

Activity in human reward-sensitive brain areas is strongly context dependent

Sander Nieuwenhuis,^{a,*} Dirk J. Heslenfeld,^a Niels J. Alting von Geusau,^a Rogier B. Mars,^{b,c} Clay B. Holroyd,^d and Nick Yeung^e

^aDepartment of Cognitive Psychology, Vrije Universiteit, Van der Boechorststraat 1, 1081 BT Amsterdam, The Netherlands

^bF.C. Donders Centre for Cognitive Neuroimaging, Nijmegen, The Netherlands

^cNijmegen Institute for Cognition and Information, Radboud Universiteit Nijmegen, The Netherlands

^dDepartment of Psychology, University of Victoria, Victoria, Canada

^eDepartment of Psychology, Carnegie Mellon University, Pittsburgh, PA 15213, USA

Received 3 November 2004; revised 13 December 2004; accepted 17 December 2004

Available online 16 February 2005

Functional neuroimaging research in humans has identified a number of brain areas that are activated by the delivery of primary and secondary reinforcers. The present study investigated how activity in these reward-sensitive regions is modulated by the context in which rewards and punishments are experienced. Fourteen healthy volunteers were scanned during the performance of a simple monetary gambling task that involved a “win” condition (in which the possible outcomes were a large monetary gain, a small gain, or no gain of money) and a “lose” condition (in which the possible outcomes were a large monetary loss, a small loss, or no loss of money). We observed reward-sensitive activity in a number of brain areas previously implicated in reward processing, including the striatum, prefrontal cortex, posterior cingulate, and inferior parietal lobule. Critically, activity in these reward-sensitive areas was highly sensitive to the range of possible outcomes from which an outcome was selected. In particular, these regions were activated to a comparable degree by the best outcomes in each condition—a large gain in the win condition and no loss of money in the lose condition—despite the large difference in the objective value of these outcomes. In addition, some reward-sensitive brain areas showed a binary instead of graded sensitivity to the magnitude of the outcomes from each distribution. These results provide important evidence regarding the way in which the brain scales the motivational value of events by the context in which these events occur.

© 2004 Elsevier Inc. All rights reserved.

Keywords: fMRI; Punishment; Feedback; Reward processing; Context dependence; Prediction error

A fundamental function of the brain is to evaluate the motivational significance of ongoing events and to use this information to adapt behavior. Neuroimaging studies of reward processing have identified a number of brain areas that are activated by the delivery of primary reinforcers such as appetitive stimuli (Berns et al., 2001; McClure et al., 2003; O’Doherty et al., 2004), and secondary reinforcement such as monetary gains and losses (Breiter et al., 2001; Delgado et al., 2004; Elliott et al., 2000; Holroyd et al., 2004b; Thut et al., 1997). However, it remains to be determined precisely how information about reward and punishment is encoded in these reward-sensitive neural circuits. The aim of the present study was to investigate one critical aspect of reward processing: the extent to which neural coding of reward value is influenced by the context in which outcomes are experienced.

We manipulated reward context by varying the range of possible outcomes from which a particular outcome was drawn. The fundamental question was whether neural activity in reward-sensitive brain regions would vary solely with objective reward value (*context independence*), or whether activity in these areas would be scaled by subjective expectations about the magnitude of the possible outcomes (*context dependence*). We used functional magnetic resonance imaging (fMRI) to explore this issue of context dependence and independence at the level of the whole brain. The experimental paradigm we used was a simple monetary gambling task (cf. Holroyd et al., 2004a; Fig. 1A). In one condition, the *win condition*, the three possible outcomes were +60¢, +30¢, and +0¢. In the other condition, the *lose condition*, the three possible outcomes were –40¢, –20¢, and –0¢. The question of interest was the extent to which fMRI activity in reward-sensitive brain regions would be modulated by the context (win or lose) in which outcomes were experienced.

One possibility is that particular brain areas show context-dependent activity (Fig. 1B, top). An important feature of a context-

* Corresponding author.

E-mail address: stn20@dds.nl (S. Nieuwenhuis).

Available online on ScienceDirect (www.sciencedirect.com).

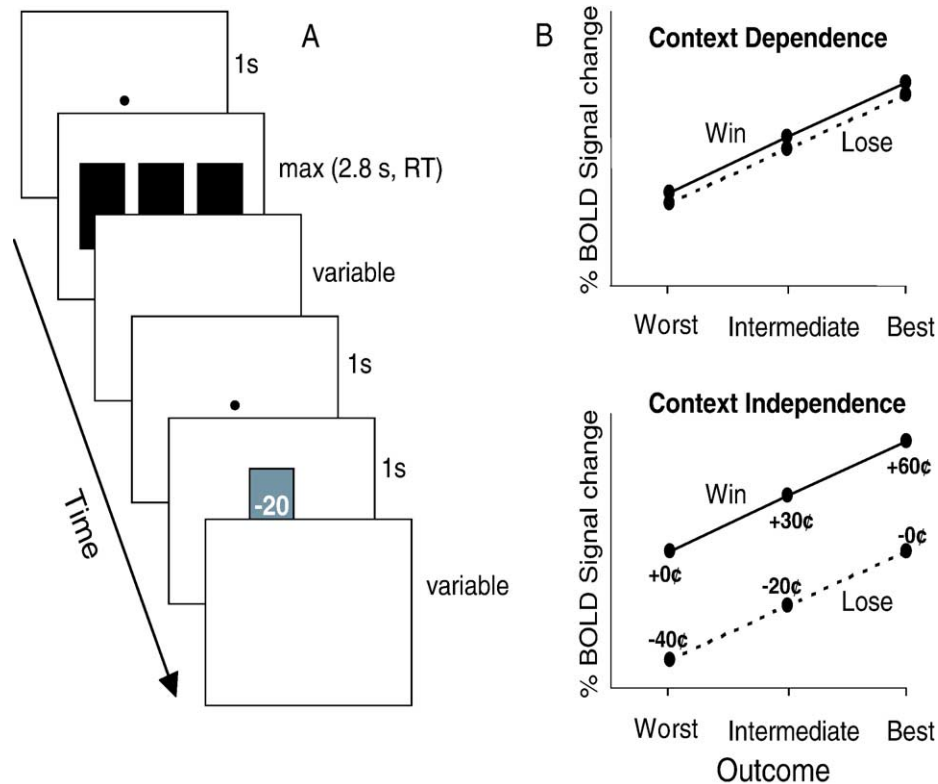


Fig. 1. (A) Illustration of task displays. See text for actual colors and size. (B) Schematic examples of context sensitivity in reward processing. Lose, win: task conditions. Worst, intermediate, best: the worst, intermediate, and best of the three possible outcomes in each condition, as given by their objective values. The ordinate corresponds to percentage BOLD signal increase associated with each outcome, and assumes that better outcomes elicit larger BOLD activity, as is the case in the current data. Note that the relationship between the outcomes and percentage BOLD signal increase might follow any monotonically increasing function; linear functions were adopted here only for ease of exposition.

dependent system is that it should show a different neural response to a 0¢ outcome when it is experienced in the win condition (in which 0¢ is the worst possible outcome) than when it is experienced in the lose condition (in which 0¢ is the best possible outcome), but should show a similar response to the +30¢ and -20¢ outcomes (the intermediate outcomes of their respective conditions). A context-independent system would show the opposite pattern: equivalent responses to 0¢ outcomes in the win and lose conditions but a different response to the +30¢ and -20¢ outcomes (Fig. 1B, bottom). The former possibility would be consistent with theories like reward prediction error theory, according to which brain areas are sensitive to deviations from expected reward rather than to the objective value of reward (Montague and Berns, 2002). The latter possibility would be consistent with previous findings of overlapping but somewhat dissociable systems for processing abstract rewards and punishments (Elliott et al., 2000; O'Doherty et al., 2001; Zalla et al., 1999).

Materials and methods

Participants

Fourteen young adults (six women), ranging in age from 22 to 31 years ($M = 25.4$) participated in the experiment. All participants were right-handed and had normal or corrected-to-normal visual acuity. They were paid € 20, plus a feedback-related bonus of € 9, as described below. All participants provided written informed

consent, and the experiment was approved by the local research ethics panel.

Task

Each trial (see Fig. 1A for an example) started with the presentation of a central fixation point that was replaced after 1 s by three colored rectangles. The participants were asked to imagine that these rectangles were playing cards presented face-down. The three cards remained on the screen (for a maximum of 2.8 s) until the participant selected one of them by pressing a spatially corresponding button with the right index, middle, or ring finger. After a variable interval, the participant was presented with a central fixation point for 1 s, followed by a gray rectangle with a number indicating the outcome of the trial (in eurocents). Participants were asked to imagine that this was their chosen card flipped over, and were informed that the number on each chosen card would be added to (or subtracted from) the total amount of bonus money awarded to them at the end of the experiment. The feedback display remained on the screen for 1 s, followed by an empty screen until the start of the next trial. The time interval between the onset of the choice display and the feedback display, and between the onset of the feedback display and the choice display for the next trial varied between three values that occurred with equal frequency: 4, 6.5, and 9 s. The interval between stimulus events was jittered in order to allow more efficient deconvolution of the hemodynamic signal (Burock et al., 1998).

The experiment consisted of two conditions, indicated by the color (red or green) of the cards in the choice display. One color was associated with the *win condition*: Participants were informed that the numbers on the back of the three cards with this color were +60, +30, and 0, and that their goal was to try to pick cards that would increase their bonus total. The other color was associated with the *lose condition*: Participants were informed that the numbers on the back of these three cards were -40, -20, and 0, and that their goal was to try to avoid picking cards that would decrease their bonus total. As in previous research (e.g., Breiter et al., 2001; Delgado et al., 2003), the gains were made larger than the losses to compensate for the tendency of participants to assign greater weight to a loss than to a gain of equal magnitude (Tversky and Kahneman, 1981). The difference in the size of gains and losses also ensured that participants tended to accumulate winnings gradually throughout the experiment at an average rate of 5¢ per trial. The assignment of card colors to experimental conditions was varied across participants.

The two possible choice displays (red and green cards) and the six possible outcomes (+60, +30, 0 in the win condition, 0 in the lose condition, -20, -40) were varied on a trial-to-trial basis and occurred equally often across the experiment. The order of outcomes was varied across participants according to a predetermined, randomized sequence that was constrained to minimize correlations between the eight predictors. The participants were not told about these computer-controlled contingencies, nor were they told that the sequence of outcomes was unaffected by their response choices. Instead, they were simply instructed to “use any strategy you want to help you maximize your bonus total”. Because in truth there was no strategy to learn, there was no meaningful performance measure in this task. Instead, the task simply provided a realistic context in which rewards and penalties and the anticipation thereof were experienced. Nevertheless, at debriefing, most participants reported that they had attempted to find a systematic pattern in the association between responses and outcomes, and that they had felt disappointed when testing of a specific hypothesis regarding this association led to unfavorable outcomes.

Participants received instructions and ten practice trials outside the scanner before entering the experimental phase. The experimental phase consisted of 180 trials altogether, divided into six equal blocks, with short breaks in between. At the end of each block, participants were informed about their current cumulative winnings. At the end of the experiment, participants had won a €9 bonus, which was then paid to them, in addition to their basic financial compensation.

Stimuli

Stimuli were presented in color on a black background projected onto a frosted screen, viewed by the participants through a tilted mirror placed on top of the head coil. The fixation point was white and subtended 0.4°. The three rectangles in the choice display were red or green with a thin white outline. They subtended 4.4° × 5.9°, and the visual angle between the centers of the rectangles was approximately 6.0°. The rectangle in the feedback display was gray with a thin white outline and subtended 4.4° × 5.9°. The outcome numerals were presented in a white, 44-size, bold Courier New font and subtended approximately 1.6° vertically.

fMRI image acquisition

Images were collected with a 1.5-T Siemens Sonata scanner equipped with a volume head coil. Prior to the experimental session, anatomical images were collected using a T1-weighted MP-RAGE sequence (TR = 2700 ms, TE = 3.95 ms, TI = 950 ms, FA = 8°, 256 × 160 coronal matrix, 1.0 × 1.0 mm in-plane resolution, 224 1.1-mm slices). Functional images were reconstructed from twenty oblique slices acquired using a T2*-weighted EPI sequence (TR = 2000 ms, TE = 60 ms, FA = 90°, 64 × 64 matrix, 3.0 × 3.0 mm in-plane resolution, 20 5.0-mm slices, 1-mm gap between slices). Image acquisition varied systematically across trials with respect to stimulus onset, yielding an effectively higher temporal sampling rate (Josephs et al., 1997). Six functional runs (202 scans each) were collected. The first two scans of each run were discarded to allow the longitudinal magnetization to reach a steady state recovery value.

fMRI image analysis

Data were preprocessed and analyzed with BrainVoyager 2000 software (Maastricht, The Netherlands). Image preprocessing consisted of: rigid-body 3D motion correction using trilinear interpolation; slice scan time correction using sinc interpolation; 3D spatial smoothing with a 4-mm fullwidth at half maximum (FWHM) Gaussian kernel; voxel-wise linear detrending, highpass filtering (above 7 cycles per time course) to remove low frequencies, and lowpass filtering with a 2.8-s FWHM Gaussian kernel to remove high frequencies. Spatial normalization was performed using the standard 9-parameter landmark method of Talairach and Tournoux (1988).

For each participant and voxel, the blood oxygen-level dependent (BOLD) responses across the scanning run were modeled with a general linear model that included nine regressors. Two regressors accounted for choice displays indicating the win condition and the lose condition. Six additional regressors accounted for the six possible outcomes. A final regressor accounted for the fixation period before each choice display and feedback display. The hemodynamic response to each event was modeled by convolving each regressor with a standard gamma function (Boynton et al., 1996). For each voxel and each event type, a parameter estimate (beta coefficient) was computed that indicated the strength of covariance between the data and the modeled response function; these estimates were corrected for temporal autocorrelation using a first-order autoregressive model. Pair-wise contrasts between parameter estimates for different events were calculated for each participant and the results were submitted to a group analysis that treated inter-subject variability as a random effect. In a first step of the analyses, a whole-brain analysis using a +60¢ vs. -40¢ contrast was carried out to identify reward-sensitive brain areas. Statistical parametric maps were derived from the resulting *t* values associated with each voxel and were thresholded at a conservative value ($P < 0.0005$, uncorrected) with a contiguity threshold of 60 mm³ as a further precaution against type-1 errors (Forman et al., 1995). The location of the peak activity associated with each cluster of activation was reported in Talairach coordinates (Talairach and Tournoux, 1988). The reward-sensitive brain areas obtained in the initial whole-brain analysis were defined as regions of interest (ROIs) for subsequent analyses aimed at distinguishing between context dependence and context independence.

Results

Behavior

There are few meaningful behavioral measures in the task, because outcomes were unrelated to the participants' behavior. Since participants had 2.8 s to choose a card, there were very few trials on which no response was recorded ($M = 1.0$ trials). The choice reaction times for the trials on which a response was recorded were faster for the win condition ($M = 884$ ms) than for the lose condition [$M = 981$ ms; $F(1,13) = 13.4$, $P = 0.003$].

fMRI

The analysis of the fMRI data consisted of two steps. Previous studies have found a number of brain areas that show strong sensitivity to the magnitude of monetary rewards and punishments. In the first step of the analysis, we wanted to identify reward-sensitive brain areas in our participants by contrasting BOLD responses to the highest (+60¢) and lowest (−40¢) possible outcomes. This contrast should reveal areas that are sensitive to differences in experienced reward value, regardless of whether they are context dependent or independent. Activation was observed in a number of brain regions previously implicated in reward processing, including several basal ganglia areas (putamen, caudate, globus pallidus), prefrontal cortex, posterior cingulate, inferior parietal lobule, and cerebellum (Table 1). Each of these regions showed greater activity for +60¢ than for −40¢; no region showed greater activity for −40¢ than for +60¢. The ten brain regions listed in Table 1 served as our reward-sensitive ROIs for subsequent analyses.

In the second step of the analysis, we examined whether the reward-sensitive ROIs showed context dependence or independence (or some intermediate pattern of activity) by contrasting, for each area, the estimated magnitude of the BOLD responses associated with the two intermediate outcomes (i.e., +30¢ vs. −20¢), and those associated with the two 0¢ outcomes (+0¢ vs. −0¢). In line with previous research (Holroyd et al., 2004a), brain areas were considered context *dependent* if they satisfied the following critical condition: a larger BOLD response to the best outcome in the lose condition (−0¢) than to the worst outcome in the win condition (+0¢). In addition, context-dependent brain areas

should also show comparable BOLD responses to the intermediate outcomes from the win (+30¢) and lose (−20¢) conditions (Fig. 1B, top). Conversely, brain areas were considered context *independent* if they satisfied the two contrary conditions: Most critically, context-independent brain areas should show differential responses to the intermediate outcomes (+30¢ vs. −20¢) of the two conditions. Furthermore, one should expect comparable responses to the +0¢ outcome in the win condition and −0¢ outcome in the loss condition (Fig. 1B, bottom). For each contrast (+0¢ vs. −0¢, and +30¢ vs. −20¢), corresponding beta coefficients derived from the general linear model were compared using 2-tailed paired t tests. The resulting P values are reported in Table 1.

As Table 1 indicates, most of the reward-sensitive ROIs showed context dependence: With the exception of two areas, the right cerebellum and the left medial frontal gyrus, all ROIs showed a significant difference in activity between the 0¢ outcomes in the win condition (in which 0¢ was the worst outcome) and the lose condition (in which 0¢ was the best outcome). In addition, although one cannot draw strong conclusions from statistical null effects, it is worth noting that of these ROIs, six showed comparable BOLD responses to the +30¢ and −20¢ outcomes, the intermediate outcomes of the two conditions. Thus, the pattern of BOLD responses for these ROIs is consistent with the definition of pure context dependence (−0¢ > +0¢, +30¢ = −20¢). Two further ROIs, the putamen and globus pallidus, demonstrated a pattern of activity that fell between pure context dependence and pure context independence: These regions showed differential responses to the 0¢ outcomes in the win and lose conditions, but also showed differential responses to the +30¢ and −20¢ intermediate outcomes. In contrast, there were no ROIs that satisfied the definition of pure context independence (+30¢ > −20¢, −0¢ = +0¢).

The results are illustrated in Fig. 2, which shows event-related averages associated with each of the six outcomes for representative reward-sensitive ROIs. In general, this figure underlines the notion that reward-sensitive brain activity is scaled in relation to the range of possible outcomes. That is, the magnitudes of the BOLD responses are ordered much like in Fig. 1B (top): The best outcomes in each condition elicited the largest BOLD response and the worst outcomes in each condition elicited the smallest BOLD response, regardless of the objective value of those outcomes. Two aspects of Fig. 2 are particularly note-

Table 1

List of reward-sensitive brain areas revealed by contrasting the highest outcome and lowest outcome (+60¢ vs. −40¢)

Area	Left/right	Cluster size (mm ³)	X	Y	Z	Max t value	+30¢ vs. −20¢ ^a	+0¢ vs. −0¢ ^a	+60¢ vs. −0¢ ^a	CD/CI
Putamen	Right	1481	19	7	1	8.92	0.03 ^b	<0.0001 ^b	0.77	CD + CI
Caudate nucleus	Left	150	−8	4	3	5.97	0.32	<0.0005 ^b	0.31	CD
Globus pallidus	Left	443	−14	1	−6	7.16	0.04 ^b	0.004 ^b	0.70	CD + CI
Inferior parietal lobule	Left	283	−48	−52	45	6.68	0.28	0.002 ^b	0.48	CD
Posterior cingulate gyrus	Left	290	−2	−35	29	7.64	0.26	0.005 ^b	0.37	CD
Posterior cingulate sulcus	Left	63	−1	−31	47	5.35	0.97	0.03 ^b	0.65	CD
Medial frontal gyrus	Left	105	−6	7	53	6.23	0.34	0.24	0.23	−
Middle frontal gyrus	Left	60	−19	4	59	5.13	0.74	0.002 ^b	0.18	CD
Cerebellum	Right	296	10	−57	−10	6.62	0.26	0.30	0.18	−
Cerebellum	Left	92	−38	−51	−36	6.98	0.20	0.002 ^b	0.30	CD

Note. All regions are $P < 0.0005$, uncorrected.

CD = context dependent; CI = context independent.

^a P values for paired t tests (statistical comparison of the beta coefficients associated with different outcomes; $df = 13$).

^b Significant test results ($P < 0.05$).




Fig. 2. Event-related averages of the BOLD response associated with each of the six possible outcomes for six representative reward-sensitive regions of interest. The abscissa indicates time relative to the onset of the outcome display. The ordinate indicates the percentage BOLD signal increase. Cond = condition.

worthy. First, as noted above, by far the largest BOLD responses were elicited by the best outcomes in each of the two conditions: +60¢ and -0¢. Interestingly, despite the large difference in the absolute value of these two outcomes, the corresponding BOLD responses were strikingly similar: In none of the ROIs was there a reliable difference between these two outcomes (Table 1); if anything, in most ROIs, the -0¢ outcome elicited a somewhat greater neural response than the +60¢ outcome. A second important feature of Fig. 2 is that for some of the ROIs—most notably the caudate and the posterior cingulate gyrus—the BOLD response to the intermediate outcomes was similar to the BOLD response to the worst outcomes. This was supported by a statistical comparison (t test) between the summed beta coefficients for the intermediate outcomes (+30¢ and -20¢) and the summed beta coefficients for the worst outcomes (+0¢ and -40¢). There were only four ROIs that showed a significant difference ($P < 0.05$): the putamen, inferior parietal lobule, middle frontal gyrus, and right cerebellum. These findings suggest that most reward-sensitive ROIs showed a binary instead of a graded sensitivity to the magnitude of the

outcomes. That is, these areas appear to classify outcomes into two categories: The best outcomes are classified as good and all other outcomes are classified as bad.

In an additional analysis, we examined *expectation*-related activation in each of the reward-sensitive ROIs by contrasting the BOLD responses to the choice displays associated with the win and lose conditions (i.e., red vs. green cards). This analysis indicated that only two of the ROIs exhibited differential BOLD responses during the expectation period: the left posterior cingulate sulcus, which showed increased activity in the win condition compared to the lose condition ($P = 0.02$), and the medial frontal gyrus, which showed the reverse pattern ($P = 0.03$). A subsequent whole-brain analysis of the expectation period established that at the statistical threshold of $P < 0.0005$ (uncorrected), there were no areas that exhibited differential activity for the two choice displays. Lowering the threshold to $P < 0.005$ (uncorrected) revealed only a few areas—including the right precuneus (14, -45, 44), cerebellum (3, -41, -10), and the left inferior frontal gyrus (-35, 15, -8)—each of which showed larger activity in the win condition.

Discussion

The results of the present study provide important insights into the way in which information about rewards and penalties is encoded in the brain. Our central finding is that reward-sensitive areas throughout the brain exhibit a high degree of context sensitivity. More specifically, the results suggest that reward processing systems determine whether an outcome is favorable or unfavorable on the basis of the range of possible outcomes encountered in a particular setting—judging the best possible outcome to be favorable and the worst possible outcome to be unfavorable, regardless of the absolute magnitudes of these outcomes. The scaling of the reward by the range of possible outcomes is consistent with reward prediction error theory, according to which brain areas are sensitive to deviations from expected reward rather than to absolute magnitude of reward (Holroyd and Coles, 2002; Montague and Berns, 2002; Schultz, 2002). Our findings are also consistent with previous research that has identified brain areas showing modulation of reward-sensitive activity by the context of the recent history of monetary rewards and punishments (Akitsuki et al., 2003; Elliott et al., 2000; Nakahara et al., 2004).

Initial evidence consistent with our finding of context dependence was reported in a study by Breiter et al. (2001). In this study, participants experienced monetary gains and losses that were drawn, with the participants' knowledge, from different but partly overlapping outcome distributions. Context dependence was found for the nucleus accumbens and amygdala, in which the neural response to winning or losing nothing (i.e., \$0) depended on the distribution from which the outcome was drawn. However, whereas Breiter et al. restricted their analysis of context dependence to a limited set of brain areas, our findings indicate that reward-sensitive regions throughout the brain demonstrate a high degree of context dependence. Another difference with the current study is that in Breiter et al.'s study, the presented outcomes were not in any way contingent upon the participants' choices, or perceived by participants as such; instead, the participants were simply attending to the various stimulus displays. Importantly, recent research has demonstrated that the perception of action-outcome contingencies has a large impact on the (degree of) involvement of various brain areas in reward processing (Tricomi et al., 2004; Walton et al., 2004; Yeung et al., *in press*).

Other evidence consistent with our results has been reported by recent studies of a reward-sensitive electrophysiological brain potential thought to be generated in the medial frontal cortex: the feedback negativity (reviewed in Nieuwenhuis et al., 2004). These studies have indicated that the evaluative system that produces the feedback negativity processes gains and losses in a context-dependent manner (Holroyd et al., 2004a). That is, the amplitude of the feedback negativity is determined by the value of the eliciting outcome relative to the range of outcomes possible, rather than by the objective value of the outcome. In the present study, an activation focus in the caudal anterior cingulate cortex (−6, 34, 29), the presumed generator of the feedback negativity, fell just below our conservative significance threshold for identifying reward sensitivity ($P = 0.0007$). Nevertheless, the pattern of BOLD responses for this area showed a similar pattern of pure context dependence as observed for most of the other ROIs examined. The lack of significant activation in the caudal part of the anterior cingulate is inconsistent with previous findings of

significant activation of the dorsal–caudal ACC to unfavorable outcomes (Ridderinkhof et al., 2004). This finding in combination with the reliable activation of the posterior cingulate appears more consistent with recent results suggesting that the feedback negativity does not originate in the caudal anterior cingulate but instead reflects the summed activity of a more rostral region of the anterior cingulate and the posterior cingulate (Müller et al., 2005; Van Veen et al., 2004).

An interesting aspect of our results was that some reward-sensitive ROIs—including the caudate and posterior cingulate gyrus—did not show a parametric ordering of the three outcomes in each condition. Instead, these regions showed comparably small BOLD responses to the intermediate and worst outcomes. To explain this finding, one might propose that reward processing, at least of abstract financial reinforcers, is rather coarse in nature, providing a discrete evaluation of events as good or bad regardless of magnitude. Similar results have been reported for the feedback negativity (Holroyd et al., 2004a; Yeung and Sanfey, 2004). In contrast, the results seem inconsistent with a recent study that found a parametric ordering (in terms of both timing and magnitude) of the BOLD responses in left caudate nucleus to the magnitude of monetary gains and losses (Delgado et al., 2003). As can be seen in Fig. 2, the left caudate nucleus did not exhibit a similar parametric ordering in our study. Instead, the BOLD responses associated with small and intermediate outcomes were largely overlapping. The precise cause of the discrepancy between these results remains to be determined. However, a notable difference between the two studies is that, in the present study, gains and losses were presented in different contexts, whereas in the Delgado et al. study, all monetary outcomes were presented in the same context.

In further analyses, we found that the brain areas that showed context dependence in the processing of outcomes did not seem to activate differentially in response to the two choice displays (that provided the context in which the outcomes were experienced). This finding might suggest that areas in which activity is modulated by context do not themselves code the relevant contextual information, but instead receive this information from elsewhere in the brain (cf. Watanabe et al., 2002). However, a whole-brain analysis using a conservative statistical threshold did not reveal a single brain area that showed differential BOLD responses to the choice displays associated with the win and lose conditions. Previous research, most notably by Knutson and colleagues, has revealed the involvement of several distinct brain areas during the anticipation of monetary gains and losses (for review, see Knutson and Peterson, *in press*). A likely reason responsible for this discrepancy concerns the nature of the task used by Knutson and colleagues: Participants could obtain a promised reward of varying magnitude by producing a speeded response to a target stimulus. The participants used the expectation period to prepare this response and were presumably more motivated to do so in case of a large promised reward. Thus, the observed reward-dependent brain activity may reflect both anticipation of the potential reward and the degree of motivation to work for that reward (Roesch and Olson, 2004). In contrast, in the current study, participants did not have to make active use of the expectation phase, and were not told the specific reward they could obtain (but instead the range of possible rewards). Together, these factors may explain the relatively weak differential activation during the expectation phase observed in the present study.

A potential caveat of our study is that the scanning protocol that we used did not allow us to measure reliable signals from the orbitofrontal cortex, a brain area critically involved in reward processing. This leaves open the question to what extent the orbitofrontal cortex exhibits context dependence and/or independence in its response to monetary gains and punishments. Another possibility is that the orbitofrontal cortex will show a pattern of activations that is entirely different from those illustrated in Fig. 1B: Several studies have reported that activation of the orbitofrontal cortex is non-linearly related to objective value, such that the largest activations are observed for the extreme (highest and lowest) outcomes (Breiter et al., 2001; Elliott et al., 2003). In any case, this pattern of results suggests a strong sensitivity of the orbitofrontal cortex to the range of possible outcomes (see also Tremblay and Schultz, 1999).

An important issue for future research is to determine whether the present findings of context dependence in reward processing generalize to situations in which experienced outcomes are primary rewards and punishments (rather than abstract monetary rewards and penalties). This question could be addressed by modifying the present design so that in the win condition, the choice display signals a range of appetitive outcomes (e.g., 2 units of juice, 1 unit of juice, or water), and in the lose condition, the choice display signals a range of aversive outcomes (e.g., water, 1 unit of saline, or 2 units saline). The current results predict that brain activity elicited by appetitive and aversive outcomes should be influenced by the current task context, independently of whether the outcome eliciting the activity is a reward or a punishment per se. Such an extension of the present research would provide further critical insights into how the reward processing system evaluates the motivational significance of events as a function of the context in which these events occur.

Acknowledgments

This research was supported by the Netherlands Organization for Scientific Research (S.N.). We thank Nisan Mol and Sanne Boesveldt for technical assistance.

References

- Akitsuki, Y., Sugiura, M., Watanabe, J., Yamashita, K., Sassa, Y., Awata, S., Matsuoka, H., Maeda, Y., Matsue, Y., Fukuda, H., Kawashima, R., 2003. Context-dependent cortical activation in response to financial reward and penalty: an event-related fMRI study. *NeuroImage* 19, 1674–1685.
- Berns, G.S., McClure, S.M., Pagnoni, G., Montague, P.R., 2001. Predictability modulates human brain response to reward. *J. Neurosci.* 21, 2793–2798.
- Boynton, G.M., Engel, S.A., Glover, G.H., Heeger, D.J., 1996. Linear systems analysis of functional magnetic resonance imaging in human V1. *J. Neurosci.* 16, 4207–4221.
- Breiter, H.C., Aharon, I., Kahneman, D., Dale, A., Shizgal, P., 2001. Functional imaging of neural responses to expectancy and experience of monetary gains and losses. *Neuron* 30, 619–639.
- 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.
- Delgado, M.R., Locke, H.M., Stenger, V.A., Fiez, J.A., 2003. Dorsal striatum responses to reward and punishment: effects of valence and magnitude manipulations. *Cogn. Aff. Beh. Neurosci.* 3, 27–38.
- Delgado, M.R., Stenger, V.A., Fiez, J.A., 2004. Motivation-dependent responses in the human caudate nucleus. *Cereb. Cortex* 14, 1022–1030.
- Elliott, R., Friston, K.J., Dolan, R.J., 2000. Dissociable neural responses in human reward systems. *J. Neurosci.* 20, 6159–6165.
- Elliott, R., Newman, J.L., Longe, O.A., Deakin, J.F., 2003. Differential response patterns in the striatum and orbitofrontal cortex to financial reward in humans: a parametric functional magnetic resonance imaging study. *J. Neurosci.* 23, 303–307.
- Forman, S.D., Cohen, J.D., Fitzgerald, M., Eddy, W.F., Mintun, M.A., Noll, D.C., 1995. Improved assessment of significant activation in functional magnetic resonance imaging (fMRI): use of a cluster-size threshold. *Magn. Reson. Med.* 33, 636–647.
- Holroyd, C.B., Coles, M.G.H., 2002. The neural basis of human error processing: reinforcement learning, dopamine, and the error-related negativity. *Psychol. Rev.* 109, 679–709.
- Holroyd, C.B., Larsen, J.T., Cohen, J.D., 2004a. Context dependence of the event-related brain potential to reward and punishment. *Psychophys* 41, 245–253.
- Holroyd, C.B., Nieuwenhuis, S., Yeung, N., Nystrom, L.E., Mars, R.B., Coles, M.G.H., Cohen, J.D., 2004b. Dorsal anterior cingulate cortex shows fMRI response to internal and external error signals. *Nat. Neurosci.* 7, 497–498.
- Josephs, O., Turner, R., Friston, K., 1997. Event-related fMRI. *Hum. Brain Mapp.* 5, 243–248.
- Knutson, B., Peterson, R., in press. Neurally reconstructing expected utility. *Games Econ. Beh.*
- McClure, S.M., Berns, G.S., Montague, P.R., 2003. Temporal prediction errors in a passive learning task activate human striatum. *Neuron* 38, 339–346.
- Montague, P.R., Berns, G.S., 2002. Neural economics and the biological substrates of valuation. *Neuron* 36, 265–284.
- Müller, S.V., Möller, J., Rodriguez-Fornells, A., Münte, T.F., 2005. Brain potentials related to self-generated and external information used for performance monitoring. *Clin. Neurophys.* 116, 63–74.
- Nakahara, H., Itoh, H., Kawagoe, R., Takikawa, Y., Hikosaka, O., 2004. Dopamine neurons can represent context-dependent prediction error. *Neuron* 41, 269–280.
- Nieuwenhuis, S., Holroyd, C.B., Mol, N., Coles, M.G.H., 2004. Reinforcement-related brain potentials from medial frontal cortex: origins and functional significance. *Neurosci. Biobehav. Rev.* 28, 441–448.
- O'Doherty, J., Kringelbach, M.L., Rolls, E.T., Hornak, J., Andrews, C., 2001. Abstract reward and punishment representations in the human orbitofrontal cortex. *Nat. Neurosci.* 4, 95–102.
- O'Doherty, J., Dayan, P., Schultz, J., Deichmann, R., Friston, K., Dolan, R.J., 2004. Dissociable roles of ventral and dorsal striatum in instrumental conditioning. *Science* 304, 452–454.
- Ridderinkhof, K.R., Ullsperger, M., Crone, E.A., Nieuwenhuis, S., 2004. The role of the medial frontal cortex in cognitive control. *Science* 306, 443–447.
- Roesch, M.R., Olson, C.R., 2004. Neuronal activity related to reward value and motivation in primate frontal cortex. *Science* 304, 307–310.
- Schultz, W., 2002. Getting formal with dopamine and reward. *Neuron* 36, 241–263.
- Talairach, J., Tournoux, P., 1988. *Co-planar Stereotaxic Atlas of the Human Brain: An Approach to Medical Cerebral Imaging*. Thieme: Stuttgart, Germany.
- Thut, G., Schultz, W., Roelcke, U., Nienhusmeier, M., Missimer, J., Maguire, R.P., Leenders, K.L., 1997. Activation of the human brain by monetary reward. *NeuroReport* 8, 1225–1228.
- Tremblay, L., Schultz, W., 1999. Relative reward preference in primate orbitofrontal cortex. *Nature* 398, 704–708.
- Tricomi, E.M., Delgado, M.R., Fiez, J.A., 2004. Modulation of caudate activity by action contingency. *Neuron* 41, 281–292.

- Tversky, A., Kahneman, D., 1981. The framing of decisions and the psychology of choice. *Science* 211, 453–458.
- Van Veen, V., Holroyd, C.B., Cohen, J.D., Stenger, V.A., Carter, C.S., 2004. Errors without conflict: implications for performance monitoring theories of anterior cingulate cortex. *Brain Cogn.* 56, 267–276.
- Walton, M.E., Devlin, J.T., Rushworth, M.F., 2004. Interactions between decision making and performance monitoring within prefrontal cortex. *Nat. Neurosci.* 7, 1259–1265.
- Watanabe, M., Hikosaka, K., Sakagami, M., Shirakawa, S., 2002. Coding and monitoring of motivational context in the primate prefrontal cortex. *J. Neurosci.* 22, 2391–2400.
- Yeung, N., Sanfey, A.G., 2004. Independent coding of reward magnitude and valence in the human brain. *J. Neurosci.* 24, 6258–6264.
- Yeung, N., Holroyd, C.B., Cohen, J.D., in press. ERP correlates of feedback and reward processing in the presence and absence of response choice. *Cereb. Cortex.*
- Zalla, T., Koechlin, E., Basso, P., Aquino, P., Sirigu, A., Grafman, J., 1999. Differential amygdala responses to winning and losing: an fMRI study in humans. *Eur. J. Neurosci.* 12, 1764–1770.