



Universiteit
Leiden
The Netherlands

The neurocognitive development of social decision-making

Bos, W. van den

Citation

Bos, W. van den. (2011, April 12). *The neurocognitive development of social decision-making*. Retrieved from <https://hdl.handle.net/1887/16711>

Version: Not Applicable (or Unknown)

License: [Licence agreement concerning inclusion of doctoral thesis in the Institutional Repository of the University of Leiden](#)

Downloaded from: <https://hdl.handle.net/1887/16711>

Note: To cite this publication please use the final published version (if applicable).

7. Better than expected or as bad as you thought? The neurocognitive development of probabilistic feedback processing

Abstract

Learning from feedback lies at the foundation of adaptive behavior. Two prior neuroimaging studies have suggested that there are qualitative differences in how children and adults use feedback by demonstrating that dorsolateral prefrontal cortex (DLPFC) and parietal cortex were more active after negative feedback for adults, but after positive feedback for children. In the current study we used functional magnetic resonance imaging (fMRI) to test whether this difference is related to valence or informative value of the feedback by examining neural responses to negative and positive feedback while applying probabilistic rules. In total, 67 healthy volunteers between ages 8 and 22 participated in the study (8–11 years, $n = 18$; 13–16 years, $n = 27$; 18–22 years, $n = 22$). Behavioral comparisons showed that all participants were able to learn probabilistic rules equally well. DLPFC and dorsal anterior cingulate cortex were more active in younger children following positive feedback and in adults following negative feedback, but only when exploring alternative rules, not when applying the most advantageous rules. These findings suggest that developmental differences in neural responses to feedback are not related to valence *per se*, but that there is an age related change in processing learning signals with different informative value.

7.1 Introduction

Learning to correctly adapt your behavior in a changing environment is an essential feature of human cognition and has been studied extensively over the past decades (for reviews, see Ridderinkhof and van den Wildenberg, 2005; Rushworth and Behrens, 2008). When adapting behavior, individuals often make use of feedback signals, which can be positive, encouraging the continuation of behavior, or negative, discouraging the continuation of behavior and signaling the need for adjustment. Prior studies have indicated that adaptive learning based on feedback signals undergoes pronounced developmental improvements between late childhood and early adulthood, as is evident from tasks in which participants need to switch between multiple rules (Crone and van der Molen, 2004; Somsen, 2007) or in which they need to infer sorting rules based on positive and negative signals (van Duijvenvoorde et al., 2008).

Early developmental improvements in adaptive behavior are observed when feedback has a direct mapping to deterministic rules (Somsen, 2007), however, when the feedback is probabilistic, changes in adaptive learning are observed until late adolescence (Hooper et al., 2004). In these situations, individuals must learn the statistical regularities between actions and outcomes, and use that information to interpret current feedback signals (see also Rangel et al., 2008). Feedback which is not directly mapped to behavior is often more complex because it requires individuals to attend to long term consequences and override the tendency to respond directly to local environmental change.

Neuroimaging studies have shown that regions previously associated with cognitive control and response selection (Miller and Cohen, 2001; Toni et al., 2002) are also active when adults receive negative performance feedback, including the dorsal anterior cingulate cortex (dACC) and the dorsolateral prefrontal cortex (DLPFC) (Klein et al., 2007; Taylor et al., 2007). The dACC is thought to monitor action outcome regularities and is important for signaling adjustment (Botvinick et al., 2001; Yeung et al., 2004). In addition, the dACC may exercise behavioral control via the engagement of the DLPFC (Kerns et al., 2004; Zanolie et al., 2008), which in turn is important for trial-to-trial adjustments of behavior (Dosenbach et al., 2008). Similar to the DLPFC, the parietal cortex is also involved in feedback processing, in particular negative feedback (Crone et al., 2008; van Duijvenvoorde et al., 2008). Finally, these regions are thought to work in close concert with the basal ganglia, specifically the caudate nucleus, which is thought to be engaged when learning action-outcome regularities (for a review see Cools, 2008).

In two prior developmental studies we have identified the developmental time course of these regions during adaptive feedback processing. In the first study (Crone et al., 2008), participants were instructed to infer rules based on positive and negative feedback which could change without warning. Following Somsen (2007), we were interested in the way children, adolescents, and adults processed negative feedback indicating a rule shift. As anticipated, adults engaged DLPFC, dACC, and the parietal cortex when processing negative feedback indicating a rule shift. A similar pattern was observed in 14- to 15-year-old adolescents, but 8- to 11-year-old children engaged these regions less following negative feedback in comparison to positive feedback or a low-level fixation baseline. In the second study (van Duijvenvoorde et al., 2008), participants were instructed to guess a correct rule. Because there were two possible rules, there was a 50% chances of receiving positive feedback, and therefore both feedback signals (negative and positive) were similarly salient and probable. Again, adults engaged DLPFC, dACC, and the parietal cortex following negative feedback, but in this study 8-year-old children engaged

DLPFC and the parietal cortex more following positive feedback relative to negative feedback. The developmental trajectory of the dACC followed a different pattern, as it slowly emerged in response to negative feedback at the age of 12, but it was not more active following negative compared to positive feedback at a younger age (see also Velanova et al., 2008). Although the caudate nucleus was involved in these tasks, these studies revealed that there were no developmental differences in activation patterns.

Together, these findings indicate that the possible meaning of positive and negative feedback signals, and the role of the associated neural circuits, changes during development. However, prior studies could not dissociate between neural activation as a result of valence versus informative value, given that negative feedback always signaled response adjustment and therefore had different informative value than positive feedback. Thus, it remains to be determined how the involvement of DLPFC and the parietal cortex is dependent on valence versus informative value of the feedback.

Prior research suggests that differences in positive and negative feedback adjustment are the result of differences in attention regulation (Somsen, 2007). Following this hypothesis, it is argued that children are less able to update the relevant feedback information and therefore they are less flexible in selecting alternative actions. We therefore reasoned that the brain regions implicated in prior feedback studies may be sensitive to the informative value of feedback, and that activation in these brain regions is indicative of feedback attendance. Furthermore, we predicted that attention to feedback may also underlie the developmental differences in brain activation. We hypothesized that DLPFC and parietal cortex would be more active following positive feedback in children and following negative feedback in adults, but only when the feedback has informative value for learning and response adjustment. Thus, we sought to test how neural responses are sensitive to informative value for learning versus valence of feedback, and the developmental trajectory of feedback processing.

We reasoned that feedback valence versus informative value could be disentangled after participants learned probabilistic feedback rules. In the probabilistic learning paradigm, participants need to learn from positive and negative feedback under different levels of probability, and therefore not all positive feedback signals response continuation and not all negative feedback signals response adjustment. The probabilistic learning (i.e., trial-and-error) task employed in this study was based on a prior study by Frank et al. (2004), but was simplified for use with children. In our version of the probabilistic learning task, two different stimulus pairs (AB or CD) were presented in random order, and participants had to learn over trials that one stimulus was more likely to result in positive feedback (70–80%) (see Figure 7.1). Over the course of the

experiment participants had to learn the statistical regularities and thus had to learn to choose the stimuli with a high probability of positive feedback (A and C) more often than those with a low probability of positive feedback (B and D).

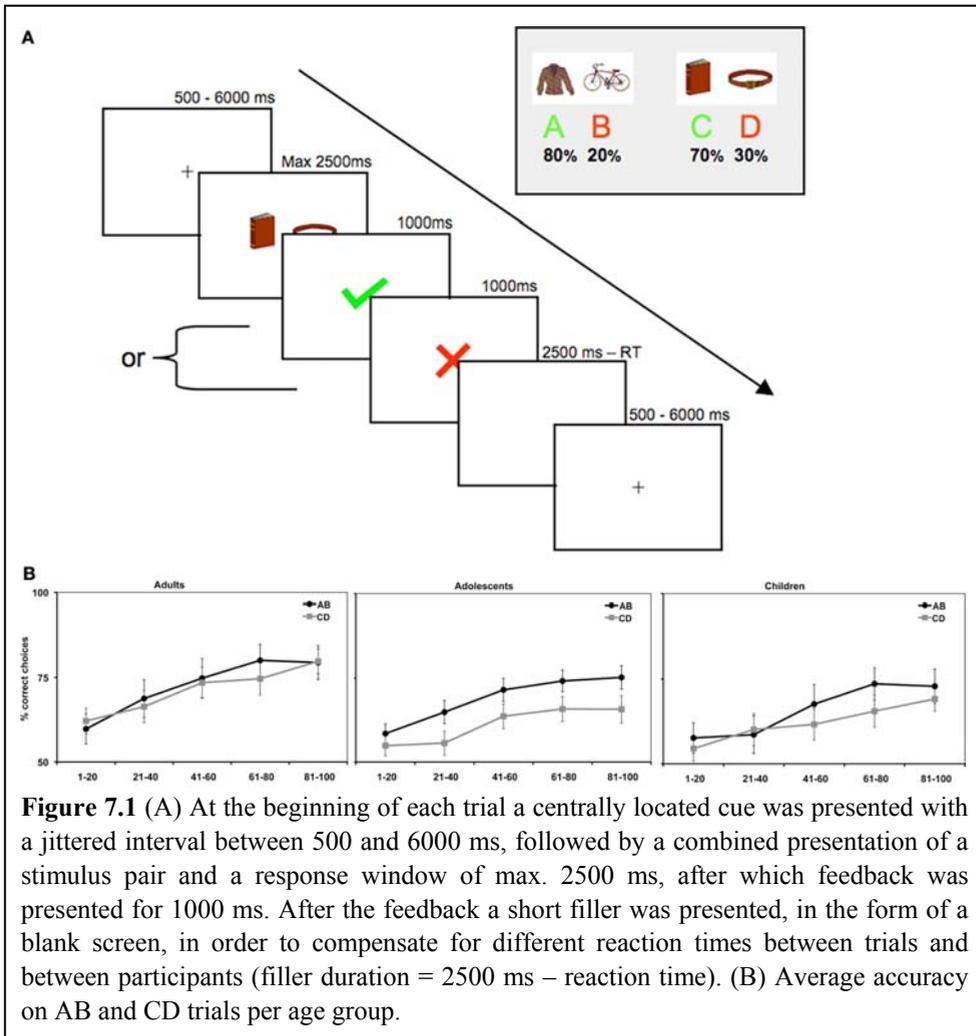


Figure 7.1 (A) At the beginning of each trial a centrally located cue was presented with a jittered interval between 500 and 6000 ms, followed by a combined presentation of a stimulus pair and a response window of max. 2500 ms, after which feedback was presented for 1000 ms. After the feedback a short filler was presented, in the form of a blank screen, in order to compensate for different reaction times between trials and between participants (filler duration = 2500 ms - reaction time). (B) Average accuracy on AB and CD trials per age group.

When participants have gained knowledge of the statistical regularities, they were expected to more often apply the correct rule. Notably, in probabilistic learning tasks individuals generally do not consistently apply the correct rule but show matching behavior; i.e., they choose the correct stimulus with a frequency that is proportional to the probability of positive feedback associated with that stimulus (Estes, 1961; Herrnstein, 1961; Shanks et al., 2002; Frank and Kong, 2008). Thus, we anticipated that participants would apply the correct rule (in this study, choosing the high probability stimuli A and C) more often,

but we also anticipated that they would remain exploring the alternative rule (choosing the low probability stimuli B and D). Therefore, this paradigm allowed us to investigate the processing of positive and negative feedback that carries different informative value. In particular, receiving negative feedback when choosing the correct rule should not be interpreted as a signal to switch to the alternative rule because the probability of positive feedback remains higher than for the alternative rule. In contrast, receiving negative feedback when choosing the alternative rule should lead to a switch to the correct rule. To be able to address the question how neural responses are sensitive to feedback signals in the context of learned rules, we only analyzed neural responses after participants had reached a learning plateau.

Based on prior studies, we expected that DLPFC and the parietal cortex would be sensitive to whether feedback signals required greater attention, and would contain greater informative value for performance adjustment on subsequent trials. Therefore, we expected that these regions would be engaged mostly after choosing the alternative rule (B or D), because this feedback contained learning signals for performance adjustment, independent of valence. We also examined the role of the dACC and the caudate as these regions have previously been implicated in feedback processing (Schultz, 2007; Cools, 2008; Rushworth and Behrens, 2008). We expected that the dACC would be most sensitive to negative feedback signals, particularly when indicating the need for behavioral adjustment (Kerns et al., 2004), whereas we expected that the caudate would be most sensitive to positive feedback which signals response continuation (Cools, 2008).

The second question concerned developmental differences in performance and neural activation. In prior research, developmental differences were observed between childhood and mid-adolescence, but differences between adolescence and adulthood remain unclear (Crone et al., 2008; van Duijvenvoorde et al., 2008). For this purpose, we compared behavioral and neural responses of three age groups; children (8–11 years), adolescents (13–16 years), and adults (18–22 years). Behaviorally, we predicted that differences in adaptive learning would be largest between childhood and adolescence, with refinement of learning between adolescence and adulthood (Luna and Sweeney, 2001; Crone and van der Molen, 2004; Somsen, 2007). In addition, we expected to find that these behavioral changes would be paralleled by changes in the areas involved in adaptive control (dACC, DLPFC, parietal cortex and caudate nucleus). For the fMRI analyses, we had three specific age related hypotheses based on prior studies. First, we expected an increase in differentiation in the dACC for positive and negative feedback processing with increasing age (van Duijvenvoorde et al., 2008; Velanova et al., 2008). Second, we expected an

attention-based shift in recruitment of DLPFC and the parietal cortex from positive to negative performance feedback with age. Third, we expected age differences in how learned probabilities would be associated with neural changes in feedback processing; in particular we predicted that feedback after exploring the alternative rule would be associated with developmental differences. Because of the children's putative focus on positive feedback, we expected that with increasing age there would be a decrease in activity related to processing positive feedback and an increase in activity related to processing negative feedback following selection of the alternative rule.

Finally, our paradigm allowed us to investigate age differences in adaptive behavior, that is, whether participants stay or shift on subsequent trials based on the received feedback. Besides behavioral analyses of sequential effects, we also employed exploratory sequential condition analyses to further understand the relation between neural activation and subsequent adjustment of behavior (see also Kerns et al., 2004).

7.2 Materials and Methods

7.2.1 Participants

Sixty-seven healthy right-handed paid volunteers (35 female, 32 male; ages 8–22 participated in the fMRI experiment. Age groups were based on adolescent development stage, resulting in three age groups: children (8- to 11-year-olds, $n = 18$; 9 female), mid-adolescents (13- to 16-year-olds, $n = 27$; 13 female) and young adults (18- to 22-year-olds, $n = 22$; 13 female). A chi square analysis indicated that the gender distribution was similar across age groups, $\chi^2(2) = 0.79$, $p = 0.67$. All participants reported normal or corrected-to-normal vision and participants or their caregivers indicated an absence of neurological or psychiatric impairments. Participants and their caregivers (for minors) gave informed consent for the study and all procedures were approved by the medical ethical committee of the Leiden University Medical Center. In accordance with Leiden University Medical Center policy, all anatomical scans were reviewed and cleared by the radiology department following each scan. No anomalous findings were reported.

7.2.2 Behavioral Assessment

Parents filled out the Child Behavior Check List (CBCL, Achenbach, 1991) for participants younger than 18 years, in order to screen for psychiatric conditions. All participants scored below clinical levels on all subscales of the CBCL, and had scores within 1 SD of the mean of a normative standardized sample.

Participants completed two subscales (similarities and block design) of either the Wechsler Adult Intelligence Scale (WAIS) or the Wechsler

Intelligence Scale for Children (WISC) in order to obtain an estimate of their intelligence quotient (Wechsler, 1991, 1997). There were no significant differences in estimated IQ scores between the different age groups, $F(2, 66) = 1.63, p = 0.20$ (see Table 7.1).

Table 7.1 Groups Measures

	IQ	RT (ms)	points	head motion Avg (mm)	Max (mm)
Adults	107 (2.4)	811(44)	118(3)	.08(.01)	1.56
Adolescents	108 (2.0)	773(39)	114(3)	.08(.01)	2.96
Children	111 (2.6)	804(42)	107(6)	.09(.01)	2.85

Displays means per age groups, standard errors between brackets. Final column represents the maximum head motion between two time points in each group

7.2.3 Task Procedure

The procedure for the probabilistic learning task (Frank et al., 2004) was as follows: The task consisted of two stimulus pairs (called AB and CD). The stimulus pairs consisted of pictures of everyday objects (e.g., a chair and a clock). Each trial started with the display of one of the two stimulus pairs and subsequently the participant had to choose one of the two stimuli (e.g., A or B), which were presented on the left or the right side of the screen. The stimulus pairs were presented in random order. Participants were instructed to choose either the left or the right stimulus by pressing a button with the index or middle finger of the right hand within a 2500 ms window, which was followed by a 1000 ms feedback display. The feedback display consisted of a green V-signal for positive feedback and a red cross for negative feedback. If no response was given within 2500 ms, the text “too slow” was presented on the screen. This occurred on less than 2% of the trials.

The feedback displayed was probabilistic. Choosing stimulus A led to positive feedback on 80% of AB trials, whereas choosing stimulus B led to positive feedback on 20% of these trials. The CD pair procedure was similar, but probability for positive feedback was lower; choosing stimulus C led to positive feedback on 70% of CD trials, whereas choosing stimulus D led to positive feedback on 30% in these trials. Thus, the correct choice in order to obtain most positive feedback was A or C, whereas the incorrect choice was B or D.

Participants were instructed to earn as many points as possible (as indicated by receiving a positive feedback signal), but were also informed that it would not be possible to receive positive feedback on every trial. Further, participants were informed that although stimuli sometimes appeared on the right side and sometimes on the left side, that laterality was an irrelevant dimension. After the

instructions and right before the scanning session, the participants played 40 practice rounds on a computer in a quiet laboratory to ensure proficiency on the task.

In total, the task in the scanner consisted of two blocks of 100 trials each: 50 AB trials and 50 CD trials per block. To ensure that participants had to learn a new mapping in both task blocks, the first and the second block consisted of different sets of pictures. The duration of each block was approximately 8.5 min. The stimuli were presented in pseudo-random order with a jittered interstimulus interval (min = 1000 ms, max = 6000 ms) optimized with OptSeq2 (Dale, 1999). During inter trial intervals, a central fixation cross was shown.

7.2.4 Data Acquisition

Participants were familiarized with the scanner environment on the day of the fMRI session through the use of a mock scanner, which simulated the sounds and environment of a real MRI scanner. Data were acquired using a 3.0T Philips Achieva scanner at the Leiden University Medical Center. Stimuli were projected onto a screen located at the head of the scanner bore and viewed by participants by means of a mirror mounted to the head coil assembly. First, a localizer scan was obtained for each participant. Subsequently, T2*-weighted Echo-Planar Images (EPI) (TR = 2.2 s, TE = 30 ms, 80×80 matrix, FOV = 220, 35 2.75 mm transverse slices with 0.28 mm gap) were obtained during two functional runs of 232 volumes each. The first two scans were discarded to allow for equilibration of T1 saturation effects. A high-resolution T1-weighted anatomical scan and a high-resolution T2-weighted matched-bandwidth high-resolution anatomical scan, with the same slice prescription as the EPIs, were obtained from each participant after the functional runs. Stimulus presentation and the timing of all stimuli and response events were acquired using E-Prime software. Head motion was restricted by using pillow and foam inserts that surrounded the head.

7.2.5 fMRI Data Analysis

Data were preprocessed using SPM5 (Wellcome Department of Cognitive Neurology, London). The functional time series were realigned to compensate for small head movements. Translational movement parameters never exceeded 1 voxel (<3 mm) in any direction for any subject or scan. There were no significant differences in movement parameters between age groups $F(2, 65) = 0.152$, $p = 0.85$, (see Table 7.1). Functional volumes were spatially smoothed using a 6 mm full-width half-maximum Gaussian kernel. Functional volumes were spatially normalized to EPI templates. The normalization algorithm used a 12 parameter affine transformation together with a nonlinear transformation

involving cosine basis functions and resampled the volumes to 3 mm cubic voxels. The MNI305 template was used for visualization and all results are reported in the MNI305 stereotaxic space (Cosoco et al., 1997), an approximation of Talairach space (Talairach and Tournoux, 1988).

Statistical analyses were performed on individual participants' data using the general linear model in SPM5. The fMRI time series data were modeled by a series of events convolved with a canonical haemodynamic response function (HRF). The presentation of the feedback screen was modeled as 0-duration events. The stimuli and responses were not modeled separately as these occurred in one prior or overlapping EPI images as feedback presentation.

In the model, feedback was further subdivided into correct vs. alternative rule and positive vs. negative feedback. These trial functions were used as covariates in a general linear model, along with a basic set of cosine functions that high-pass filtered the data, and a covariate for run effects. The least-squares parameter estimates of height of the best-fitting canonical HRF for each condition were used in pair-wise contrasts. The resulting contrast images, computed on a participant-by-participant basis, were submitted to group analyses. At the group level, contrasts between conditions were computed by performing one-tailed *t*-tests on these images, treating participants as a random effect. We further performed voxelwise ANOVAs to identify regions that showed age related differences in relation to feedback processing. We tested for linear increases (-1 0 1) and decreases (1 0 -1) in the contrasts specified below.

We applied AlphaSim (Ward, 2000) to calculate the appropriate threshold significance level and cluster size for the whole-brain analyses. A significance threshold of $p < 0.05$, corrected for multiple comparisons was calculated by performing 10,000 Monte Carlo simulations in AlphaSim resulting in an uncorrected threshold of $p < 0.001$, requiring a minimum of 24 voxels in a cluster. This threshold was used for all whole-brain analyses.

We used the Marsbar toolbox for use with SPM5 (Brett et al., 2002) to perform Region of Interest (ROI) analyses to further characterize patterns of activation. We created ROIs of the regions that were identified in the functional mask of whole-brain analyses. The masks used to generate functional ROIs was based on the general (positive vs. negative feedback) contrasts ($p < 0.001$, > 24 voxels) across all participants, which was unbiased for effects of probability rule or age. Because this statistical image spanned several distinct functional brain regions in the striatum, we used Marsbar anatomical masks for the caudate nucleus to further specify our ROIs.

For all ROI analyses, effects were considered significant at an α of 0.0125, based on Bonferroni correction for multiple comparisons, $p = 0.05/4$ ROIs (caudate, DLPFC, parietal cortex and dACC), unless reported otherwise.

7.4 Results

7.4.1 Performance

To investigate the age differences in learning performance for the different stimulus pairs we calculated the percentage of correct choices (choosing the high probability stimulus) per block of 20 trials for each participant, resulting in five blocks in total. Because the two runs in the scanner consisted of new stimulus pairs, the two runs were collapsed.

As expected, the age (8–11 years, 13–16 years, 18–22 years) \times probability (AB, CD) \times task block (5) ANOVA showed that participants learned to make more correct choices over time, as indicated by a main effect of task block, $F(4, 260) = 40.44$, $p < 0.001$, (See Figure 7.1B). There was a significant difference in accuracy between the two probabilities; participants were more accurate on the AB (80%–20%) trials than the CD (70%–30%) trials, $F(1, 65) = 11.58$, $p < 0.001$. Contrary to predictions, there were no age differences in learning (age \times task block interaction, $F(8, 260) = 1.38$, $p = 0.11$), no age differences in accuracy on the two pairs (age \times probability interaction, $F(2, 65) = 0.941$, $p = 0.393$), and no age \times probability \times task block interaction ($p > 0.10$). A similar ANOVA for reaction times revealed no differences for age, probability, or task block (all p 's > 0.10) (see Table 7.1).

The task block factor allowed us to obtain the point in learning where participants reached a plateau. By selecting the task phase in which there were no longer differences in learning, we could examine how feedback was processed in the context of applying the correct (choosing the stimuli with a high probability of positive feedback) or alternative rule (choosing the stimuli with a low probability of positive feedback). Follow up comparisons showed that the last 60 trials were appropriate for this purpose, as performance stabilized and participants showed probability matching behavior (Shanks et al., 2002). That is, both the AB and the CD pairs showed no effects of block (learning) on accuracy in the last three blocks, $F(2, 130) = 3.47$, $p = 0.08$ and $F(2, 130) = 1.81$, $p = 0.52$, respectively. When we reanalyzed these last 60 trials, we still found a significant effect of stimulus pair, $F(1, 65) = 16.51$, $p < 0.001$, $\eta_p^2 = 0.203$, and again no significant interactions with age (all p 's > 0.3).

To summarize, the behavioral results showed that all participants learned to perform more accurately over time and they learned faster on the easier AB trials than the more difficult CD trials. Performance stabilized in the last 60 trials, at which point participants showed probability matching behavior (Shanks et al., 2002).

The fMRI analyses focused on the last 60 trials. In order to have enough trial numbers in each condition, we collapsed across probabilities in the

analyses below. Thus, we differentiated between over-learned high probabilities (A and C collapsed) and alternative low probabilities (B and D trials collapsed). These will be referred to as the correct and alternative rules. Each of these rules could result in positive and negative feedback.

Table 7.2. : Brain Regions revealed by whole brain contrasts.

Anatomical region	L/R	voxel volume	Z	MNI coordinates		
				x	y	z
Positive > Negative						
Striatum (ventral and dorsal)	L/R	774	7.49	-6	12	-3
Dorsolateral prefrontal cortex	L	71	4.61	-27	24	51
Superior parietal cortex	L	170	4.23	-30	-75	48
Precuneus	L/R	137	4.07	-3	-36	33
Ventral Medial PFC	L/R	26	4.03	3	54	-12
Visual Cortex	L/R	332	4.50	27	-93	-9
Negative > Positive						
Dorsal Anterior Cingulate Cortex	L/R	63	4.43	9	21	36

MNI coordinators for main effects, peak voxels reported at $p < .001$, at least 24 contiguous voxels.

7.4.2 fMRI Results Positive Versus Negative Feedback

Whole-brain comparisons across age groups

First, we identified the neural correlates of feedback processing by comparing the (positive feedback vs. negative feedback) contrast across all participants. This analysis revealed increased BOLD responses for positive feedback > negative feedback in several regions including the left and right caudate, left DLPFC and left parietal cortex (see Figure 7.2A). The opposite contrast (negative > positive feedback) resulted in increased activation in the dACC. The coordinates for these comparisons (positive feedback vs. negative feedback) are reported in Table 7.2.

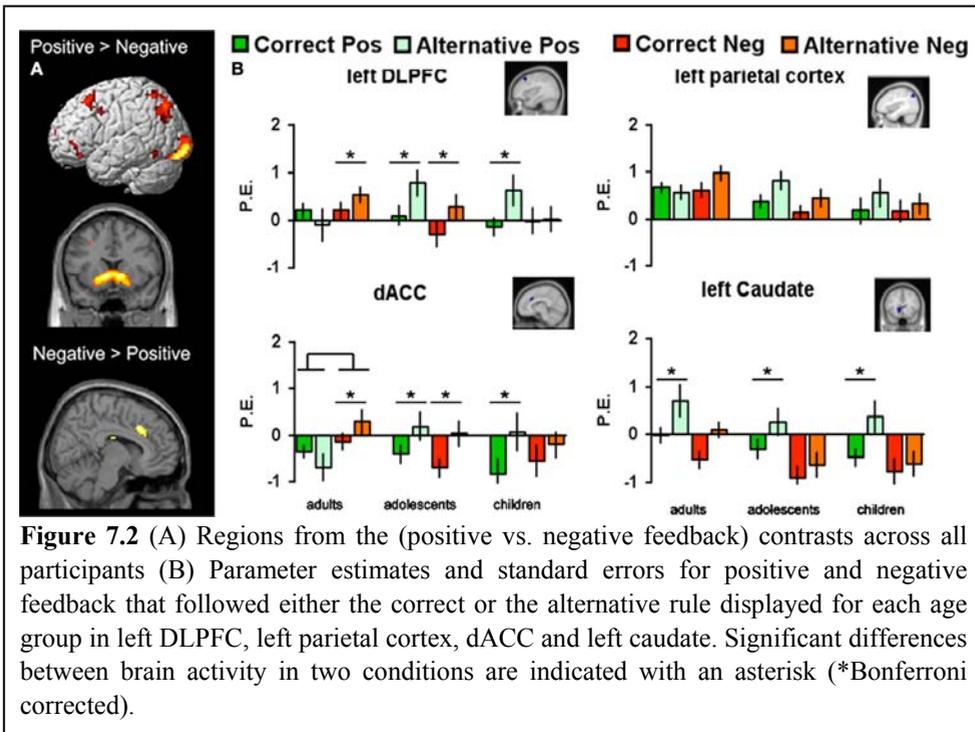


Figure 7.2 (A) Regions from the (positive vs. negative feedback) contrasts across all participants (B) Parameter estimates and standard errors for positive and negative feedback that followed either the correct or the alternative rule displayed for each age group in left DLPFC, left parietal cortex, dACC and left caudate. Significant differences between brain activity in two conditions are indicated with an asterisk (*Bonferroni corrected).

7.4.3 fMRI Region of Interest Results for Feedback \times Rule \times Age Group Interactions

Next, we tested for age differences and rule sensitivity in these regions by performing region of interest (ROI) analyses. The ROI analyses were restricted to the four *a priori* defined regions which emerged in the (positive vs. negative) contrast across participants: bilateral caudate, left DLPFC, left parietal cortex and dACC. In order to investigate whether there were age differences in how the statistical regularities learned by the participants had an effect on how feedback was processed we performed $3 \times 2 \times 2$ ANOVAs testing for the interaction between valence (positive vs. negative) and rule (correct vs. alternative) as within-subjects factors and age (children, adolescents, adults) as the between-subjects factor for each ROI (see Figure 7.2B).

Left DLPFC. The (age group \times valence \times rule) ANOVA for left DLPFC resulted in an interaction between valence and rule, $F(2, 64) = 6.32, p < 0.01$, showing that left DLPFC was more active for both negative and positive feedback after choosing the alternative rule compared to the correct rule, but this difference was larger for positive than negative feedback. In addition, there was an interaction between rule (AC vs BD) and age group, $F(2, 64) = 3.87, p =$

0.02, and a three-way interaction between rule, valence, and age group, $F(2, 64) = 6.77, p < 0.01$.

As can be seen in Figure 7.2B, children and adolescents showed more activity for positive feedback after choosing the alternative rule compared to the correct rule ($t(17) = 2.64, p < 0.01$ and $t(26) = 3.18, p < 0.004$, respectively), whereas this difference was not present in adults. In addition, adults and adolescents showed more activity for negative feedback after choosing the alternative rule compared to the correct rule, ($t(21) = -2.49, p = 0.02$ and $t(23) = -2.81, p < 0.01$ respectively), but this difference was not present in children.

Left parietal cortex. The (age group \times valence \times rule) ANOVA for the left parietal cortex revealed a similar three-way interaction which approached significance, $F(2, 64) = 3.16, p = 0.05$ (see Figure 7.2B). Although the pattern of activation for the different conditions in the left parietal cortex appears similar to the pattern for left DLFPC, it did not survive Bonferroni correction and none of the *post hoc* comparisons resulted in significant effects.

dACC. The (age group \times valence \times rule) ANOVA for the dACC resulted in a rule \times valence interaction, $F(2, 64) = 14.14, p < 0.001$, an age \times valence interaction, $F(2, 64) = 4.11, p < 0.01$, and an age \times rule interaction, $F(2, 64) = 4.81, p = 0.03$, but the three-way interaction failed to reach significance $F(2, 64) = 0.28, p = 0.75$.

As can be seen in Figure 7.2B, adults showed more activation in dACC after negative feedback than after positive feedback, $F(1, 21) = 8.25, p < 0.01$, but this was not found for the younger age groups. Children and adolescence, in contrast, showed more dACC activation after positive feedback for the alternative rule relative to the correct rule ($t(17) = 2.51, p < 0.01$ and $t(26) = 3.44, p < 0.01$ respectively). In addition, adults and adolescents showed more activity for negative feedback after choosing the alternative rule compared to the correct rule, ($t(21) = -2.89, p < 0.01$ and $t(26) = -3.32, p < 0.003$ respectively), but this difference was not present in children.

Left and right caudate. Finally, we performed an (age group \times valence \times rule) ANOVA for the left caudate nucleus. This analyses did not reveal any age effects, but a main effect for feedback, $F(1, 64) = 33.17, p < 0.001$, and a feedback \times rule interaction $F(2, 64) = 17.21, p < 0.01$. All age groups showed more activity for the alternative (low probability) compared to the correct rule (high probability) positive feedback (all p 's < 0.001), but there were no additional main or interaction effects (Figure 7.2B). Similar analyses for right

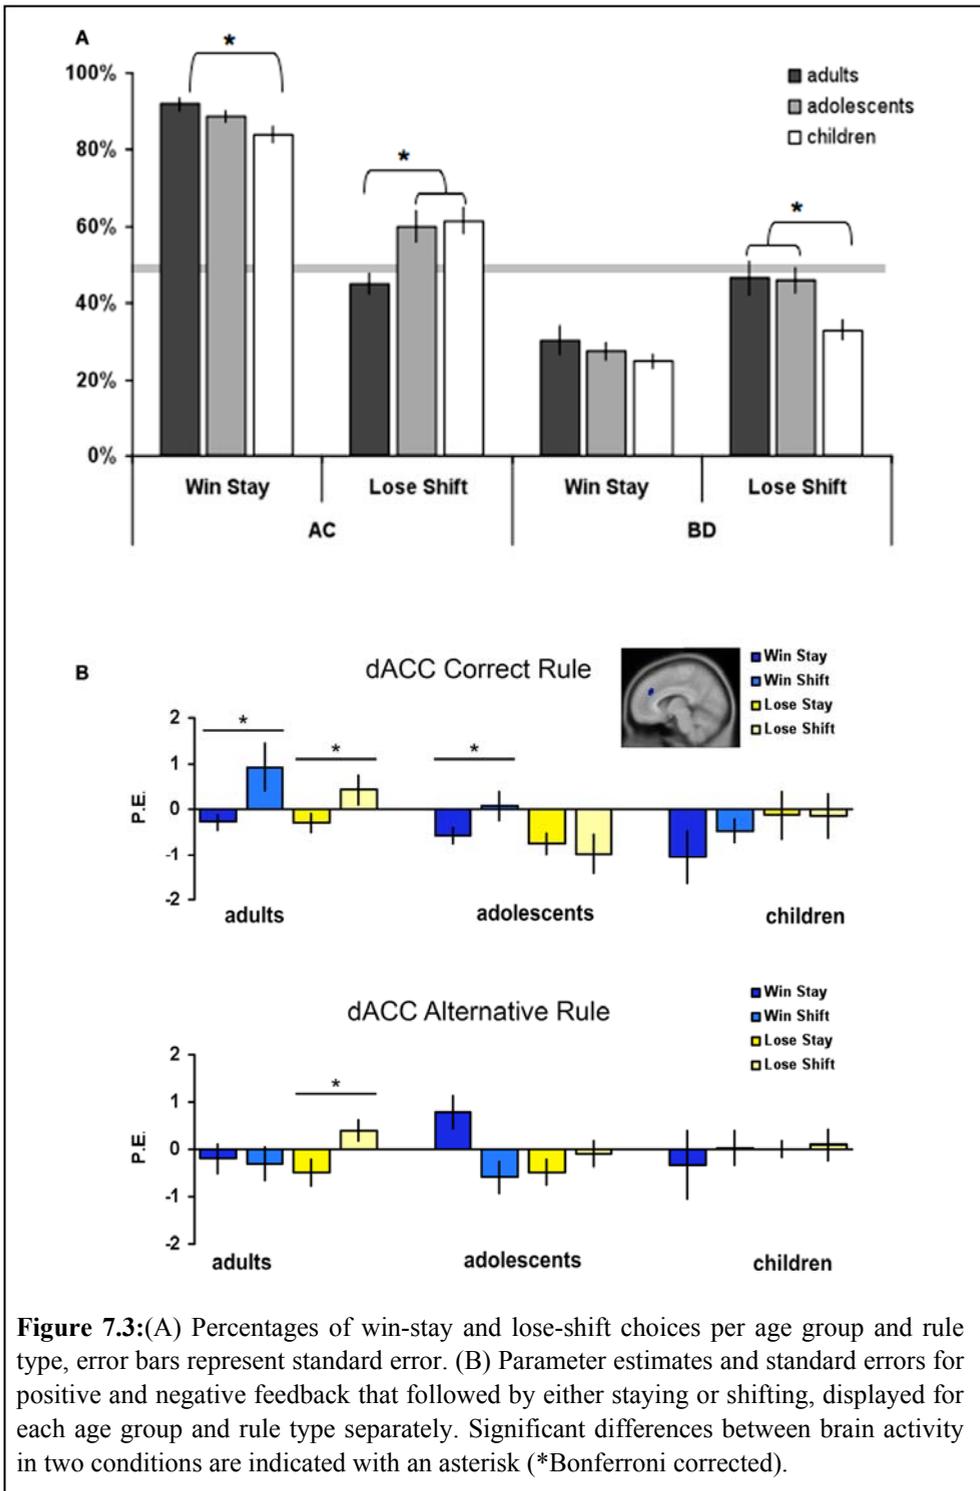
caudate yielded the same results; a main effect of feedback, $F(1, 64) = 28.16$, $p < 0.005$, and a feedback \times rule interaction $F(2, 64) = 19.33$, $p < 0.01$.

7.4.4 Win Stay – Lose Shift Strategies: Behavior and Brain Analyses

Finally, to further investigate differences in feedback processing we explored developmental changes in decision-making strategies on the behavioral and neural level. In order to investigate the strategy used on the task we examined how often participants chose either the same stimulus after positive feedback (win-stay) or the other stimulus after negative feedback (lose-shift). For this set of analyses we further broke down the trials based on the subsequent choice when presented with the same stimulus pair; win-stay, win-shift, lose-stay and lose-shift. The factor ‘win-stay’ was computed by calculating the proportion of choice repetitions following positive feedback as a function of the total number of positive feedback events. Likewise, the factor ‘lose-shift’ was computed by calculating the proportion of choice shifts following negative feedback as a function of the total number of negative feedback events. Because previous analyses revealed that positive and negative feedback were processed differently dependent on rule type we analyzed the sequential effects for the correct and alternative rule separately.

Task Strategy. For correct rules, the univariate ANOVAs with age group as the between-subjects factor revealed a significant age difference in lose-shift strategies, $F(2, 64) = 4.04$, $p < 0.02$ as well as in win-stay strategies, $F(2, 64) = 4.51$, $p < 0.02$ (see Figure 7.3A). These results illustrate that adults showed more optimizing behavior than adolescents and children; they stayed more often with the correct rule after positive feedback and shifted less often after negative feedback.

For the alternative rules, the univariate ANOVAs revealed no age differences for win-stay strategies, $F(2, 64) = 0.85$, $p = 0.43$, but there was a significant age difference in lose-shift strategies, $F(2, 64) = 3.91$, $p < 0.03$. In the latter case, children showed less optimal behavior compared to the adolescents and adults; surprisingly, they stayed more often with the alternative (incorrect) rule after negative feedback.



ROI analyses. In order to explore the relation between brain activity and behavior on the subsequent trial, we compared brain activity after positive and negative feedback that resulted in staying or shifting for the two rule types separately. We explored the same ROIs as reported above. These analyses revealed significant shift and age effects only in the dACC and left DLPFC, but not in the caudate or the parietal cortex. In general, the ANOVAs showed that in adults, dACC and DLPFC were more active when participants shifted on the next trial. There were some differences in significance levels, but overall this effect seemed generally independent of feedback valence or rule. The analyses are described in more detail below.

The dACC showed the strongest relation between brain activity and subsequent behavioral change. When applying the correct rule, the shift \times age group ANOVA for positive feedback revealed a main effect of shifting, $F(1, 65) = 6.27, p < 0.01$ but no interaction with age, $F(2, 64) = 2.29, p = 0.11$ (see Figure 7.3B). There was more dACC activity when shifting after positive feedback. The same ANOVA for negative feedback revealed an age \times shift interaction, $F(2, 64) = 3.62, p = 0.03$. *Post hoc* comparisons revealed that there was more dACC activity when shifting compared to staying after negative feedback for adults ($t(21) = -2.76, p < 0.01$) but not for the adolescents and children (both p 's > 0.1).

When applying the alternative rule, the shift \times age group ANOVA for positive feedback revealed no significant effects of age or shifting. However, the same ANOVA for negative feedback revealed an age \times shift interaction ($F(2, 63) = 5.31, p < 0.01$). *Post hoc* comparisons revealed that there was more dACC activity when shifting after negative feedback for adults ($t(21) = -3.01, p < 0.01$) but not for adolescents and children (both p 's > 0.2).

Finally, the pattern of activation in the left DLPFC appeared similar to that of the dACC (Figure S7.2 in Supplementary Material). The shift \times age ANOVAs for the correct rule resulted in significant shift \times age interactions for both positive and negative feedback ($F(2, 63) = 4.46, p = 0.03$ and $F(2, 64) = 4.91, p = 0.02$, respectively). *Post hoc* test revealed that there was more left DLPFC activity when shifting on the next trial after positive and negative feedback, but this was only significant for the adults ($t(21) = -2.54, p < 0.01$ and $t(21) = -2.32, p = 0.03$, respectively). There were no significant effects for the alternative rule (all p 's > 0.2).

7.4 Discussion

The goal of this study was to examine the neural developmental changes when processing positive and negative feedback signals in a probabilistic

decision-making task. As predicted, all participants learned to choose the correct rules (high probability stimuli A and C) more often than the alternative rules (low probability stimuli B and D) (Frank et al., 2004; Klein et al., 2007). After approximately 40 trials, participants adapted a performance pattern consistent with ‘probability matching behavior’, and this behavioral phase was the focus of our further analyses.

Behavioral analyses showed two important patterns: (1) probability matching behavior occurred in all age groups, but there were no age differences in overall learning rate, and (2) task adaptive win-stay, lose-shift strategies were observed, but age differences in adaptive behavior indicated more task-adaptive optimizing behavior in adults. These task and age differences in decision-making strategy were paralleled by changes in functional brain activity; (1) neural responses in DLPFC, dACC, and caudate were sensitive to rule \times feedback interactions and an age related difference was observed in DLPFC and dACC, and (2) activity in DLPFC and dACC predicted behavioral change on subsequent trials more strongly in adults than in adolescents and children. These behavioral data and their neural correlates provide important new insights in feedback processing in general and across development. The discussion will be organized according to these themes.

7.4.1 Feedback processing in adults

Our analysis of positive and negative feedback processing in a probabilistic environment demonstrated that feedback-related activity in the DLPFC, dACC and caudate was dependent on valence and information value. We started out with a general whole-brain comparison for positive versus negative feedback and used ROI analyses to explore the areas identified in this contrast. This analysis revealed that especially left DLPFC, dACC and bilateral caudate were sensitive to feedback \times rule context interactions. Before interpreting age differences in these activation patterns, we start out with the interpretation of feedback sensitivity observed in adults, which will set the stage for interpreting the developmental effects.

When exploring the data for adults separately, the results showed increased recruitment of DLPFC after receiving negative feedback following the alternative compared to the correct rule. Given that negative feedback after choosing the alternative, but not the correct, rule indicates the need for a switch in behavior, the adult findings are consistent with previous studies demonstrating negative feedback-related sensitivity in DLPFC for feedback that is important for subsequent behavioral adjustment (Kerns, 2006; van Duijvenvoorde et al., 2008; Zanolie et al., 2008) and not for negative feedback *per se*.

Besides DLPFC, the parietal cortex has previously been implicated in feedback processing (Crone et al., 2008, van Duijvenvoorde et al., 2008) and implementing cognitive control as part of the fronto-parietal network (Brass et al., 2005; Bunge et al., 2002; Dosenbach et al., 2008). In support of this hypothesis our whole-brain analyses revealed that the left superior parietal cortex was involved in feedback processing. However, in contrast with previous studies (van Duijvenvoorde et al., 2008), our subsequent *post hoc* analyses could not confirm a strong contribution of the superior parietal cortex. Possibly, the parietal cortex was more engaged in prior studies because these involved trial-to-trial learning, whereas in the current study we investigated feedback processing when rules were already learned. Future research is necessary to elucidate the role of the superior parietal cortex in feedback processing in relation to learning.

The analyses of dACC revealed a very similar activation pattern as DLPFC, however the dACC activation pattern in adults was more supportive of a general increase in activity after negative feedback regardless of rule type. Possibly, this finding indicates that, at least in adults, the dACC has a more general role in processing negative feedback; both in terms of detecting general conflict (Brown and Braver, 2005) and signaling the need for behavior change (Holroyd and Coles, 2008; Rushworth, 2008).

Finally, the caudate nucleus also showed sensitivity to feedback and rule type, but this region was more active after positive compared to negative feedback when participants chose the alternative rule. Given that this effect was specific for positive feedback, and that the probability for positive feedback for the alternative rule was low, the signal in the caudate could reflect a positive prediction error; i.e., signaling that the outcome is better than predicted (for review see Schultz, 2007).

Together, analysis of the adult activation pattern confirms prior findings showing that DLPFC and dACC are sensitive to negative feedback and the caudate is sensitive to positive feedback, but the findings further elucidate that these neural responses are dependent on the extent to which these feedback signals provide a learning signal of future performance. That is, DLPFC and caudate responses were more pronounced after selecting the incorrect rule which had a low probability of resulting in positive feedback, but which may have been important to explore. In contrast, when applying over-learned high probability rules, DLPFC and caudate were less involved, possibly because the informative value was smaller.

7.4.2 Feedback Processing: Developmental Comparisons

The neural activation patterns described above were differentially sensitive to age modulations. The first notable finding is that of differential activation patterns in the DLPFC. All participants, regardless of age, showed increased recruitment of DLPFC when choosing the alternative rule compared to the correct rule. However, children, but not adults, showed more activation in DLPFC after positive feedback when choosing the alternative rule. In contrast, adults, but not children, showed more activation in DLPFC after negative feedback when choosing the alternative rule. Adolescents seemed to be in a transition phase, because their neural response to positive feedback was similar to that observed in children, but their neural response to negative feedback was similar to that observed in adults. Thus, consistent with prior studies, these developmental differences indicate a shift from focus on positive to a focus on negative feedback with age (Somsen, 2007; Crone et al., 2008; van Duijvenvoorde et al., 2008), which appears to continue across adolescence. In addition, the current results extend previous findings by showing that developmental differences in neural responses to feedback are not related to valence *per se*, but suggest an age related change in processing learning signals with different informative value.

In contrast, for all age groups the caudate nucleus was more active for positive compared to negative feedback, in particular when participants chose the alternative rule. This finding indicates that part of the feedback processing network, which is implicated in processing statistical regularities of reward (Schultz, 2007) matures already at an early age, whereas the part of the network that is involved in processing negative feedback and the subsequent control of behavior has a more protracted developmental time course. These findings are consistent with prior reports using cognitive tasks, as these studies have also reported early maturation of subcortical regions and protracted development of cortical brain areas (Casey et al., 2004; van Duijvenvoorde et al., 2008; Velanova et al., 2008). It should be noted that other developmental studies have reported increased sensitivity of the striatum in early adolescence, however, these studies have employed paradigms with a more affective content, such as gambling tasks with real monetary rewards or emotion recognition (Ernst et al., 2005; Galvan et al., 2006; McClure-Tone et al., 2008; van Leijenhorst et al., 2009). In future studies, it will be of interest to examine whether the caudate activation can be modulated by the use of affective task modulations when learning rules or processing performance feedback.

7.4.3 Adaptive Behavior and Brain Activation across Development

One of the challenging questions for future studies is how the neural activation is associated with trial-to-trial learning. For example, we did not observe age differences in general learning performance, despite differences in neural activation. This was unexpected, and again demonstrates that differences in neural activation can be present without differences in observable behavior (Ladouceur et al., 2004). However, consistent with prior studies, the sequential analyses revealed that with age, participants became better at using the negative feedback signals to adjust their behavior on subsequent trials (Crone and van der Molen, 2004). As expected, when receiving positive feedback after having applied the correct rule, participants were more likely to stay and select the same stimulus on the subsequent trial. Likewise, when receiving negative feedback after having applied the incorrect alternative rule, participants were more likely to shift and select the correct stimulus on the subsequent trial. Overall, adults appeared better at optimizing than adolescents, and adolescents performed better than children. Based on these findings, in combination with the developmental differences in neural activation, the data are supportive of a linear increase across adolescence. Although these findings differ from earlier reports which have showed larger differences in early adolescence than in later adolescence (e.g. Ladouceur et al. 2004) the findings are consistent with prior fMRI results showing late changes in brain activation and behavior (e.g. Scherf et al., 2006; van Duijvenvoorde et al., 2008).

Intriguingly, even though children were more likely than adults to shift after receiving negative feedback when applying the correct rule, they were also more likely to stay after receiving negative feedback when applying the incorrect alternative rule. The reason for this behavioral pattern is still unclear, but it is possible that children waited with shifting when applying the incorrect alternative rule until they received positive feedback (20%). Future research should use task manipulations that allow for further investigation of this hypothesis.

We performed exploratory analyses to investigate the relation between brain activity and win-stay, lose-shift behavior, although it should be noted that these analyses are preliminary as our study design was not optimized to test for these differences. The analyses on the ROIs identified in the main analyses revealed that, consistent with prior research, dACC and left DLPFC activity predicted behavioral adjustment on the subsequent trial in adults (Kerns et al., 2004; Jocham et al., 2009). However, this pattern was observed for both rule types and appeared independent of feedback valence. Possibly, the dACC and left DLPFC were important for trial-by-trial adjustment (Kerns et al., 2004). We found a similar pattern in adolescents, but only when applying the correct rule. We

failed to find similar relations in children, which may indicate that the neural mechanisms that facilitate future behavioral adjustment are still immature or that they employed different strategies to perform the task. These interpretations are consistent with an ERP study showing increased error related negativity across adolescence (Ladouceur et al., 2007). Furthermore, the same study showed that only in adults the ERN amplitude was related to task performance.

The current study is limited by the relatively small number of trials for some of the contrasts examining the neural correlates of shifting behavior. Future studies should make use of tasks that are optimized for studying these developmental differences in more detail.

In addition, a challenging direction for future research will be to investigate the developmental differences in the learning phase. The combined use of computational reinforcement learning models (Klein et al., 2007) with imaging techniques could be a promising endeavor to parse out the developmental changes in different phases of learning (e.g. learning rate) and their neural correlates. These methods could be combined with trial-to-trial data categorization to understand how the observed developmental change in sensitivity from positive to negative feedback hinders or facilitates learning locally versus oriented towards future goals.

Conclusion

Taken together, the current findings confirm that DLPFC, dACC and caudate are important for probabilistic feedback processing, and show that they have dissociable roles as reflected in differential sensitivity to feedback valence and rule types. The DLPFC and dACC were sensitive to information value in response to negative feedback, but the caudate was sensitive to information value in response to positive feedback. These findings are consistent with previously suggested computational models of feedback learning (Cohen, 2008; Frank and Kong, 2008).

The results of this study replicate the previously reported developmental shift in sensitivity from positive to negative feedback as reflected in neural activation in the DLPFC, with a transition phase in adolescence. Using probabilistic feedback stimuli, we could dissociate between two competing hypotheses with respect to this developmental change. The results confirm the hypothesis that this shift is associated with different attention focus on learning signals and disconfirm the hypothesis that this shift reflects a simple valence effect. Further understanding of the age related changes in strategy differences, and how to influence decision-making strategies by guiding attention regulation, promise to be useful sources to improve learning behavior of children and adolescents.

7.5 Supplementary Material

7.5.1 Additional tests for Feedback x Rule and Rule x Age groups interactions

The ROI analyses presented in the manuscript suggest that neural responses to feedback valence are modulated by rule selection and age. Additionally, we performed whole-brain ANOVAs testing for interactions between valence (positive and negative), rule (correct high probability vs. alternative low probability) and age in order to explore whether additional regions were sensitive to these interactions. The whole brain ANOVA and subsequent ROI analyses of age related changes on effects of rule choice further supported the hypothesis of a shift in focus from positive feedback to negative feedback from childhood to adulthood in the DLPFC when choosing the alternative rule (see Figure S7.1).

Rule & Valence

The first ANOVA was performed to test for regions that were sensitive to the rule x valence interaction across participants. This analysis revealed a single region in the right DLPFC (BA 9, MNI: [45, 39, 30], see Figure S7.1A). This region was further explored by extracting the ROI and was found to be more active for positive feedback following the alternative rule compared to positive feedback following the correct rule, $t(58) = 4.30$, $p < .001$, and was also more active for negative feedback following the alternative rule compared to negative feedback following the correct rule, $t(63) = -3.93$, $p < .001$. In addition, a comparison of positive and negative feedback for the alternative rule indicated that the neural response was enlarged for positive feedback signals, $t(58) = 2.08$, $p < .05$.

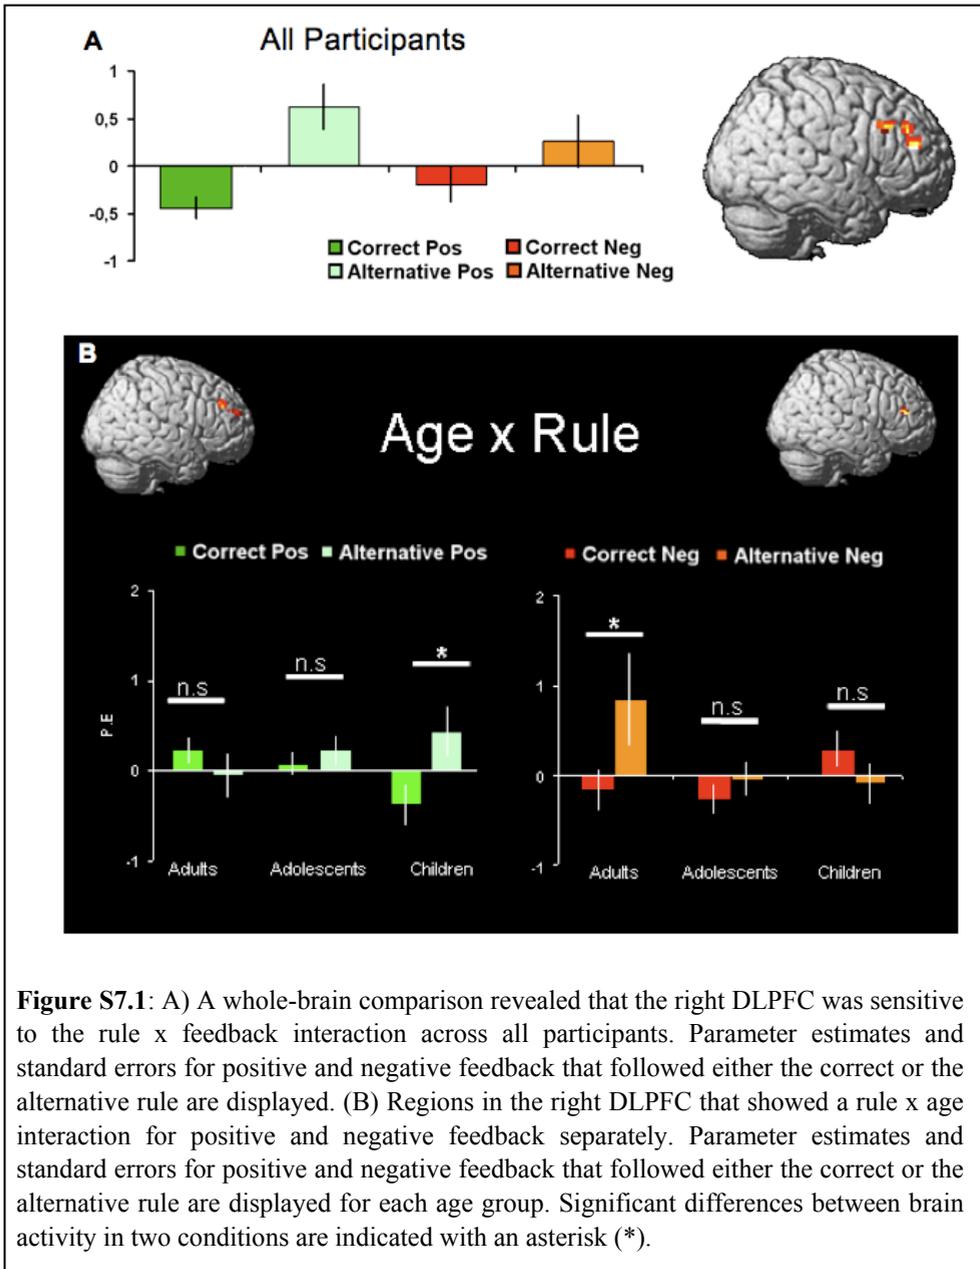


Figure S7.1: A) A whole-brain comparison revealed that the right DLPFC was sensitive to the rule x feedback interaction across all participants. Parameter estimates and standard errors for positive and negative feedback that followed either the correct or the alternative rule are displayed. (B) Regions in the right DLPFC that showed a rule x age interaction for positive and negative feedback separately. Parameter estimates and standard errors for positive and negative feedback that followed either the correct or the alternative rule are displayed for each age group. Significant differences between brain activity in two conditions are indicated with an asterisk (*).

Age x Rule for positive and negative feedback

Next, we tested for age differences by performing whole-brain ANOVAs with age group as between participants factor, testing for both linear increases [-1 0 1] as well as decreases [1 0 -1] with age. Given that differences in feedback processing were expected to differentiate based on the rule that was applied (correct vs. alternative), we tested for age differences in processing positive and

negative feedback separately. In the rule x age ANOVA for positive feedback, the decrease contrast [1 0 -1] revealed an age related change in right DLPFC (BA 9, MNI: [39, 27, 17] Figure S7.1B). The ROI of this region was extracted to test for pattern differences. Post hoc comparisons showed more activation for positive feedback after the alternative rule compared to the correct rule for children only, $t(17)=2.93$, $p<.01$. In contrast, adults and adolescents did not show differences in processing positive feedback following the two rules (both p 's $>.2$). No regions were detected for the increasing age contrast [-1 0 1].

The same whole-brain rule x age ANOVA was performed for negative feedback. Here, the increasing age contrast [-1 0 1] revealed a slightly lower area in the right DLPFC (BA 46, MNI: [39, 27, 17], Figure 7.3B). Subsequent post hoc comparisons for the ROI which was extracted of this region revealed increased activity for negative feedback after the alternative compared to the correct rule for adults only, $t(19)=-2.45$, $p<.01$. The adolescents and children did not show any effect of rule choice on negative feedback (both p 's $>.1$). No regions were detected for the decreasing age contrast [1 0 -1].

In sum, the whole brain ANOVA analyses of age related changes in effects of rule choice further supported the hypothesis of a shift in focus on positive feedback to negative feedback from childhood to adulthood in the DLPFC when choosing the alternative rule. Notably, the regions which were identified in this set of ANOVAs were right lateralized. Even though the post hoc comparisons of left and right DLPFC resulted in similar activation patterns, we interpreted this difference as right DLPFC being relatively more sensitive to rule context, and left DLPFC to feedback valence.

