processing in visual and auditory choice reaction tasks. *Journal of Psychophysiology*, **3**, 32.
- HUNT, M. K., HOPKO, D. R., BARE, R., LEJUEZ, C. W., & ROBINSON, E. V. (2005). Construct validity of the balloon analog risk task (BART): Associations with psychopathy and impulsivity. *Assessment*, **12**, 416-428.
- HUSAIN, M., PARTON, A., HODGSON, T. L., MORT, D., & REES, G. (2003). Self-control during response conflict by human supplementary eye field. *Nature Neuroscience*, **6**, 117-118.
- ITO, S., STUPHORN, V., BROWN, J., & SCHALL, J. D. (2003). Performance monitoring by anterior cingulate cortex during saccade countermanding. *Science*, **302**, 120-122.
- JOHANSEN, J. P., & FIELDS, H. L. (2004). Glutamatergic activation of anterior cingulate cortex produces an aversive teaching signal. *Nature Neuroscience*, **7**, 398-403.

- JONES, A. D., CHO, R., NYSTROM, L. E., COHEN, J. D., & BRAVER, T. S. (2002). A computational model of anterior cingulate function in speeded response tasks: Effects of frequency, sequence, and conflict. *Cognitive, Affective, & Behavioral Neuroscience*, **2**, 300-317.
- KAHNEMAN, D., & TVERSKY, A. (1979). Prospect theory: An analysis of decision under risk. *Econometrica*, **47**, 263-291.
- KENNERLEY, S. W., WALTON, M. E., BEHRENS, T. E., BUCKLEY, M. J., & RUSHWORTH, M. F. (2006). Optimal decision making and the anterior cingulate cortex. *Nature Neuroscience*, **9**, 940-947.
- KERNS, J. G., COHEN, J. D., MACDONALD, A. W., III, CHO, R. Y., STENGER, V. A., & CARTER, C. S. (2004). Anterior cingulate conflict monitoring and adjustments in control. *Science*, **303**, 1023-1026.
- KIM, H., SHIMOJO, S., & O'DOHERTY, J. P. (2006). Is avoiding an aversive outcome rewarding? Neural substrates of avoidance learning in the human brain. *PLoS Biology*, **4**, e233.
- KNUTSON, B., TAYLOR, J., KAUFMAN, M., PETERSON, R., & GLOVER, G. (2005). Distributed neural representation of expected value. *Journal of Neuroscience*, **25**, 4806-4812.
- KUHNE, C. M., & KNUTSON, B. (2005). The neural basis of financial risk taking. *Neuron*, **47**, 763-770.
- LEJUEZ, C. W., AKLIN, W. M., ZVOLENSKY, M. J., & PEDULLA, C. M. (2003). Evaluation of the balloon analogue risk task (BART) as a predictor of adolescent real-world risk-taking behaviours. *Journal of Adolescence*, **26**, 475-479.
- LEJUEZ, C. W., READ, J. P., KAHLER, C. W., RICHARDS, J. B., RAMSEY, S. E., STUART, G. L., ET AL. (2002). Evaluation of a behavioral measure of risk taking: The balloon analogue risk task (BART). *Journal of Experimental Psychology: Applied*, **8**, 75-84.
- LESIEUR, H. R., & BLUME, S. B. (1987). The south oaks gambling screen (SOGS): A new instrument for the identification of pathological gamblers. *American Journal of Psychiatry*, **144**, 1184-1188.
- LIDDLE, P. F., FRISTON, K. J., FRITH, C. D., HIRSCH, S. R., JONES, T., & FRACKOWIAK, R. S. J. (1992). Patterns of cerebral blood flow in schizophrenia. *British Journal of Psychiatry*, **160**, 179-186.
- 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., COWAN, W. B., & DAVIS, K. A. (1984). On the ability to inhibit simple and choice reaction time responses: A model and a method. *Journal of Experimental Psychology: Human Perception & Performance*, **10**, 276-291.
- MACCALLUM, R. C., ZHANG, S., PREACHER, K. J., & RUCKER, D. D. (2002). On the practice of dichotomization of quantitative variables. *Psychological Methods*, **7**, 19-40.
- MACDONALD, A. W., COHEN, J. D., STENGER, V. A., & CARTER, C. S. (2000). Dissociating the role of the dorsolateral prefrontal cortex and anterior cingulate cortex in cognitive control. *Science*, **288**, 1835-1838.
- MAGNO, E., FOXE, J. J., MOLHOLM, S., ROBERTSON, I. H., & GARAVAN, H. (2006). The anterior cingulate and error avoidance. *Journal of Neuroscience*, **26**, 4769-4773.
- MENON, V., ADLEMAN, N. E., WHITE, C. D., GLOVER, G. H., & REISS, A. L. (2001). Error-related brain activation during a go/nogo response inhibition task. *Human Brain Mapping*, **12**, 131-143.
- NIEUWENHUIS, S., SCHWEIZER, T., MARS, R. B., BOTVINICK, M. M., & HAJCAK, G. (2007). Error-likelihood prediction in the medial frontal cortex: A critical evaluation. *Cerebral Cortex*, **17**, 1570-1581.
- NORDAHL, T. E., CARTER, C. S., SALO, R. E., KRAFT, L., BALDO, J., SALAMAT, S., ET AL. (2001). Anterior cingulate metabolism correlates with Stroop errors in paranoid schizophrenia patients. *Neuropsychopharmacology*, **25**, 139-148.
- 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. 1-18). New York: Plenum.
- OLLINGER, J. M., CORBETTA, M., & SHULMAN, G. L. (2001). Separating processes within a trial in event-related functional MRI. *NeuroImage*, **13**, 218-229.
- OLSON, C. R., & GETTNER, S. N. (2002). Neuronal activity related to rule and conflict in macaque supplementary eye field. *Physiological Behavior*, **77**, 663-670.
- PADOA-SCHIOPPA, C., & ASSAD, J. A. (2006). Neurons in the orbitofrontal cortex encode economic value. *Nature*, **441**, 223-226.
- PAULUS, M. P., & FRANK, L. R. (2006). Anterior cingulate activity modulates nonlinear decision weight function of uncertain prospects. *NeuroImage*, **30**, 668-677.
- PAULUS, M. P., HOZACK, N., FRANK, L., & BROWN, G. G. (2002). Error rate and outcome predictability affect neural activation in prefrontal cortex and anterior cingulate during decision-making. *NeuroImage*, **15**, 836-846.
- RUDEBECK, P. H., WALTON, M. E., SMYTH, A. N., BANNERMAN, D. M., & RUSHWORTH, M. F. (2006). Separate neural pathways process different decision costs. *Nature Neuroscience*, **9**, 1161-1168.
- SCHIFFERS, M. K., & COLES, M. G. (2000). Performance monitoring in a confusing world: Error-related brain activity, judgments of response accuracy, and types of errors. *Journal of Experimental Psychology: Human Perception & Performance*, **26**, 141-151.
- SHIDARA, M., & RICHMOND, B. J. (2002). Anterior cingulate: Single neuronal signals related to degree of reward expectancy. *Science*, **296**, 1709-1711.
- SHIMA, K., & TANJI, J. (1998). Role of cingulate motor area cells in voluntary movement selection based on reward. *Science*, **282**, 1335-1338.
- STUPHORN, V., TAYLOR, T. L., & SCHALL, J. D. (2000). Performance monitoring by the supplementary eye field. *Nature*, **408**, 857-860.
- TALAIRACH, J., & TOURNOUX, P. (1988). *Co-planar stereotaxic atlas of the human brain*. New York: Thieme.
- TREMBLAY, L., & SCHULTZ, W. (1999). Relative reward preference in primate orbitofrontal cortex. *Nature*, **398**, 704-708.
- TVERSKY, A., & KAHNEMAN, D. (1981). The framing of decisions and the psychology of choice. *Science*, **211**, 453-458.
- 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 VEEN, V., COHEN, J. D., BOTVINICK, M. M., STENGER, V. A., & CARTER, C. S. (2001). Anterior cingulate cortex, conflict monitoring, and levels of processing. *NeuroImage*, **14**, 1302-1308.
- VON NEUMANN, J., & MORGANSTERN, O. (1944). *Theory of games and economic behavior*. Princeton, NJ: Princeton University Press.
- WALTON, M. E., BANNERMAN, D. M., & RUSHWORTH, M. F. (2002). The role of rat medial frontal cortex in effort-based decision making. *Journal of Neuroscience*, **22**, 10996-11003.
- WALTON, M. E., DEVLIN, J. T., & RUSHWORTH, M. F. (2004). Interactions between decision making and performance monitoring within prefrontal cortex. *Nature Neuroscience*, **7**, 1259-1265.
- WEBER, E., BLAIS, A., & BETZ, N. (2002). A domain-specific risk-attitude scale: Measuring risk perceptions and risk behaviors. *Journal of Behavioral Decision Making*, **15**, 263-290.
- WOODS, R. P., CHERRY, S. R., & MAZZIOTTA, J. C. (1992). Rapid automated algorithm for aligning and reslicing PET images. *Journal of Computer Assisted Tomography*, **16**, 620-633.
- WOODS, R. P., GRAFTON, S. T., HOLMES, C. J., CHERRY, S. R., & MAZZIOTTA, J. C. (1998). Automated image registration: I. General methods and intrasubject, intramodality validation. *Journal of Computer Assisted Tomography*, **22**, 139-152.
- ZUCKERMAN, M. (1994). *Behavioral expressions and biosocial bases of sensation-seeking*. Cambridge: Cambridge University Press.

(Manuscript received January 26, 2007;
revision accepted for publication July 8, 2007.)