One learns by experience

A Life-span View on Learning and Error Processing

Dissertation
zur Erlangung des Grades eines
Doktors der Philosophie
der Philosophischen Fakultäten
der Universität des Saarlandes

vorgelegt von

Benjamin Eppinger
aus Sindelfingen

Saarbrücken, 2008

Dekan:

Prof. Dr. Rainer Krause, Universität des Saarlandes

Berichterstatter:

Prof. Dr. J. Kray, Universität des Saarlandes

Prof. Dr. G. Kerkhoff, Universität des Saarlandes

Tag der Disputation: 20. 02. 2008

To Emilia and Caramel, my beloved monsters in heaven. I miss you.

Overview

Figure Index

Table Index

I Theoretical Part

1 Introduction

"One learns by experience." What this proverb implies is that we learn from positive, as well as negative experiences, from our mistakes or errors, as well as from our correct choices. Whether, under certain conditions, we might learn more from the negative or the positive outcomes of our decisions (or vice versa) is a matter of current debate (Frank, Seeberger, & O'Reilly, 2004; Holroyd & Coles, 2002). Recent findings point to the view that there are considerable individual differences in the way participants learn from their errors and correct choices, and that these learning biases seem to be associated with alterations in dopamine levels (Frank, Woroch, and Curran, 2005). Moreover, there are several findings that indicate that the role of errors for learning changes as a function of lifespan development. On the one hand, there is data that point to the view that children are more sensitive to errors and error feedback during learning than younger adults, (see Crone, Jennings, & van der Molen, 2004; Crone, Somsen, Zanolie, & van der Molen, 2006; Crone & van der Molen, 2007; van Meel, Oosterlaan, Heslenfeld, & Sergeant, 2005). On the other hand, there is considerable evidence for the view that older adults are impaired in error processing and in the processing of error feedback during learning (Band & Kok, 2000; Falkenstein Hoormann, & Hohnsbein, 2001; Mathewson Dywan, Segalowitz, 2005, Nieuwenhuis et al., 2002; Themanson Hillman, & Curtin, 2006).

The basic theoretical idea of the present work is to combine neurophysiological models of reward and reinforcement learning (for reviews see Schultz, 2000; 2002; 2006; 2007) with theoretical accounts and empirical findings on error processing (Holroyd & Coles, 2002; Nieuwenhuis et al., 2002). In order to investigate the role of error processing for learning the present work will adopt an event-related potential (ERP) approach, which allows to investigate brain activity during error commission and the processing of external error feedback during learning with a high temporal resolution.

The aim of this work is to investigate how children and older adults differ from younger adults in the way they process errors and negative feedback during learning. The empirical part of this work consists of three ERP experiments. In the first experiment, a probabilistic learning task has been applied to study the effects of aging on learning and error processing. In the second experiment, this learning task was adapted, so that it could be used to study learning and error processing in children. The third experiment was based on the findings of the first experiment, but takes a slightly different approach. For this experiment, a learning task was developed that allowed to investigate individual differences in the way younger and older adults learn from their errors and correct choices.

The theoretical background of the present research could be broadly separated into two modules. The first module consists of neurophysiological models of reinforcement learning, which have been developed based on electrophysiological findings in primates and functional imaging findings in humans (for reviews see Montague, Hyman, & Cohen, 2004, Schultz, 2007, Wise, 2004). These models assume that reinforcement learning depends on learning signals from the dopamine system, which are projected to subcortical and particularly to prefrontal target areas. Moreover, there are several theoretical accounts and empirical findings in this field that show that the dopamine system and especially its projections to the prefrontal cortex are subject to pronounced changes during childhood development, as well as aging (Braver & Barch, 2002; Diamond, 1996; Diamond, Briand, Fossella, & Gehlbach, 2004; Goldman-Rakic & Brown, 1982; Bäckman, Nyberg, Lindenberger, Li, & Farde, 2006).

The second module consists of electrophysiological findings and neuro-computational models on error processing in humans (Botvinick, Braver, Barch, Carter, & Cohen, 2001; Coles, Scheffers, & Holroyd, 1998; Holroyd & Coles, 2002; Yeung, Botvinick, & Cohen, 2004). This field of research is mainly based on the observation of the error-related negativity (ERN), a component of the event-related potential (ERP) that can be observed when participants commit errors on reaction time tasks (Falkenstein Hohnsbein, & Hoormann, 1990; Gehring, Goss, Coles, Meyer, & Donchin, 1993). The amplitude of the ERN has been suggested to increase with learning (Holroyd & Coles,

2002) and has been shown to be affected by changes in dopamine levels (de Bruijn, Hulstijn, Verkes, Ruigt, & Sabbe, 2004; Zirnheld et al., 2004). Furthermore, there is strong evidence that error processing and its ERP correlate the ERN change as a function of lifespan development (Band & Kok, 2000; Davies, Segalowitz, & Gavin, 2004; Falkenstein, Hoormann, & Hohnsbein, 2001; Nieuwenhuis et al., 2002; Santesso, Segalowitz, & Schmidt, 2006; West, 2004).

2 Literature review

2.1 *Overview*

According to the two modules of the theoretical background that were outlined above, the literature review is structured into two main sections. In the first section, I will review the literature on the basic properties of the dopamine system and how dopamine is implicated in reinforcement learning. One focus will be on the subcortical dopamineric structures involved in reinforcement learning and their projections to the prefrontal cortex (PFC). The second focus will be on how the dopamine system and its projections to the PFC change from childhood to older age. In the second section I will review the current literature on error processing and its ERP correlate, the error-related negativity (ERN). The main emphasis in this section will be on how the error processing changes during lifespan development. The two sections will end in the description of a neurocomputational model that aims at integrating the role of dopamine for learning with the error processing system that elicits the ERN (Holroyd & Coles, 2002).

Coming from this theoretical position, in the synopsis, I will outline how transformations in the dopamine system during lifespan development might be associated with age-related changes in learning and error processing as reflected in the ERN.

2.2 *The Neurophysiological Basis of Reinforcement Learning*

2.2.1 *A Definition of Reinforcement Learning*

In order to better understand the theoretical background of the neurophysiological models on which this work is based, it is necessary to briefly recourse to the early theories on this issue and to define reinforcement learning. One early definition of reinforcement learning has been formulated by Edward Thorndike (1911) in the law of effect:

> *Of several responses made to the same situation, those which are accompanied or closely followed by satisfaction to the animal will,…, be more likely to recur; those which are accompanied or closely followed by discomfort to the animal,…, will be less likely to occur. (p. 244)*

In other word, the law of effect proposes that a certain behavior is learned if it is followed by a positive, rewarding outcome, whereas it will be avoided if a negative, punishing outcome follows it. One of the classical paradigms that have been used to study reinforcement learning is the operant (or instrumental) conditioning task (see Skinner, 1938). In this task a certain behavior (e.g., a lever press) is paired with a reward (e.g., a drop juice in case of a monkey). After training the monkey will show the now conditioned response (lever press) even in absence of the reward. In contrast to the operant conditioning task, in which reward is contingent upon a certain behavior, in the classical (or Pavlovian) conditioning task (Pavlov, 1898) reward is not contingent upon behavior, but with respect to environmental events.

Thorndike’s and Skinner’s work on instrumental conditioning formed the basis of the Behaviorism, which was probably the most influential psychological research paradigm from the 1920s to the 1960s. However, with the emergence of the cognitivism in the 1960s the behaviorist learning theories became less influential and psychological research on learning focused more on how information is processed and represented (Piaget, 1996; Rumelhart, 1981). Yet, the basic principles of reinforcement learning were picked up in the field of computer science and machine learning and led to the development of computational models of reinforcement learning (see Rescorla & Wagner, 1972; Sutton & Barto, 1990). The central assumption of these models is that reinforcement learning is driven by expectancy violations. Thus, these models suggest that actions are learned if they lead to an outcome that is better than predicted, whereas actions will be avoided if they are followed by outcomes that are worse than expected.

The integration of these models with findings on the role of dopamine in reward processing and reinforcement learning led to the neurophysiological models on which this study is based on (see Schultz, Dayan, & Montague, 1997 Waelti, Dickinson, & Schultz, 2001).

2.2.2 The Role of Dopamine for Reinforcement Learning

In the last 15 years tremendous progresses have been made in the research on the neurophysiological basis of reinforcement learning. The groundbreaking work in this field has been performed by Wolfram Schultz and colleagues (for reviews see Schultz, 2000; 2002; 2006; 2007). Using electrophysiological recordings from subcortical dopamine neurons in primates, they showed that the neurotransmitter dopamine plays a central role during reinforcement learning (Mirenowicz & Schultz, 1996). Moreover, they were able to integrate their findings with computational models on reinforcement learning (Sutton & Barto, 1990) and showed that the firing rates of the subcortical dopaminergic neurons could be predicted on the basis of these models (see Schultz et al., 1997; Waelti et al., 2001). The objective of this section is to give an overview on these neurophysiological models of reinforcement learning and to review the literature on lifespan developmental changes in the mesencephalic dopamine system.

The mesencephalic dopamine system (MDS) refers a network of different cortical and subcortical areas (involving the basal ganglia) that receive projections from dopamine neurons, which are mainly located in the pars compacta of the substantia nigra and the ventral tegmental area (VTA) (see Ungless, 2004 and see also Figure 1). The dopaminergic neurons in the VTA seem to operate on different time scales ranging from phasic short-latency responses (100 - 300 ms) to more tonic (continuous) activity (see Schultz, 2002; 2007). In the present research the focus will be on the phasic responses of these neurons, which have been shown to play an important role in reward-driven learning (for reviews, see Schultz, 2000; 2002; 2007; Wise, 2004). For a discussion of the more tonic aspects of dopamine see Niv, Daw, Joel, and Dayan (2007) and Schultz, (2007).

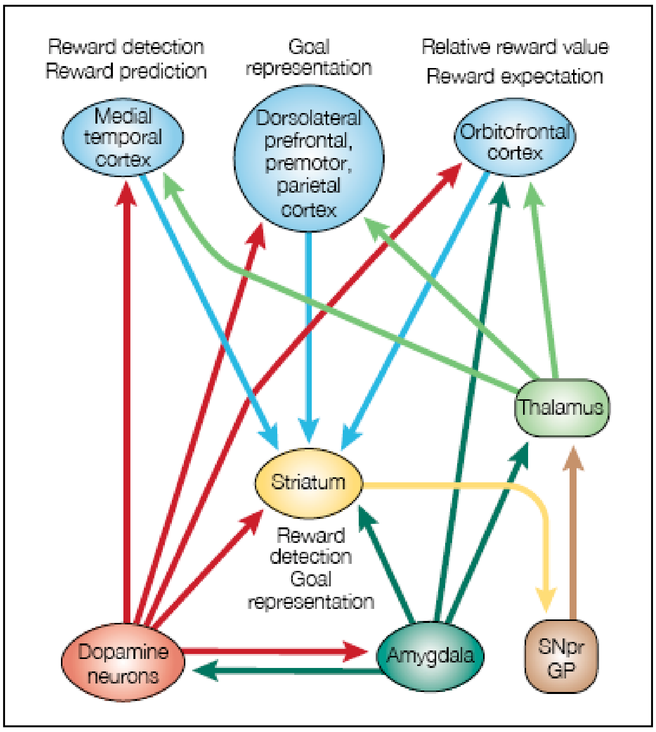

Figure 1: The mesencephalic dopamine system (MDS): Dopamine neurons in the pars compacta of the substantia nigra and ventral tegmental area project to sub-cortical and cortical areas including the basal ganglia (striatum) and different parts of the prefrontal cortex. The present work focuses on the role of these projections (depicted in red) for learning and error processing. Figure adapted from Schultz, 2000.

A long-held hypothesis regarding the function of dopamine during learning was the so-called hedonia hypothesis proposed by Wise and colleagues (Wise, Spindler, deWit, & Gerber, 1978; Wise, 2004). According to the hedonia hypothesis dopamine contributes to reinforcement learning by mediating feelings of pleasure and satisfaction when receiving a reward. However, in more recent times especially electrophysiological findings in primates challenged this view and pointed to a broader and subtler role of dopamine for learning (see Schultz, 2000; 2002; 2007). One of the first studies that investigated the behavior of the dopamine neurons in the VTA during reinforcement learning in monkeys was performed by Schultz and colleagues, (1997). They used an instrumental conditioning task in which a monkey learned to touch a lever after the presentation of a visual stimulus and received reward (a drop of juice). During learning they recorded from dopamine neurons in the VTA and showed that the phasic activity of these dopamine neurons undergoes systematic changes during learning (see Figure 2). At the beginning of learning phasic bursts of dopaminergic activity are found for rewarding outcomes (see Figure 2, top). With learning the phasic dopaminergic activity propagates back in time and is then elicited by the conditioned stimulus that is, the monkey is now able to internally predict the occurrence of the reward (see Figure 2, middle). However, if the reward is not delivered at the time to which it is expected, a phasic decrease of dopaminergic activity is observed that reflects the violation of the expectation of the monkey (see Figure 2, bottom).

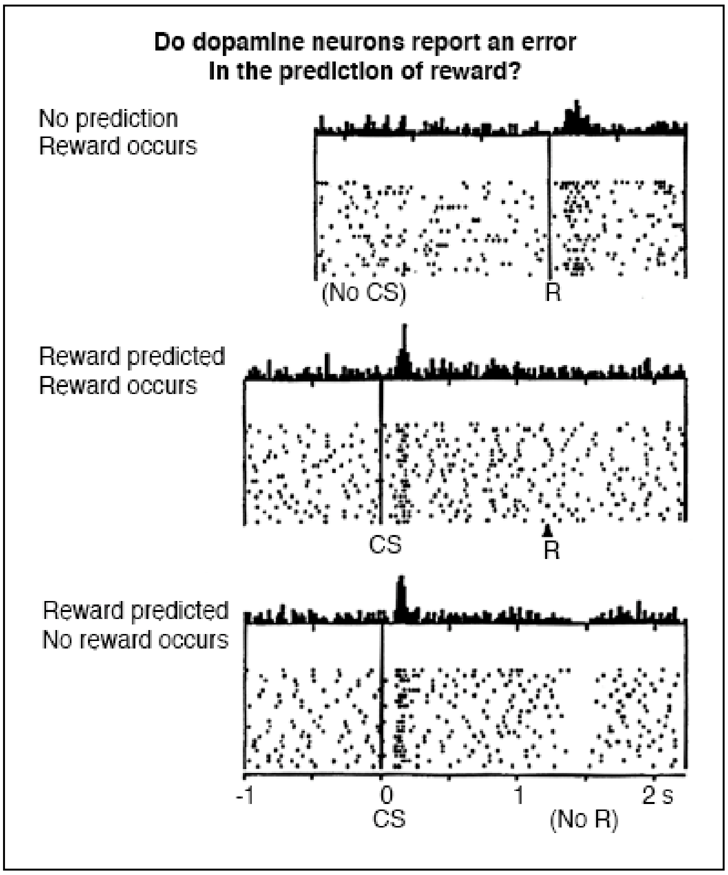

Figure 2: Firing rates of dopamine neurons from the ventral tegmental area (VTA) during instrumental conditioning in monkeys. Top: At the beginning of learning phasic increases of the activity of the dopamine neurons are found for the reward. Middle: After learning the conditioned stimulus elicits the phasic increase of dopaminergic activity. Bottom: The predicted reward is omitted and the dopaminergic signal is suppressed at the time at which the reward was expected.

Based on these findings Schultz and colleagues (1997) proposed that the learning-related changes in the phasic dopaminergic activity could be formalized using a temporal difference learning model (see Rescorla & Wagner, 1972; Sutton & Barto, 1990). The idea is that dopaminergic neurons from the ventral tegmental area signal the extent to which a rewarding outcome deviates from a prediction during learning. That is, they code prediction errors that reflect changes in the value of ongoing events, when events are suddenly better or worse than expected (Schultz, et al., 1997). According to this model, learning is induced when a reinforcer occurs that is better than predicted (positive prediction error), as at the beginning of learning. In contrast, a reinforcer that is worse than predicted or omitted (negative prediction error) leads to extinction of learned behavior. The electrophysiological data in monkeys show that positive prediction errors are reflected in phasic bursts of activity in mesencephalic dopamine neurons, whereas negative prediction errors are reflected in a phasic depression of activity of these neurons.

However, in order to determine that there is causal relationship between behavioral learning and the phasic dopaminergic activity obtained in monkeys, it needs to be shown that learning depends on the presence of predictions errors as coded in phasic changes of dopaminergic activity. To do so, Waelti and colleagues (2001) used a blocking paradigm and recorded from dopaminergic neurons in the substantia nigra and VTA of monkeys. In the blocking paradigm a new, unconditioned stimulus is paired with a pretrained, reward-predicting stimulus. During compound learning both stimuli are then paired with a reward. In a subsequent learning test both stimuli are tested on whether they produce the trained behavior (licking) and the expected neuronal responses. The behavioral, as well as the electrophysiological results showed that learning of a second stimulus is blocked when the other stimulus already predicts the occurrence of reward. These results nicely show that reinforcement learning depends on errors in reward prediction, and that these predictions errors are reflected in phasic changes of the firing rates of the dopaminergic neurons. Moreover, these findings support the view that activity of the MDS could be formalized using temporal difference learning models (see Rescorla & Wagner, 1972; Sutton & Barto, 1990).

2.2.3 Dopaminergic Projections to the Basal Ganglia

The primary subcortical projection area of the mesencephalic dopamine neurons is the basal ganglia and here particularly the striatum (see Schultz, 2000; 2002; 2007). The basal ganglia refer to a group of subcortical structures that involve the striatum (caudate nucleus, putamen and nucleus accumbens) and the globus pallidus. Neuronanatomically, the striatum could be subdivided along a dorsal-ventral axis into the dorsal striatum, involving caudate nucleus and putamen and the ventral striatum, which is comprised of the nucleus accumbens (Heimer & Van Hoesen, 1979). The basal ganglia have since long been implicated in motor behavior since damage to these regions (as in Parkinson's disease, for example) causes dramatic motor and movement disorders (see Ahlskog, 2007; Samii, Nutt, & Ransom, 2004). However, in more recent times the 'cognitive' aspects of the basal ganglia and especially the role of the dorsal striatum for the learning of stimulus-response (S-R) associations, or habits, have received more and more attention (for a review, see Packard & Knowlton, 2002).

In line with this, results from a recent functional imaging (fMRI) study in humans' points to the view that different structures within the basal ganglia are involved in distinct aspects of reinforcement learning (O'Doherty et al., 2004). The idea of this study was to compare classical and instrumental learning and to investigate whether different structures in the basal ganglia are involved in these different forms of learning. The hypothesis was that the functional organization of the basal ganglia might comply with the architecture of so-called actor critic models (see Barto, 1995; Schultz, 2000; 2002; 2007). These computational models are based on the assumption that two distinct components are involved in reinforcement learning. A critic, which learns to predict future rewards based on reward history and by this carries the reward prediction error and an actor, which uses these signals to learn stimulus-actions associations, so that actions associated with higher rewards are more likely to be chosen again. O'Doherty and colleagues (2004) showed that activity in the ventral striatum (mainly nucleus accumbens) is correlated with prediction errors during instrumental as well as classical conditioning. This is in line with the idea that the ventral striatum plays the role of an adaptive critic, which learns to predict future rewards. In contrast, activity in the dorsal striatum showed correlations with prediction errors only during instrumental, but not during classical conditioning. This is consistent with its role as an actor, which maintains information about the rewarding outcomes of actions, in order to enable better ones to be chosen more frequently.

2.2.4 Dopaminergic Projections to the Prefrontal Cortex (PFC)

The lateral prefrontal cortex and the anterior cingulate cortex are two of the main cortical projection areas of the mesencephalic dopamine system (Braver et al., 2001; Cohen, Braver, & Brown, 2002; Durstewitz, Seamans, & Seijnowski, 2000). In the following I will briefly review findings from computational models (Braver et al., 2001; Cohen, Braver, & Brown, 2002; Hazy, Frank, & O'Reilly, 2007) and empirical studies (Hampton & O'Doherty, 2007; O'Doherty et al, 2004; Shima & Tanji, 1998; Shidara & Richmond, 2002) that propose that dopamine has neuromodulatory effects on the cognitive functions that are associated with the lateral PFC and the ACC.

The lateral prefrontal cortex has since long been suggested to be implicated in cognitive control and especially in the active representation of task goals (for reviews see Miller, 2000; Miller & Cohen, 2001; Smith & Jonides, 1999). One critical aspect regarding the representation of task goals in the PFC is how they are built and how they are adaptively updated. Evidence from computational modeling (Braver et al., 2001; Cohen et al., 2002; Hazy et al., 2007) suggests that projections from mesencephalic dopamine system might implement this mechanism by controlling the gating of afferent information into the PFC. That is, phasic increases in dopaminergic activity as elicited for example by unpredicted rewards are suggested to open the gate and allow afferent signals to establish a new goal representation (Montague et al., 2004). Hence, the idea is that rewarding outcomes during reinforcement learning as reflected in phasic dopaminergic learning signals are implicated in the formation and updating of task goals in the lateral PFC. According to these task goals processing in other parts of the brain could be modulated in order to perform a task at hand (see Miller & Cohen, 2001).

The anterior cingulate cortex (ACC) is the second main cortical target area of the MDS and shows the highest density of dopamine fibers in the human cortex (see Paus, 2001). The ACC lies on the medial surface of the PFC and can be broadly subdivided into two parts. Its ventral part is located at the surface of the cingulate gyrus and contains the Brodman area 24. This region has been referred to as the caudal cingulate zone (CCZ) (see Picard and Strick, 1996). The dorsal part of the ACC lies in the cingulate sulcus and contains mainly Brodman area 32. The dorsal part of the ACC has also been termed the rostral cingulate zone (see Ridderinkhof, Ullsperger, Crone, & Nieuwenhuis, 2004) (see Figure 3).

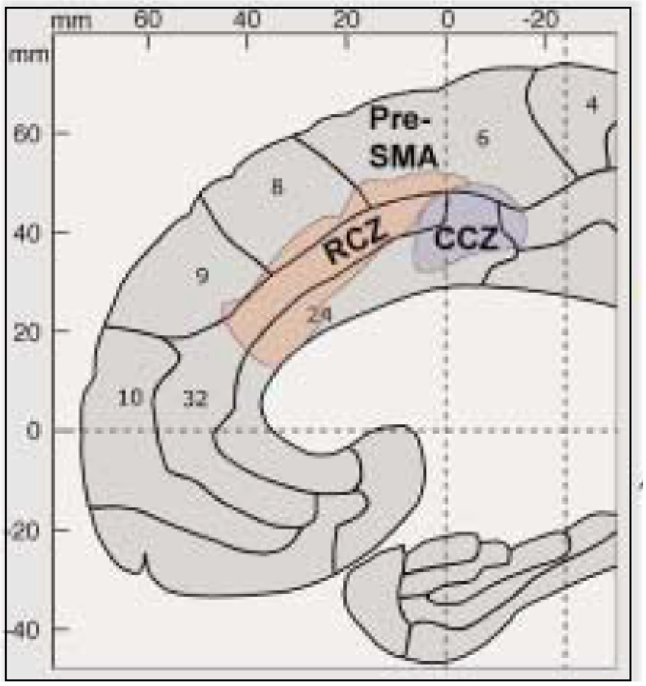

Figure 3: Anatomical map of the medial frontal cortex. The dorsal part of the ACC contains the Brodman area 32 and overlaps with what is referred to as the rostral cingulate zone (RCZ), which is shaded in red. The more ventral part of the ACC has been termed caudal cingulate zone (CCZ) and comprises Brodman area 24. The pre-SMA comprises Brodman area 6 and is the most rostral part of the supplementary motor area which is involved in establishing motor programs (Figure adapted from Ridderinkhof et al., 2004)

The dorsal part of the ACC has been suggested to be involved in several cognitive control functions such as in conflict monitoring (see Botvinick et al., 2001; Yeung et al., 2004) and error processing (Gehring, Goss, Coles, Meyer, & Donchin, 1993; Scheffers & Coles, 2000) or generally in performance monitoring (Ridderinkhof et al., 2004). What these theoretical accounts have in common is that they propose that the ACC is involved in detecting unfavorable outcomes and signaling the need for performance adjustments. This view is consistent with the results from an fMRI study by Bush and colleagues (2001). In this study a reward-based decision making task was applied in which participants had to switch their response strategy if a reduction of reward occurred. This study showed increased activation in the ACC for unexpected reductions of reward, which indicates that the ACC is sensitive to expectancy manipulations. Moreover, O'Doherty, Critchley, Deichmann and Dolan (2003) showed that this punishment-related activity is most pronounced when it precedes a switch in choice behavior. This suggests that the ACC not simply detects a change in reward value, but rather signals the need for a shift in response strategy after reward contingencies have changed (see also Hampton & O'Doherty, 2007; O'Doherty et al, 2003). This is intriguingly in line with data from single neuron recordings in humans, showing that the ACC is implicated in linking reward-related information with appropriate actions, especially when reward is reduced (Williams, Bush, Rauch, Cosgrove, & Eskander, 2004).

However, electrophysiological recordings in monkeys have shown that the ACC is not only activated by unfavorable outcomes but also by positive events such as unexpected rewards (Ito, Stuphorn, Brown, & Schall, 2003) and during reward expectancy (Matsumoto et al., 2003; Shidara & Richmond, 2002). These findings point to a broader role of the ACC in integrating reward information with action selection (Shima & Tanji, 1998; Shidara & Richmond, 2002). Given the strong dopaminergic innervations of the ACC, it seems reasonable to suggest that ACC activity during reward processing might be modulated by input from the MDS. According to such a view, the ACC detects changes in reward value as signaled by the MDS and uses these learning signals to evaluate whether a shift in response strategy is necessary.

2.2.5 Age-related Changes in Dopamine System Across the Lifespan

In the previous sections, it has been shown that the mesencephalic dopamine system has a fundamental role in driving reinforcement learning, as well as in the formation and updating of goal representations and in performance monitoring. However, the structure and the functional efficiency of the dopamine system and its target areas in the prefrontal cortex show substantial changes during life-span development. The following section aims at providing an overview on how the MDS develops during childhood and how it declines in older age and how these (lifespan-) developmental changes in dopamine might be related to age differences in learning and cognitive control.

Up to now little is known about the ontogenetic development of the MDS. In part this is due to the fact that the use of radioactive tracers, as necessary in PET studies, is problematic in children. Therefore, the direct evidence on how dopamine receptor density or dopamine levels change as a function of age during childhood development is limited. However, indirect evidence for the role of dopamine during development comes from studies in nonhuman primates, genetic studies as well as neuropsychological studies in children with deficits in the function of the MDS. For example, research in non-human primates has shown that the density of dopamine receptors in the PFC and in the striatum increases dramatically from childhood to late adolescence (Goldman-Rakic & Brown, 1982; Lambe, Krimer, & Goldman-Rakic, 2000). This has been further confirmed by findings from a post-mortem study in humans, which suggest that pre- and postsynaptic markers of dopamine increase during childhood.

Further indirect evidence for the role of dopamine during childhood development comes from genetic studies (Diamond et al., 2004; Rueda, Rothbart, McCandliss, Saccimanno, & Posner, 2005). Diamond and colleagues (2004) have genotyped children with respect to three polymorphisms of the COMT gene, which are associated with different levels of dopamine in the PFC. Subsequently they tested these children on tasks that are known to involve the prefrontal cortex. They found that children with the Met-Met polymorphism, which is associated the highest dopamine levels in the PFC, perform significantly better in these tasks than children with polymorphisms that lead to lower dopamine levels. Strikingly similar evidence has been obtained in a neuropsychological study on children with phenylketonuria (Diamond, 1996). Phenylketonuria is a genetic disorder that is associated with disturbances in the dopaminergic projections to the PFC. Children with phenylketonuria have been found to be specifically impaired in working memory and inhibitory control tasks, which are assumed to depend on dopaminergic input to the PFC (Diamond, 1996).

Another very interesting line of evidence comes from event-related potential (ERP-) studies in children with attention deficit hyperactivity disorder (ADHD) (Liotti et al., 2005; van Meel et al., 2005). ADHD is a developmental disorder that seems to be due to alterations in the dopamine system and its target areas in the PFC (see Biederman & Faraone, 2005). Liotti and colleagues investigated the effects of ADHD on inhibitory control and the ERP correlate of internal (self-generated) errors, the error-related negativity (ERN). To do so, they recorded ERPs while participants performed a Stop signal paradigm, which is a two-choice reaction time task in which participants have withdraw their response if an infrequent ‘Stop signal’ is provided. They found that children with ADHD showed reduced ERN components and higher error rates compared to healthy controls, suggesting that these children are impaired in processing internal (self-generated) errors.

However, changes in error-related ERPs in children with ADHD are not only found during the processing of internal error information, but also when external error information (negative feedback) has to be processed. Van Meel and colleagues (2005) investigated the processing of external error feedback and its ERP correlate the feedback-related negativity (FRN) in children with ADHD and healthy controls. They applied a guessing task, in which participants had to decide for one of four stimuli and received positive or negative feedback. Unbeknown to the participants, feedback was provided independently of their choice. Interestingly, they found that in contrast to the ERN (see Liotti et al., 2005) the FRN is enhanced in ADHD children compared to healthy controls, which suggests that they are more sensitive to unfavorable outcomes. Taken together, these findings suggest that ADHD is characterized by an enhanced sensitivity to external error feedback, as well as impairments in the processing of internal error information.

In contrast to the limited direct evidence on the role of dopamine for childhood development, we know much more about how the dopamine system changes during older age (for reviews see Bäckman et al., 2006; Braver & Barch, 2002). Most of this evidence comes from studies using positron emission tomography (PET), with different radioactive tracers. These studies show that presynaptical (e.g., dopamine transporter, DAT, see Erixon-Lindroth et al., 2005) as well as the postsynaptical (e.g., D1- and D2 receptor densities (see Volkow et al., 1996; Wang et al., 1998) biochemical markers of the MDS decline as a function of aging. Moreover, these age-related changes in the MDS are associated with performance impairments on tasks that are assumed to involve the PFC. For example, Volkow and colleagues (1998) showed that the declining availability of dopamine D2 receptors in the striatum with age is correlated with performance in the Wisconsin Card Sorting Test and the Stroop task. These tasks have been shown to rely on performance monitoring functions, associated with the PFC (Miller & Cohen, 2001; Ridderinkhof et al., 2004).

Based on these and other findings (Bäckman et al., 2000; Volkow et al., 2000), several theoretical accounts have been proposed on how age-related impairments in learning and cognitive control might be associated with age-related changes in dopamine function. One of these accounts is the context-processing hypothesis of aging as proposed by Braver and colleagues (2001). This model suggests that older adults are impaired in the active representation and maintenance of task-relevant information and that this impairment is due to an age-related decline of the dopaminergic projections to the lateral PFC. This computational model is based on the idea that dopamine controls the gating of afferent information into the PFC and by this allows the updating or maintenance of representations in the prefrontal cortex (see Cohen et al., 2002; Hazy et al., 2007, Montague et al., 2004). More specifically, the context processing hypothesis supposes that an age-related reduction of the dopaminergic projections to the PFC results in less reliable representations of task context and age-related impairments in the maintenance of task-relevant information (see Braver et al., 2001; Braver, Satpute, Rush, Racine, & Barch, 2005).

Another line of modeling work proposes that age-related deficits in dopaminergic function lead to less distinctive neural representations, resulting in age-related impairments of episodic memory and attentional functions (Li, Lindenberger, & Sikström, 2001). In the computational model by Li and colleagues (2001) dopaminergic modulation is conceptualized as altering the signal-to-noise ratio of neural information processing. This is modeled by adjusting the gain parameter of the neural network models. The results of the simulations are in line with a range of behavioral data on age-related impairments in learning and memory, suggesting that age-related deficits in dopaminergic neuromodulation might indeed cause less distinctive cortical representations.

A more recent neurocomputational account assumes that age-related reductions in the phasic activity of the mesencephalic dopamine system might explain the impairments of older adults in reinforcement learning and error processing (Nieuwenhuis et al., 2002). This account is an extension of the reinforcement learning (RL-) theory of the error-related negativity (ERN) by Holroyd and Coles (2002) (for details, see section 2.4). The R-L theory proposes that reinforcement learning is driven by phasic dopaminergic signals that lead to the generation of the ERN in the anterior cingulate cortex. According to the extended model by Nieuwenhuis et al. (2002), these learning signals are weakened in older adults, which leads to reduced ERN components, as well as impaired learning in the elderly. Hence, this model suggests a direct link between age-related deficits in the dopamine system and impaired activity of the ACC during error processing. Convincing support for such a link has been provided by a PET-study by Volkow and colleagues, (2000). They have shown that age-related reductions in striatal dopamine D2 receptor availability are associated with a decrease of glucose metabolism in the ACC. This points to the view that age-related impairments in learning and error processing, as reflected in ACC activity, might be a consequence of an age-related decline in the availability of dopamine D2 receptors in the striatum. Hence, there is good evidence that age-related decline in the MDS and the dopamine receptors in the prefrontal cortex are associated with age-related impairments in cognitive control and learning.

2.2.6 Summary

The electrophysiological findings in monkeys suggest that dopaminergic neurons in the VTA signal the extent to which an actual outcome deviates from a prediction during learning. That is, they show phasic increases of their firing rates if an outcome is better than predicted (positive prediction error) and phasic decreases of their firing rates if an outcome is worse than predicted (negative prediction error) (Schultz et al., 1997; Montague et al., 1996). By doing so, these dopaminergic neurons provide learning signals to their sub-cortical and cortical projection areas, which could then be used for associative learning and performance adaptations (Waelti et al., 2001)[1]. The basal ganglia and particularly the striatum are the primary subcortical target areas of the learning signals provided by the mesencephalic dopamine neurons (see Schultz, 2000; 2002; 2007). In line with the architecture of the actor-critic models of reinforcement learning (Barto, 1995; Sutton & Barto, 1990) the ventral striatum seems to play the role of an adaptive critic, which learns to predict future rewards based on the dopaminergic learning signals. In contrast, the dorsal striatum seems to be implicated in action selection based on these computations (see O'Doherty et al., 2004). Together, these structures are building a sub-cortical learning network, which is referred to as the mesencephalic dopamine system (MDS) in the following.

The primary cortical projection areas of the MDS are the lateral prefrontal cortex (PFC) and the anterior cingulate cortex (ACC). Computational models on the neuromodulatory effect of dopamine in the lateral PFC suggest that dopamine regulates the gating of afferent information into the PFC and by this allows the formation and updating of task goals (Braver et al., 2001; Cohen et al., 2002; Hazy et al., 2007). The second main cortical target area of the dopaminergic projections of the MDS is the ACC (Paus, 2001). Neurophysiological data in monkeys (Ito et al., 2003; Matsumoto et al., 2003; Shidara & Richmond, 2002) and functional imaging (fMRI-) data in humans suggest

---

[1] Pease note that the dopaminergic learning signals have different modifying effects on postsynaptic neurons in their projection areas. First, dopaminergic prediction error signals may change synaptic plasticity when they are coincident with cortical inputs to the same postsynaptic neuron. Second, dopamine may have an enhancing and focusing effect on postsynaptic neurons by altering signal-to-noise ratio. Third, dopamine has been shown to induce long-term potentiation in the PFC, which might be a mechanism by which dopamine could contribute to learning without mediating synaptic plasticity directly. For a review on the mechanisms by which dopamine contributes to asssociative learning, please refer to Schultz, (2002).

that the ACC receives input from the MDS and uses these learning signals to evaluate whether the outcome of an action deviates from a prediction and then indicates, whether a shift in response strategy is necessary (O'Doherty, et al., 2003; Hampton & O'Doherty, 2007). This speaks for the view that the ACC is implicated in performance monitoring by linking reward-related information with appropriate actions, especially when the outcome of an action deviates from a prediction (Shima & Tanji, 1998; Shidara & Richmond, 2002; Williams et al., 2004).

The MDS and its projection areas in the prefrontal cortex show considerable changes during childhood development (see Goldman-Rakic & Brown, 1982; Weickert et al., 2007), as well as during healthy aging (see Bäckman et al., 2006; Braver & Barch, 2002). Although the direct evidence for the role of dopamine in cognitive development during childhood is limited, there is a huge number of studies with different approaches that indirectly show that dopamine and particularly dopaminergic projections to the PFC play an essential role during cognitive development (Diamond, 1996; Diamond et al., 2004; Liotti et al., 2005; Rueda et al., 2005; van Meel et al., 2005). Studies on adult age differences in dopamine metabolism have shown that the pre- and postsynaptical markers of the mesencephalic dopamine system decline in older age (Erixon-Lindroth et al., 2005; Wang et al., 1998; Volkow et al., 1996). This decline is correlated with impairments on tasks that have been shown to rely on PFC function (see Volkow et al., 1998). Moreover, it has been shown that age-related reductions in striatal dopamine D2-receptor availability are associated with a decrease of glucose metabolism in the ACC (Volkow et al., 2000). This points to the view that age-related impairments in tasks that rely on ACC function might be due to a decline of dopamine metabolism in the striatum.

Most of the modeling work (Braver et al., 2001; Li et al., 2002) has focused on how age-related reductions in the activity of the dopamine system might modulate representations in the PFC. Yet, the empirical support for these models is largely restricted to behavioral data and there is less evidence for these modulations on the level of brain activity. However, the recent neurocomputational models by Holroyd and Coles

(2002) and Nieuwenhuis et al. (2002) go beyond this, by providing a direct link between changes in phasic dopaminergic activity and the activity of the ACC, as reflected in the error-related negativity (ERN). The core idea of these models is that the ERN is driven by learning signals from the MDS that indicate that an event is worse than expected. Hence, these models provide a testing ground for the investigation of age-related impairments in reinforcement learning and dopaminergic function on the behavioral and electrophysiological level.

The purpose of the following sections will be to give an overview on the ERP correlates of error processing and how they change as a function of lifespan development, as well as to provide a more detailed view on the R-L theory of the ERN (Holroyd & Coles, 2002) and its extension to older age by Nieuwenhuis and colleagues (2002).

2.3 Neuropsychological Theories of Error Processing

2.3.1 Theoretical Accounts to the Error-Related Negativity (ERN)

When participants commit errors in reaction time tasks a phasic negative deflection can be observed at around 80 milliseconds after the onset of the erroneous response at fronto-central electrodes (see Figure 4).

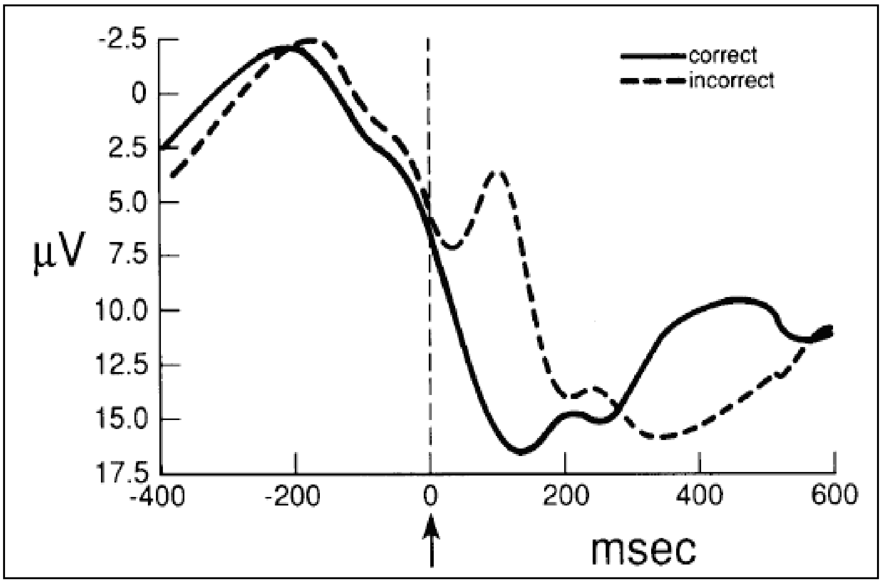

Figure 4: The ERN can be observed at around 80 ms after the onset of an incorrect response (dashed line). (Figure adapted from Gehring et al. 1993).

The two research groups that first observed this component termed it error-related negativity (ERN) (Gehring et al., 1993), or error negativity (Ne) (Falkenstein et al., 1990). Converging results from imaging studies (Carter et al., 1998; Holroyd et al., 2004), dipole analyses (Miltner et al., 2003; Van Veen & Carter, 2002), and neuropsychological studies (Swick & Turken, 2002), point to the view that the ERN is generated in the dorsal part of the ACC (for a review see Ridderinkhof et al., 2004).

One of the most influential functional interpretations of the ERN is that it reflects error detection (Bernstein, Scheffers, & Coles, 1995; Coles, Scheffers, & Holroyd, 2001; Scheffers & Coles, 2000). This view is supported by data that shows that the ERN increases if participants are instructed to focus on accuracy rather than speed. In contrast, the ERN decreases when response speed is emphasized at the expense of accuracy (Gehring et al., 1993). Furthermore, it has been shown that the ERN also increases with the number of incorrectly chosen response parameters (Bernstein et al., 1995), and that it is irrespective of stimulus modality (Falkenstein, Hoorman, Christ, & Hohnsbein, 2000) or response modality (Holroyd, Dien, & Coles, 1998). Together, these findings suggest that the ERN reflects the detection of a mismatch between the representation of a correct response and the actual (incorrect) response (Bernstein et al., 1995; Coles, et al., 2001; Scheffers & Coles, 2000). The larger this mismatch, the larger the ERN. However, what this theoretical interpretation implies is that the amplitude of the ERN depends on the intact representation of the correct response, since otherwise no mismatch could be detected. This is an important issue with respect to the relation between the ERN and learning because it shows that the ERN is related to the participants' ability to represent the correct response.

In opposition to the mismatch theory of the ERN, there is another line of research that proposes that the ERN reflects the monitoring of response conflict (Botvinick et al., 2001; Yeung, Botvinick, & Cohen, 2004). According to this view, the ERN reflects conflict that develops after an error as a consequence of continued stimulus processing. Most of the evidence for this view comes from fMRI-studies on conflict processing using the

Stroop or the Flanker task (Carter et al., 1998; MacDonald, Cohen, Stenger, Carter, 2000; Milham, Banich, & Barad, 2003; van Veen et al., 2001). However, these studies have mainly focused on ACC activations on correct trials and have shown that ACC activity is enhanced when response conflict is present. Moreover, ACC activity has been found to co-vary with the degree of conflict, being larger on high conflict compared to low conflict trials (Carter et al., 2000). This is in line with the results from recent ERP studies using similar tasks (Bartholow et al., 2005; Kray, Eppinger, & Mecklinger, 2005; Eppinger, Kray, Mecklinger, & John, 2007). These studies have shown that response conflict processing, as well as the adaptation to changing demands on conflict processing, is reflected in a negativity on correct trials (CRN). This negativity shares the temporal and topographical characteristics of the ERN and hence might reflect ACC activity during response conflict processing.

In contrast to these 'classic' accounts to the functional significance of the ERN, more recent models suggest that the ERN is associated phasic changes of activity of the MDS, and by this is implicated in reinforcement learning (Holroyd & Coles, 2002; Nieuwenhuis et al., 2002, see section 2.4). In line with these models, Frank and colleagues (2005) have shown that the amplitude of the ERN seems to be associated with individual differences in whether participants learn from positive or negative feedback. In Parkinson patients these learning biases have been shown to depend on whether patients are on or off dopaminergic medication (Frank et al., 2004)[2], which makes it quite plausible to assume that the ERN is associated with individual differences in dopamine levels.

---

[2] Frank and colleagues (2004) investigated learning biases in Parkinson patients on and off medication compared to age-matched healthy controls. They used a probabilistic learning task, in which participants had to learn to choose one of two stimuli based on feedback. Subsequently they tested the participants with novel combinations of these stimuli. Based on the performance in this test, participants were defined as positive learners, when they chose for the stimuli that were rewarded during learning. In contrast, participants were defined as negative learners, when they avoided choosing for stimuli for which they received negative feedback during learning. The results showed that when Parkinson patients are on medication, they are biased towards learning more positive feedback, whereas they tend to learn more from negative feedback, when they are off medication.

Further evidence for the view that the ERN might be related to the function of the MDS comes from pharmacological studies on error processing (de Bruijn et al., 2004; Zirnheld et al., 2004). Results of these studies suggest that dopamine antagonists (e.g., haloperidol) lead to a reduction of the ERN and impair learning (Zirnheld et al., 2004). In contrast, indirect dopamine agonists (e.g., D-amphetamine) lead to an enhancement of the ERN (de Bruijn et al., 2004). This is nicely in line with studies on the effects of dopaminergic medication on error processing in Parkinson patients. Studies that tested the Parkinson patients off medication found comparable ERN amplitudes as in healthy controls (Holroyd, Praamstra, Plat, & Coles, 2002), whereas reduced ERNs were found when patients were on dopaminergic medication (Falkenstein et al., 2001). Hence, these findings consistently show that the amplitude of the ERN is affected by individual differences in dopamine levels.

2.3.2 Theoretical Accounts to the Feedback-Related Negativity (FRN)

Interestingly, ERN-like components are not only elicited by internally generated errors, but also in response to external error feedback. The so-called feedback-ERN (called FRN in the following) was first observed by Miltner et al. (1997) and shows the same medial-frontal topography as the ERN. It can be observed between 200 – 300 ms after the onset of a negative feedback stimulus, indicating that an error has occurred. Based on the similarities with the ERN, Miltner and colleagues (1997) proposed that the FRN might as well reflect the activity a generic error detection system. According to this view, the FRN reflects error detection based on external error information, in contrast to the ERN, which reflects the detection of internally generated errors. The idea that the ERN and FRN reflect the activity of a similar underlying network is further supported by findings from an fMRI-study by Holroyd and colleagues (2004), which showed that internal and external errors activate the same region in the dorsal ACC. This underlines the view that ERN and FRN reflect activity of an error processing system that detects internal and external error information and involves the dorsal ACC.

However, the conclusion that the FRN reflects the processing of external error information has been recently questioned by the results of a study on the processing of monetary gains and losses by Gehring and Willoughby (2002). In this study a gambling task was applied, in which participants had to perform a two-choice decision and received feedback for their choice, as well as for the alternative choice. The results showed that the FRN seems to reflect the processing of monetary losses rather than incorrect choices. That is, losses elicited a FRN even if the alternative outcome would have yielded a greater loss. In contrast, gains did not elicit a FRN even if the alternative choice would have resulted in a greater gain. Gehring and Willoughby (2002) concluded that the FRN might not reflect the evaluation of performance per se, but the motivational impact of the outcome.

Yet, a recent reexamination of these results by Nieuwenhuis, Yeung, Holroyd, Schurger, and Cohen, (2004) showed that whether the FRN is sensitive to monetary losses, or the correctness of the choice, depends on which aspect is most salient in the feedback. Hence, their results revealed that the FRN depends on whether the color of the feedback stimulus indicates the utilitarian (gain/loss) or the performance (correct/incorrect) aspect of the feedback. In line with the R-L theory (Holroyd & Coles, 2002), Nieuwenhuis and colleagues (2002) argue that, similar to the ERN, the FRN indicates that a current outcome is worse than expected.

2.3.3 Age-related Changes in Error Processing across the Lifespan

Developmental differences in the processing of internally generated errors and its correlate, the ERN, have received more and more attention in the recent literature. Several studies have investigated error-related ERPs using the Flanker task in children of different age groups (Davies et al., 2004; Ladouceur et al., 2004; Santesso et al., 2006). The results of these studies consistently showed that the ERN increases with increasing age until late adolescence (see Figure 5). However, in all of these studies developmental differences in the ERN were accompanied by performance differences between age groups (Davies et al., 2004; Ladouceur et al., 2004; Santesso et al., 2006).

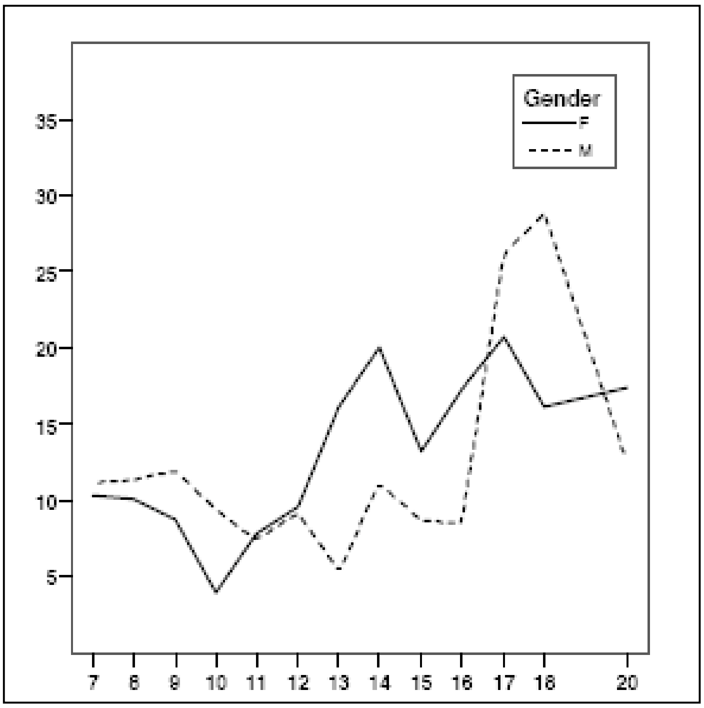

Figure 5: ERN amplitude for boys (solid line) and girls (dashed line) as a function of age. Figure adapted from Segalowitz et al. (2005).

This indicates that in these studies developmental changes in the ERN might have been confounded by age differences in task performance. In an attempt to address this issue, Hogan, Vargha-Khadem, Kirkham, and Baldeweg (2005) compared adolescents (12 - 18 years) to adults (18 - 22 years) using forced-choice visual reaction-time tasks of different complexity. They observed a reduced ERN and behavioral impairments for adolescents only in the more complex version of the task, indicating that controlling for task performance is critical when comparing ERN amplitudes in different age groups. This view is further supported by findings from a study by Kim et al. (2007), who used a Go/No-Go task to examine response inhibition and the ERN in children (7-11 years) and adults (21 - 25 years). They did not obtain a significant reduction of the ERN for children compared to adults and similar to Hogan et al. (2005) concluded that developmental differences in the ERN might depend on task complexity. Hence, there is ample evidence for the view that whether or not the ERN is reduced in children compared to adults might depend on performance differences between age groups. The reason for this performance dependence of the ERN could be that the mismatch signal that is reflected in the ERN depends on an appropriate representation of the correct response. If this representation is impaired, the mismatch signal and hence the ERN is reduced.

Less attention has been paid to the question whether children differ from adults in how they process external error information. In a recent study, Crone and colleagues (2004) examined changes in heart rate to positive and negative feedback during probabilistic learning in 8 - 12 year old children and younger adults. They showed that adult heart rate was slowed following negative feedback in a learning condition, whereas it was not slowed in a non-learning condition (see also Somsen, van der Molen, Jennings, & van Beek, 2000). In contrast, heart rate in children was slowed for both conditions, indicating that children were less able to distinguish informative from uninformative feedback during learning. To my knowledge, the only developmental ERP study on external error processing has been performed in children with ADHD (van Meel, Oosterlaan, Heslenfeld, & Sergeant, 2005). In this study, children with ADHD and age-matched controls performed a gambling task and the ERPs for positive and negative feedback were compared. The results of this study suggested that children with ADHD are more sensitive to negative feedback, as reflected in a larger FRN compared to age-matched controls. This indicates that deficits in the MDS in ADHD children might have resulted in an enhanced sensitivity to negative feedback in these children. However, this study does not provide evidence on how the FRN is affected by age.

Hence, there is a strong need for more research on the question how healthy children differ from younger adults in the way they process external error feedback and the question how they use this information for learning. On the basis of the scarce literature on this issue it seems reasonable to expect that the FRN should be reduced for children compared to adults. However, as for the ERN the question emerges whether this is still the case when performance levels are equated between age groups (see Hogan et al., 2005; Kim et al., 2007). Moreover, as described above, there are some data points that even indicate the opposite pattern, suggesting an enhanced sensitivity to negative feedback for children compared to adults (Crone et al., 2004; van Meel et al., 2005).

Similar to the increasing interest in developmental differences in error processing there is an extending literature on the question how processes of error detection and performance monitoring change as a function of older age. Most of these studies have applied the Erikson Flanker task and their results quite consistently indicate that older adults are impaired in the processing of internal error information, as reflected in a reduced ERN amplitude (see Band & Kok, 2000; Falkenstein et al., 2001; Mathewson et al., 2005, Nieuwenhuis et al., 2002; Themanson et al., 2006; West, 2004) (see Figure 6).

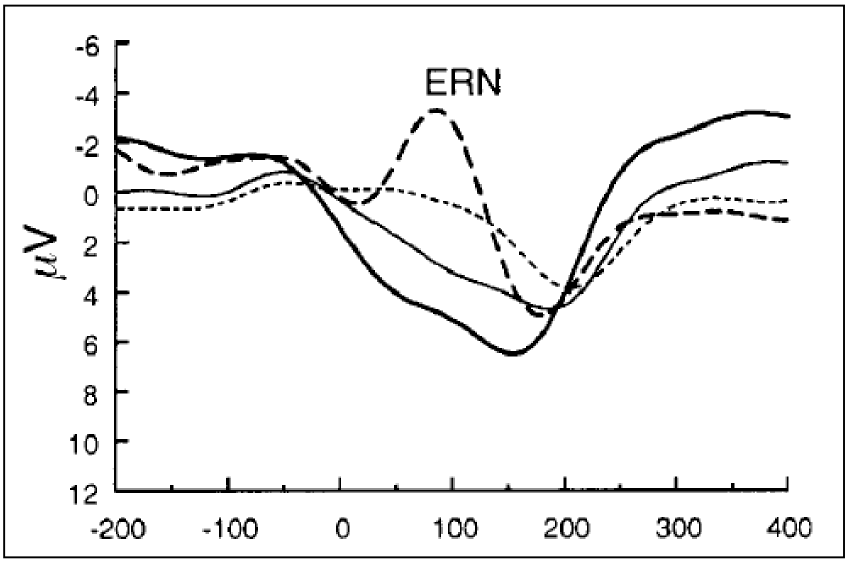

Figure 6: ERN in younger adults (thick dashed line) and ERN in older adults (thin dashed line). Figure ed from Nieuwenhuis et al. (2002).

This amplitude reduction seems to be specific to the ERN since other stimulus-locked components do not seem to be affected by age (Falkenstein et al., 2001; Nieuwenhuis et al., 2002; West, 2004). Moreover, the reduction of the ERN seems to be independent of the type of task used and the type of error that is committed (see Falkenstein et al., 2001; Mathewson et al., 2005). These results point to a general impairment of older adults in the processing of internal error information. However, similar to children, reduced ERNs in older adults were accompanied by performance impairments in the elderly (Band & Kok, 2000; Mathewson et al., 2005; Nieuwenhuis et al., 2002; West, 2004). This suggests that as in children, the reduced mismatch signal in older adults might be due to the fact that they are impaired in representing the correct response.

This idea is supported by findings that suggest that older adults have impairments in response-related conflict processing, as reflected in an increased CRN component (see Eppinger et al., 2007; Kray et al., 2005; Pietschmann, Endrass, & Kathmann, 2007). This suggests that age differences may not be restricted to error processing but may be generally observed during performance monitoring. Interestingly, this pattern of a reduced ERN and enhanced CRN in older adults is strikingly similar to that obtained in patients with lesions in lateral PFC, suggesting that similar to PFC patients older adults might be impaired in representing the correctness of the response. Together, these findings point to a more general impairment of older adults in performance monitoring.

To my knowledge, to date, there is only one study that has investigated the effects of aging on external error processing (Nieuwenhuis et al., 2002). In this study, a probabilistic learning task was applied and age differences in the ERN and the FRN were investigated over the course of learning. Nieuwenhuis ad colleagues (2002) found that similar to the ERN, the FRN is reduced in older compared to younger adults. Moreover, in contrast to younger adults, in whom the FRN increased the larger the mismatch between the actual and the expected feedback, this was not the case in older adults. Together, these findings point to the view that older adults might not only be impaired in internal, but also in external error processing.

2.3.4 Summary

Taken together, there is substantial evidence for the view that the ERN reflects the processing of internally generated errors (Bernstein et al., 1995; Coles et al., 2001; Scheffers & Coles, 2000), whereas the FRN reflects the processing of external error information (Miltner et al., 1997). The ERN has been shown to increase the larger the mismatch between the actual and the intended response, which suggests that it depends on the participants' ability to appropriately represent the correct response. Recent findings point to the view that the ERN is associated with individual differences in dopamine levels (de Bruijn et al., 2004; Zirnheld et al., 2004) and biases towards learning more from negative compared to positive feedback (Frank et al., 2005). Similar to the ERN the FRN has been suggested to be generated in the ACC (Holroyd et al., 2004). Recent data shows that the FRN is sensitive to the first aspect of a feedback stimulus that signals that the outcome of an action is worse than expected (Nieuwenhuis et al., 2004).

The reviewed literature on age-related changes in internal error processing indicate that the ERN increases during development until late adolescence (Davies et al., 2004; Ladouceur et al., 2004; Santesso et al., 2006) and decreases in older age (Band & Kok, 2000; Falkenstein et al., 2001; Mathewson et al., 2005, Nieuwenhuis et al., 2002; Themanson et al., 2006). However, the findings by Hogan et al. (2005) and Kim et al. (2007) also show that whether or not the ERN is reduced in children depends on task complexity and whether the children are able to appropriately represent the correct response. When the task is less complex and children are able to perform comparably to younger adults no reduction of the ERN is found.

A similar argumentation could be applied regarding the reduction of the ERN in older adults. In most of the studies on the effects of aging on error processing (Band & Kok, 2000; Mathewson et al., 2005; Nieuwenhuis et al., 2002; West, 2004) older adults performed worse than younger adults, suggesting that age-related changes in the ERN might have been confounded by performance differences between age groups. Regarding age differences in the processing of external error information there is a clear need for more research in children as well as in older adults. The few studies that have been performed to date point to an enhanced sensitivity to external error feedback in children (Crone et al., 2004; van Meel et al., 2005) and to a reduction of the FRN in older adults (Nieuwenhuis et al., 2002)

2.4 The Reinforcement Learning (R-L) Theory of the ERN

In order to integrate the findings on the role of dopamine for learning with the error processing function associated with the ACC, Holroyd and Coles (2002) proposed the reinforcement learning (R-L) theory of error processing. According to this theory, internal errors (response errors) and external errors (feedback indicating an error) represent negative prediction errors that are elicited when an event is worse than expected. More precisely, the model states that internal and external errors induce phasic decreases in mesencephalic dopaminergic activity. The ERN/FRN is generated when such a dip in dopaminergic input disinhibits neurons in the ACC (see Figure 7).

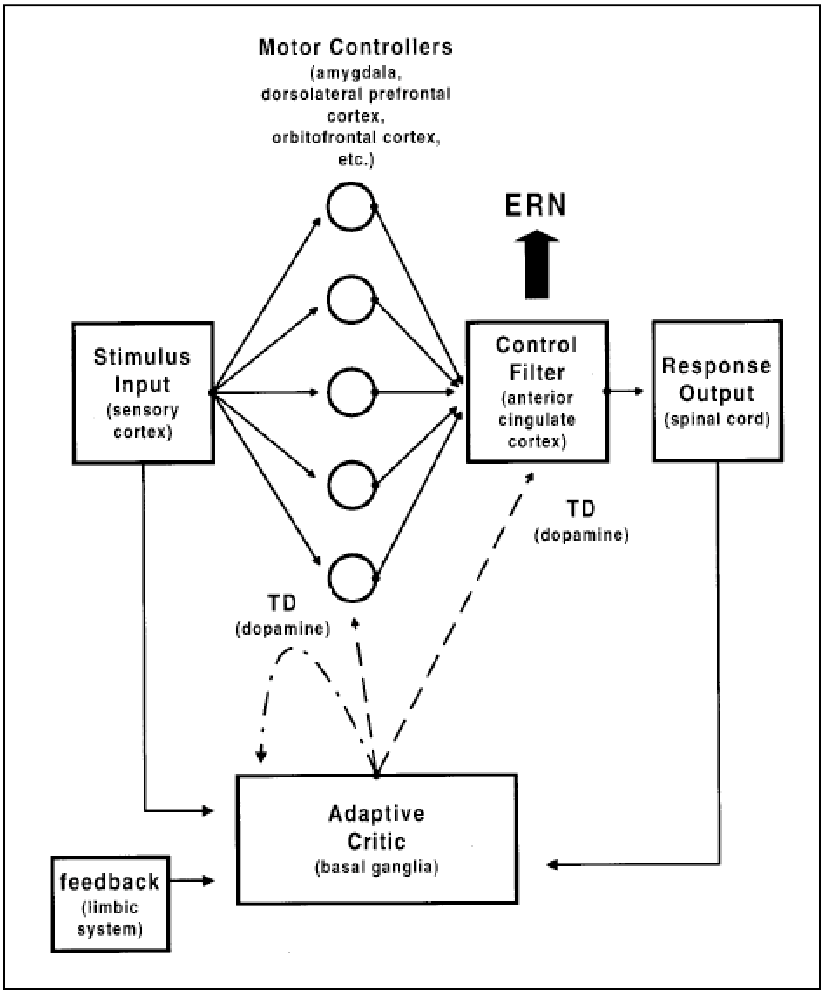

Figure 7: The reinforcement learning theory of the error-related negativity (ERN) (Holroyd & Coles, 2002). TD corresponds to temporal difference error. Figure adapted from Holroyd & Coles, (2002).

In other words, their model suggests that the ERN/FRN reflects a negative prediction error and is generated in the ACC when an event is worse than expected[3]. On the basis of these assumptions the model predicts a trade-off between the ERN and the FRN during reinforcement learning. The ERN is expected to increase over the course of learning since the negative prediction error increases as a function of learning. This is

---

[3] Note that in recent publications from the proponents of this model (Holroyd, Nieuwenhuis, Yeung, & Cohen, 2003; Holroyd, 2004) it has been suggested that positive predictions errors might inhibit the ACC and lead to a positive deflection in the ERP.

because the better participants have learned a certain stimulus-response mapping the more they expect themselves to be correct and hence the larger the prediction error when they commit an error. In turn, the FRN is expected to decrease as a function of learning since the more participants learn the more they are able to internally predict the outcome of their response and the smaller the negative prediction error that is elicited by negative feedback.

To study the relation between the ERN/FRN and reinforcement learning, Holroyd and Coles (2002) used a probabilistic learning task in which participants learned stimulus-response assignments by trial and error based on feedback information. The results of this study showed that the ERN indeed increased with learning and that in the condition in which learning was possible the ERN was larger than the FRN. Thus, for the most part the empirical data supported the R-L theory by showing that the ERN is getting larger with learning, which is in line with the idea that it reflects negative learning signals from the MDS and is generated when an event is worse than expected. Moreover, the data showed the expected trade-off between the ERN and the FRN. That is, the more participants learn and hence are able to internally predict the correctness of the feedback the larger the ERN and the smaller the FRN. Taken together, this data supports the view that the ERN is driven by learning signals from the MDS and is generated when the outcome of an action is worse than expected.

2.5 A Computational Account to Altered Error Processing in Older Age

On the basis of the R-L theory Nieuwenhuis and colleagues (2002) proposed a computational account that suggests that the impairments of older adults in learning and error processing might be due age-related deficits in the function of the MDS. The central assumption of this account is that an age-related reduction of phasic dopaminergic activity leads to a reduced ERN and learning impairments in older adults. Nieuwenhuis and colleagues (2002) tested this hypothesis by measuring performance and ERNs in younger and older adults during probabilistic learning. Their learning task included three learning conditions in which feedback validity was manipulated (100%, 80% and 50% feedback validity). That is, they varied the probability to which a certain response was predictive with respect to the outcome and hence allowed participants to learn more or less from feedback (for a similar paradigm see Figure 8). The results for younger adults generally replicated the findings from Holroyd and Coles (2002) and further supported the R-L theory. For older adults, they found reduced ERNs and less pronounced differences in the ERN between learning conditions, which is in line with the predictions of their model. The FRN was also reduced in the elderly, but did not vary as function of learning condition.

2.6 Synopsis

To summarize, the literature reviewed above shows that reinforcement learning depends on phasic changes in dopaminergic activity that signal the extent to which an outcome deviates from a prediction. Phasic increases of dopaminergic activity are found if the outcome of an action is better than expected, whereas phasic decreases of dopaminergic activity are observed if an outcome is worse than expected (see Schultz 2000; 2002; 2007). By this the MDS provides learning signals that modulate activity in cortical areas like the lateral PFC (see Braver et al., 2001; Cohen et al., 2002; Durstewitz et al., 2000) and particularly the ACC (see Frank et al., 2005; Holroyd & Coles, 2002; Nieuwenhuis et al., 2002). Neurophysiological data in monkeys point to the view that the ACC evaluates these signals and indicates whether a change in response strategy is necessary (Hampton & O’Doherty, 2007; O’Doherty et al, 2003). Hence, the ACC is implicated in performance monitoring, by linking reward-related information with appropriate actions (Shima & Tanji, 1998; Shidara & Richmond, 2002).

There is considerable evidence for age-related changes in the MDS during childhood development (Diamond et al., 2004; Goldman-Rakic & Brown, 1982) as well as during older age (see Bäckman et al., 2006; Braver & Barch, 2002). These age-related changes in dopamine metabolism are associated with impairments in tasks that rely on dopaminergic projections to the ACC (Volkow et al., 1998; 2000). This is supported by electrophysiological data that shows that functions like internal and external error processing, which are associated with activity in the ACC, seem to be particularly affected by disturbances in dopamine metabolism (de Bruijn et al., 2004; Liotti et al., 2005; van Meel et al., 2005; Zirnheld et al., 2004;).

The present research is based on a neurocomputational model (Holroyd & Coles, 2002) that integrates the role of dopamine for learning with the error processing function associated with the ACC and its ERP-correlates, the error-related negativity (ERN) and the feedback-related negativity (FRN), respectively. The reinforcement learning (RL-) theory proposes that the ERN is generated when a negative reinforcement learning signal is conveyed from the MDS to the ACC. By this, the model suggests that the ERN reflects a negative prediction error that is generated when the outcome of an action is worse than expected. The R-L theory predicts that the amplitude of the ERN should increase with learning, since the negative prediction error increases with learning. In contrast, the amplitude of the FRN is suggested to decrease with learning since participants rely less on external error feedback the more they are able to internal predict the outcome of the response.

In order to explain the impairments of older adults in error processing and learning Nieuwenhuis and colleagues (2002) have proposed an extension of the reinforcement learning theory. The central assumption of this account is that age-related reductions of dopaminergic activity lead to impairments in learning and error processing as reflected in reduced ERN and FRN amplitudes. Although the study by Nieuwenhuis and colleagues (2002) provides some important insights into the role of error processing for reinforcement learning in older age there are two major concerns regarding this account and the data that is thought to support it. The first issue relates to the fact that in Nieuwenhuis et al. (2002) study older adults performed much worse than younger adults, which suggests that age differences in the ERN might have been confounded by performance differences between age groups[4]. Due to these performance impairments older adults might have

---

[4] Please note that there is already some evidence for this view from developmental studies on error processing (Hogan et al., 2005, Kim et al., 2007). In this study the authors showed that whether or not the ERN is reduced in children compared to adults depends on task complexity and hence on whether there are performance differences between age groups.

been less able represent the correct response and hence perceived less mismatch and showed a reduced ERN when they committed an error. The second issue that is not addressed in this study is how the ERN changes over the course of learning in younger and older adults. In the study by Nieuwenhuis et al. (2002), learning was investigated by comparing different learning conditions in which the validity of feedback was manipulated and by this more or less learning was possible.

The first experiment intends to address these issues by using a probabilistic learning task that allows older adults to perform comparably to younger adults. If the ERN is performance- rather than age-sensitive, comparable ERN amplitudes for both age groups should be obtained, when performance levels are equated. Moreover, if the ERN co-varies with learning, this should be the case for younger, as well as for older adults. In order to investigate age differences in the time course of these learning effects, changes in the ERN will be examined over the course of learning.

The second Experiment investigates developmental differences in learning and error processing, using a very similar version of the learning paradigm that was applied in the first experiment. This experiment aims at giving insights into the question whether children differ from younger adults in the ability to use internal and external error information for learning. Similar to the first experiment, the objective was to equate performance levels between age groups in order to compare the ERN between children and adults in the absence of performance differences.

The third experiment is based on the findings from the first Experiment and explores an alternative account on learning and error processing that has recently been proposed by Frank and colleagues (2004). This model suggests that individual differences in dopamine levels are related to learning biases and are reflected in the amplitude of the ERN. This experiment aims at answering the question whether older adults differ from younger adults with respect to their learning biases. That is, whether they have a tendency towards learning more form negative or more from positive feedback.

II Empirical part

3 *Experiment 1*

3.1 Statement of Problem

The objective of the first experiment was to replicate and extend recent findings on age differences in learning and error processing. The idea of this study was to address two important issues that have not been resolved in a recent study on the effects of aging on learning and error processing (Nieuwenhuis et al., 2002). The first issue relates to the question how performance differences between age groups affect age differences in the ERP correlates of learning and error processing (the ERN and FRN, respectively). This is an important point because the R-L theory (Holroyd & Coles, 2002) suggests that the ERN is performance-dependent and increases the larger the mismatch between the actual and the intended response. Hence, it could be argued that the finding of a reduced ERN during learning in older adults (see Nieuwenhuis et al., 2002) could be explained by the fact that older adults performed much worse than younger adults. In other words, this suggests that in the aforementioned study age differences in the ERN were confounded by differences in performance between age groups. Thus, the major aspect of this experiment was to design a learning task that would allow us to compare the ERNs of younger and older adults in the absence of performance differences between age groups.

The second important aspect that has not been addressed in Nieuwenhuis et al. (2002) study is how the ERN changes over the course of learning in younger and older adults. In this study, learning was investigated by comparing different learning conditions, in which feedback validity was manipulated, and by this more or less learning was possible. Thus, one important further goal of the first experiment was to precisely track age differences in error processing during the time-course of learning.

3.2 Design

The first experiment addresses these issues by using a probabilistic learning task. In this task the participants were asked to make a two-choice decision upon presentation of an imperative stimulus and received positive or negative feedback. Feedback validity was manipulated in three conditions (100%, 80%, or 50% validity). In the 100% validity condition, feedback was always valid and participants were able to learn from the feedback. In the 80% validity condition, feedback was valid in 80% of the trials but also invalid in 20% of the trials. Hence, learning was impaired in this condition since invalid information occurred during learning. In the 50% validity condition, which served as control condition, feedback was delivered randomly so that no learning was possible. In order to obtain similar performance levels in younger and older participants, an algorithm was implemented in the learning task that adaptively adjusted the response deadline (for details, see Method). This was done because it is well known that aging is accompanied by a substantial general slowing that accounts for several age-related impairments in cognitive tasks (see e.g., Salthouse, 1996; 2000). Using an equal response deadline for both age groups, as it was done in the Nieuwenhuis et al. (2002) study, would have led to a disproportional time pressure for older adults, thereby impairing their ability to learn. Hence, the adaptive algorithm allows each individual to take time for responding by maintaining moderate time pressure.

3.3. Hypotheses

Based on the neurocomputational models reviewed above, the following hypotheses were derived. First, the use of an adaptive response deadline should increase learning rates in older adults. Thus, older adults were expected to perform comparably to younger adults, at least in the 100% validity condition. Second, based on the dopamine hypothesis of aging (Nieuwenhuis et al., 2002) older adults were expected to show reduced ERN components during learning. However, if age differences in the ERN were confounded by performance differences between age groups in the previous studies, an equation of performance levels can be expected to result in comparable ERN amplitudes for younger and older adults. Third, the R-L theory (Holroyd & Coles, 2002) predicts that the ERN should increase with learning in younger adults and the Nieuwenhuis et al. (2002) account suggests that this increase should be smaller for older compared to younger adults. Given that learning rates are equated between age groups, and given that the ERN is indeed performance-rather than age-sensitive, comparable increases of the ERN with learning should be obtained for both age groups. Finally, the R-L theory (Holroyd & Coles, 2002) predicts that the FRN should decrease with learning, since participants rely less on the feedback. However, since in the Nieuwenhuis et al. (2002) study the effects of learning on the FRN were rather small and no age differences in the learning effects were obtained, it is necessary to replicate these findings and to explore whether the FRN indeed changes as a function of learning.

3.4 Method

Participants. Forty-two adults participated in the study. The experimental procedure lasted about 3 hours and the subjects received 22.5 Euro for participation. One younger adult had to be excluded from data analysis due to technical problems during data acquisition. Three younger and two older adults had to be excluded because they did not commit enough errors to analyze the error-related ERP components over the course of learning. The effective sample consisted of eighteen younger adults and eighteen older adults (see Table 1). According to self-report, all participants were healthy, had a right-hand preference, no color blindness, and no history of neurological or psychiatric problems. The participants performed two psychometric tests, one from the domain of fluid intelligence (the Digit-Symbol Substitution test; adapted from Wechsler, 1982) and one from the domain of crystallized intelligence (the Spot-a-Word test; adapted from Lehrl, 1977). As expected on the basis of prior findings (Verhaeghen & Salthouse, 1997) and the two component model of intelligence (Baltes, Staudinger, & Lindenberger, 1999), younger adults reached a higher score than older adults on the Digit-Symbol Substitution test, $\underline{F}(1, 34) = 43.66$, $\underline{p} < .0001$, $\underline{\eta}^2 = .56$ (see Table 1), which reflects the age-related decline in perceptual speed of processing. In contrast, in the Spot-a-Word test both age groups reached comparable scores, $\underline{F}(1, 34) = 0.42$, $\underline{p} < .52$, $\underline{\eta}^2 = .01$ (see Table 1), which speaks for age-related stability in semantic knowledge.

	Younger adults	***Older adults***
N / Gender	18 / 9 female	18 / 9 female
Age Range	19 – 26	65 – 75
Mean Age (M / SD)	20.8 (1.8)	68.5 (2.8)
Digit-Symbol Substitution test (M / SD)	61.9 (7.7)	43.4 (9.1)
Spot-a Word test (M / SD)	25.2 (3.1)	26.0 (4.1)

Table 1: Demographic characteristics of the sample in Experiment 1. Digit-Symbol Substitution test, adapted from Wechsler, (1982), Spot-a-Word test; adapted from Lehrl, (1977)

Stimuli and Task. Stimuli were presented in color against a dark grey background on a 17-inch computer screen. The stimulus set consisted of 36 colored images of objects from the Snodgrass and Vanderward (1980) picture database. The objects belonged to one of the following six categories: Clothes, vehicles, fruit, vegetables, furniture, and domestic appliances. The German words ‘RICHTIG’ (‘correct’) printed in green and ‘FALSCH’ (‘incorrect’) printed in red served as feedback stimuli. When the response deadline was missed, the German words ‘ZU LANGSAM’ (‘too slow’) were presented.

The subjects were asked to make a two-choice decision upon presentation of the imperative stimulus and to press one of two response keys (C and M on a standard computer keyboard). They were instructed to infer the stimulus-response mappings by trial and error based on the feedback. In order to increase the motivation of the participants, they were told that they could win between 50 Euro Cents and 450 Euro Cents per block, depending on their performance. At the end of each block, they received feedback about the amount of money they had won during the block. This monetary feedback depended on the mean performance in the 100% validity condition.

Experimental design. The design involved three learning conditions, in which the validity of feedback was manipulated (see Figure 8).

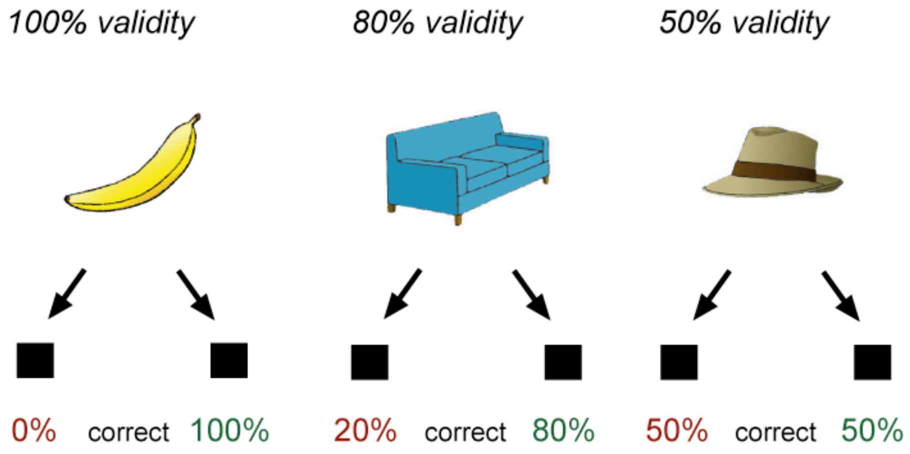

Figure 8: Feedback validity was manipulated in three validity conditions (100%, 80%, and 50% validity).

In the 100% validity condition, in which the feedback was always valid, one stimulus (A) was mapped to the right response key and the other stimulus (B) to the left response key. If participants responded to A with a right button press, they always received positive feedback, and they always received negative feedback if they responded with a left button press (and vice versa for stimulus B). Two other stimuli (C and D) were associated with the 80% validity condition. If participants responded to C with a left button press, they received positive feedback in 80% and negative feedback in 20% of the button presses. If they responded with a right button press, they received negative feedback in 80% of the button presses and positive feedback in 20% of the button presses (and vice versa for Stimulus D). In the 50% validity condition, positive and negative feedback for responses to the stimuli E and F was delivered randomly. The assignment of stimuli and responses was randomized across subjects. For all validity conditions feedback was drawn with replacement, thus the percentage of feedback validity was equal for each bin.

Trial Procedure. At the beginning of each trial a fixation cross was displayed for 500 ms, which was followed by the imperative stimulus for again 500 ms. The response deadline was adapted in 100 ms steps in a range of 600 to 1000 ms, depending on the proportion of time-out trials relative to performed trials (see Figure 9).

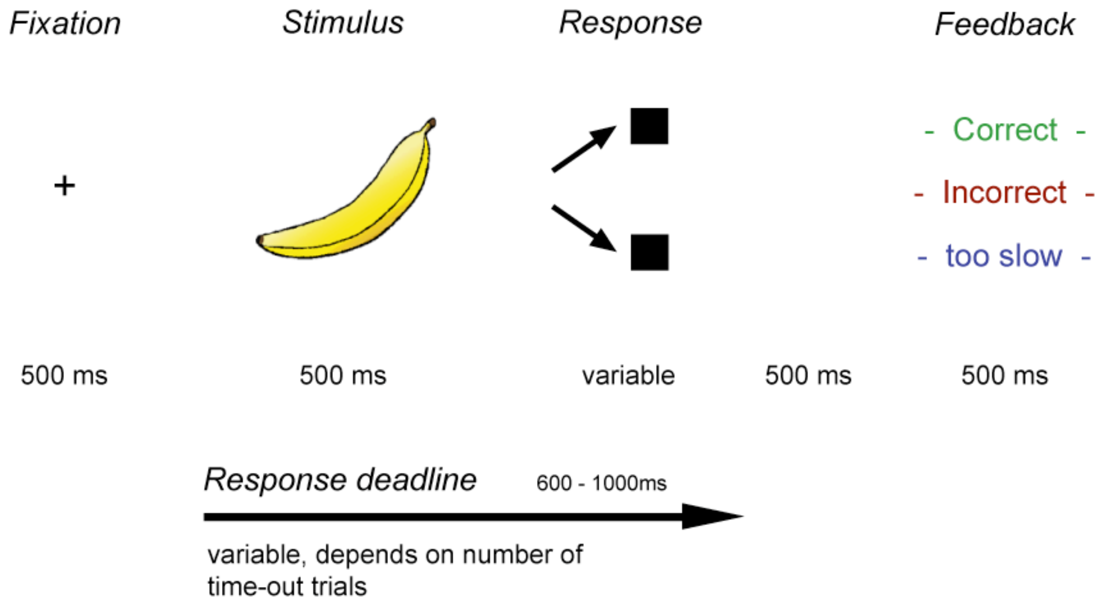

Figure 9: Trial procedure of the probabilistic learning task.

Each participant started with a response deadline of 800 ms. After the first trial the algorithm kept track of the proportion of time-outs (number of time-out trials relative to the trials performed). If the proportion of time-outs was smaller than two percent, a response deadline of 600 ms was applied. With steps of two percent, the response deadline increased for 100 ms and reached a maximum deadline of 1000 ms with over eight percent of time-out trials. This was done in order to make sure that all subjects produced a similar proportion of time-outs (M = .02, SD = .01, for younger adults, M = .06, SD = .04, for older adults), and thereby had a similar opportunity to learn from feedback. For similar deadline procedures see Light et al. (2006) and Rinkenauer et al. (2004). Following the key press, a blank screen was displayed for 500 ms and then the feedback stimulus appeared for again 500 ms. Then participants entered the next trial.

Procedure. First, each participant filled out an informed consent and a short demographic questionnaire. Then, they performed the two psychometric tests. The experiment consisted of one practice block and five experimental blocks. Each block involved a new set of six imperative stimuli, which were drawn randomly (without replacement) from the six stimulus categories (see Stimuli). In a practice block (150 trials) the participants were familiarized with the experimental setting. Finally, they performed the five experimental blocks. In the experimental blocks, each of the six imperative stimuli was presented 50 times in random order. Thus, each participant performed 300 trials per experimental block, yielding in a total number of 1500 trials.

Data Recording. An IBM compatible computer was used for collecting reaction times (RTs) and accuracy data. The stimuli were presented on a CTX 17-inch color monitor with a dark grey background. Responses were registered using the response keys C and M on a standard computer keyboard. The experiment was controlled by the Software E-Prime. EEG and EOG activity were recorded continuously (Brain Amp DC Recorder and Brain Vision Recorder acquisition software) from 64 Ag/AgCl electrodes (10 - 10 system) using EasyCaps recording caps. The left mastoid was used as reference and the right mastoid was recorded as an active channel. The EEG and EOG signals were filtered online from DC - 70 Hz and digitized at 500 Hz. Vertical and horizontal EOG was recorded from two electrode pairs placed on the infra- and supraorbital ridges of the right eye and on the outer canthi of the two eyes. Impedances were kept below 10 kΩ. To increase S-R ratio, the EEG data were offline low-pass filtered with 30 Hz prior to statistical analyses.

Behavioral Data Analyses. Responses faster than 167 ms (more than two standard deviations from the mean reaction time in both age groups) and responses that exceeded the response deadline (younger adults: M = 706 ms, SD = 117 ms; older adults: M = 851 ms, SD = 126 ms) were excluded from data analysis. The accuracy data was analyzed by averaging mean accuracy rates individually for each subject and validity condition into four bins (of 75 trials), reflecting the four quarters of the learning blocks (see Table 2, Appendix). The mean accuracy rates (% correct) were then subjected to an analysis of variance (ANOVA). In order to quantify the learning-related changes in the accuracy data, the learning curves were fitted individually for each subject and for the three validity conditions using a linear ($Y = b0 + (b1 * t)$) and an inverse function ($Y = b0 + (b1 / t)$), as implemented in SPSS. The slope (b1 or β-) parameters of the functions that fitted the data most adequately (inverse learning function for the 100% and the 80% validity condition and linear function for the 50 % validity condition, see Table 2) were then subjected to the analyses of variance. The mean fit parameters (R^2) and the mean slope parameters (β) of the learning functions are displayed separately for the two age groups and the three validity conditions in Table 3 (see Appendix).

ERP-Data Analyses. The EEG epochs were averaged with respect to response and feedback onset to obtain response-locked and feedback-locked ERPs. The response-locked EEG data was baseline corrected by subtracting the average activity during the 200 ms preceding the imperative stimulus. For the feedback-locked EEG data, the average activity from -100 ms to feedback onset served as baseline. Prior to averaging, trials containing eye-movement artifacts or other artifacts were excluded from further analysis using a threshold criterion (standard deviations greater than 30 µV within a sliding window of 200 ms). Remaining vertical and horizontal eye movements were corrected using a modified version of the linear regression approach developed by Gratton et al. (1983), as it is implemented in EEProbe software (ANT Software).

In a first step, response- and feedback-locked ERP components were analyzed separately for correct and incorrect responses (positive and negative feedback). The response-locked components were measured as the mean amplitudes in a 0 – 100 ms time window post-response at the electrode FCz. The feedback-locked components were measured as the mean amplitudes within a 100 ms time window centered on the peak of the FRN at the electrode FCz (260 ms in younger adults and 300 ms in older adults).

In the second step, the ERN, CRN and FRN were specifically analyzed by means of peak-to-peak measurements (see Frank, Woroch, & Curran, 2005; Yeung & Sanfey 2004). For the peak-to-peak analyses, response-locked, as well as feedback-locked EEG data were filtered using a 15Hz low-pass filter in order to obtain more reliable peak amplitude measures. Following Frank et al., (2005) and Yeung and Sanfey (2004), the ERN and the CRN (in older adults) were defined as the peak-to-peak voltage difference between the most negative peak between -50 and 150 ms around the response and the preceding positive peak. The FRN was defined as the difference between the most negative peak within 200 to 400 ms and the preceding positive peak. Scalp potential topographic maps of selected ERP results were generated using all electrode positions by means of a two-dimensional spherical spline interpolation (Perrin et al., 1989) and a radial projection from CZ, which respects the length of the median arcs. Whenever necessary the Geisser-Greenhouse correction was applied (Geisser & Greenhouse, 1958). In these cases the original F-value, the adjusted p-values, and the Epsilon values are reported. In addition, effect sizes (eta squared, η^2) are reported, which reflect the proportion of variance that is accounted for by the experimental manipulations (see Cohen, 1973).

As for the behavioral data, the ERPs were averaged into four bins reflecting the four quarters of the learning blocks. To quantify the learning-related changes, each individual's learning curves were fitted separately for the three validity conditions using a linear ($Y = b0 + (b1 * t)$) and an inverse function ($Y = b0 + (b1 / t)$), as for the analysis of the behavioral data. The slope (b1 or β-) parameters that were estimated using these functions were then subjected to the analyses of variance. The mean fit parameters (R^2) and the mean slope parameters (β) of the learning functions are displayed separately for the two age groups and the three validity conditions in Table 3 (see Appendix).

3.5 Results

3.5.1 Accuracy data

The accuracy data (see Figure 10) was analyzed using an ANOVA design with the factors Age group (young, old), Validity (100%, 80% and 50% validity), and Bin (Bin1, Bin2, Bin3, Bin4). The ANOVA revealed a significant effect of validity, $\underline{F}(2, 68) = 91.24$, $\underline{p} < .0001$, $\underline{\varepsilon} = .88$, $\underline{\eta}^2 = .71$. Contrasts for each of the levels of the factor validity showed a higher accuracy for the 100% compared to the 80% condition and for the 80% compared to the 50% condition ($\underline{p}$'s < .0001, $\underline{\eta^2\text{'s}}$ > .45). Moreover, a marginally significant effect of age group, $F(1, 34) = 3.50$, $p < .07$, $\underline{\eta}^2 = .38$ and a marginally significant interaction between age group and validity, $\underline{F}(2, 68) = 2.55$, $\underline{p} < .09$, $\underline{\varepsilon} = .88$, $\underline{\eta}^2 = .02$ were obtained. Separate ANOVAs for each of the validity conditions revealed significant age differences only in the 80% condition, $\underline{F}(1, 34) = 4.94$, $\underline{p} < .03$, $\underline{\eta}^2 = .12$, suggesting that older adults performed worse than younger adults when feedback was partially invalid (see Figure 10)[5]. These findings show that accuracy increased with feedback validity. However, age differences were only obtained in the 80% validity condition, indicating that older adults were impaired when invalid information occurred during learning.

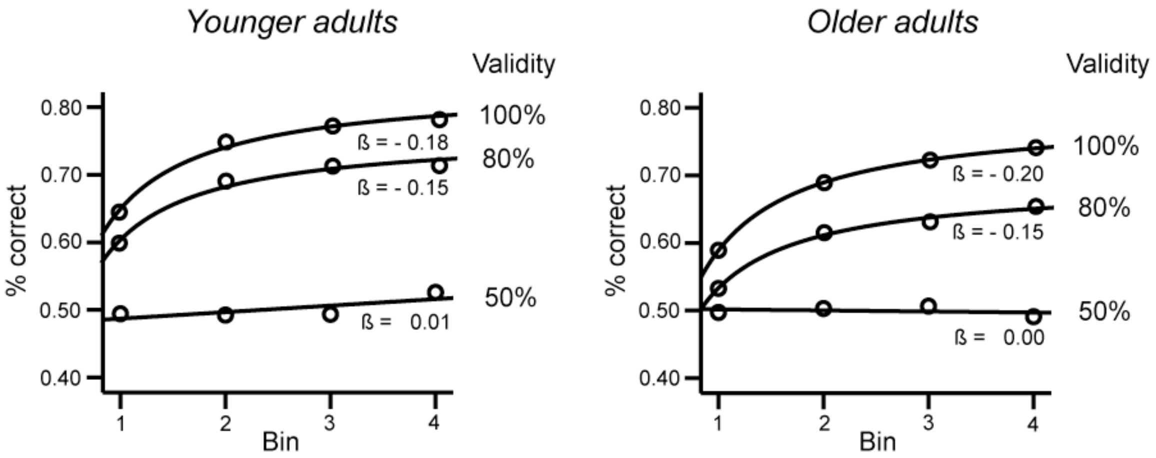

Figure 10: Accuracy learning curves for the three validity conditions (100%, 80%, and 50% validity) displayed separately for younger and older adults.

Learning effects. Of most interest for the present work were the learning effects in the different validity conditions. The analysis showed a significant effect of bin, $\underline{F}(3, 102) = 41.69$, $\underline{p} < .0001$, $\underline{\varepsilon} = .78$, $\underline{\eta}^2 = .55$ and a significant interaction between validity and bin $\underline{F}(6, 204) = 15.69$, $\underline{p} < .0001$, $\underline{\varepsilon} = .75$, $\underline{\eta}^2 = .30$. Separate ANOVAs for each of the validity conditions revealed significant effects of bin for the 100% and the 80% conditions ($\underline{p}$'s $< .0001$, $\underline{\eta}^2$'s $> .41$). As expected, no significant effect of bin was found for the 50% condition ($\underline{p} = .63$), indicating that the accuracy increased over the course of learning only in the 100% and 80% validity condition (see Figure 10). In order to investigate age differences in accuracy over the course of learning, pair-wise comparisons for the levels of the factor

---

[5] The accuracy rates in the 80% validity condition reflect the mean accuracy for the 80% valid trials of this condition. For the 20% invalid trials mean accuracy is lower than chance ($\underline{M} = .33$, $\underline{SD} = .10$ for younger adults; $\underline{M} = .39$, $\underline{SD} = .14$ for older adults) since participants learned to respond according to the dominant (but here incorrect) mapping. For the analysis of the response-locked ERPs valid and invalid trials were aggregated in the 80% condition since there should be no difference between these trial types at the level of the response. For the feedback-locked ERPs only valid trials were averaged in the 80% condition.

validity were performed separately for each bin. These contrasts revealed significant differences between all of the validity conditions in all of the bins (p's $< .001$, η^2's $> .21$). However, these contrasts did not reveal significant age differences for the 100% condition compared to the 80% and 50% validity conditions in any of the four bins (p's $> .09$). In contrast, in line with the age differences in overall accuracy in the 80% condition, significant age differences for the 80% compared to the 50% condition were found for the first, second, and third bin (p's $< .02$). However, at the end of learning (in the fourth bin), no significant age differences in the 80% condition were obtained ($p = .60$).

These findings show that in the 100% validity condition no age differences in accuracy were obtained over the course of learning. Age differences in the 80% validity condition were most pronounced at the beginning of learning, but absent at the end of learning (see Figure 10). In order to examine age differences in the learning functions, an ANOVA on the slope parameters of the learning curves (for details, see Method) was performed. The ANOVA involved the factors Age group and Validity. Results revealed a significant effect of validity, $F(2, 68) = 43.92$, $p < .0001$, $\varepsilon = .94$, $\eta^2 = .56$. Contrasts for each of the levels of the factor validity showed higher slope parameters for the 100% and 80% validity conditions compared to the 50% validity condition (p's $< .0001$, η^2's $> .54$). However, only a marginally significant difference was obtained between the 100% and the 80% validity condition ($p = .08$). Importantly, no age differences in the slope parameters were obtained (see Figure 10). Thus, these findings show that the adaptive adjustment of the response deadlines leads to similar learning rates in younger and older adults.

3.5.2 ERP data

In the following, analyses of response-locked and feedback-locked ERPs will be presented. The ERPs were analyzed separately for correct and incorrect responses, as well as for positive and negative feedback. This was done because as Figures 11 and 12 show, the ERPs varied as a function of validity for correct as well as for incorrect responses (positive and negative feedback).[6] In the first part of the analyses, the mean amplitudes of the ERP components to correct and incorrect responses (positive and negative feedback) were investigated. In the second part, peak-to-peak measurements were used for an additional quantification of the ERN and FRN. The additional peak-to-peak measurements were necessary because the mean amplitude measures of these components are confounded by an overlapping positivity (see Figures 11 and 12). As for the accuracy data, learning-related effects in the ERP components were investigated by analyzing the slope parameters of the learning functions (for details, see Methods).

Response-locked ERPs.

Figure 11 shows the ERPs for correct and incorrect responses in the three validity conditions (100%, 80% and 50% validity), separately for younger and older adults at electrode FCz. In both age groups incorrect responses were followed by a phasic negativity, the error-related negativity (ERN) that seemed to be larger the more valid the feedback. However, as also apparent from Figure 11, correct responses were followed by a positivity that also varied as function of the feedback validity, being largest for the 100%, intermediate for the 80% and smallest for the 50% validity condition. This component will be termed response-locked positivity in the following. In older adults, superimposed on this response-locked positivity, a small negativity for correct trials (CRN) can be observed that seemed to get larger the more invalid the feedback. Figure 11 also displays the topographical distribution of the difference between correct and incorrect responses for all validity conditions and the two age groups. As can be seen in the topographical maps the difference wave is maximal at fronto-central electrodes, which is in line with ERN topographies reported in previous studies (Holroyd & Coles, 2002; Nieuwenhuis et al., 2002).

---

[6] Please note, that this procedure is in contrast to previous studies (Holroyd & Coles, 2002; Nieuwenhuis et al., 2002), in which the ERN and FRN were investigated by means of a difference wave approach.

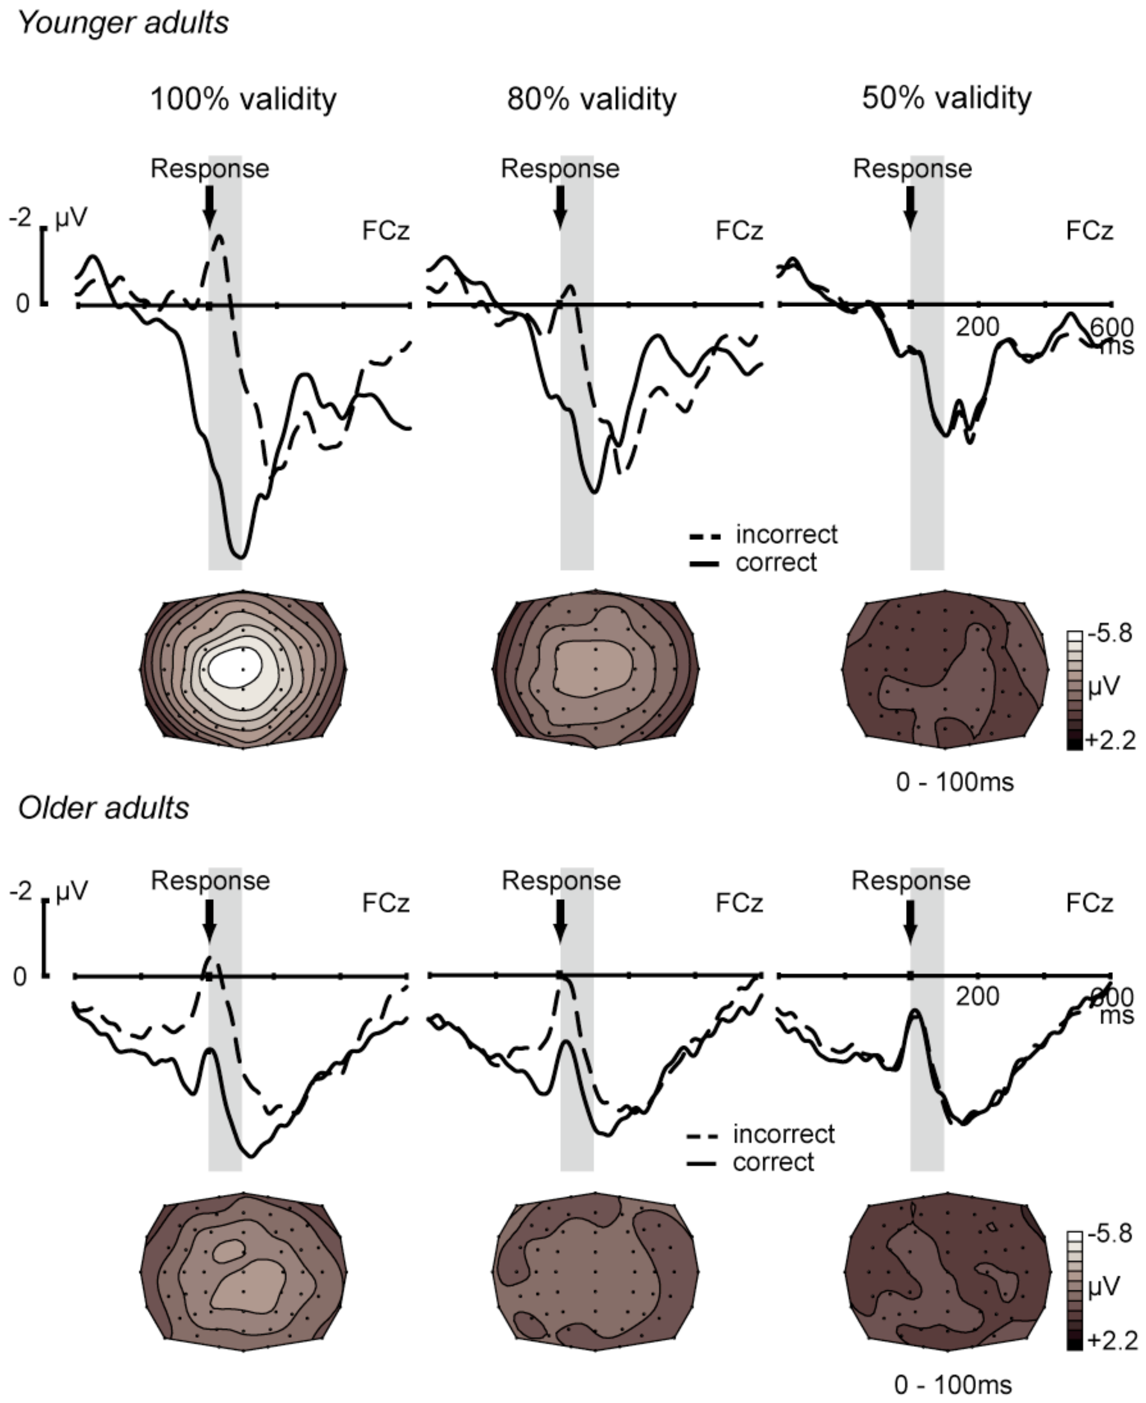

Figure 11: Response-locked ERPs for the three validity conditions, displayed separately for correct (solid lines) and incorrect (dashed lines) responses, for younger and older adults at the electrode FCz. Topographical distribution of the ERP difference wave for correct and incorrect responses displayed separately for the three validity conditions and the two age groups.

Analysis of correct and incorrect responses. The mean amplitude measures of the ERPs to correct and incorrect responses were analyzed using an ANOVA design involving the factors Age group (younger, older), Validity (100%, 80% and 50% validity), Response type (correct, incorrect) and Bin (Bin1 – Bin 4). The analysis showed a significant main effect of response type, F(1, 34) = 87.03, p < .0001, η^2 = .66 and an interaction between age group and response type, F(1, 34) = 11.19, p< .002, η^2 = .08. Moreover, a reliable interaction between response type and validity, F(2, 68) = 56.58, p < .0001, ε = .89, η^2 = .57 was obtained. Separate ANOVAs for the factor response type revealed significant main effects of validity for correct, as well as incorrect responses (p's < .0001, η^2's > .41). These findings show that the response-locked positivity for correct trials, as well as the error-related negativity (ERN), both got larger with increasing feedback validity (see Figure 11). Furthermore, a significant three-way interaction between age group, validity, and response type, F(2, 68) = 8.20, p < .001, ε = .90, η^2 = .08 was obtained. Yet, a significant interaction between age group and validity was only found for correct responses, F(2, 68) = 7.74, p < .002, ε = .81, η^2 = .11, but not for incorrect responses (p = .23). To further investigate the age differences in the response-locked positivity three post-hoc contrasts comparing each of the levels of the factor validity were performed. This analysis revealed significant age differences for correct responses in the 100% - 50% and the 100% - 80% contrasts (p's < .02, η^2 > .12). As can be seen in Figure 11, these results reflect the fact that the increase of the response-locked positivity with feedback validity is less pronounced in older compared to younger adults, suggesting that the elderly may have been less able to differentiate between the validity conditions (see Figure 11).

Learning-related effects in the response-locked positivity. Since the focus of this study was on the time course of learning, interactions involving the factor Bin were of most interest. Indeed, significant interactions between response type and bin, F(3, 102) = 10.09, p < .0001, ε = .90, η^2 = .22 and between validity, response type, and bin, F(6, 204) = 2.81, p < .03, ε = .71, η^2 = .08 were obtained. Separate analyses for the factor response type showed a significant main effect of bin, F(3, 102) = 13.89, p < .0001, ε = .83, η^2 = .29 and a significant interaction between validity and bin, F(6, 204) = 3.94, p < .002, ε = 83, η^2 = .10, only for correct trials. For incorrect trials, neither the main effect of bin, nor the interaction between validity and bin (p's > .29) was significant. Post-hoc tests for the factors response type and validity showed significant main effects of bin on correct trials for the 100% and 80% validity conditions (p's < .0001, η^2 > .19), but not for the 50% validity condition (p = .20). Hence, the present results do not provide evidence for learning-related changes in the ERN, as measured using mean amplitude values. In contrast, as shown in Figure 12, changes over the course of learning were only observed for the response-locked positivity on correct trials (see Figure 12).

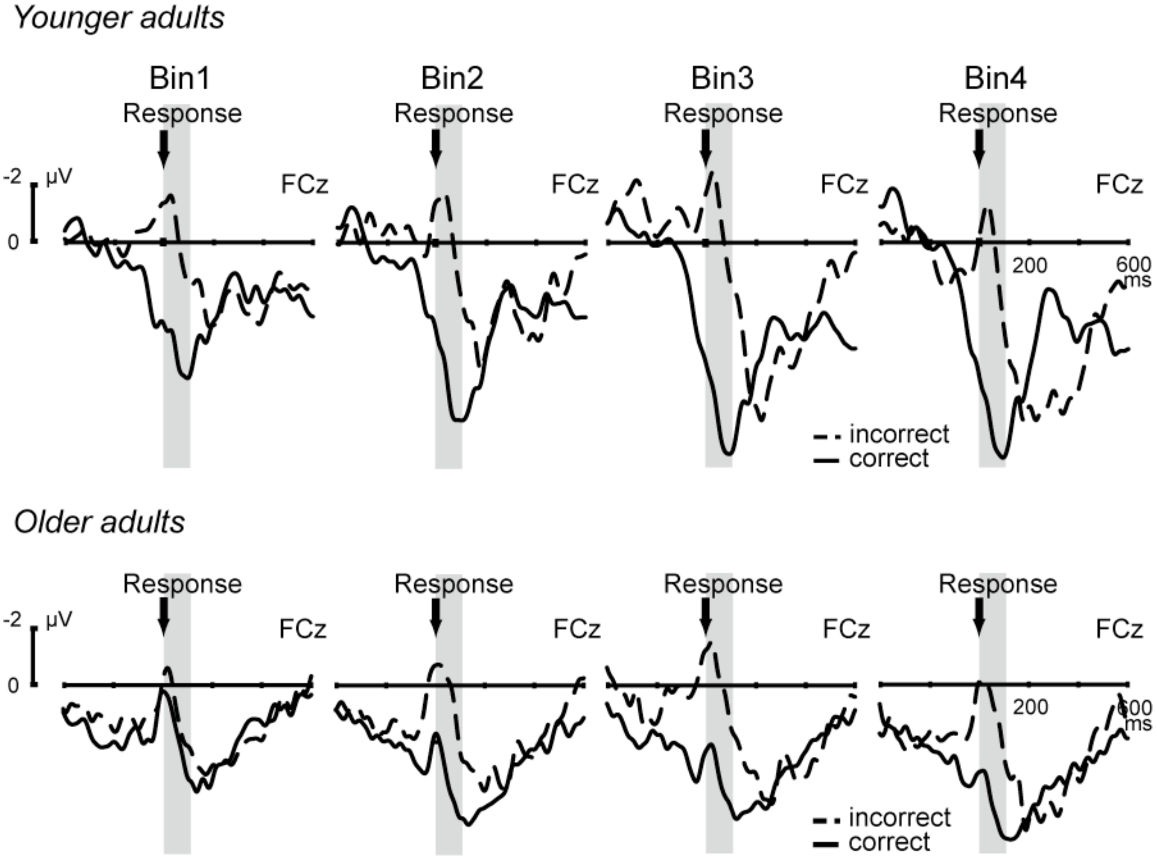

Figure 12: Response-locked ERPs over the course of learning (averaged into four bins) for the 100% validity condition displayed separately for correct (solid) and incorrect (dashed) trials for younger (top) and older (bottom) adults at the electrode FCz.

In order to further investigate the learning-related effects in the response-locked positivity, the slope parameters of the learning functions (for details, see Methods) were subjected to an ANOVA involving the factors Age group and Validity. The analysis revealed a significant main effect of validity, $F(2, 68) = 21.21$, $p < .0001$, $\varepsilon = .93$, $\eta^2 = .38$. However, neither a significant main effect of age group (p = .21) nor a significant interaction between age group and validity (p = .43) was obtained. Post-hoc contrasts for each of the levels of the factor validity revealed significantly larger slope parameters for

100% and the 80% compared to the 50% validity condition ($\underline{p}$'s < .0001, $\underline{\eta}^2$ = .37) (see Figure 13).

Response-locked positivity

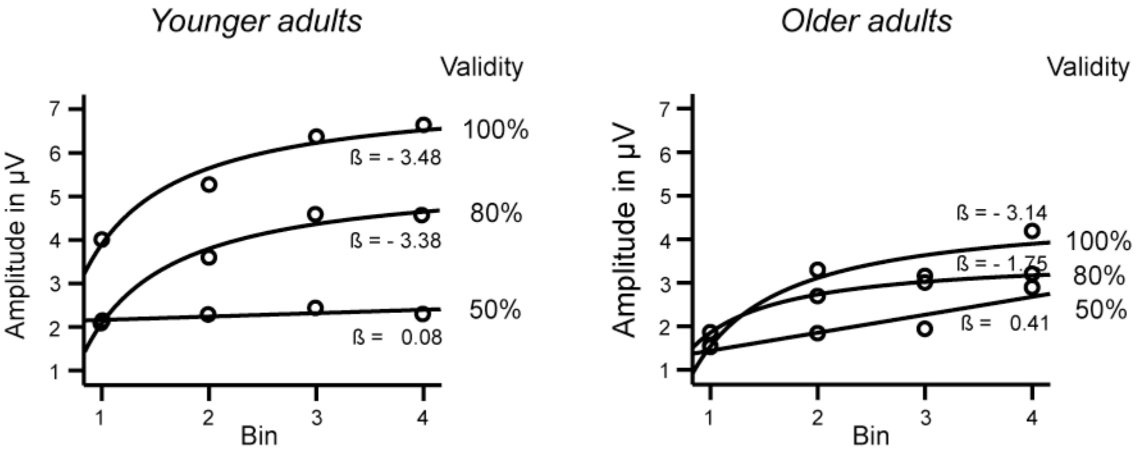

ERN

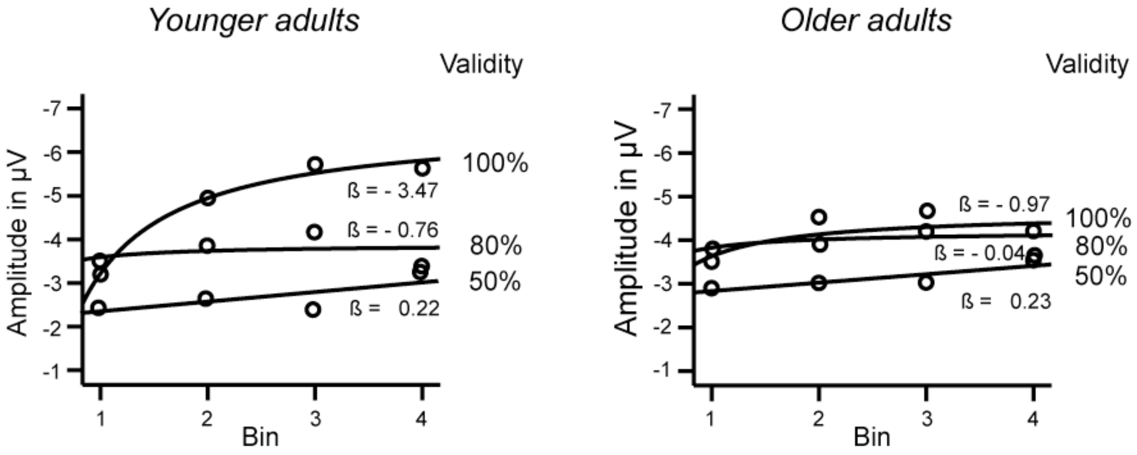

Figure 13: Learning curves for the response-locked positivity (top) and the ERN (bottom) for the three validity conditions, displayed separately for younger (left) and older (right) adults. The y-axis indicates the amplitude in µV, the x-axis shows the course of learning averaged into four bins of trials. The β-parameters indicate the steepness of the learning functions (for details, see Method).

Yet, the comparison between the 100% and the 80% condition was not significant ($\underline{p}$ = .14). The fact that this pattern of results was obtained for younger adults ($\underline{p}$'s < .002, $\underline{\eta}^2$'s = .46 for the 100% and 80% validity conditions), as well as for older adults ($\underline{p}$'s < .02, $\underline{\eta}^2$ = .29), indicates that both age groups showed comparable learning-related effects in the response-locked positivity for the two learning conditions (see Figure 13).

Peak-to-peak analysis of the ERN. The peak-to-peak measures of the ERN were analyzed using an ANOVA with the factors Age group, Validity, and Bin. The ANOVA showed a significant main effect of validity, $F(2, 68) = 9.10$, $p < .002$, $\varepsilon = . 70$, $\eta^2 = .20$, a significant main effect of bin, $F(3, 102) = 2.94$, $p < .04$, $\varepsilon = .87$, $\eta^2 = .08$, and a significant interaction between validity and bin, $F(6, 204) = 2.45$, $p < .04$, $\varepsilon = .84$, $\eta^2 = .07$. In separate analyses for the factor validity, a significant main effect of bin only was observed for the 100% validity condition, $F(3, 102) = 4.31$, $p < .01$, $\varepsilon = .79$, $\eta^2 = .11$. No significant main effect of bin was obtained for the 80% or 50% validity conditions (p's > .22) (see Figure 13). This indicates that in contrast to the analysis of the mean amplitude measures, the peak-to-peak analysis of the ERN showed a significant learning-related increase, however, only in the 100% validity condition (see Figure 13). Taken together, the present data suggest that learning-related effects, though present in the ERN (when measured peak-to-peak), are much more pronounced in the response-locked positivity for correct trials. Moreover, it is important to note, that the peak-to-peak analysis did neither reveal a significant main effect of age ($p = .93$) nor any significant interactions involving the factor age group (p's > .26). Hence, there is no evidence for a reduction of the ERN in older adults in the present data, neither with peak-to-peak nor with mean amplitude measures (see Figure 11). This result is in contrast to several recent findings, which pointed to an age-related reduction of the ERN (Band & Kok, 2000; Falkenstein et al., 2001; Mathewson et al., 2005; Nieuwenhuis et al., 2002; West, 2004).

Peak-to-peak analysis of the CRN. Since the correct response negativity (CRN) could not be measured reliably in younger adults (see Figure 11), the analysis was focused on the CRN in older adults. The CRN was analyzed using an ANOVA involving the factors Validity and Bin. The analysis revealed a significant main effect of validity, $F(2, 34) = 4.41$, $p < .02$, $\varepsilon = .98$, $\eta^2 = .21$. Contrasts for each of the levels of the factor validity showed that the CRN was increased for the 50% compared to the 100% validity condition ($p < .01$, $\eta^2 = .32$). This finding suggests that the CRN in older adults was larger the more invalid the feedback, indicating that the elderly were less certain about the appropriate response (see Figure 11). Moreover, the analysis revealed a significant interaction between validity and bin, $F(6, 102) = 3.14$, $p < .01$, $\varepsilon = .78$, $\eta^2 = .16$ and separate ANOVAs for the factor validity showed a significant effect of bin only for the 50% validity condition ($p < .02$, $\eta^2 = .20$). Post-hoc contrasts for each of the bins in the 50% validity condition showed that the CRN was reduced at the end of the learning blocks in the fourth bin, ($M = -3.99$ µV, $SD = 1.78$ µV) compared to the third bin, ($M = -2.94$ µV, $SD = 1.84$ µV), ($p < .006$, $\eta^2 = .36$).

Feedback-locked ERPs.

Figure 14 displays the ERPs for positive and negative feedback and the topographical distribution of the difference between both feedback types in the three validity conditions (100%, 80%, and 50% validity), separately for younger and older adults. For younger adults, a pronounced feedback-related negativity (FRN) for negative compared to positive feedback can be observed for all validity conditions. In contrast, for older adults the FRN is strongly reduced for all validity conditions (see Figure 14). As also illustrated in Figure 14, in younger adults the difference between positive and negative feedback shows a fronto-central distribution and gets larger the more invalid the feedback. In contrast, for older adults no such effect can be observed (see Figure 14). Similar to the response-locked ERPs, learning-related effects were most evident in a positivity for positive feedback, which will be called feedback-locked positivity in the following. In contrast, the FRN remained stable over the course of learning (see Figure 15).

Analysis of positive and negative feedback. For the analysis of the ERPs to positive and negative feedback an ANOVA design involving the factors Age group (younger, older), Validity (100%, 80% and 50% validity), Feedback type (positive, negative) and Bin (Bin1 – Bin 4) was applied. The analysis revealed a significant main effect of age, $F(1, 34) = 6.68$, $p < .01$, $\eta^2 = .16$, a significant main effect of feedback type, $F(1, 34) = 32.29$, $p < .0001$, $\eta^2 = .39$, and a significant interaction between age and feedback type, $F(1, 34) = 15.98$, $p < .0003$, $\eta^2 = .19$. Separate ANOVAs for the two age groups showed a significant main effect of feedback type for younger adults, $F(1, 17) = 42.93$, $p < .0001$, $\eta^2 = .72$, but not for older adults (p = .23). As illustrated in Figure 14, this finding suggests a differential sensitivity to negative and positive feedback between younger and older adults.

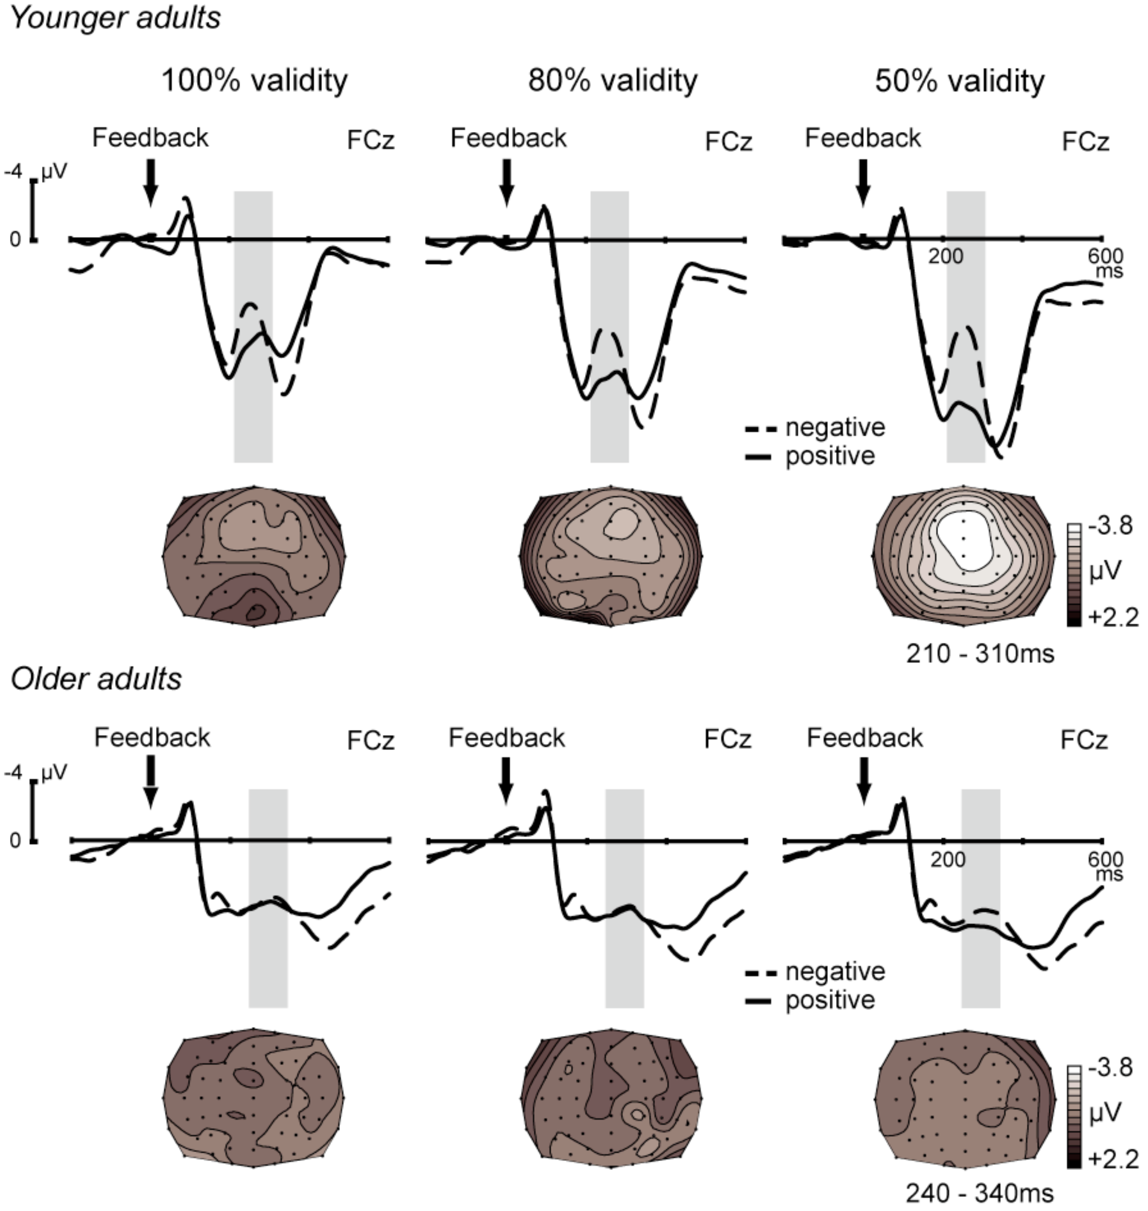

Figure 14: Feedback-locked ERPs for the three validity conditions displayed separately for positive (solid lines) and negative (dashed lines) feedback for younger (top) and older (bottom) adults at the electrode FCz. Topographical distribution of the ERP difference wave for positive and negative feedback displayed separately for the three validity conditions and the two age groups.

Moreover, a significant interaction between feedback type and validity, $F(2, 68) = 7.85$, $p < .001$, $\varepsilon = .91$, $\eta^2 = .18$ and a marginally significant interaction between age group feedback type and validity $F(2, 68) = 2.96$, $p < .06$, $\varepsilon = .91$, $\eta^2 = .07$ was obtained. Separate analyses for the two age groups revealed a significant interaction between feedback and validity for younger adults ($p < .003$, $\eta^2 = .29$), but not for older adults ($p = .23$). As shown in Figure 14, this pattern of results reflects the fact that for younger adults the effects of feedback type were larger the more invalid the feedback, which was not the case for older adults.

Learning-related effects in the feedback-locked ERPs. Again, the time course of learning in the feedback-locked ERPs was of most interest. The analysis revealed a significant main effect of bin, $F(3, 102) = 8.44$, $p < .0003$, $\varepsilon = .73$, $\eta^2 = .19$, a significant interaction between validity and bin, $F(6, 204) = 7.42$, $p < .0001$, $\varepsilon = .90$, $\eta^2 = .17$ and a significant three-way interaction between feedback type, validity, and bin, $F(6, 204) = 2.94$, $p < .02$, $\varepsilon = .81$, $\eta^2 = .08$. Most interestingly, separate ANOVAs for the factor feedback type revealed a significant interaction between validity and bin only for positive feedback, $F(6, 204) = 11.62$, $p < .0001$, $\varepsilon = .78$, $\eta^2 = .25$, but not for negative feedback ($p = .31$). Post-hoc tests for the factors feedback type and validity revealed significant effects of bin for positive feedback for the 100% and 80% validity conditions (p's $< .0005$, η^2's $> .18$), but not for the 50% validity condition ($p = .49$).

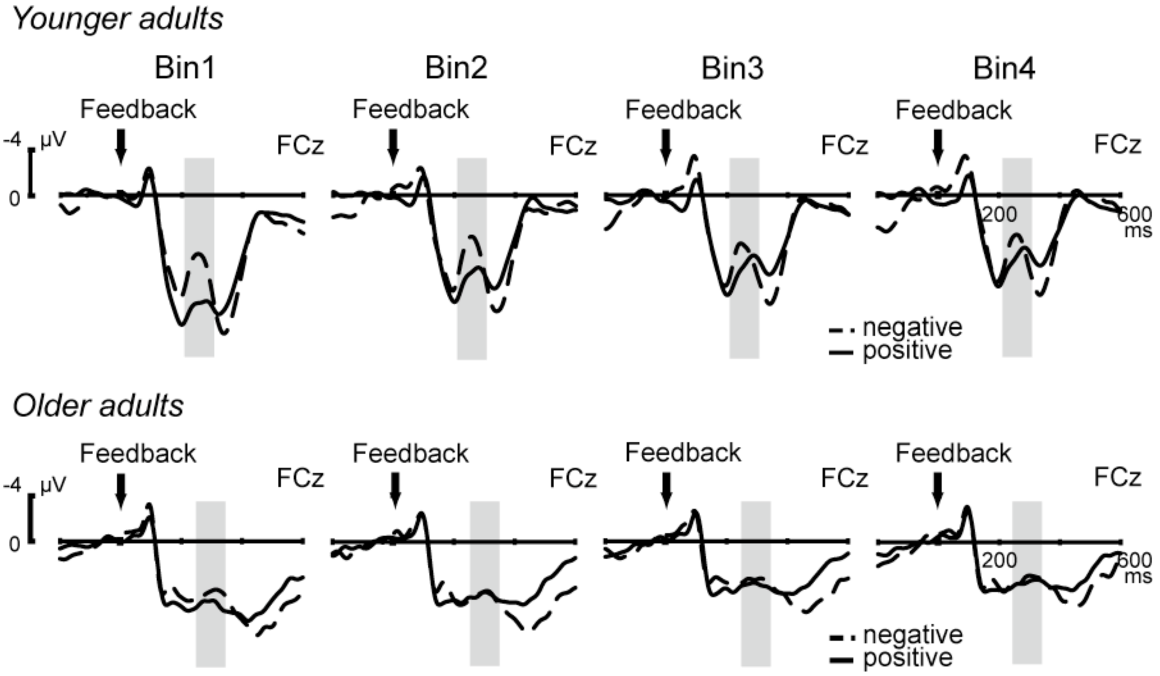

Figure 15: Feedback-locked ERPs over the course of learning (averaged into four bins) for the 100% validity condition displayed separately for positive (solid) and negative (dashed) feedback for younger (top) and older (bottom) adults at the electrode FCz.

Thus, the feedback-locked positivity decreased with learning in the both learning conditions, whereas no learning effect was obtained for the FRN (see Figures 15 and 16). Hence, similar to the response-locked ERPs learning-related effects were only obtained for positive, but not for negative feedback (see Figure 15).

In order to quantify the learning-related effects in the feedback-locked positivity the slope parameters of the learning functions (for details, see Methods) were subjected to an ANOVA involving the factors Age group and Validity. The analysis showed a significant main effect of validity, $\underline{F}(2, 68) = 20.65$, $\underline{p} < .0001$, $\underline{\varepsilon} = .96$, $\underline{\eta}^2 = .36$ and a significant interaction between age group and validity $\underline{F}(2, 68) = 3.11$, $\underline{p} < .05$, $\underline{\varepsilon} = .96$, $\underline{\eta}^2 = .05$. Separate ANOVAs for the two age groups showed significant effects of validity for younger adults ($p < .0001$, $\underline{\eta}^2 = .48$), as well as older adults ($p < .009$, $\underline{\eta}^2 = .26$). Post-hoc contrasts for each of the levels of the factor validity showed that the slope parameters were larger for the 100% compared to the 80% validity condition ($\underline{p} < .0008$, $\underline{\eta}^2 = .26$), as well as for the 80% compared to the 50% validity condition ($\underline{p} < .009$, $\underline{\eta}^2 = .19$) (see Figures 15 and 16). These findings point to the view that the learning-related effects in the feedback-locked positivity were the larger the more valid the feedback, and were more pronounced in younger compared to older adults.

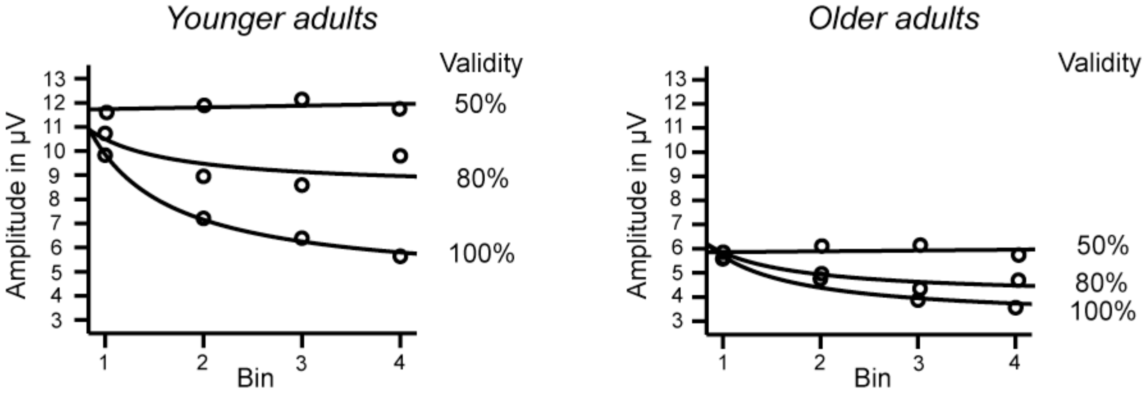

Figure 16: Learning curves for the feedback-locked positivity for the three validity conditions displayed separately for younger (left) and older (right) adults. The y-axis indicates the amplitude in µV, the x-axis shows the course of learning averaged into four bins of trials. The β-parameters indicate the steepness of the learning functions (for details, see Methods).

Peak- to-peak analysis of the FRN. The peak-to-peak measures of the FRN (for details, see Methods) were subjected to an ANOVA involving the factors Age group, Validity, and Bin. This analysis only revealed a significant main of age, $\underline{F}(1, 34) = 10.33$, $\underline{p} < .003$, $\underline{\eta}^2 = .23$, which reflects the strongly reduced FRNs for older compared to younger adults (see Figure 14). However, neither the main effects of validity or bin, nor their interaction or interactions with age turned out to be significant ($\underline{p}$'s > .15).

3.6. Summary

In line with previous results (Holroyd & Coles, 2002; Nieuwenhuis et al., 2002), the analysis of the accuracy data revealed that response accuracy increased with feedback validity, being largest for the 100% validity condition. Age differences were only obtained in the 80% validity condition, in which older adults showed a reduced overall accuracy. An analysis of the time course of learning in the 80% condition showed that age differences were most pronounced at beginning of learning, but absent at the end of learning. In contrast, in the 100% validity condition no age differences were obtained, not even at the beginning of learning. These findings show that in contrast to the study by Nieuwenhuis et al. (2002) in the present experiment performance levels were equated in the 100% condition, and in the 80% validity condition at the end of learning. However, these findings also indicate that older adults are impaired when invalid information interferes with learning (in the 80% condition) and that this impairment is most pronounced at the beginning of learning (see Figure 10). In contrast to overall accuracy, no age differences were obtained for the slope parameters of the learning functions, which were comparable for the two age groups (see Figure 10). This finding indicates that the adaptive adjustment of the response deadlines yields similar learning rates in younger and older adults.

Consistent with the increase of response accuracy with feedback validity, the difference between the ERPs to correct and incorrect responses also increased the more valid the feedback (see Figure 11). This is in line with previous results (Holroyd & Coles, 2002; Nieuwenhuis et al., 2002) and suggests that the more participants learn the more they are able to internally represent the correctness of the response. In contrast to the aforementioned studies, which used a difference wave approach to study learning-related changes in the ERN, in the present study the ERPs were analyzed separately for correct and incorrect responses. As illustrated in Figure 11, the results of this analysis show that the negativity to incorrect responses (ERN), as well as the positivity to correct responses (response-locked positivity) both increase with feedback validity. In contrast to the R-L theory, these findings point to the view that the reward-related variance in the response-locked ERPs is driven by correct as well as incorrect responses. Moreover, age differences were only obtained for the response-locked positivity, which showed a less pronounced increase with feedback validity for older compared to younger adults (see Figure 11). This suggests that the elderly may have been less able to differentiate between the validity conditions. However, this finding was in part due to the fact that in older adults there was a CRN superimposed on the response-locked positivity (see Figure 11). The CRN in older adults was larger the more invalid the feedback, which suggests that older adults were less certain about the appropriate response when feedback was invalid.

Interestingly, the analysis of the mean amplitude measures for correct and incorrect responses also revealed learning-related changes. However, in contrast to the predictions of the R-L theory (Holroyd & Coles, 2002) this was only the case for the response-locked positivity on correct trials, but not for the ERN. A learning-related increase in the ERN was only found in the peak-to-peak analysis, in which the component is captured most precisely. Moreover, in this analysis changes in the ERN with learning were only obtained when feedback was fully valid (see Figure 13). These findings show that learning-related changes in the response-locked ERPs are more pronounced on correct compared to incorrect trials and by this provide an important extension to recent theoretical accounts that focused on the role of errors and negative feedback for learning (Holroyd & Coles, 2002; Nieuwenhuis et al., 2002).

Most important, the present data does not show evidence for a reduction of the ERN in older adults, neither with peak-to-peak nor with mean amplitude measures (see Figure

11). This stands in contrast to several recent findings (Band & Kok, 2000; Falkenstein et al., 2001; Mathewson et al., 2005; Nieuwenhuis et al., 2002; West, 2004) and suggests that older adults show similar ERN amplitudes as younger adults if performance levels are equated between age groups.

In line with previous findings (Holroyd & Coles, 2002; Nieuwenhuis et al., 2002) the analysis of the feedback-locked ERP data showed that older adults, in contrast to younger adults, did not differentiate between positive and negative feedback. Moreover, whereas in younger adults the difference between positive and negative feedback increased the more invalid the feedback, no such effect was obtained for older adults (see Figure 14). The peak-to-peak analysis showed that these effects were due to the FRN, which was strongly reduced for the elderly in the present study. This suggests that older adults are less sensitive to negative feedback than younger adults and points to an age-related asymmetry in feedback processing (see Figure 14). Similar to the response-locked ERPs, a separate analysis for positive and negative feedback showed that learning-related effects were only obtained for positive, but not for negative feedback. Thus, it was the feedback-locked positivity and not the FRN that decreased with learning (see Figure 15). The learning-related decrease of the feedback-locked positivity indicates that the more participants are able to internally predict the correctness of their response (with learning) the less they rely on the external feedback. An analysis of the ERP learning functions revealed that the decrease of the feedback-locked positivity with learning was the larger the more valid the feedback. This effect was less pronounced for older compared to younger adults, indicating that the elderly were less able to disengage from processing positive feedback during learning (see Figure 16).

4 Experiment 2

4.1 Statement of problem

The second Experiment aims to provide insights into the question, which role the processing of internal and external error information plays for reinforcement learning during childhood development. More precisely, the aim of this study is to investigate, whether children differ from younger adults in their ability to use internal and external error signals for learning and whether this shows up in the ERP correlates of error processing (the ERN and FRN, respectively). Again, one major aspect of this study was to avoid potentially confounding effects of performance differences between children and younger adults (see Hogan et al., 2005; Kim et al., 2007).

4.2 Design

Apart from some minor changes that were necessary to adapt the learning task for the children (for details see Methods), a highly similar version of the learning paradigm that has already been used in the first experiment was applied. That is, the learning task again involved the three validity conditions and the same algorithm was used to adaptively adjust the response deadline (for details, see Methods).

4.3 Hypotheses

The adaptive response deadline procedure should enable children to learn comparably to younger adults, at least in the 100% validity condition. However, similar to older adults (see Experiment 1), it was expected that children would be impaired in accuracy when invalid feedback interferes with learning in the 80% validity condition. On the basis of recent findings (Davies et al., 2004, Ladouceur et al., 2004; Santesso et al., 2006) and the predictions of the RL-theory (Holroyd & Coles 2002), children should be impaired in the processing of internal error information, as reflected in reduced ERN components during learning. However, given the results from the first Experiment and the findings by Hogan et al. (2005), it is seems questionable whether there are still age differences in the ERN when performance levels are equated between age groups. Little is known about age differences in the FRN between children and adults. However, according to the findings on developmental differences in heart rate during feedback processing (Crone et al., 2004) and the FRN in children with ADHD (van Meel et al., 2006), children could be expected to be less able to disengage from external error information and show increased FRN components for all validity conditions, as well as no change of the FRN with learning.

4.4 Method

Participants. Twenty-one younger adults[7] and twenty-one children participated in the study. All subjects received 22.5 Euro for their participation. One younger adult had to be excluded from data analysis due to technical problems during data acquisition. Two younger adults were excluded because they did not commit enough error trials to analyze the error-related ERPs over the course of learning. One child felt so uncomfortable with the EEG setting that the experimental session had to be stopped. Three children had to be eliminated from further analyses because they performed at chance level even in the 100% validity condition (M = 0.48, SD = 0.03) and responded much faster (M = 278 ms, SD = 22) than the mean of the children group (M = 404 ms, SD = 128). These children probably pressed the button before they were able to fully perceive the stimulus.

The effective sample consisted of eighteen younger adults and seventeen children (see Table 4). According to self-report all participants were healthy, had a right-hand preference, no color blindness, and no history of neurological or psychiatric problems. The participants performed two psychometric tests, one from the domain of fluid intelligence (the Digit-Symbol Substitution test; adapted from Wechsler, 1982) and one from the domain of crystallized intelligence (the Spot-a-Word test; adapted from Lehrl, 1977). Adults reached a higher score than children on the Digit-Symbol Substitution test, $F(1, 33) = 20.57$, $p < .0001$, $\eta^2 = .38$ (see Table 4). A similar pattern was obtained for the Spot-a-Word test, in which adults performed better than children, $F(1, 33) = 102.07$, $p < .0001$, $\eta^2 = .75$ (see Table 4). Consistent with several other studies (e.g., Cepeda, Kramer, & Gonzalez de Sather, 2001; Kray, Eber, & Lindenberger, 2004) and the two-component model of intelligence (Baltes, Staudinger, & Lindenberger, 1999), these findings suggest

---

[7] Note that the sub-sample of younger adults already served as a control group in the first experiment.

that the speed of information processing, as well as semantic knowledge increase from childhood to adulthood.

	Younger adults	***Children***
N / Gender	18 / 9 female	17 / 9 female
Age Range	19 – 26	10 – 12
Mean Age (M / SD)	20.8 (1.8)	11.4 (0.8)
Digit-Symbol Substitution test (M / SD)	61.9 (7.7)	48.1 (10.3)
Spot-a Word test (M / SD)	25.2 (3.1)	13.2 (3.9)

Table 4: Demographic characteristics of the sample in Experiment 2. Digit-Symbol Substitution test, adapted from Wechsler, (1982), Spot-a-Word test; adapted from Lehrl, (1977).

Stimuli and Task. As in the first experiment, stimuli were presented in color against a dark grey background on a 17-inch computer screen. The stimuli were identical with those used in Experiment 1. The subjects were asked to make a two-choice decision upon presentation of the imperative stimulus and to press one of two response keys (C and M on a standard computer keyboard). Further, they were instructed to infer the stimulus-response mappings by trial and error, based on the feedback information. To motivate the children, a cover story similar to a Donald Duck comic was constructed. The children were told that they should help Scrooge McDuck to sort objects into two safes (represented by the two response buttons) to protect them from the “Beagle Boys”. They were instructed to use the feedback to learn, which object belongs to which safe (response button). To further motivate the children, there was a short break of 15 seconds in the middle of each learning block. In this break a monetary feedback was displayed that indicated what they had already won (the feedback was independent of their performance). At the end of each learning block, monetary feedback was displayed to all subjects. This monetary feedback depended on the mean performance in the 100% condition and participants could win between 50 Euro Cents and 450 Euro Cents. Children additionally received the amount of money they had won in form of chocolate coins.

Experimental Design. The experimental design was the same as in first experiment and involved three learning conditions in which feedback validity was manipulated (for further information, see Methods of Experiment 1 and Figure 8).

Trial Procedure. The trial procedure was similar to the one that was used in the first experiment (see Methods of Experiment 1 and Figure 9). Furthermore, a similar adaptive procedure was applied in order to individually adjust the response deadlines (see Method Experiment 1). However, in contrast to younger adults, for whom the response deadline was adapted in a range of 600 to 1000 ms, for children response deadlines in the range of 800 to 1200 ms were applied[8]. This was done in order to account for their larger variability in response times (see Leth-Steensen, Elbaz, & Douglas, 2000; Williams, Hultsch, Strauss, Hunter, & Tannock, 2005). As in the previous experiment the purpose of the adaptive procedure was to make sure that all subjects produced a similar proportion of time outs (M = .02, SD = .01, for younger adults, M = .03, SD = .03, for children), and thereby had a similar opportunity to learn from feedback.

Procedure. The experimental procedure lasted approximately three hours. First, each participant (the parents in case of the children) filled out an informed consent and a short demographic questionnaire. Then they performed the two psychometric tests. The experiment consisted of one practice block and five experimental blocks. Each block involved a new set of six imperative stimuli, which were drawn randomly (without replacement) from the six stimulus categories. In a practice block (150 trials) the participants were familiarized with the experimental setting. Then they performed five

---

[8] Each participant started with a response deadline of 800 ms. After the first trial the algorithm kept track of the proportion of time-out trials (number of time-out trials relative to the trials performed) and adjusted the response deadline in steps of 100 ms. If the proportion of time-out trials was smaller than two percent, a response deadline of 600 ms (adults) or 800 ms (children) was applied. With steps of two percent, the response deadline increased for 100 ms and reached a maximum deadline of 1000 ms (adults) or 1200 ms (children) with over eight percent of time-out trials.

experimental blocks. In the experimental blocks each of the six imperative stimuli were presented 50 times in random order. Thus, each participant performed 300 trials per validity condition, which corresponds to an overall trial number of 1500 trials.

Data Recording. The recording parameters of the behavioral data and the EEG data were identical to Experiment 1 (see Data Recording Experiment 1).

Data Analysis. Accuracy data. Responses faster than 140 ms (more than two standard deviations from the mean reaction time in both age groups) and responses that exceeded the response deadline (mean response deadline in younger adults: M = 706 ms, SD = 117 ms; mean response deadline in children: M = 897 ms, SD = 119 ms) were excluded from data analysis. The analysis of variance (ANOVA) was based on accuracy (% correct)[9]. To analyze the behavioral learning effects, mean accuracy was averaged individually for each subject and validity condition into four bins (of 75 trials), reflecting the four quarters of the learning blocks. Mean accuracy rates are displayed as a function of age group, validity condition, response type, and bin in Table 5 (see Appendix).

In order to quantify the learning-related changes in the accuracy data, learning curves were fitted separately for each subject and the three validity conditions using a linear function ($Y = b0 + (b1 * t)$), as implemented in SPSS. A linear function was used, since it fitted the data in children more adequately than the inverse function used in the first experiment. The slope (b1- or β-) parameters of these learning functions were then subjected to the analyses of variance. The mean fit parameters (R^2) and the mean slope parameters (β) of the learning functions are displayed separately for the two age groups and the three validity conditions in Table 6 (see Appendix).

ERP data. The EEG epochs were averaged with respect to response and feedback onset to obtain response-locked and feedback-locked ERPs. Similar to previous studies

---

[9] Note that the accuracy rates in the 80% validity condition reflect the mean accuracy for the 80% valid trials of this condition. For the 20% invalid trials mean accuracy was lower than chance (M = 34, SD = .11 for younger adults; M = 43, SD = .12 for children) since participants learned to respond to the dominant (but here incorrect) mapping.

on developmental differences in the ERN (Hogan et al., 2005; Ladouceur et al., 2004; Wiersema et al., 2007) the response-locked EEG data were baseline corrected by subtracting the average activity during -200 and -50 ms preceding the response. For the feedback-locked EEG data the average activity from -100 ms to feedback onset served as baseline. As in the previous study ocular artifacts or other artifacts were excluded from further analysis and remaining eye movements were corrected using a modified version the approach developed by Gratton, Coles, and Donchin (1983) (see Data Analysis of Experiment 1).

The response-locked ERPs were measured as the mean amplitudes in a 0 - 100 ms time window following the response at electrode FCz. The feedback-locked components were measured as the mean amplitudes within a 100 ms time window centered on the peak of the FRN at the electrode FCz (260 ms in younger adults and 290 ms in children). For the peak-to-peak analyses response-locked, as well as feedback-locked EEG data were filtered using a 15Hz low-pass filter in order to obtain more reliable peak amplitude measures. As in the first experiment the ERN was defined as the peak- to-peak voltage difference between the most negative peak between -50 and 150 ms around the response and the preceding positive peak. The FRN was defined as the difference between the most negative peak within 200 to 400 ms and the preceding positive peak (for a similar procedure see Frank et al. (2005) and Yeung & Sanfey, (2004)).

Whenever necessary, the Geisser-Greenhouse correction was applied (Geisser & Greenhouse, 1958). In these cases the original F-value, the adjusted p-values, and the Epsilon values (ε) are reported. Furthermore, effects sizes (eta squared, η^2) are reported, which reflect the proportion of variance that is accounted for by the experimental manipulations (see Cohen, 1973). As for the behavioral data, ERPs were averaged into four bins reflecting the four quarters of the learning blocks. In order to quantify the learning-related changes, a linear ($Y = b0 + (b1 * t)$) learning function (as implemented in SPSS) was fitted to each individual's learning curves, separately for the three validity conditions. The slope (b1- or β-) parameters of these learning functions were then subjected to the analyses of variance. The mean fit parameters (R^2) and the mean slope parameters (β) of the learning functions are displayed separately for the two age groups and the three validity conditions in Table 6 (see Appendix).

4.5 Results

4.5.1 Accuracy data

The accuracy data (see Figure 17) was analyzed using the same ANOVA design as in Experiment 1, involving the factors Age group, Validity, and Bin.

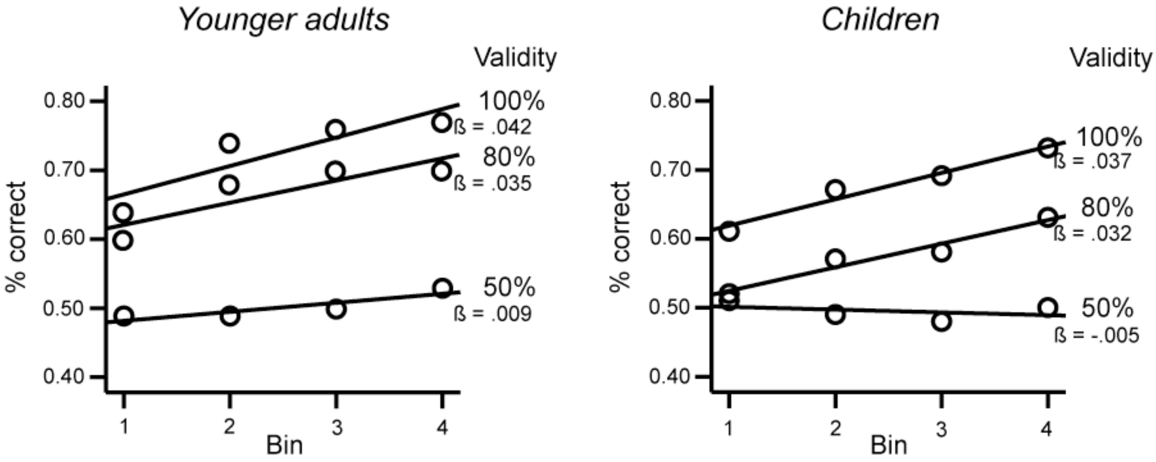

Figure 17: Accuracy learning curves for the three validity conditions (100%, 80%, and 50% validity) displayed separately for adults and children.

The analysis showed a significant main effect of age group, $F(1, 33) = 8.37$, $p < .007$, $\eta^2 = .20$, and a significant main effect of validity, $F(2, 66) = 102.02$, $p < .0001$, $\varepsilon = .96$, $\eta^2 = .73$. Contrasts for each of the levels of the factor validity showed that participants performed better in the 100% compared to the 80% validity condition and in the 80% compared to the 50% validity condition (p's $< .0001$, η^2's $> .45$). Hence, response accuracy increased with feedback validity. Moreover, the ANOVA showed a significant two-way interaction between age group and validity, $F(2, 66) = 4.59$, $p < .01$, $\varepsilon = .96$, $\eta^2 = .03$. Separate analyses for the factor validity revealed that children performed worse than adults in the 80% validity condition ($p < .0009$, $\eta^2 = .28$). However, neither for the 100%, nor for the 50% validity condition significant age differences were obtained (p's $>.12$).

Thus, similar to older adults, children showed impaired overall accuracy in the 80% validity condition, when invalid information occurred during learning.

Learning Effects. Again, as in the first experiment, age differences in the learning effects and hence interactions involving the factor bin were of most interest. The analysis showed a significant main effect of bin, $\underline{F}(3, 99) = 30.83$, $\underline{p} < .0001$, $\underline{\varepsilon} = .73$, $\underline{\eta}^2 = .47$, and a significant interaction between validity and bin, $\underline{F}(6, 198) = 14.32$, $\underline{p} < .0001$, $\underline{\varepsilon} = .70$, $\underline{\eta}^2 = .30$. Separate analyses for the factor validity indicated that learning took place in the 100% and the 80% validity conditions ($\underline{p\text{'s}} < .0001$, $\underline{\eta^2\text{'s}} > .37$). As expected, in the 50% validity condition no significant main effect of bin was obtained ($\underline{p} = .08$). In order to investigate developmental differences in accuracy over the course of learning, pair-wise comparisons for each of the levels of the factor validity were performed separately for the four bins. These contrasts revealed significant differences between all of the validity conditions in all of the bins ($\underline{p}$'s $< .0009$, $\underline{\eta^2\text{'s}} > .28$). However, these contrasts did not reveal significant age differences for the 100% condition compared to the 80% and 50% validity conditions in any of the four bins ($\underline{p}$'s $> .08$). In contrast, in line with the age differences in overall accuracy in the 80% condition, significant developmental differences were found for the 80% compared to the 50% condition for the first, second, and third bin ($\underline{p}$'s $< .01$, $\underline{\eta^2\text{'s}} > .07$). However, at the end of learning (in the fourth bin), no significant differences between children and adults in the 80% condition were obtained ($\underline{p} = .17$). These findings show that in the 100% validity condition no developmental differences in accuracy were obtained over the course of learning. In contrast, in the 80% condition age differences were most pronounced at the beginning of learning, but absent at the end of learning (see Figure 17).

To analyze the learning effects over the course of the four bins, an ANOVA on the slope parameters of the individual learning functions was performed (for details, see Method). This analysis revealed a significant main effect of validity, $\underline{F}(2, 66) = 25{,}57$, $\underline{p} < .0001$, $\underline{\varepsilon} = .78$, $\underline{\eta}^2 = .42$. Contrasts for each of the levels of the factor validity showed higher slope parameters for the 100% and the 80% validity condition compared to the 50% validity condition ($\underline{p\text{'s}} < .0001$, $\underline{\eta^2\text{'s}} > .40$). However, the comparison between the

100% and the 80% validity condition was not significant (p = .11), indicating that learning effects were comparable for both conditions. Most importantly, neither a significant main effect of age group, nor a significant interaction between age group and validity (p's > .29) was obtained, suggesting that children and adults showed comparable learning functions (see Figure 17).

4.5.2 ERP Data

Similar to the first experiment, response-locked and feedback-locked ERPs will be presented. In a first step the ERPs for correct and incorrect responses (positive and negative feedback) were analyzed using mean amplitude measures. In the second step peak-to-peak measurements were used for an additional quantification of the ERN and FRN. As for the accuracy data, learning-related effects in the ERP components were examined by analyzing the slope parameters of the individual learning functions (for details, see Method).

Response-locked ERPs

Figure 18 shows the response-locked ERPs for correct and incorrect responses in the three validity conditions (100%, 80% and 50% validity), separately for younger adults and children at the electrode FCz. In both age groups incorrect responses were followed by an error-related negativity (ERN) that seemed to be larger the more valid the feedback. However, as also apparent from Figure 18, correct responses elicited a positivity that also varied as a function of the feedback validity. Similar to the first experiment, this component will be termed response-locked positivity in the following. Generally, children showed a similar ERP pattern as younger adults. However, in the 80% validity condition, the difference between correct and incorrect responses in children was less pronounced than in younger adults. This can also be observed in the topographic maps for ERP difference between correct and incorrect responses (see Figure 18).

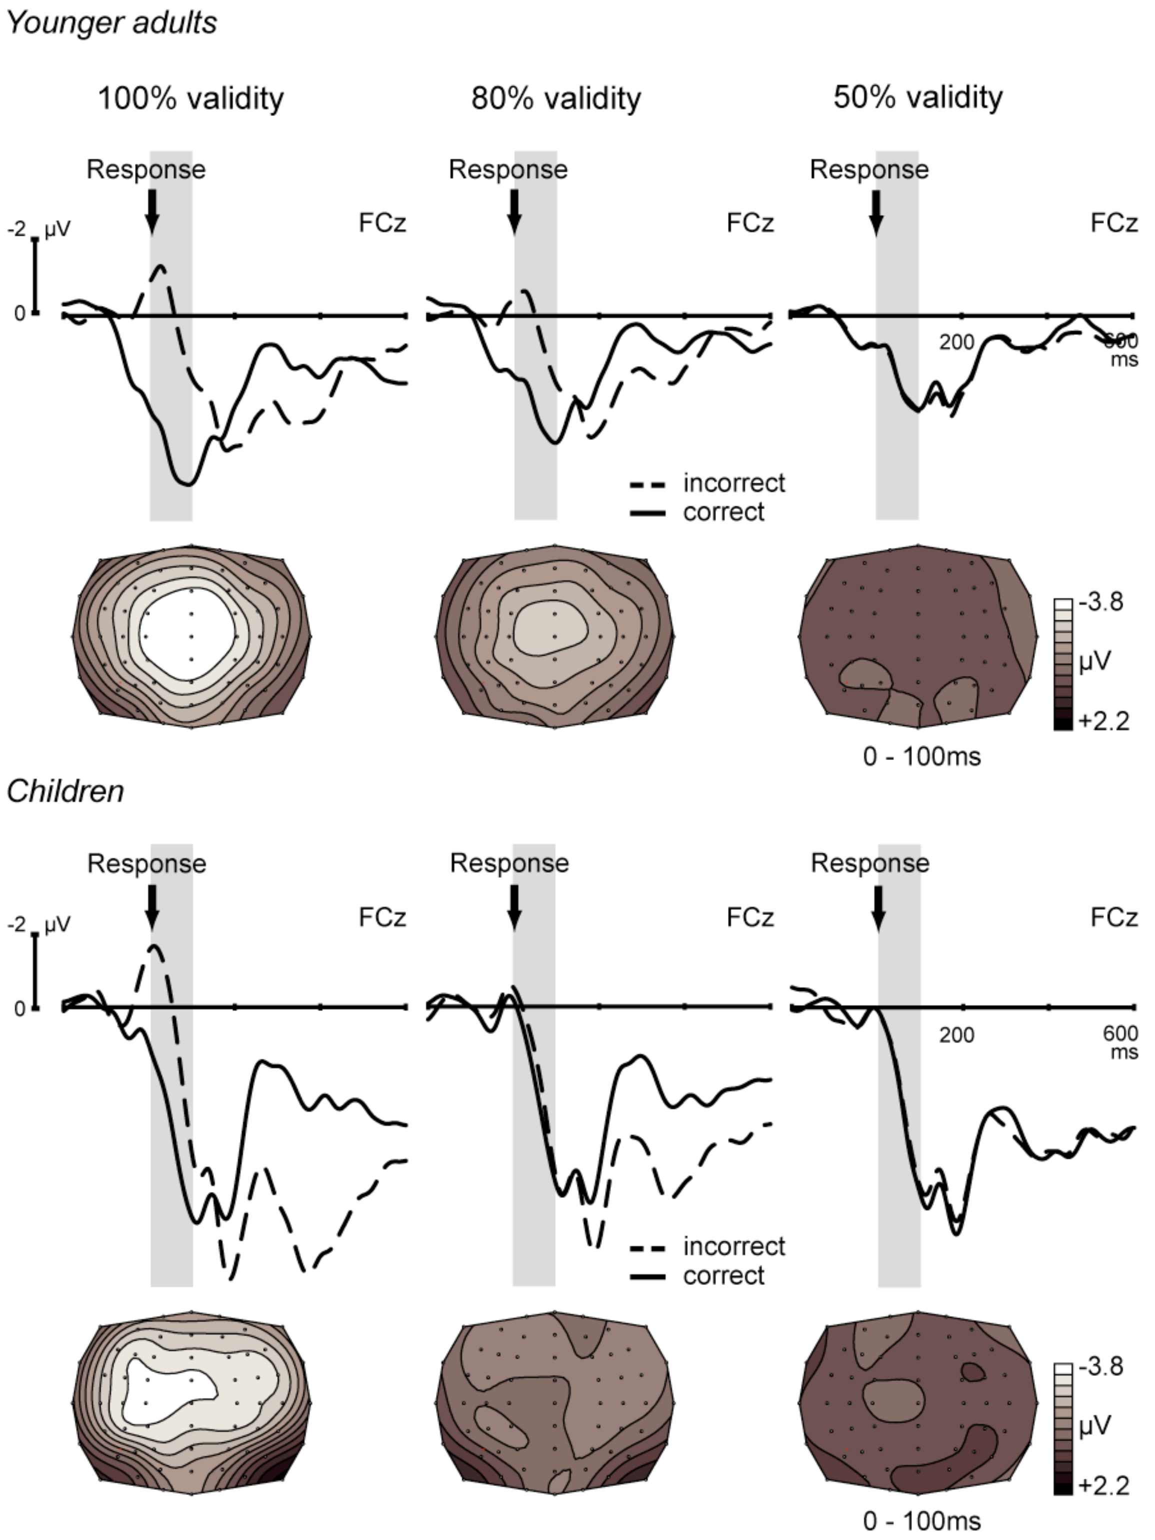

Figure 18: Response-locked ERPs for the three validity conditions, displayed separately for correct (solid lines) and incorrect (dashed lines) responses, for adults and children. Topographical distribution of the ERP difference wave for correct and incorrect responses displayed separately for the three validity conditions and the two age groups.

The ANOVA for the response-locked components involved the factors Age group, Validity, Response type, and Bin. The analysis showed a significant main effect of response type, $F(1, 33) = 33.17$, $p < .0001$, $\eta^2 = .48$, a significant interaction between response type and validity, $F(2, 66) = 36.41$, $p < . 0001$, $\varepsilon = 83$, $\eta^2 = .50$ and a marginally significant interaction between age group, response type, and validity, $F(2, 66) = 2.77$, $p < .08$, $\eta^2 = .04$. Separate analyses for the factors age group and validity revealed significant main effects of response type for younger adults in the 100% and the 80% validity conditions (p's $< .0003$, η^2's $> .55$). In contrast, for children a significant main effect of response type was only obtained for the 100% validity condition ($p < .001$, $\eta^2 = .50$), but not for the 80% validity condition ($p = .20$). As expected, no significant main effect of response type was obtained in the 50% condition for the two age groups (p's $> .72$). These findings show that children are as well able as adults to internally represent correct and incorrect responses in the 100% validity condition, whereas this representation is impaired when feedback is partially invalid in the 80% validity condition (see Figure 18). Separate ANOVAs for the factors age group and response type showed significant effects of validity for correct, as well as incorrect trials in both age groups (p's $< .003$, η^2's $> .30$), indicating that the response-locked positivity and the ERN both increased the more valid the feedback.

Learning-related Effects in the ERN and the Response-locked Positivity. Since the purpose of the second experiment was to study developmental differences in learning-related ERPs, the focus of this analysis was on interactions with the factor bin. The ANOVA showed a significant main effect of bin, $F(3, 99) = 6.30$, $p < .002$, $\varepsilon = .79$, $\eta^2 = .16$, a significant two-way interaction between response type and bin, $F(3, 99) = 3.33$, $p < .03$, $\varepsilon = .85$, $\eta^2 = .09$ and a significant three-way interaction involving the factors response type, validity, and bin, $F(6, 198) = 4.50$, $p < .001$, $\varepsilon = .77$, $\eta^2 = .12$. Separate analyses for the factor response type showed significant interactions between validity and bin for correct ($p < .02$, $\eta^2 = .07$), as well as incorrect trials ($p < .04$, $\eta^2 = .07$). Post-hoc tests for the factors response type and validity revealed a significant main effect of bin for correct trials in the 100% and 80% validity conditions (p's $< .0001$, η^2's $> .23$). In contrast, for incorrect trials a significant main effect of bin was obtained in the 50% validity condition (p

< .01, η^2 = .11). These findings reflect the fact that for correct trials the response-locked positivity increases over the course of learning for the 100% and 80% validity conditions (see Figure 19).

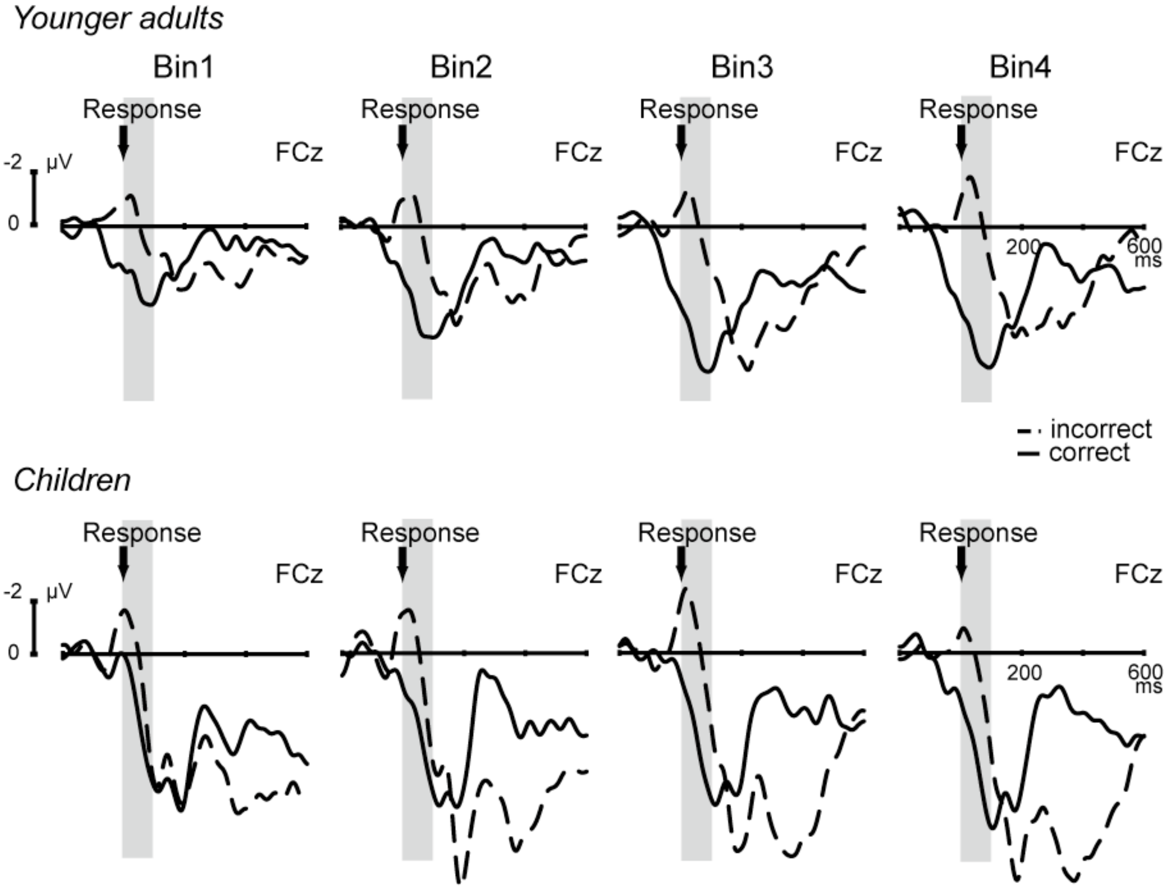

Figure 19: Response-locked ERPs over the course of learning (averaged into four bins) for the 100% validity condition displayed separately for correct (solid) and incorrect (dashed) trials for adults (top) and children (bottom) at the electrode FCz.

However, for incorrect trials an increasing positivity was found for the 50% validity condition, whereas for the two learning conditions (100% and 80% validity) no significant changes with learning were obtained (p's > .42). Hence, it appears that in contrast to the R-L theory (Holroyd & Coles, 2002), the response-locked positivity on correct trials, rather than the ERN, increases as a function of learning.

To analyze the learning-related effects in the response-locked positivity over the course of the four bins, an ANOVA on the slope parameters of the individual learning functions was performed (for details, see Method). This analysis revealed a significant main effect of validity, $F(2, 66) = 7.21$, $p < .002$, $\varepsilon = .97$, $\eta^2 = .17$. Contrasts for each of the levels of the factor validity showed higher slope parameters for the 100% and the 80% validity condition than for the 50% validity condition ($\underline{p}$'s < .008, $\underline{\eta}^2$'s > .19). However, the comparison between the 100% and the 80% validity condition was not significant ($\underline{p}$ = .62), suggesting that for the two learning conditions comparable learning effects were obtained in the response-locked positivity (see Figure 20).

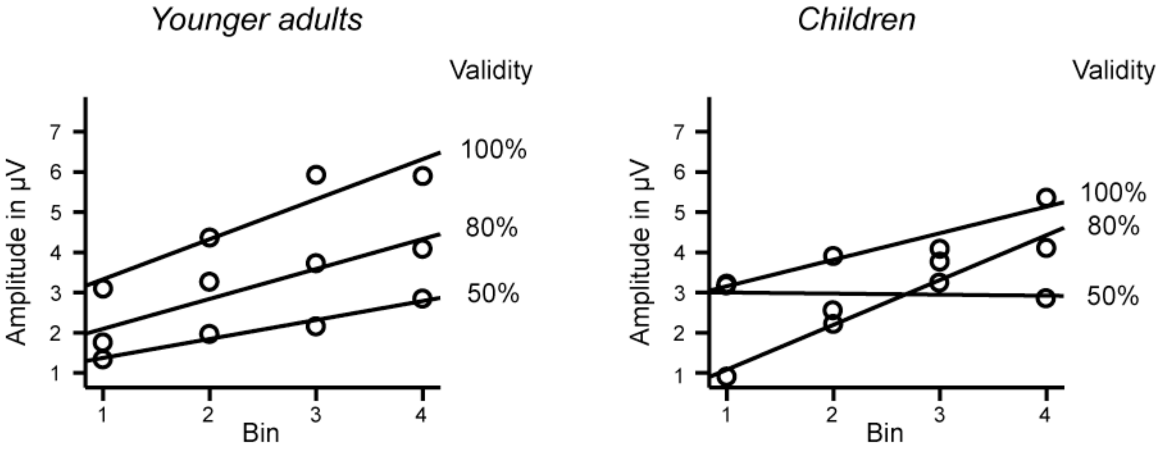

Figure 20: Learning curves for the response-locked positivity for the three validity conditions displayed separately for adults (left) and children (right). The y-axis indicates the amplitude in µV, the x-axis shows the course of learning averaged into four bins of trials. The β-parameters indicate the steepness of the learning functions (for details, see Method).

Peak-to-peak Analysis of the ERN. To obtain a more precise measure of the ERN that is less confounded by the overlapping positive component (see Figure 19), a peak-to-peak analysis for the ERN was performed. The peak-to-peak measures were analyzed using an ANOVA design with the factors Age group, Validity, and Bin. This analysis revealed a significant main effect of age group, $\underline{F}(1, 33) = 6.20$, $\underline{p} < .02$, $\underline{\eta}^2 = .16$, suggesting that the ERN was larger in children than in younger adults (see Figure 19). Moreover, a significant main effect of validity, $\underline{F}(2, 66) = 19.29$, $\underline{p} < .0001$, $\underline{\varepsilon} = .74$, $\underline{\eta}^2 = .36$ was obtained. Post-hoc contrasts revealed that the ERN was larger for the 100% than for the 80% validity condition and for the 80% than for the 50% validity condition ($\underline{p}$'s < .006, $\underline{\eta}^2$'s > .20) (see Figure 19). In line with the predictions of the R-L theory (Holroyd & Coles, 2002), this finding shows that the more valid the feedback the better the internal representation of an incorrect response. However, in contrast to this theory no evidence for an increase of the ERN with learning was found ($\underline{p}$'s > .13).

Correlation Analysis. To investigate the relation between the amplitude of the response-locked positivity and response accuracy a correlation analysis using Pearson's correlation coefficients was performed. As can be observed in Figure 21, the amplitude of the response-locked positivity was significantly correlated with the overall accuracy in the 100% validity condition, r(35) = .57, p < .0001, as well as in the 80% validity condition, r(35) = .34, p < .05. In contrast, no significant correlation was obtained for the 50% validity condition, r(35) = .19, p = .27. Moreover, the correlations in the 100% validity condition turned out to be reliable for younger adults, r(18) = .45, p < .06, as well as for children, r(17) = .67, p < .003 (see Figure 21).

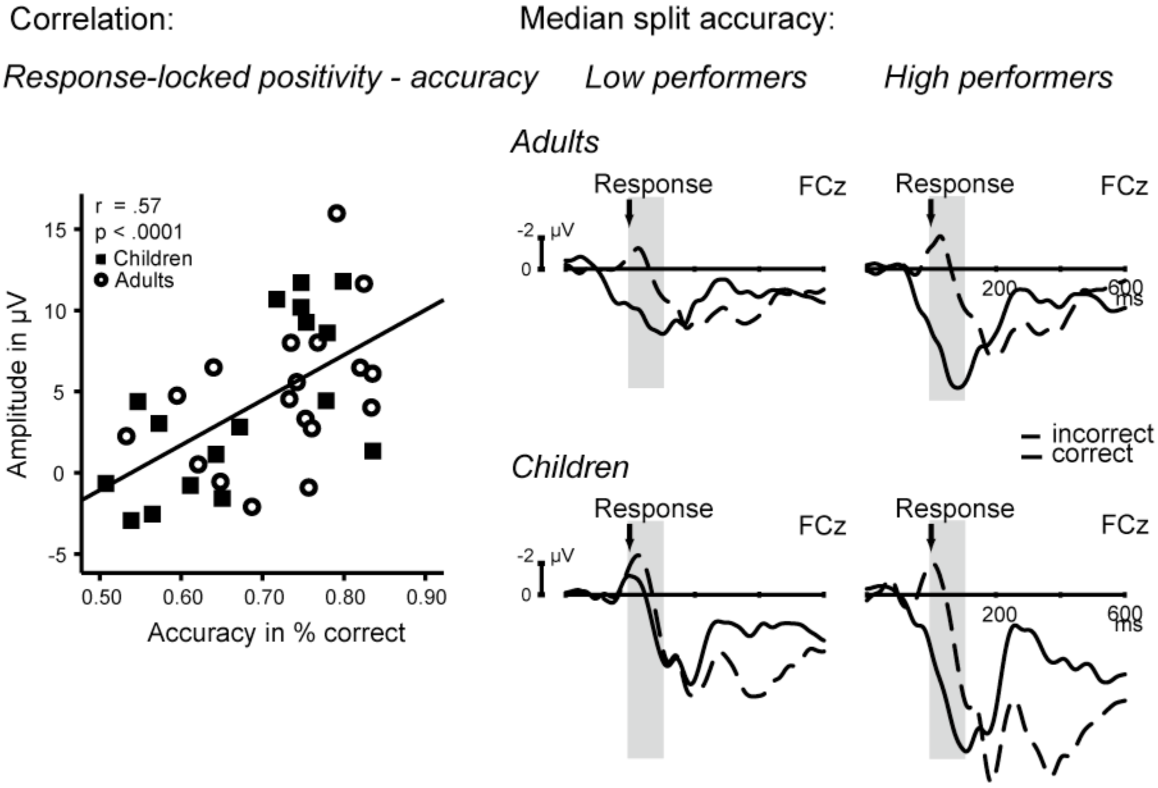

Figure 21: Left: Scatter plot illustrating the correlation between accuracy and the amplitude of the response-locked positivity in the 100% validity condition. Children are displayed by squares and adults are displayed by circles. Right: Median split for the mean accuracy in the 100% condition. The Figure displays the ERPs for correct (solid) and incorrect (dashed), separately for adults (top) and children (bottom) at the electrode FCz.

Thus, better performance with learning is related to a larger amplitude of the response-locked positivity. In contrast, no reliable correlations were found between the amplitude of the ERN and overall accuracy in either of the learning conditions (r's < .24,

p's $> .17$). Hence, the response-locked positivity rather than the ERN co-varied with response accuracy during learning.

Feedback-locked ERPs

Figure 22 displays the feedback-locked ERPs for positive and negative feedback in the three validity conditions (100%, 80%, and 50% validity), separately for children and adults. For both age groups a pronounced feedback-related negativity (FRN) can be observed for all validity conditions. In line with previous findings the ERP difference between positive and negative feedback shows a fronto-central topography and seems to get larger the invalid the feedback (see also Figure 22). Learning-related effects seem to be most pronounced in ERP component for positive feedback, which will be termed feedback-locked positivity in the following. In contrast, the FRN seems to remain stable over the course of learning (see Figure 23).

The ANOVA for the feedback-locked components included the factors Age group, Validity, Feedback type, and Bin. The analysis revealed a main effect of feedback type, $F(1, 33) = 91.12$, $p < .0001$, $\eta^2 = .70$, and a significant interaction between age group and feedback type, $F(1, 33) = 6.28$, $p < .02$, $\eta^2 = .05$, which reflects the larger feedback effects for children than for adults (see Figure 22). Separate analyses for the factor feedback type showed a marginally significant main effect of age group for negative feedback ($p < .09$, $\eta^2 = .09$), but not for positive feedback ($p = .58$) (see peak-to-peak analysis for a more precise measurement of the FRN). Moreover, the analysis revealed a main effect of validity, $F(2, 66) = 28.77$, $p < .0001$, $\varepsilon = .79$, $\eta^2 = .46$, and an interaction between feedback type and validity, $F(2, 66) = 4{,}73$, $p < .01$, $\varepsilon = .95$, $\eta^2 = .12$. Separate analyses for the factor validity showed main effects of feedback type for all validity conditions (p's $< .0001$, η^2's $> .47$). As depicted in Figure 22, the effects of feedback type were larger the more invalid the feedback.

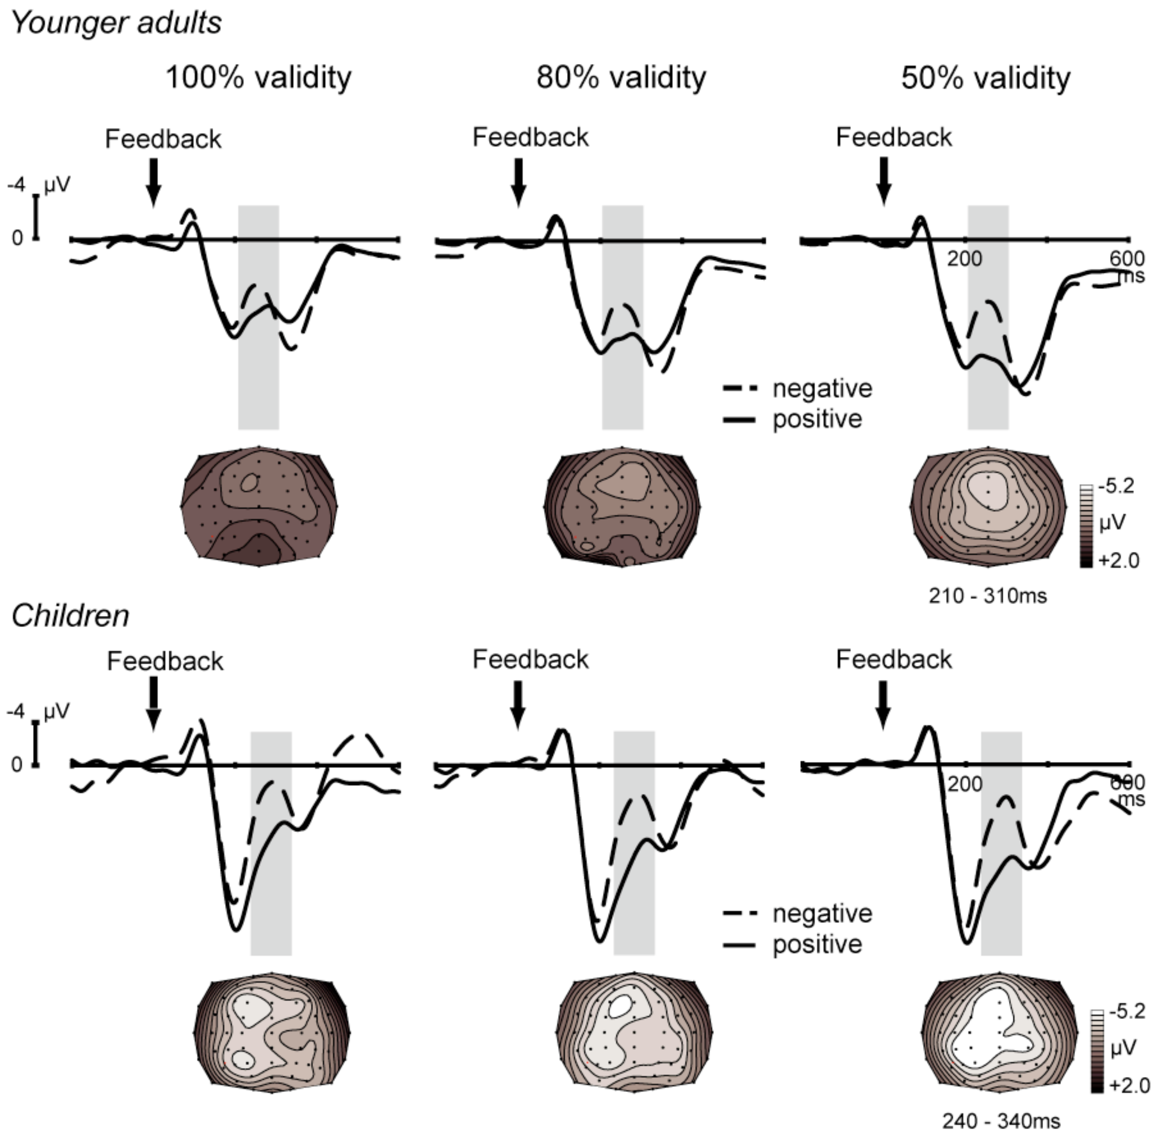

Figure 22: Feedback-locked ERPs for the three validity conditions displayed separately for positive (solid lines) and negative (dashed lines) feedback for adults (top) and children (bottom) at the electrode FCz. Topographical distribution of the ERP difference wave for correct and incorrect responses displayed separately for the three validity conditions and the two age groups.

Learning-related Effects in the Feedback-locked Positivity. Again, the focus of the analysis was on the effects of learning on the feedback-locked ERPs. The ANOVA showed a significant main effect of bin, $F(3, 99) = 8.39$, $p < .0001$, $\varepsilon = .92$, $\eta^2 = .20$, and an interaction between the factors validity and bin, $F(6, 198) = 2.89$, $p < .01$, $\varepsilon = .87$, $\eta^2 = .08$. Separate analyses for the factor validity showed main effects of bin for 100% and the 80% validity conditions (p's $< .02$, η^2's $> .10$), but not for the 50% validity condition ($p = .36$). These findings indicate that the feedback-locked ERP components for positive and for negative feedback decrease over the course of learning (see Figure 23). However, as can be also observed in Figure 23, for negative feedback this effect seems to be due to fact that especially in children the ERPs got generally more negative over the course of the four bins.

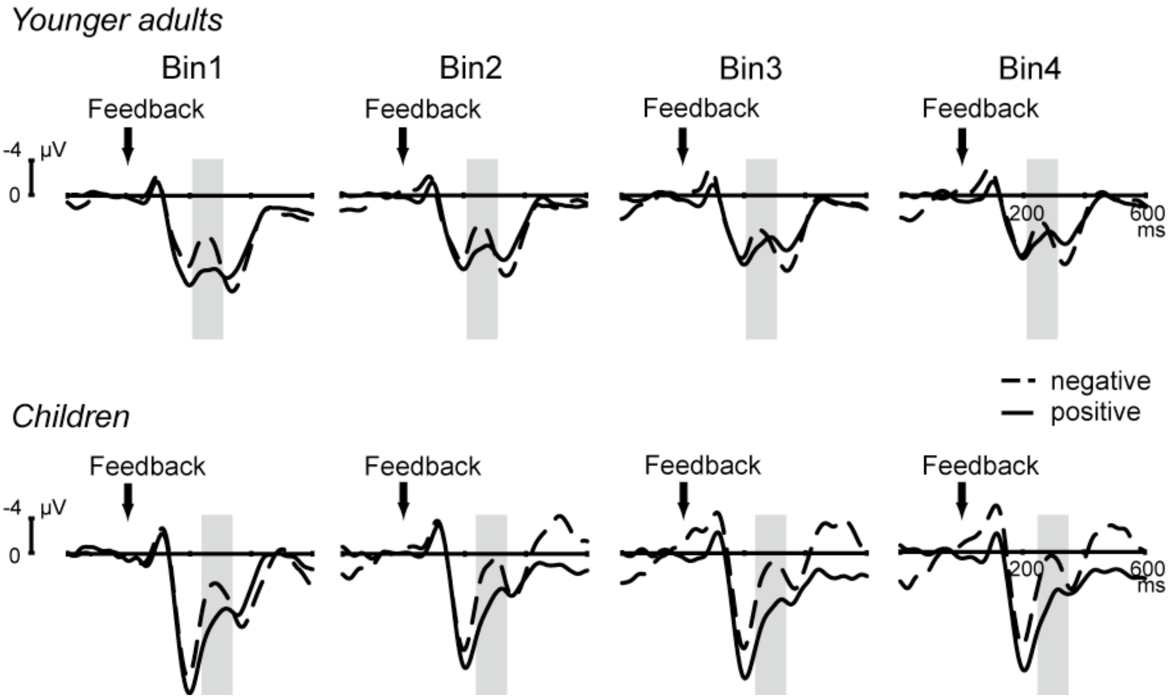

Figure 23: Feedback-locked ERPs over the course of learning (averaged into four bins) for the 100% validity condition displayed separately for positive (solid) and negative (dashed) feedback for adults (top) and children (bottom) at the electrode FCz.

In order to quantify the learning-related effects in the feedback-locked positivity, the slope parameters of the individual learning functions (for details, see Method) were subjected to an ANOVA involving the factors age group and validity. This analysis showed a main effect of validity, $F(2, 66) = 5.54$, $p < .009$, $\varepsilon = .87$, $\eta^2 = .13$, and a marginally significant interaction between age group and validity, $F(2, 66) = 2,79$, $p < .08$, $\varepsilon = .87$, $\eta^2 = .07$. Contrasts for each of the levels of validity showed significantly higher slope parameters for the 100% than for the 80% and the 50% validity condition (p's $< .008$, $\eta^2 = .18$). However, slope parameters were not significantly larger for the 80% than for the 50% validity condition ($p = .59$), indicating that learning-related effects were restricted to the 100% validity condition. Separate analyses for the two age groups revealed that the effects of feedback validity on the slope parameters were much more pronounced for adults (p's $< .0001$, $\eta^2 = .47$) than for children (p's $> .53$, $\eta^2 = .01$) (see Figure 24). To summarize, these findings suggest that the feedback-locked positivity decreased with learning. Moreover, these learning-related effects were less pronounced for children compared to adults, indicating that children were less able to disengage from feedback during learning (see Figure 24).

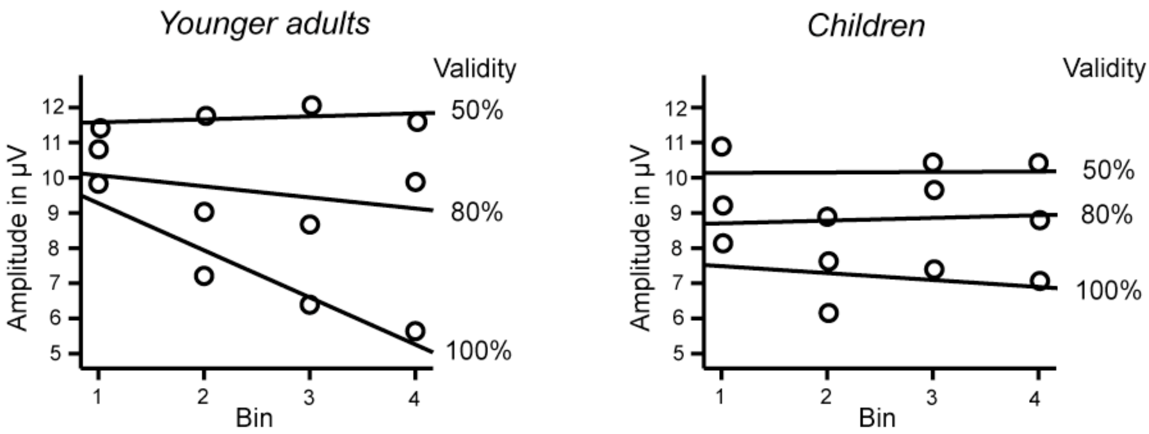

Figure 24: Learning curves for the feedback-locked positivity for the three validity conditions displayed separately for adults (left) and children (right). The y-axis indicates the amplitude in µV, the x-axis shows the course of learning averaged into four bins of trials. The β-parameters indicate the steepness of the learning functions (for details, see Method).

Peak-to-peak Analysis of the FRN. The peak-to-peak measures of the FRN were subjected to an ANOVA involving the factors Age group, Validity, and Bin. The analysis only revealed a main effect of age group, $F(1, 33) = 27.10$, $p < .0001$, $\eta^2 = .45$, which reflects the larger FRN amplitudes for children than for adults (see Figure 22). However, neither significant main effects of validity or bin (p's > .12), nor a significant interaction between these factors ($p = .14$) was obtained, indicating that the FRN did not vary as a function of feedback validity or learning.

4.6 Summary

The analysis of the accuracy data revealed that similar to first experiment, response accuracy increased with feedback validity. Moreover, significant age differences were obtained in the 80% validity condition, which suggests that similar to older adults, children were impaired when invalid information interfered with learning. Also similar to older adults, children performed worse than adults at the beginning of learning in the 80%

validity condition, but reached a similar performance level at the end of learning. In contrast, no developmental differences in response accuracy were obtained for the 100% validity condition, not even in the beginning of learning. This indicates that performance levels were equated in this condition (see Figure 17). Again, no significant age differences in the learning rates were obtained, suggesting that children learned to a similar extent as younger adults (see Figure 17).

The analysis of the response-locked ERPs revealed that children were as well able as younger adults to internally represent correct and incorrect responses in the 100% validity condition. However, children did not show a significant difference between the ERPs to correct and incorrect responses in the 80% validity condition (see Figure 18). In line with the accuracy data this points to the view that children were impaired in representing the correctness of the response when feedback was partially invalid. Consistent with the results of the first experiment learning-related changes were only found for the response-locked positivity, but not for the ERN. Moreover, an analysis of the correlations between response accuracy and the amplitude of the response-locked positivity showed that better performance with learning was related to a larger amplitude of the response-locked positivity. This was the case for adults, as well as for children (see Figure 21). However, no association between the ERN and accuracy was obtained, which is inconsistent with the predictions of the R-L theory (Holroyd & Coles, 2002). Hence, these findings again point to the view that the response-locked positivity, rather than the ERN varies with response accuracy during learning. Consistent with the results from the first experiment and the findings by Hogan et al. (2005), the peak-to-peak analysis did not reveal a reduction of the ERN in children compared to adults. This stands in contrasts to previous findings (Davies et al., 2004, Ladouceur et al., 2004; Santesso et al., 2006) and points to the importance of controlling for performance levels when comparing the ERN between age groups.

Similar to previous findings in adults (Holroyd & Coles, 2002; Nieuwenhuis et al., 2002) and the results of the first experiment, the analysis of the feedback-locked ERPs revealed that the difference between positive and negative feedback increased the more invalid the feedback. However, the peak-to-peak analysis showed that the FRN was not affected by feedback validity, suggesting that it was the feedback-related positivity, which increased the more invalid the feedback (see Figure 22). Moreover, the FRN was increased for children compared to adults, whereas the amplitude of the feedback-locked positivity was comparable for the two age groups. This points to the view that children are more sensitive to negative compared to positive feedback during learning. In line with the first experiment the feedback-locked positivity but not the FRN decreased with learning, indicating that participants relied less on the feedback the more they were able to internally represent the correctness of the response. Moreover, these effects were less pronounced for children compared to adults, suggesting that children are less able to disengage from feedback during learning.

5. Interim Discussion

The first experiment investigated the effects of aging on reinforcement learning and error processing. The focus of this experiment was on the role of error processing for learning and the question whether the ERP-correlates of internal and external error processing, the ERN and the FRN, respectively reflect learning-related changes in younger and older adults. The study was based on a recent neurocomputational account to altered error processing in older age (Nieuwenhuis et al., 2002) that suggests that the impairments of older adults in error processing and learning result from age-related changes in the mesencephalic dopamine system (MDS).

The second experiment examined developmental differences in the processing of internal and external error information during reinforcement learning. Similar to older adults, it was expected that changes in the MDS during childhood development should be associated with developmental differences in learning and error processing. The purpose of this experiment was to investigate whether children differ from younger adults in the way they use error information for learning and whether this is reflected in the ERP correlates of internal and external error processing.

The objective of the two experiments was to provide an integrative view on how error processing and learning, and the ERP correlates of these processes, change over the lifespan. In both experiments a probabilistic learning task was applied that involved three learning conditions, in which the validity of feedback was manipulated (100%, 80%, and 50% validity). By this, the possibility to learn the stimulus-response assignments on the basis of feedback (100% and 80% validity condition) was varied relative to a control condition, in which no learning was possible (50% validity condition). In order to equate performance levels between age groups an adaptive procedure was implemented in the learning task. This algorithm individually adjusted the response deadlines, depending on the number of time outs (for details, see Method).

5.1. Accuracy Data

Results of the first experiment revealed that a) older adults had a similar overall accuracy in the 100% validity condition as younger adults b) older adults performed overall worse than younger adults in the 80% validity condition (see Figure 10) c) these age differences in the 80% condition were most pronounced at beginning of learning, but absent at the end of learning d) no age differences were obtained in the learning functions. These findings suggest that the adaptive responses deadline procedure led to similar accuracy levels in both age groups in the 100% validity condition and the 80% condition at the end of learning. Moreover, the absence of age differences in the learning rates points to the view that older adults were not impaired in learning per se. Nevertheless, older adults performed overall worse at the beginning of learning in the 80% condition, in which feedback was partially invalid.

Hence, the behavioral findings stand in contrast to the results of Nieuwenhuis and colleagues (2002), who found age differences for all learning conditions. Based on their findings Nieuwenhuis et al. (2002) argued that older adults might be generally impaired in reinforcement learning and error processing. However, one has to keep in mind that in this study both age groups were treated using the same response deadlines (700 milliseconds), which produces a disproportionate time pressure on the elderly and impairs their ability to learn. Moreover, the time pressure on older adults can be expected to result in an increased number of time-outs, which lowers the number of trials in which they can learn from feedback. On the other hand, one could argue that the adaptive deadline procedure might have masked age differences in learning, since it allows older adults to respond slower than younger adults[10]. However, it should be noted that slower reaction times in older adults are not a surprising phenomenon that per se points to a more conservative response bias in the elderly. In contrast, general slowing has been shown to be one of the hallmarks of cognitive aging (see Birren & Fisher, 1995; Salthouse, 1996, 2000) and the present study aims at accounting for these general age effects by individually adjusting the response deadlines. Yet, the first experiment differed from that of Nieuwenhuis et al. (2002) not only with this respect, but also in the kind of feedback provided to the subjects. Nieuwenhuis et al. (2002) used rather ambiguous feedback stimuli (head of a lettuce and a carrot). In contrast, in the first experiment unambiguous feedback stimuli (German words for ‘correct’, printed in green and ‘incorrect’ printed in red) were used, which are easy to encode and process and might have helped older adults in learning.

The second experiment focused on developmental differences in learning and error processing. In order to allow children to perform the learning task successfully and in order to enhance their motivation, the learning task had to be slightly modified. First, a cover story was constructed to increase the children’s motivation. Second, there was a short break of 15 seconds in the middle of each learning block, in which a fake monetary feedback was displayed. Third, in order to account for their larger variability of response times the response deadlines were increased.

The analysis of the accuracy data in the second experiment revealed a similar pattern of behavioral results for children, as it was obtained for older adults. Consistent

---

[10] Note that overall accuracy was similar for the two age groups in the 100% validity condition, and in the 80% condition at the end of learning. Hence there are no age differences in accuracy in the first experiment. This indicates that although there are age differences in reaction times (older adults M = 520 ms, SD = 132 ms responded slower than younger adults M = 407 ms, SD = 95 ms) the present behavioral findings are not confounded by age differences in speed-accuracy trade-offs.

with the results in the elderly, there were no significant differences in response accuracy in the 100% validity condition for children compared to younger adults, not even at the beginning of learning. Moreover, similar to older adults, children did not differ from younger adults with respect to the steepness of their learning functions, suggesting that they learned comparably to younger adults. These findings again validate the success of the adaptive deadline procedure in equating performance levels between age groups. Nevertheless, similar to older adults, children were impaired in overall accuracy when invalid feedback occurred in the 80% validity condition. Moreover, as for older adults, the age differences in the 80% condition were most pronounced at the beginning of learning, but absent at the end of learning.

To summarize, age differences in overall accuracy were only obtained in the 80% condition, suggesting that older adults and children are impaired in accuracy when invalid information interferes with learning. For both age groups these impairments were most pronounced at the beginning of learning, but absent at the end of learning. This indicates that for children and older adults invalid information interferes most at the beginning of learning, whereas these effects appear to be compensated at the end of learning. In line with the absence of age differences in the learning functions, this points to the view that the basic reinforcement learning mechanisms are similar across the lifespan. However, invalid information seems to impair the ability of children and older adults to acquire the stimulus-response mappings, especially at the beginning of learning. This suggests that both age groups are particularly impaired when control requirements are enhanced due to interference by invalid feedback, as it is the case in the 80% condition at the beginning of learning. This view is supported by a recent study that examined the effects of aging on reversal learning (Mell et al., 2005). In this study participants had to flexibly learn and relearn stimulus-response associations depending on feedback information. The results of this study showed that older adults are particularly impaired in relearning stimulus-response mappings that is, when the previously learned but now invalid mappings interfere. Further evidence for this view comes from studies on age differences in the

Wisconsin Card Sorting Test (WCST) (Hartman et al., 2001; Ridderinkhof et al., 2002; for a meta-analytic review see Rhodes, 2004). These studies showed that older participants have the tendency to perseverate when a change in the response rule is indicated. Taken together, these finding suggest that older adults are particularly impaired when invalid information interferes with learning. Interestingly, similar results have been obtained in developmental studies on discrimination learning (Moran & McCullers, 1979; Offenbach, 1973). Results of these studies suggest that especially younger children are impaired in learning when invalid feedback is provided.

5.2. Response-locked ERPs

The analysis of the response-locked ERPs in the first experiment revealed that the difference between the ERPs to correct and incorrect responses increased with feedback validity. This suggests that the more participants learned, the more they were able to internally represent the correctness of the response. This is in line with the predictions that were derived based on the results of previous studies (Holroyd & Coles, 2002; Nieuwenhuis et al., 2002). However, in these studies, a difference wave approach was adopted to investigate the effects of learning on the ERN. The general problem of a difference wave approach is that the variance of correct, as well as incorrect trials contributes to the difference wave. Hence, on the basis of this approach it cannot be decisively concluded whether it was the ERN, or some component on correct trials that increased as a function of feedback validity. Indeed, as illustrated in Figure 11, a separate analysis for correct and incorrect responses revealed that the response-locked positivity for correct trials, as well as the ERN increased the more valid the feedback. Age differences were only obtained for the response-locked positivity, which showed a less pronounced increase with feedback validity for older compared to younger adults. This suggests that older adults may have been less able to differentiate between the validity conditions.

Consistent with these results, changes over the course of learning were most pronounced in the response-locked positivity for correct trials, which increased with learning for the 100% and 80% validity conditions. The analysis for the learning functions of the response-locked positivity showed that the learning-related effects were larger for the two learning conditions compared to the 50% condition. Yet, they were not significantly different between the two learning conditions, which is nicely in line with the findings in the accuracy data. In contrast, a significant increase of the ERN with learning was only found for the 100% validity condition, when the ERN was captured most precisely using peak-to-peak measurements. Hence, these findings suggest that learning-related changes were much more pronounced for correct compared to incorrect trials, an effect, which may have been overlooked in previous studies that mainly focused on the effects of errors and negative feedback for learning (Holroyd & Coles, 2002; Nieuwenhuis et al., 2002).

The second experiment revealed that similar to younger and older adults, children were able to differentiate correct from incorrect responses in the 100% validity condition (in which no performance differences between age groups were obtained). That is, they showed a pronounced ERN on incorrect trials and a response-locked positivity on correct trials (see Figure 18). However, in the 80% validity condition, in contrast to younger and older adults, children did not differentiate between correct and incorrect responses (see Figure 18). This pattern of results is nicely consistent with the accuracy data, which showed that children performed worse than younger adults only in the 80% condition (see Figure 17). Thus, these findings suggest that children are impaired in response accuracy and in the ability to represent the correctness of the response if feedback is partially invalid and interferes with learning.

In line with the results in younger and older adults the response-locked ERP data in the second experiment does not provide evidence for the view that the ERN varies considerably with learning. Similar to younger and older adults, children showed a learning-related increase in the response-locked positivity for correct trials in the two learning conditions (100% and 80% validity condition). Moreover, for children, as well as for younger adults a significant positive correlation between the amplitude of the response-locked positivity and response accuracy was obtained in both learning conditions. As illustrated in Figure 21, this correlation reflects the fact that the higher the response accuracy, the larger the response-locked positivity. This further supports the view that the response-locked positivity, rather than the ERN varies with response accuracy.

The central idea of the first experiment was that an equation of performance levels between age groups should reveal similar ERN amplitudes for younger and older adults. On the one hand, this prediction was based on the observation that in most of the studies on age-related impairments in error processing age differences in the ERN were confounded with age differences in performance (Band & Kok, 2000; Mathewson et al., 2005; Nieuwenhuis et al., 2002; West, 2004). On the other hand, the mismatch model of the ERN (Bernstein et al., 1995; Coles et al., 2001; Scheffers & Coles, 2000) and the R-L theory (Holroyd & Coles, 2002) itself, suggest that the ERN depends on an intact internal representation of the correct response. From this it follows that if performance is impaired, participants are less able represent the correct response and perceive less mismatch and show a reduced ERN when they commit an error. Hence, in order to fairly compare the ERN between younger and older adults it seems rather important to avoid accuracy differences between age groups. In line with this prediction, no differences in the ERN between younger and older adults were found, suggesting that it is indeed not age per se, but differences in performance level (in the expectation on the correctness of the response) that drive the ERN (for a similar finding see Pietschmann, Endrass, & Kathmann, 2007). Moreover, in contrast the prediction of the dopamine hypothesis of aging, there is also no evidence that the ERN develops differentially over the course of learning for older compared to younger adults (see Figure 13). This is in line with the absence of age differences in the behavioral learning functions.

Similar to the findings in older adults, it was predicted that if performance levels were similar in children and younger adults, comparable ERN amplitudes should be obtained for the two age groups. As the findings in older adults, this result stands in contrast to data from several recent developmental ERN studies (Davies et al., 2004; Ladouceur et al., 2004; Santesso et al., 2006; Wiersema, van der Meere, & Roeyers, 2007), which pointed to the view that the ERN is reduced in children compared to adults. Hence, the data in children strengthens the view that the ERN is driven by performance, rather than age per se. Further evidence for this idea comes from a recent study on developmental differences in error processing between adolescents and younger adults (Hogan et al., 2005). In this study, task complexity was manipulated and age differences in the ERN were only obtained in the more complex task version, in which performance was also impaired in adolescents. In contrast, in the less complex task version, no significant age differences in the ERN or performance were obtained (see also Kim et al., 2007). Taken together, these findings suggest that in previous studies age differences in ERN may have been confounded by performance differences between age groups. When performance is equated between children, younger adults and older adults no age differences in internal error processing as reflected in the ERN are obtained. Hence, the ERN seems to be performance- rather than age-sensitive.

5.3 Feedback-locked ERPs

The analysis of the feedback-locked ERPs in the first experiment revealed that for younger adults the difference between positive and negative feedback increased the more invalid the feedback (see Figure 14), which is in line with previous data (Holroyd & Coles, 2002; Nieuwenhuis et al., 2002). For older adults there was no significant difference between positive and negative feedback and no effect of feedback validity on this difference was obtained. This is inconsistent with the findings of the Nieuwenhuis et al. (2002) study, which revealed that older adults, in contrast to younger adults, showed an increase of the difference the more valid the feedback. Nieuwenhuis et al. (2002) explained this effect, by assuming that the amount of attention that older adults pay to a feedback stimulus might depend on their subjective probability of committing an error. However, the present data if at all, showed a small increase of the difference between positive and negative feedback with feedback validity for older adults, which is more in line with the data obtained in younger adults (see Figure 14). On the one hand, this result might be due to the fact that in the present study older adults performed similarly as younger adults and therefore, most likely, did not differ from younger adults with respect to their subjective probability of committing an error. On the other hand, this effect could also be a result of the different types of feedback stimuli that were used in the two studies. As mentioned above, Nieuwenhuis et al. (2002) used feedback stimuli that might have been difficult to disambiguate for older adults. In contrast, in the present study, unambiguous feedback stimuli were used, which were easy to encode and process.

Hence, apart from the inconsistencies in the findings in older adults, the increase of the difference wave with feedback validity is consistent with the data from Holroyd and Coles (2002) and Nieuwenhuis et al. (2002). However, in contrast to the predictions of the R-L theory (Holroyd & Coles, 2002), the peak-to-peak analysis showed that the FRN was not modulated by feedback validity. This suggests that it was not the FRN, but the feedback-locked positivity (on positive feedback trials), which was affected by the feedback validity. This view receives further support from a comparison of the 20% invalid trials with the 80% valid trials of the 80% validity condition[11]. For this comparison the R-L theory would suggest that the FRN should be much larger for the 20% invalid trials since in that condition a strong expectation for positive feedback is violated. Indeed, the analysis showed that there is a marginally significant difference between the FRN for valid compared to invalid negative feedback. However, the effect was much more pronounced in the positivity for positive feedback, which was significantly larger for the 20% invalid

---

[11] In order to investigate the effects of expectancy violations on the feedback-locked ERPs, the 20% invalid and the 80% valid trials of the 80% condition were compared. Separate analyses for positive and negative feedback revealed a larger feedback-locked positivity for invalid compared to valid positive feedback ($\underline{p}$ < .004). This was not the case for negative feedback ($\underline{p}$ = .90). The peak-to-peak analysis showed a marginally significant difference between invalid and valid trials in the FRN ($\underline{p}$ < .07). In neither of these analyses significant interactions with age were obtained ($\underline{p}$'s > .13). These findings are consistent with the other results of the Experiment 1 in showing that differences between validity conditions are most pronounced on positive feedback trials. Since there is no interaction with age, there is no reason to assume that older adults differed from younger adults in the amount of attention they paid to the feedback stimulus on valid compared to invalid trials of the 80% condition.

trials compared to the 80% valid trials. Moreover, this effect did not interact with age, suggesting that older adults did not differ from younger adults in the way they attended to valid or invalid positive feedback in the 80% condition.

According to the R-L theory (Holroyd & Coles, 2002), it was expected that the FRN should decrease with learning, since the participants rely less on the external error feedback. However, similar to the response-locked ERPs there was no evidence for learning-related changes in the FRN, not even in the peak-to-peak analysis. In contrast, a significant decrease of the feedback-locked positivity with learning was obtained, indicating that the more participants are able to internally represent the correctness of the response the less they have to rely on external feedback. An analysis of the learning functions revealed that the learning effects in the feedback-locked positivity were larger for the 100% compared to the 80% validity condition and for both learning conditions compared to the 50% condition (see Figure 16). This fits nicely to the learning effects in the accuracy data. Moreover, these learning effects were more pronounced for younger compared to older adults, suggesting that the elderly may have been less able to disengage from processing positive feedback during learning (see Figure 16).

Similar to younger and older adults, children showed a significant increase of the difference between positive and negative feedback the more invalid the feedback. However, as in the first experiment, the peak-to-peak analysis of the FRN revealed no significant effect of feedback validity. This indicates that it was not the FRN but the feedback-locked positivity on positive feedback trials that varied with feedback validity. Also consistent with the results of the first experiment, no significant learning-related changes were obtained for the FRN in children. In contrast, the feedback-locked positivity on positive feedback trials decreased over the course of learning for the two learning conditions, but not for the 50% validity condition. This supports the view that the feedback-locked positivity, rather than the FRN varies as a function of learning. Similar to the first experiment, the learning-related changes in the feedback-locked were less pronounced for children compared to adults, suggesting that children were less able to disengage from positive feedback during learning.

One of the most interesting and surprising results of the analysis of the feedback-locked ERPs was the finding of an age-related asymmetry in the FRN across the lifespan. The first experiment revealed that the FRN is strongly reduced for older adults compared to younger adults. Thus, the present data suggest that although older adults learned comparably to younger adults, they showed reduced activity of the structures involved in the processing of negative feedback (presumably the ACC, but also the orbitofrontal cortex; see O'Doherty et al., 2001; Rolls, 2000). This result is somewhat surprising given the absence of age differences in the ERN in the present study and may point to a functional dissociation of both components (see Nieuwenhuis et al., 2005). Similar findings on an asymmetry in the processing of valence information in older adults have been obtained in research on episodic memory and decision making (Charles, Mather, & Carstensen, 2003; Mather & Johnson, 2000). These findings have been interpreted within the framework of the socio-emotional selectivity theory of aging, which proposes that the ratio between positive and negative affect improves through adulthood and leads to what is called a "positivity effect" (Carstensen, 2006; Mather & Carstensen, 2005). The idea is that older adults focus more on emotion regulation and implement cognitive control mechanisms that enhance positive and diminish negative information. Interestingly, recent fMRI findings from Larkin et al. (2007) using a gain and loss anticipation task support this view and suggest that older adults are less affected by potential losses than younger adults, whereas both age groups are equally excited by potential gains. The present data underlines these findings and provides the first electrophysiological evidence for an age-related asymmetry in feedback processing and by this supports the idea of a positivity effect in older adults.

In contrast to the results in older adults, a larger FRN was found for children compared to younger adults, whereas no age differences in the ERPs for positive feedback (feedback-locked positivity) were obtained. These findings suggest that children are more sensitive to negative feedback during learning than adults, whereas both age groups seem to be similarly affected by positive feedback. Since this is one of the first studies that investigated developmental differences in the ERP correlates of feedback processing an integrative interpretation of the present results in the light of previous findings is difficult. However, there is some evidence from developmental studies using heart-rate measures that speaks for a similar asymmetry in feedback processing between children and adults (Crone et al., 2004; Somsen, van der Molen, Jennings, & van Beek, 2000). These studies showed that in younger adults heart rate is slowed for negative compared to positive feedback, when the feedback is informative with respect to learning or performance adaptation. In contrast, 8 - 10 year-old children showed heart-rate slowing for informative, as well as uninformative negative feedback (Crone et al., 2004). This suggests that younger children may be more sensitive to negative feedback and less able to use external error information for learning. Moreover, in contrast to older children (12-14 year-old), younger children (8-10 year-old) did not show heart rate slowing following performance errors, suggesting that they are less able to internally represent incorrect responses (Crone, Somsen, Zanolie, & van der Molen, 2006). Consistent with these findings, the present data suggest that on the one hand children may be are more sensitive to negative feedback during learning, whereas on the other hand they seem to be impaired in representing the correctness of a response, especially when invalid information occurs during learning.

6 Experiment 3

6.1 Statement of Problem

The third experiment is mainly based on the findings of the first experiment that pointed to an age-related asymmetry in feedback processing, as reflected in the FRN. In the first experiment older adults showed a strongly reduced FRN, which suggests that they are less sensitive to negative feedback and might have focused more on positive feedback during learning. Moreover, the elderly did not show a reduction of the ERN when compared to younger adults, suggesting that they were not impaired in error processing per se. Since older adults learned similarly to younger in this experiment, the question arises whether the elderly might have strategically focused more on positive compared to negative feedback during learning. In other words, older adults may have a tendency towards learning more from positive compared to negative outcomes that is, they might be positive learners rather than negative learners.

Further support for such an asymmetry in valence processing comes from a recent fMRI study by Larkin et al. (2007). Larkin and colleagues (2007) used a reward anticipation task and showed similar activations in the striatum and the insula for younger and older adults during gain anticipation. However, the activations during loss anticipation were reduced for older compared to younger adults. They concluded that as proposed by the socioemotional selectivity theory (Carstensen, 2006, Mather & Carstensen, 2005), older adults may have focused more on emotion regulation as younger adults and hence implemented cognitive control mechanisms that enhance positive and diminish negative information.

Interestingly, a recent neurocomputational model (Frank et al., 2004, Frank, 2005) suggests that whether participants learn more from positive or negative feedback depends on individual differences in dopamine levels. Hence, it could be the case that age-related asymmetries in feedback processing are due to age differences in dopamine levels. In the following, I will briefly introduce this alternative neurocomputational account to reinforcement learning.

6.2 An Alternative Account to Reinforcement Learning

The reinforcement learning model by Frank and colleagues (2004) was developed based on the neurophysiological findings by Schultz and others (for reviews see Schultz, 2000; 2002; 2007). The objective of this computational model was to make predictions on how changes in dopamine levels affect the way participants learn form positive and negative feedback. These predictions were subsequently tested in Parkinson patients who were on or off medication, that is, who had either high or low dopamine levels. In line with the work by Schultz and colleagues (1997), the model suggests that positive prediction errors, which are elicited when an event is better than predicted, lead to phasic increases in dopaminergic activity. These bursts of dopamine are conceived of as Go signals that facilitate the execution of the most appropriate response (Go learning). In contrast, negative prediction errors, which lead to phasic dips in dopamine, are suggested to represent NoGo signals that suppress competing responses (NoGo learning) (see Figure 25). According to this terminology, I will refer to the model as Go-NoGo model in the following. The Go-NoGo model proposes that low levels of dopamine can lead to a bias towards learning more from negative outcomes because low levels of dopamine should impair Go learning, but should support NoGo learning. In contrast, high dopamine levels should support Go learning but hinder NoGo learning because dips in dopamine are less likely to occur when dopamine levels are generally high (see Frank, 2005).

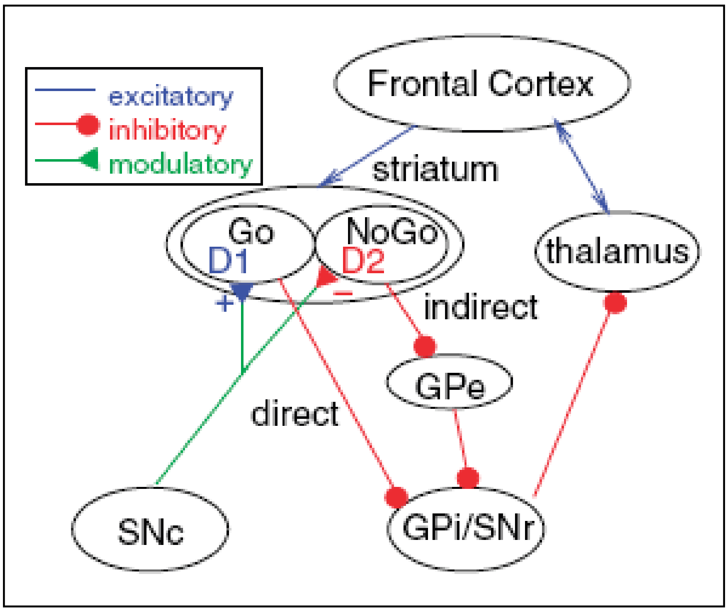

Figure 25: The Go-NoGo model (Frank, 2004) suggests that phasic bursts of activity of the dopamine neurons in the substantia nigra (SNc) during positive reinforcement activate D1 receptors in the striatum and drive Go learning. In contrast, phasic decreases in dopamine during negative reinforcement drive NoGo learning. Figure adapted from Frank et al., (2004).

These predictions were tested in patients with Parkinson’s disease, which is characterized by a decline of the dopaminergic neurons in the substantia nigra that leads to reduced tonic and phasic levels of dopamine (Frank et al., 2004; Frank, 2005). Frank and colleagues (2004) showed that Parkinson patients off medication (low dopamine level) indeed learn more from negative compared to positive outcomes. In contrast, when these patients are medicated with L-dopa (which typically leads to a so called dopamine overdose), they showed the opposite pattern and learned more from positive outcomes compared to negative outcomes. Hence, a central assertion of the ‘Go-No-Go’ model has been confirmed, showing that learning biases in Parkinson patients seem to depend on their dopamine levels.

In the following, Frank, Woroch and Curran (2005) investigated, whether a similar effect could be obtained in healthy younger subjects, and whether these learning biases might be reflected in the error-related negativity (ERN). Based on the findings by Holroyd and Coles (2002), it was predicted that if the ERN reflects negative dopaminergic learning signals, it should be affected by the degree to which participants tend to learn more from positive or negative outcomes. Indeed, Frank and colleagues (2005) were able to show that larger ERNs were associated with a bias towards learning more from negative compared to positive outcomes. This finding supports the view that the ERN is related to dips in dopamine during error commission (see Holroyd & Coles, 2002). Moreover, this result also points to the view that even in healthy young participants dopamine levels (as reflected in the ERN) might be associated with learning biases.

Thus, this alternative account to the role of dopamine in learning suggests that individual differences in dopamine levels affect the learning biases and are reflected in the relative size of the ERN. According to the Go-No-go model reduced dopamine levels, as it is the case in older adults, should lead to a tendency towards learning more from negative compared to positive outcomes. This stands in contrasts to the findings from the first experiment and the findings from Larkin et al. (2007), which suggest that older adults might focus more on positive compared to negative feedback during learning.

The aim of the third experiment was to test these competing hypotheses and to investigate whether age differences in the learning biases are reflected in the ERN. To do so, a probabilistic learning task was created that allowed to directly compare positive and negative learning in two learning conditions. As in the previous experiments, participants have to learn stimulus-response assignments based on feedback information. In the positive learning condition, participants can win 50 Euro Cents if they press the correct response button, whereas they get a neutral outcome when they press the incorrect response button. Hence, participants should learn to choose the button that leads to the positive outcome. In contrast, in the negative learning condition participants can loose 50 Euro Cents if they press the incorrect button, whereas they get a neutral outcome if they press the correct button. In this condition, participants should learn to avoid pressing the button that leads to the negative outcome.

In order to examine the stability of these learning biases and to investigate age differences in the relearning of stimulus response assignments, a reversal phase was introduced in the learning blocks. Hence, each learning block involved a learning phase, in which participants learned the stimulus response mappings until they reached a certain performance criterion (for details, see Methods). When they reached this criterion, the mappings were reversed and stimuli that were previously associated with the positive learning condition were then associated with the negative learning condition, and vice versa.

6.3 Hypotheses

In line with the data from the first experiment, which showed that older adults are impaired when invalid information interferes with learning, it was expected that in the present experiment the elderly should be particularly impaired during the reversal periods, when the previously learned stimulus-response assignments interfere with learning the new S-R mappings.

Regarding the learning biases there are two competing hypotheses: Based on the findings of the first experiment and the findings by Larkin et al. (2007), it was expected that older adults should learn more from positive outcomes compared to negative outcomes. However, the opposite pattern of results is predicted by the model of Frank and colleagues (2004), which suggests that given the age-related reductions in dopamine levels (see Bäckman et al., 2006) older adults can be expected to be negative learners.

According to previous findings by Frank and colleagues (2005), the learning biases in younger adults should be reflected in the amplitude of the ERN. That is, the ERN can be predicted to be larger for negative compared to positive learners. However, the Frank et al., (2004) model does not make predictions on how age differences in the learning biases should be reflected in the ERN. Based on the data of the first experiment it could be expected that when performance levels are equated between age groups the ERN should be of similar size for older compared to younger adults. Similar to the age-related asymmetry in feedback processing that was found in the feedback-related ERPs in the first experiment, it could be expected that age differences in the learning biases should be reflected in the FRN.

6.4 Methods

Participants. Thirty-one younger adults and 30 older adults participated in Experiment 3. Two younger adults had to be excluded from further data analysis because they did not commit enough errors to analyze the ERN over the course of learning. One older adult had to be excluded due to technical problems during data acquisition.

Positive learners were defined as participants, who showed a higher mean accuracy in the positive compared to the negative learning condition of the learning phase. In contrast, negative learners were defined as participants, who showed a higher mean accuracy in the negative compared to the positive learning condition of the learning phase. From the remaining 29 younger adults 13 were negative learners and 16 were positive learners. In contrast, from the 29 older adults 17 were negative learners, whereas 12 were positive learners. Since younger positive learners performed overall better than younger negative learners the learner groups were matched with respect to their overall accuracy[12]. This was done because the first experiment showed that the ERN is sensitive to performance differences between groups.

---

[12] Younger positive learners (M = .71, SD = .11) performed overall better than younger negative learners (M = .67, SD = .13), whereas the older learner groups showed a similar overall accuracy (negative learners: M = .56, SD = .11, positive learners: M = .55, SD = .09). In order avoid potentially confounding effects of

The effective sample consisted of 26 younger adults (mean age = 22.3, SD = 2.2) and 24 older adults (mean age = 69.6, SD = 2.8) (see Table 7). Regarding the psychometric tests (Digit-Symbol Substitution test, adapted from Wechsler, 1982 and Spot-a-Word test, adapted from Lehrl, 1977, for a description see previous experiments) a similar pattern as in the first Experiment was obtained. Younger adults reached a substantially higher score than older adults on the Digit-Symbol Substitution test, F(1, 46) = 26.94, $p < .0001$, $\eta^2 = .37$ (see Table 7). In contrast, in the Spot-a-Word older adults reached higher scores than younger adults, F(1, 46) = 15.35, $p < .0001$, $\eta^2 = .25$ (see Table 7). Neither in younger adults, nor in older adults the learner groups differed significantly with respect to age or psychometric measures (p's > .70) (see Table 7).

	Younger adults		***Older adults***	
	Positive learners	Negative learners	Positive learners	Negative learners
n / Gender	13 / 4 female	13 / 7 female	12 / 6 female	12 / 7 female
Age Range	19 - 27	20 - 27	65 - 74	66 - 75
Mean Age (M / SD)	22.2 (2.0)	22.5 (2.5)	69.7 (2.8)	69.5 (2.9)
Digit-Symbol Substitution test (M / SD)	62.6 (13.3)	63.2 (13.1)	45.4 (9.1)	47.0 (8.9)
Spot-a Word test (M / SD)	31.2 (2.1)	30.4 (3.0)	33.3 (2.1)	33.5 (2.0)

Table 7: Demographic characteristics of the sample in Experiment 3. Digit-Symbol Substitution test, adapted from Wechsler, (1982), Spot-a-Word test, adapted from Lehrl, (1977).

Stimuli and Task. Stimuli were presented in color against a dark grey background on a 17-inch computer screen. The stimulus set consisted of 32 colored images of objects from the Snodgrass and Vanderward (1980) picture database. The objects belonged to

---

performance differences between the learner groups, the individuals of the learner groups were matched with respect to their overall accuracy. This matching procedure resulted in equal sample sizes (n = 13 for the younger learner groups, n = 12 for the older learner groups) and similar overall accuracy levels (younger positive: M = .69, SD = .11, younger negative: M = .67, SD = .13, older positive: M = .55, SD = .09, older negative: M = .56 SD = .12).

one of the following four categories: clothes, fruit, vegetables, and furniture. The feedback stimuli (see Figure 26) indicated a loss of 50 Euro Cents (-50), a gain of 50 Euro Cents (+50), or a neutral outcome (*00), and were displayed in black. If the response deadline was missed, the German words ‘ZU LANGSAM’ (‘too slow’) were presented in blue color.

The participants were asked to make a two-choice decision upon presentation of the imperative stimulus and to press one of two response keys of the response pad. They were informed about the two learning conditions and about the fact that in the one condition they could either win 50 Euro Cents or get a neutral outcome, whereas in the other condition they could loose 50 Euro Cents or get a neutral outcome. They were instructed to learn the stimulus-response assignments by trial and error based on the feedback and were motivated to maximize their profit. Participants were not informed about the fact that the stimulus-response mappings were reversed within the learning blocks. Each subject received 22.5 Euro for participation and could win an additional bonus of 7.50 Euros depending on their mean performance.

Experimental design. The design involved two learning conditions, the positive learning condition, in which the participants could either win 50 Cents or get a neutral outcome and the negative learning condition, in which they could loose 50 Cents or get a neutral outcome (see Figure 26). Two stimuli (A and B) of each learning block were associated with the positive learning condition. If participants responded with a right button press to stimulus A (e.g., the pullover in Figure 26), they won fifty Euro Cents, whereas if they responded with a left button press, they received a neutral outcome (and vice versa for stimulus B). The other two stimuli (C and D) were associated with the negative learning condition. If participants responded with a right button press to stimulus C (e.g. the onion in Figure 26) they received neutral feedback, whereas if they responded with a left button press, they lost 50 Euro Cents (and vice versa for stimulus D).

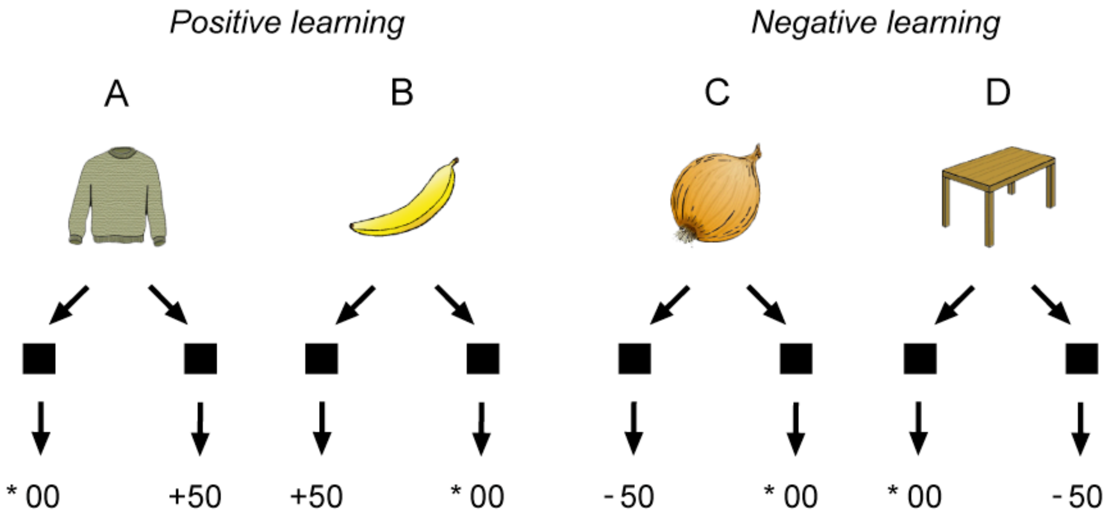

Figure 26: In the positive learning condition (stimuli A and B) participants could either win 50 Cents or get a neutral feedback. In the negative learning condition (stimuli C and D) they could either loose 50 Cents or get a neutral feedback.

Feedback was valid in 90% of the trials but also invalid in 10% of the trials. This probability of valid to invalid feedback trials was applied in order to prevent ceiling effects in the accuracy data of younger adults. Furthermore, each experimental block involved two phases, a learning phase in which the stimulus-response mappings were acquired and a reversal phase, in which the assignments were reversed. The time point at which the learning phase switched to the reversal phase depended on a performance criterion (a mean accuracy of .65 averaged across learning conditions) (see Figure 27). If participants reached this criterion after 70 trials of the learning phase the mappings were reversed. If this was not the case, the subjects had to perform additional trials, until they reached the criterion, or performed a maximum trial number of 100 trials. The same procedure was applied for the reversal phase. The reason for using this performance dependent criterion was to equate performance levels in younger and older adults at the time of the switch from the learning to the reversal phase[13]. The assignment of stimuli and responses was randomized across subjects.

---

[13] As expected, older adults performed more trials than younger adults in the learning phase (older adults: M = 92, SD = 8, younger adults: M = 79, SD = 8), as well as the reversal phase (older adults: M = 96, SD = 6, younger adults: M = 84, SD = 9). This is because it took them longer to reach the performance criterion.

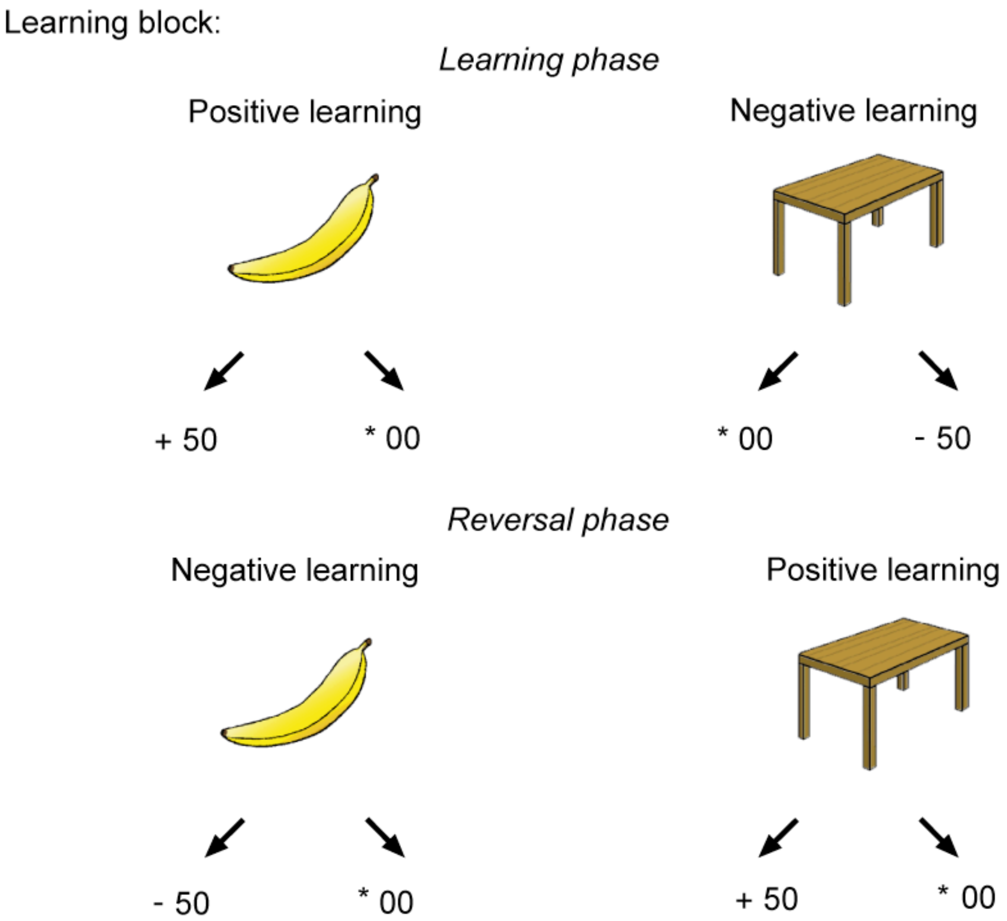

Figure 27: A learning block consists of two block phases. In the learning phase participants learn the stimulus-response assignments in the two learning conditions. In the reversal phase the stimulus-response assignments are reversed and participants have to relearn the mappings.

Trial Procedure. At the beginning of each trial, a fixation cross was displayed for 500 ms, which was followed by the imperative stimulus for again 500 ms. Similar to the previous experiments, the response deadline was adapted in 100 ms steps in a range of 700 to 1100 ms, depending on the proportion of time-out trials relative to performed trials. Each participant started with a response deadline of 800 ms. After the first trial the algorithm kept track of the proportion of time-outs (number of time-out trials relative to the trials performed). If the proportion of time-outs was smaller than two percent, a response deadline of 700 ms was applied. With steps of two percent, the response deadline increased for 100 ms and reached a maximum deadline of 1100 ms, with over eight percent of time-out trials. The deadline procedure was applied in order to make sure that all subjects produced a similar proportion of time-out trials (M = .01, SD = .01, for younger adults, M = .05, SD = .05, for older adults), and thereby had a similar opportunity to learn from feedback. Following the key press, a blank screen was displayed for 500 ms and then the feedback appeared for again 500 ms. Then participants entered the next trial.

Procedure. First, each participant filled out an informed consent and a short demographic questionnaire. Then, they performed the two psychometric tests. The experiment consisted of one practice block and eight experimental blocks. Each block involved a new set of four imperative stimuli, which were drawn randomly (without replacement), from the four stimulus categories (see Stimuli). In a practice block (100 trials), the participants were familiarized with the experimental setting. Finally, they performed the eight experimental blocks, which consisted of a learning phase and a reversal phase. In the learning, as well as the reversal phase, each of the four imperative stimuli was presented 18 – 25 times in random order, depending on whether the participants reached the performance criterion (see Experimental design).

Data Recording. An IBM compatible computer was used for collecting reaction times (RTs) and accuracy data. The stimuli were presented on a CTX 17-inch color monitor with a dark grey background. Responses were registered on a response pad (Cedrus Corporation) and the experiment was controlled by the Software E-Prime. EEG and EOG activity were recorded continuously (Brain Amp DC Recorder and Brain Vision Recorder acquisition software) from 64 Ag/AgCl electrodes (10 - 10 system) using EasyCaps recording caps. The left mastoid was used as reference and the right mastoid was recorded as an active channel. The EEG and EOG signals were filtered online from DC - 70 Hz and digitized at 500 Hz. Vertical and horizontal EOG was recorded from two electrode pairs placed on the infra- and supraorbital ridges of the right eye and on the outer canthi of the two eyes. Impedances were kept below 10 kΩ. To increase S-R ratio and to obtain more reliable mean amplitude measures, the response-locked, as well as feedback-locked EEG data were filtered using a 15Hz low-pass (see also Frank et al., 2005).

Behavioral Data Analysis. Responses faster than 149 ms (more than two standard deviations from the mean reaction time in both age groups) and responses that exceeded the response deadline (younger adults: M = 775 ms, SD = 108 ms; older adults: M = 859 ms, SD = 147 ms) were excluded from data analysis. The accuracy data was analyzed by averaging mean accuracy rates individually, for each subject, learning phase, and learning condition into two block halves (Table 8, see Appendix). The mean accuracy rates (% correct) were then subjected to an analysis of variance (ANOVA).

ERP Data Analysis. The EEG epochs were averaged with respect to response and feedback onset to obtain response-locked and feedback-locked ERPs. Similar to previous studies on age differences in the ERN (Nieuwenhuis et al., 2002; Themanson et al., 2006), the response-locked EEG data was baseline corrected by subtracting the average activity during -200 and -50 ms preceding the response. For the feedback-locked EEG data, the average activity from -200 ms to feedback onset served as baseline. As in the previous study, ocular artifacts or other artifacts were excluded from further analysis and remaining eye movements were corrected using a modified version of the approach developed by Gratton et al. (1983).

The response-locked ERPs were measured as mean amplitudes in a 0 - 100 ms time window following the response at electrode FCz. The feedback-locked components were measured as the mean amplitudes within a 100 ms time window centered on the peak of the feedback ERN at the electrode FCz (260 ms in younger adults and 290 ms in older adults). As in the first experiment, the ERN was defined as the peak-to-peak voltage difference between the most negative peak between -50 and 150 ms around the response and the preceding positive peak. The feedback ERN was defined as the difference between the most negative peak within 200 to 400 ms and the preceding positive peak (for a similar procedure, see Frank et al. (2005), Yeung & Sanfey, (2004)).

Whenever necessary the Geisser-Greenhouse correction was applied (Geisser & Greenhouse, 1958). In these cases the original F-value, the adjusted p-values, and the Epsilon values (ε) are reported. Additionally effects sizes (eta squared, η^2) are reported, which reflect the proportion of variance that is accounted for by the experimental manipulations (see Cohen, 1973). Similar to the behavioral data, learning-related effects were investigated by comparing the ERPs for the first and the second half of the learning and reversal phase.

6.5 Results

6.5.1 Accuracy Data

Response accuracy was analyzed with an ANOVA design, involving the factors Age group (younger, older), Learners (positive, negative), Block phase (learning, reversal), Learning condition (positive, negative), and Block half (first, second). The analysis revealed a significant main effect of age group, $\underline{F}(1, 46) = 41.78$, $\underline{p} < .0001$, $\underline{\eta}^2 = .47$, which reflects the fact that older adults performed overall worse than younger adults. Furthermore, a significant main effect of block phase, $\underline{F}(1, 46) = 32.24$, $\underline{p} < .0001$, $\underline{\eta}^2 = .39$, and a significant interaction between age group and block phase, $\underline{F}(1, 46) = 4.32$, $\underline{p} < .04$, $\underline{\eta}^2 = .05$ was obtained. Separate analyses for the two age groups revealed a significant effect of block phase in younger ($\underline{p\text{'s}} < .0001$, $\underline{\eta}^2 = .57$) and older adults ($\underline{p\text{'s}} < .02$, $\underline{\eta}^2 = .21$). As can be seen in Figure 28 both age groups performed better in the learning compared to the reversal phase and this effect was more pronounced for younger than for older adults.

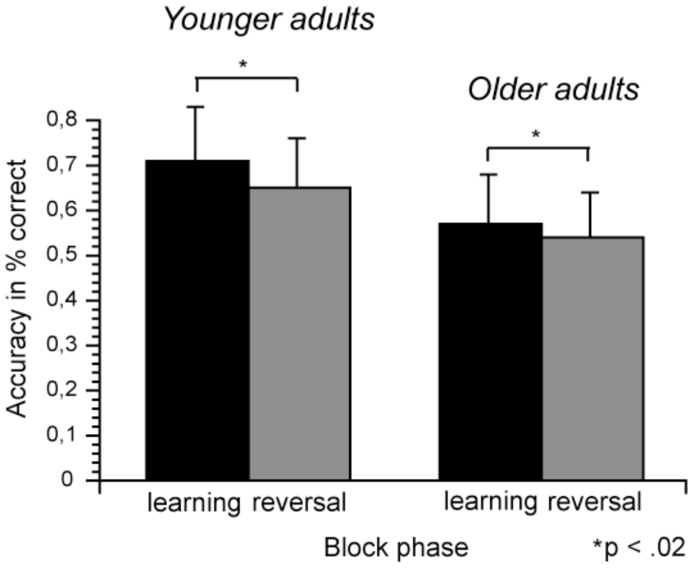

Figure 28: Mean accuracy for younger and older adults in the learning and the reversal phase, aggregated across learners, learning condition and block half.

Moreover, the analysis showed a significant main effect of learning condition, $F(1, 46) = 7.72$, $p < .008$, $\eta^2 = .06$ and an interaction between block phase and learning condition, $F(1, 46) = 17.86$, $p < .0001$, $\eta^2 = .27$. Separate analyses for the factor block phase showed a significant effect of learning condition only for the reversal phase ($p < .0001$, $\eta^2 = .18$), but not for the learning phase ($p = .99$). This finding reflects the fact that in the reversal phase participants performed better in the negative learning condition ($M = .62$, $SD = .11$) than in the positive learning condition ($M = .57$, $SD = .12$). In contrast, no effect of learning condition was obtained in the learning phase (negative learning: $M = .64$, $SD = .13$, positive learning: $M = .64$, $SD = .14$). This finding suggests that in the reversal phase, in which the previously learned mappings produce interference, participants tend to adopt a more conservative response strategy and decide for responses that lead to neutral outcomes in order to avoid negative feedback.

Learning effects: The analysis also revealed significant age differences in the learning effects that is, it showed a significant effect of block half, $F(1, 46) = 137.31$, $p < .0001$, $\eta^2 = .64$, and a significant interaction between age group and block half, $F(1, 46) = 30.29$, $p < .0001$, $\eta^2 = .14$. Separate analyses for the two age groups revealed significant effects of block half for younger adults ($p < .0001$, $\eta^2 = .86$), as well as older adults ($p < .0003$, $\eta^2 = .46$). As can be observed in Figure 29, a larger increase of accuracy from the first to the second block half was obtained for younger compared to older adults.

Analysis of learning biases: Of most interest in the third experiment were the learning biases in younger and older adults. As could be expected, according to the definition of positive and negative learners in the this experiment, the analysis revealed a significant interaction between learners, and learning condition, $F(1, 46) = 66.28$, $p < .0001$ $\eta^2 = .52$, which reflects the fact that positive learners showed a higher accuracy for the positive compared to the negative learning condition, and vice versa for negative learners. Moreover, a significant interaction between age group, learners and learning condition, $F(1, 46) = 7.41$, $p < .009$, $\eta^2 = .06$ was obtained. Post-hoc tests for the two age groups and the two learner groups showed significant learning biases in overall accuracy in the expected direction for negative and positive learners in younger adults (p's < .02, η²'s > .39), as well as older adults (p's < .003, η²'s > .56). As can be seen in Figure 29, the learning biases in overall accuracy were more pronounced for older compared to younger adults, suggesting that the asymmetry in overall performance between learner groups was more pronounced for the elderly. The fact that this pattern of results was obtained for the learning (p < .0001, η² = .57), as well as for reversal phase (p < .0001, η² = .31), suggests that participants responded according to their biases irrespective of interference in the reversal phase. This result validates the distinction between positive and negative learners.

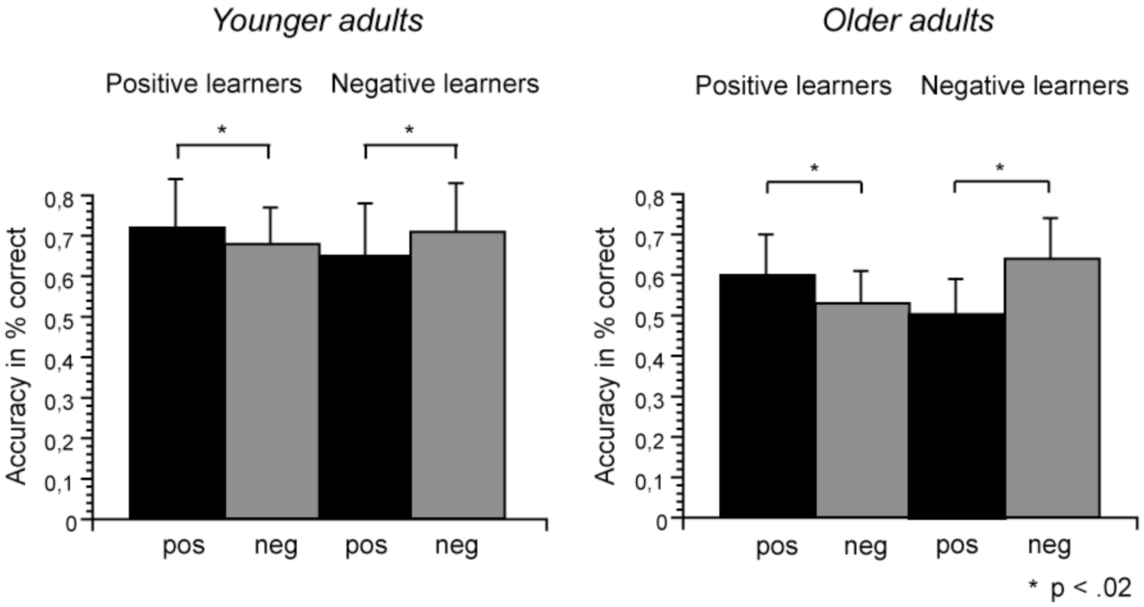

Figure 29: Mean accuracy in the positive (black) and the negative (grey) learning condition, displayed separately for younger and older adults and for positive and negative learners. Mean accuracy is averaged across block phase and block half.

Moreover, the analysis not only revealed significant learning biases in overall accuracy but also with respect to the learning effects, as reflected in a significant interaction between learners, learning condition, and block half, F(1, 46) = 15.56, p < .0003, η² = .23. Separate analyses for the two learner groups showed significant interactions between learning condition and block half for positive learners (p < .0001, η² = .28), but not for negative learners (p < .13).

However, as can be also observed in Figure 30 the absence of an interaction between valence and block half in negative learners is due to the fact that the older negative learners showed a similar increase of accuracy with learning for the positive and the negative learning condition. Post-hoc tests for the two age groups and the two learner groups showed significant interactions between learning condition and block half for younger positive learners ($p < .002$, $\eta^2 = .56$) and younger negative learners ($p < .04$, $\eta^2 = .32$). In contrast, for the elderly such an interaction was only obtained for positive learners ($p < .009$, $\eta^2 = .49$), but not for negative learners ($p = .92$).

These findings suggest that younger positive learners are biased towards learning better from positive outcomes, whereas younger negative learners are biased towards learning to avoid negative outcomes. For older adults, this pattern of results was only obtained for positive learners, whereas older negative learners were not biased towards better learning in either of the learning conditions.

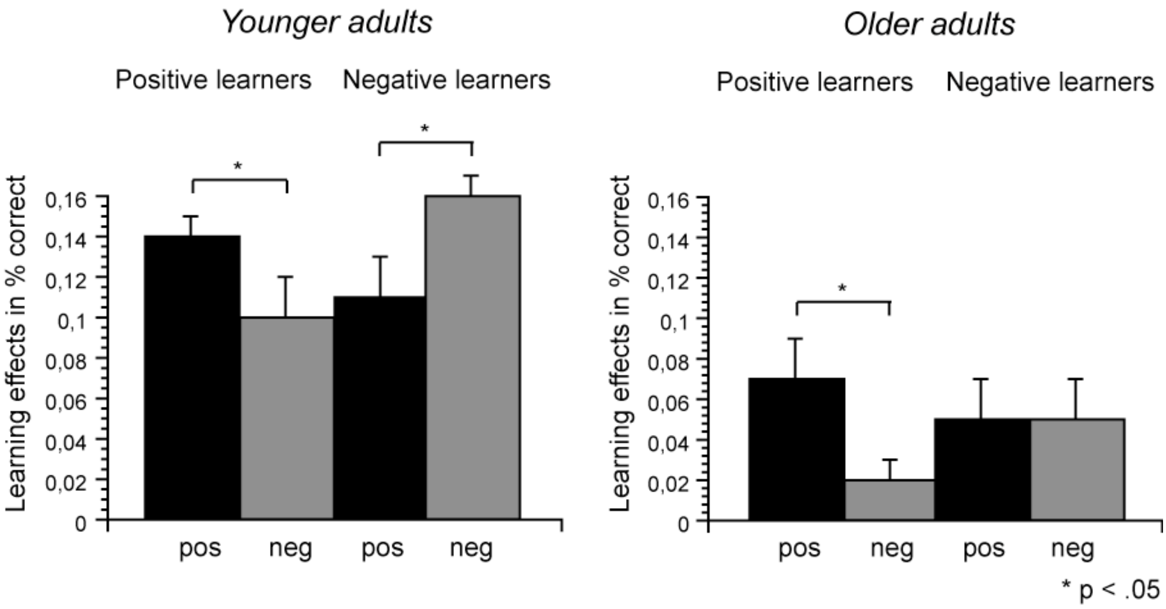

Figure 30: Mean accuracy learning effects in the positive and the negative learning condition, displayed separately for younger and older adults and positive and negative learners. Learning effects reflect the difference between block half two and block half one and are averaged across the factor block phase.

6.5.2 ERP Data

Response-locked ERPs

Figure 31 displays the ERPs to correct and incorrect responses, separately for the two age groups, the two block phases (learning and reversal) and the two learning conditions (positive and negative). As can observed in Figure 31, younger adults showed a pronounced negativity to incorrect responses (ERN), as well as a positivity to correct responses (response-locked positivity), and a small negativity for correct responses that is superimposed on the response-locked positivity. In younger adults, the ERN seems to decrease from the learning to the reversal phase. In older adults, the ERN seems to be reduced and appears to be larger for the negative compared to the positive learning condition. Most obviously, younger adults seem to be able to differentiate correct from incorrect responses in both block phases. In contrast, older adults only show a small effect of response type in the learning phase but do not seem to differentiate between correct and incorrect responses in the reversal phase. The analysis will first focus on the ERP components to correct and incorrect responses, using mean amplitude measures. In the following the ERN and the CRN will investigated by means of peak-to-peak measurements.

The mean amplitude measures of the ERPs to correct and incorrect responses were analyzed using an ANOVA design with the factors Age group (young, old), Learners (positive, negative), Block phase (learning, reversal), Learning condition (positive, negative), Response type (correct, incorrect) and Block half (first, second).

The analysis revealed a significant main effect of response type, $\underline{F}(1, 46) = 77.81$, $\underline{p} < .0001$, $\underline{\eta}^2 = .46$, an interaction between age group and response type $\underline{F}(1, 46) = 46.02$, $\underline{p} < .0001$, $\underline{\eta}^2 = .27$, as well as a significant interaction between age group, block phase, and response type, $\underline{F}(1, 46) = 5.69$, $\underline{p} < .02$, $\underline{\eta}^2 = .08$. Separate analyses for the two age groups and the factor block phase revealed that for younger adults, there was a significant main effect of response type for the learning and for the reversal phase ($\underline{p\text{'s}} < .0001$, $\underline{\eta^2\text{'s}} > .65$). In contrast, for older adults a significant main effect of response type was only observed for the learning phase ($p < .01$, $\eta^2 = .26$), but not for the reversal phase ($p = .18$) (see Figure 31).

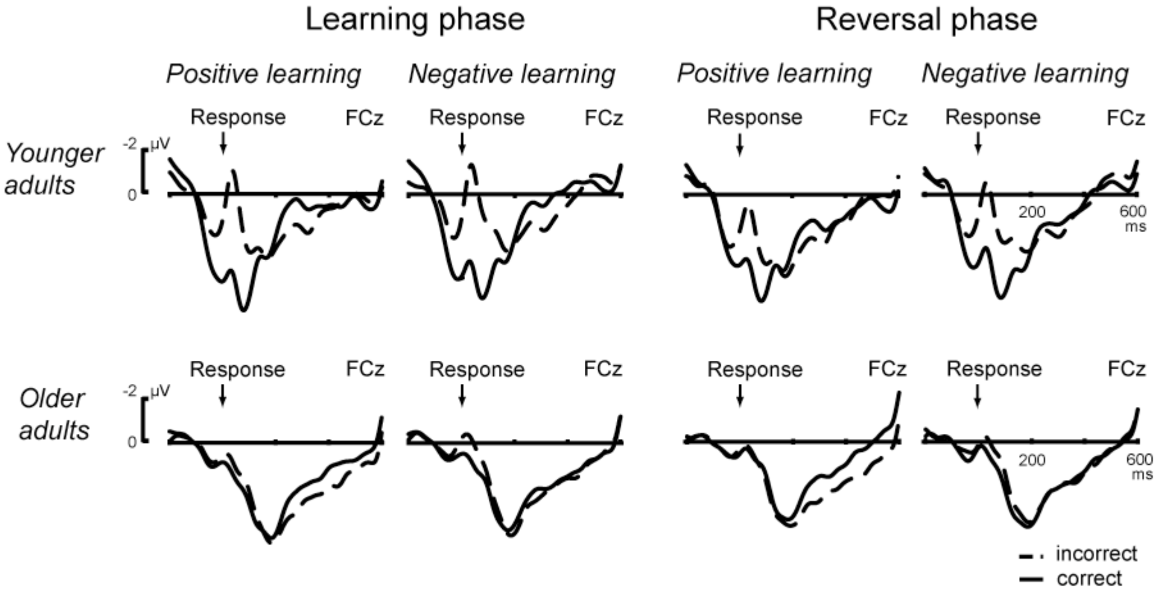

Figure 31: ERPs to correct and incorrect responses for younger and older adults, displayed separately for the learning and the reversal phase and the two learning conditions (positive learning and negative learning).

These findings suggest that younger adults were able to differentiate correct from incorrect responses in the learning and the reversal phase. In older adults, this was only the case in the learning phase, whereas they were impaired in representing the correctness of the response when stimulus-response assignments were reversed.

Leaning-related effects: The analysis showed a significant main effect of block half, $F(1, 46) = 7.43$, $p < .009$, $\eta^2 = .12$, as well as significant interactions between response type and block half, $F(1, 46) = 36.75$, $p < .0001$, $\eta^2 = .41$, and between age group, response type, and block half, $F(1, 46) = 9.50$, $p < .004$, $\eta^2 = .10$. Separate analyses for the two response types revealed a significant effect of block half and a significant interaction between age group and block half only for correct responses (p's < .0001, η^2's > .14), but not for incorrect responses (p's > .13). This finding suggests that the response-locked positivity for correct responses increases as a function of learning. Post-hoc test revealed a significant effect of block half on correct trials for younger adults ($p < .0001$, $\eta^2 = .71$) and older adults ($p < .02$). As illustrated in Figure 32, the learning-related effects in the response-locked positivity were larger for younger compared to older adults.

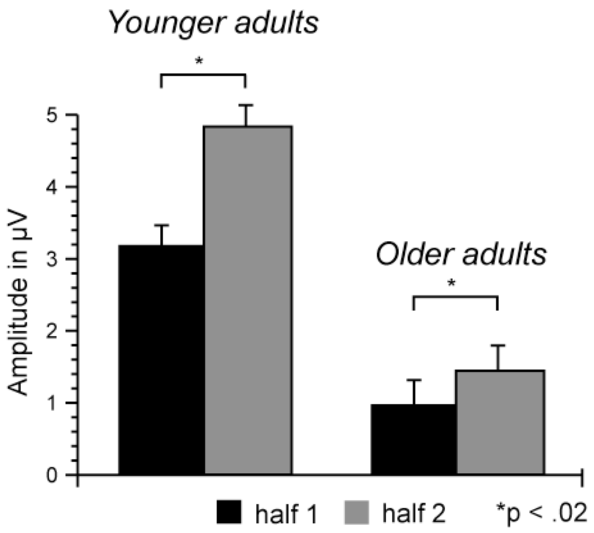

Figure 32: Mean amplitude of the response-locked positivity in the first and the second block half, displayed separately for younger and older adults. The mean amplitudes were averaged across learners, block phase and learning condition.

Learning biases. Moreover, the analysis showed a significant effect of learning condition, F(1, 46) = 15.96, p < .0002, η² = .25, and a significant three-way interaction between learners, learning condition, and response type, F(1, 46) = 4.81, p < .03, η² = .09. Post-hoc tests showed that a significant effect of learning condition was only obtained on incorrect trials for negative learners (p < .0005, η² = .42), but not for positive learners (p = .11). These findings suggest that the ERN for negative learners was larger for the negative compared to the positive learning condition, which was not the case for positive learners (see Figure 35).

Response – outcome relations. One of the most interesting aspects of the ERP data of the third experiment is the fact that older adults seem to be less able than younger adults to differentiate between correct and incorrect responses, when these responses both lead to ambiguous neutral feedback (see Figure 33). As can be observed in Figure 33, younger adults seem to clearly differentiate between correct and incorrect responses, irrespectively of whether these responses are followed by ambiguous (neutral, *00) or unambiguous (positive, +50 or negative, -50) outcomes. Older adults only seem to differentiate between the two response types if they are followed by unambiguous (positive, negative) outcomes. In contrast, there seems to be no difference in the ERPs to correct and incorrect responses in older adults, when these responses lead to neutral outcomes.

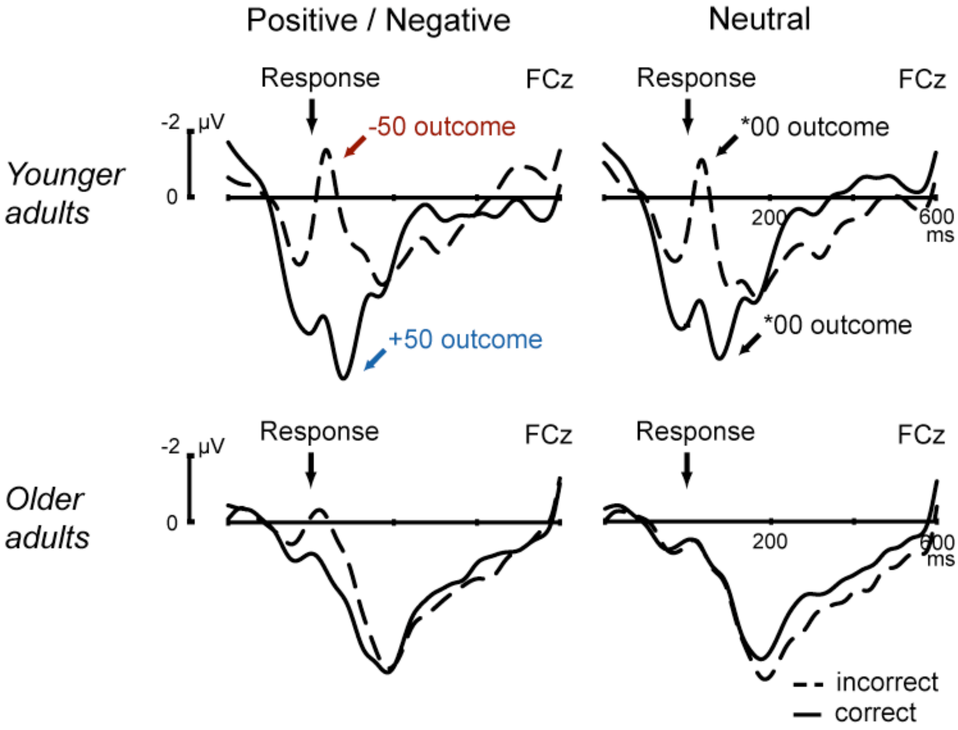

Figure 33: ERPs to correct and incorrect responses, which lead to positive (+50), negative (-50) and neutral (*00) outcomes, displayed separately for younger and older adults.

In order to investigate these effects an additional analysis was performed, in which the four response - outcome relations (correct positive, incorrect negative, correct neutral and incorrect neutral) were directly compared. This analysis revealed a main effect of response outcome, F(1, 46) = 51.63, p < .0001, η² = .41, as well as an interaction between age group and response outcome, F(1, 46) = 28.22, p < .0001, η² = .22. Post-hoc contrasts for each of the levels of the factor response outcome that were performed separately for the two age groups revealed that for younger adults there was a significant difference between correct and incorrect responses that lead to ambiguous neutral (*00) outcomes (p < .0001, η² = .58). In contrast, for older adults no significant difference between these responses was obtained (p = .87) (see Figure 33). In contrast, when comparing correct to incorrect responses that lead to unambiguous positive (+50 Cents) or negative (-50 Cents) outcomes a significant difference was obtained for younger adults (p < .0001, η² = .80), as well as for older adults (p < .005, η² = .31) (see Figure 33).

These findings show that younger adults are able to differentiate correct from incorrect responses irrespectively of whether these response lead to unambiguous (positive, negative) or ambiguous (neutral) outcomes. This suggests that younger adults are able to represent that in the negative learning condition the neutral feedback is better than the alternative (negative) outcome, whereas in the positive learning condition the neutral feedback is worse than the alternative (positive) outcome. That is, they are able to build up a relational representation of the correctness of the response. In contrast, older adults were not able to differentiate correct from incorrect responses that led to neutral outcomes. This suggests that they are impaired in representing the correctness of responses when the outcome of the response is ambiguous and must be processed in relation to the alternative outcome.

Peak-to-peak measures of the ERN. The peak-to-peak measures of the ERN were analyzed using an ANOVA design with the factors Age group, Learners, Block phase, Learning condition, and Block half. The analysis revealed a significant main effect of age group, $F(1, 46) = 12.43$, $p < .001$, $\eta^2 = .20$, indicating that the ERN was generally reduced for older compared to younger adults (see Figure 31). Furthermore, a main effect of block phase, $F(1, 46) = 18.00$, $p < .0001$, $\eta^2 = .16$, and a significant main effect of block half $F(1, 46) = 10.06$, $p < .003$, $\eta^2 = .16$ was obtained. Significant interactions were obtained between the factors age group and block phase, $F(1, 46) = 19.38$, $p < .0001$, $\eta^2 = .17$, age group and block half, $F(1, 46) = 8.05$, $p < .007$, $\eta^2 = .13$, and between age group, block phase, and block half, $F(1, 46) = 11.45$, $p < .002$, $\eta^2 = .18$. Separate analyses for the two age groups revealed a significant main effect of block phase and a significant interaction between block phase and block half for younger adults (p's $< .002$, η^2's $> .34$), but not for older adults (p's $> .42$). Post-hoc tests showed that for younger adults a significant main effect of block half was obtained in the learning phase ($p < .0003$, $\eta^2 = .42$) but not in the reversal phase ($p = .27$) (see Figure 34).

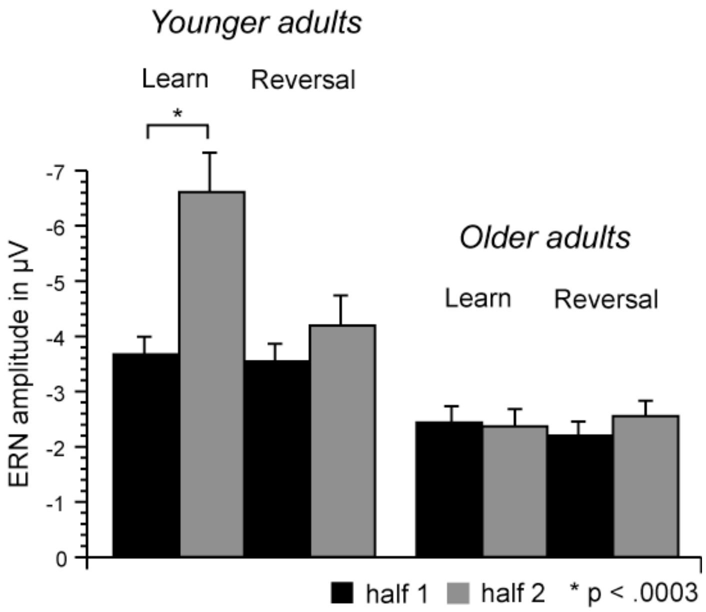

Figure 34: ERN amplitude for younger and older adults in the learning and the reversal phase, displayed separately for the first and the second block half. Note, that the mean amplitudes were averaged across the factors learners and learning condition.

These findings show that for younger adults the ERN was increased in the learning phase compared to the reversal phase, which was not the case for older adults. This effect seems to be due to the fact that in younger adults the ERN increased with block half in the learning phase, but not in the reversal phase. In contrast, no such learning-related effects in the ERN were obtained for older adults.

Learning biases in the ERN. Most interestingly, the analysis also revealed significant learning biases in the ERN (see Figure 35). The ANOVA showed a significant interaction between learners and block phase, $\underline{F}(1, 46) = 17.26$, $\underline{p} < .0001$, $\underline{\eta}^2 = .15$, as well as a three-way interaction between age group, learners, and block phase, $\underline{F}(1, 46) = 12.11$, $\underline{p} < .001$, $\underline{\eta}^2 = .11$. Separate analyses for the two age groups showed a significant interaction between learners and block phase for younger adults ($\underline{p} < .0002$, $\underline{\eta}^2 = .28$), but not for older adults, ($\underline{p} = .40$). Post-hoc contrasts revealed a significantly larger ERN for younger negative learners compared to younger positive learners in the learning phase ($\underline{p} < .03$, $\underline{\eta}^2 = .18$), but not in the reversal phase ($\underline{p} = .94$). In contrast, no learning biases in the ERN were obtained for older adults (see Figure 35).

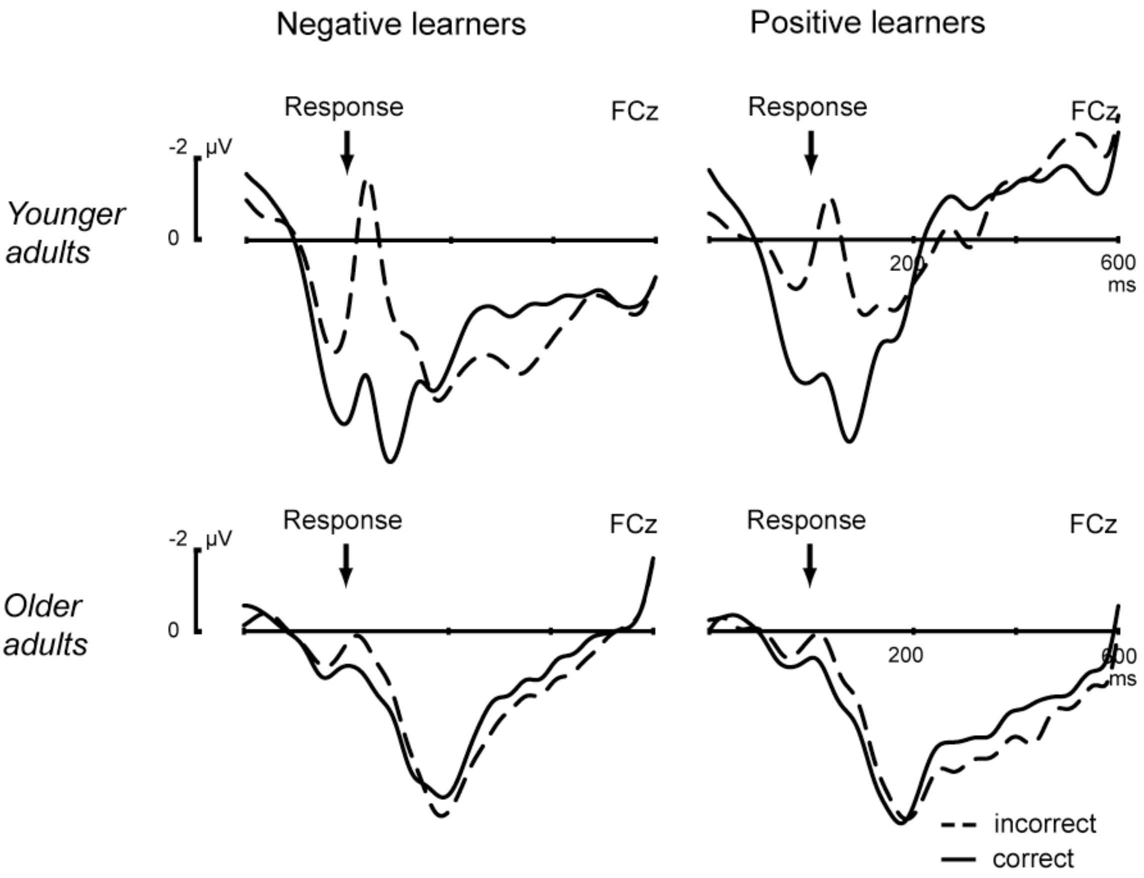

Figure 35: ERPs to correct and incorrect responses for younger and older adults in the learning phase, displayed separately for negative and positive learners. The ERN was averaged across the factor block half.

Peak-to-peak measures of the CRN: The same ANOVA design as for the ERN was applied for the analysis of the peak-to-peak measures of the CRN. The analysis revealed a significant main effect of block phase, F(1, 46) = 10.66, p < .002, η² = .16, indicating that the CRN was increased for the reversal compared to the learning phase (see Figure 31). Furthermore, a significant main effect of block half was obtained, F(1, 46) = 8.71, p < .005, η² = .15, which reflects the fact that the CRN decreases with learning from the first to the second block half. Similar to the peak-to-peak analysis of the ERN a significant interaction between age group, learners and block phase, F(1, 46) = 3.94, p < .05, η² = .06 was obtained. Separate analysis for the two age groups revealed a significant interaction between block phase and learners only for younger adults (p < .03, η² = .18), but not for older adults (p = .58). Post-hoc tests for the factors age group and block phase revealed a marginally significant main effect of learners for younger adults in the learning phase (p < .07, η² = .13), but not in the reversal phase (p < .52). Hence, as can be observed in Figure

35 similar to the ERN, the CRN was larger for younger negative learners compared to younger positive learners.

Feedback-locked ERPs

The mean amplitude measures of the feedback-locked ERP components were analyzed using an ANOVA with the factors Age group (young old), Learners (positive, negative), Block phase (learning, reversal), Outcome (positive, negative, positive zero and negative zero) and Block half (first, second). The analysis revealed a significant main effect of outcome $F(3, 138) = 176.66$, $p < .0001$, $\varepsilon = .82$, $\eta^2 = .64$ and an interaction between block phase and outcome, $F(3, 138) = 4.62$, $p < .01$, $\varepsilon = .95$, $\eta^2 = .08$. Post-hoc contrasts for each of the levels of the factor outcome that were performed separately for the two block phases revealed that for both block phases there was a significantly larger feedback-locked postivity for negative compared to positive outcomes, (p's $< .0005$, η^2's $> .23$). Moreover, the positivity was larger for positive and negative outcomes than neutral outcomes, (p's $< .0001$, η^2's $> .50$) (see Figure 36).

However, in contrast to the learning phase, in which there was no significant difference between positive and negative zero outcomes ($p = .57$) in the reversal phase the feedback-locked positivity was larger for positive ($M = 3.65$ mV $SD = 1.91$ mV) than negative zero outcomes ($M = 2.97$ mV $SD = 1.81$ mV), ($p < .005$, $\eta^2 > .16$). This suggests that positive and negative outcomes were processed similarly across the block phases. However, in contrast to the learning phase, in which participants did not differentiate between neutral outcomes, in the reversal phase participants showed a larger feedback-locked positivity to positive compared to negative zero outcomes.

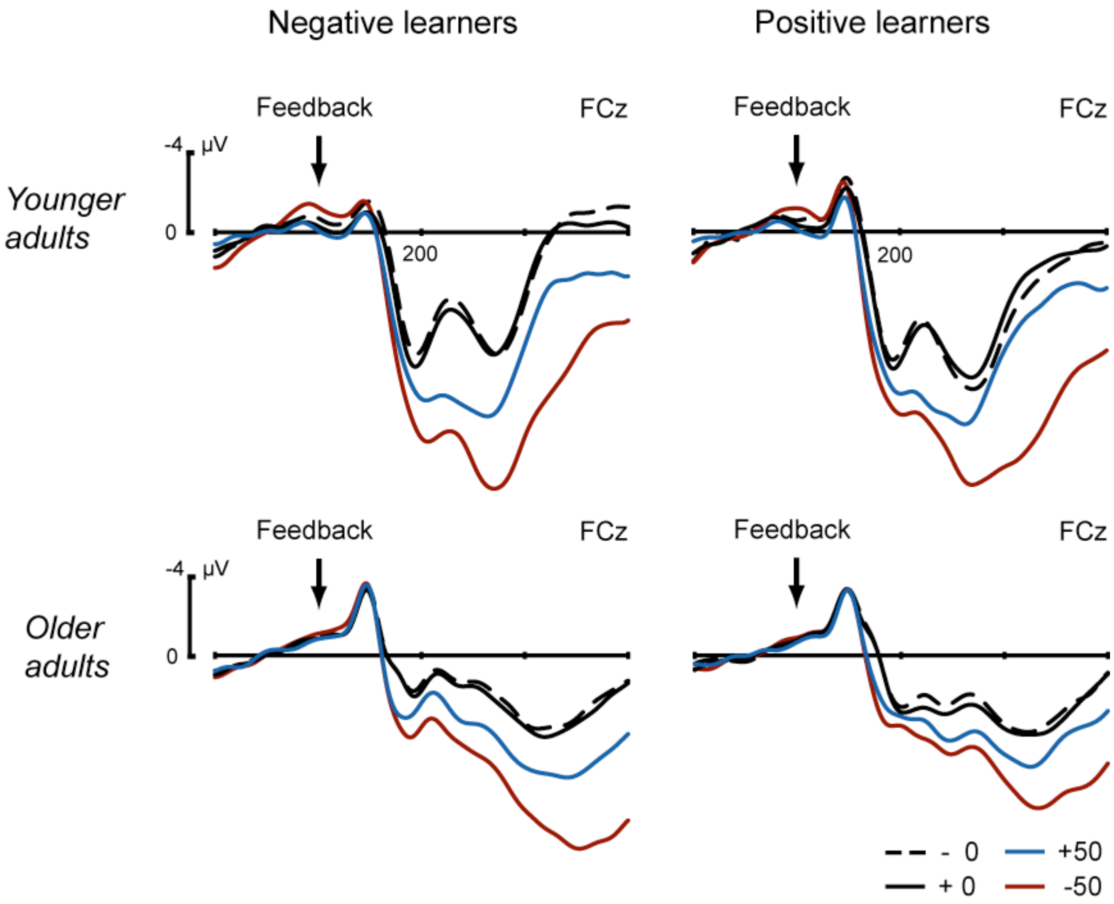

Figure 36: Feedback-locked ERPs for the four types of outcomes, displayed separately for younger and older adults and positive and negative learners.

Moreover, the analysis revealed a significant main effect of age group, $\underline{F}(1, 46) = 21.87$, $\underline{p} < .0001$, $\underline{\eta}^2 = .32$, and an interaction between age group and outcome, $\underline{F}(3, 138) = 15.77$, $\underline{p} < .0001$, $\underline{\varepsilon} = .82$, $\underline{\eta}^2 = .09$. Post-hoc contrasts for each of the levels of the factor outcome showed significant age differences when comparing neutral outcomes (positive and negative zero outcomes) with positive (+50) or negative (-50) outcomes ($\underline{p\text{'s}} < .0003$, $\underline{\eta^2\text{'s}} > .10$). However, no age differences were obtained when comparing positive with negative neutral (*00) outcomes ($\underline{p} = .71$) or when comparing positive (+50) with negative (-50) outcomes ($\underline{p} = .31$). These findings suggest that in contrast to younger adults, older adults were less sensitive to positive or negative outcomes in relation to neutral outcomes (see Figure 36).

Learning-related effects. The ANOVA also showed significant learning-related changes in the feedback-locked positivity (see Figure 37).

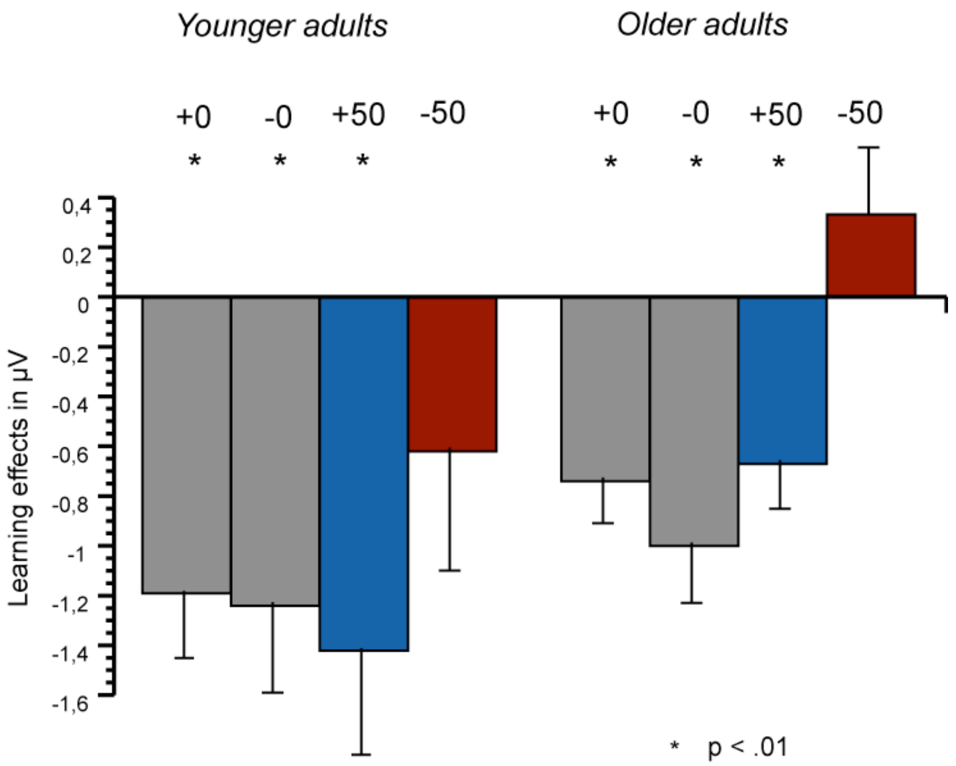

Figure 37: Learning-related reductions of the feedback-locked positivity for the four types of outcomes, displayed separately for younger and older adults.

The analysis revealed a significant main effect of block half, F(1, 46) = 25.41, $p < .0001$, $\eta^2 = .34$, and a significant interaction between outcome and block half, F(1, 46) = 4.82, $p < .005$, $\varepsilon = .90$, $\eta^2 = .09$. Separate analyses for each type of outcome showed significant main effects of block half for neutral and positive outcomes (p's < .0001, η²'s > .28). No significant effect of block half was obtained for negative outcomes ($p = .61$, $\eta^2 = .00$). As shown Figure 37, these findings suggest that the feedback-locked positivity decreases for neutral and positive outcomes, whereas no learning-related changes are found for negative outcomes.

Learning biases. Interestingly, the analysis also revealed a marginally significant interaction between learners and outcome, F(3, 138) = 2.71, $p < .06$, $\varepsilon = .82$, $\eta^2 = .01$. Contrasts for each of the levels of the factor outcome revealed a significant main effect of learners when comparing negative outcomes to positive zero outcomes ($p < .04$, $\eta^2 = .02$). In contrast, no significant effect of learners was obtained when comparing positive outcomes to negative zero outcomes ($p = .20$, $\eta^2 > .00$). As can be seen in Figure 36 this finding reflects the fact that negative learners differentiated more between negative and positive zero outcomes that is, they showed a larger difference between negative and positive feedback in the negative learning condition.

Peak-to-peak analysis of the FRN. The analysis of the peak-to-peak measures of the FRN revealed a significant main effect of block phase, $\underline{F}(1, 46) = 7.71$, $\underline{p} < .008$, $\underline{\eta}^2 = .14$, which reflects the fact that overall the FRN decreases from the learning (younger adults: $\underline{M} = -3.55$ µV, $\underline{SD} = 2.61$ µV, older adults: $\underline{M} = -3.07$ µV, $\underline{SD} = 2.69$ µV) to the reversal phase (younger adults: $\underline{M} = -3.16$ µV, $\underline{SD} = 2.30$ µV, older adults: $\underline{M} = -2.84$ µV, $\underline{SD} = 2.75$ µV). Moreover, the analysis revealed a significant main effect of outcome, $\underline{F}(3, 138) = 9.91$, $\underline{p} < .0001$, $\underline{\varepsilon} = .79$, $\underline{\eta}^2 = .14$, and a significant interaction between age group and outcome, $\underline{F}(3, 138) = 15.39$, $\underline{p} < .0001$, $\underline{\varepsilon} = .79$, $\underline{\eta}^2 = .21$. Separate analyses for the two age groups showed a significant main effect of outcome for younger adults ($\underline{p} < .0001$, $\underline{\eta}^2 = .41$), but not for older adults ($\underline{p} = .46$). Hence, in contrast to younger adults, in older adults the FRN was not sensitive to the type of outcome they obtained. In order to analyze the outcome effects in younger adults post-hoc contrasts for each of the levels of the factor outcome were performed. This analysis showed that the FRN was larger for negative (-50) compared to the positive (+50) outcomes ($\underline{p} < .0006$, $\underline{\eta}^2 = .38$). However, the FRN was found to be largest for neutral outcomes ($\underline{p\text{'s}} < .01$, $\underline{\eta^2\text{'s}} > .23$, for comparisons with negative and positive outcomes). No significant difference was obtained between positive and negative zero outcomes ($\underline{p} = .35$). Hence, in line with previous results (Holroyd & Coles, 2002; Nieuwenhuis et al., 2002), the FRN in younger adults was found to be larger for negative compared to positive outcomes. However, the component was largest for neutral outcomes, irrespective of whether they were better (positive zero outcome) or worse (negative zero outcome) with respect to the alternative outcome (see Figure 36).

6.6. Summary and Discussion

Accuracy data

The third experiment revealed significant differences in overall accuracy between younger and older adults (see Figure 29). This finding is in contrast to the results of the first experiment, which showed age differences in accuracy only in the 80% validity condition. Since feedback validity was increased in the third experiment (90% valid trials, 10% invalid trials) compared to the 80% condition in Experiment 1 there must be another factor than feedback validity, which accounts for the pronounced age differences in the third experiment. Indeed, for younger adults accuracy in the learning phase of the third experiment was increased, compared to the 80% validity condition in the first experiment (Experiment 3: M = 71, SD = .12, Experiment 1: M = .68, SD = .11). This indicates that the increase in feedback validity has supported performance in younger adults. In contrast, for older adults mean accuracy was reduced in the third experiment compared to the first experiment (Experiment 3: M = 57, SD = .11, Experiment 1: M = .61, SD = .12). Thus, although feedback validity was increased, older adults performed even worse in the third experiment.

The most apparent changes between the two experiments relate to the feedback participants received. First, the number of feedback stimuli was increased, which may have made it more difficult for the elderly to differentiate the feedback and hence may have impaired learning. However, the most likely reason for these impairments is that older adults had difficulties in disambiguating the neutral feedback. What this means is that in order to learn the stimulus-response assignments in the present task participants need to be able to differentiate the situations, in which the neutral feedback is better in relation to the alternative outcome (-50 in the negative learning condition), from those situations, in which the neutral feedback is worse in relation to the alternative outcome (+50 in the positive learning condition). What this suggests is that older adults might have had problems in building up a relational representation of feedback value. This idea is supported by findings in the response-locked ERPs that show that older adults were impaired representing the correctness of a response if the feedback that follows this response is ambiguous and needs to be processed in relation to the alternative feedback.

In contrast to the age differences in overall accuracy, a larger difference in accuracy between the learning and the reversal phase was obtained for younger, compared to older adults (see Figure 28). This suggests that although there was an effect of block phase for older adults as well, for younger adults the reversal produced more interference. These makes sense given that younger adults performed overall much better in the learning phase than older adults and thus perceived more interference when relearning the mappings. Moreover, the analysis revealed that in contrast to the learning phase, in the reversal phase participants performed better in the negative compared to the positive learning condition. This may suggest that due to the reversal participants were uncertain about the stimulus-response mappings and adopted a more conservative response strategy. That is, in order to avoid negative outcomes they decided for the neutral outcomes, but did so, irrespectively of the learning conditions and thus tended to decide for the neutral outcome even if the alternative outcome was better (in the positive learning condition).

Learning effects. Similar to overall performance larger learning effects were found for younger compared to older adults. This again is in contrast to the results in the first experiment, in which younger and older adults learned comparably. This finding indicates that older adults not only performed generally worse, but were indeed impaired in learning. In line with the argumentation raised above, it seems reasonable to assume that the ambiguity of the neutral feedback impaired the elderly in learning the stimulus-response assignments. This is because they were less able than younger adults to represent, in which cases the neutral outcomes were better, and in which they were worse with respect to the alternative outcome.

Learning biases. The focus of the third experiment was on the learning biases. Consistent with the predictions, the analysis revealed that negative learners performed significantly better in the negative learning condition, whereas positive learners showed a significantly higher overall accuracy in the positive learning condition (see Figure 29). However, one might argue that this is not a surprising pattern of results, given that participants were grouped into the learner groups based on whether they performed better in the positive or negative learning condition. Yet, there are two important additional findings that show that the learning biases in the present study indeed represent meaningful individual differences in the way participants learn from feedback. First, the learning biases were stable across the block phases, suggesting that participants, who were positive learners in the learning phase, also were positive learners in the reversal phase. What this means is that participants did not simply have a preference in responding to two stimuli, but switched their response preference according to their biases when the response outcome relations were reversed. Second, these learning biases were not only obtained for mean accuracy, but also for the learning effects (see Figure 30). Hence, younger participants not only performed better, but also learned better in the learning condition they preferred according to their learning bias. For older adults, such a learning bias was found for positive learners, whereas in negative learners small learning effects were found for both learning conditions. Thus, it could be argued that older adults differentiated less between the learning conditions, which would be in line with the predictions of the socioemotional selectivity theory (Carstensen, 2006; Mather & Carstensen, 2005).

However, a closer inspection of the mean accuracy in the first block half questions this interpretation[14]. The results of this analysis show that older negative learners performed below chance level in the positive learning condition in the first block half. This suggests that older negative learners decided for the response that leads to the neutral outcome more often than for the response that leads to the positive outcome. Hence, the

---

[14] The mean accuracy of older negative learners in the positive condition in the first block half was M = .47, SD = .07. The mean accuracy in the negative learning condition in the first block half was M = .60, SD = .09. In contrast, the mean accuracy of older positive learners in the positive learning condition in the first block half was M = .55, SD = .09, whereas the mean accuracy older positive learners in the negative condition in the first block half was M = .51, SD = .06

learning effect for older negative learners in the positive learning condition only means that coming from an accuracy below chance level in the first block half, participants perform somewhat above chance level in the second block half. Given these results it does not seem to be justified to propose that older adults have a tendency towards learning more from positive outcomes. In contrast, it seems as if older negative learners had problems in disambiguating the neutral outcomes and tended to adopt a more conservative response strategy, especially at the beginning of learning.

Thus, the present behavioral data supports recent findings that pointed to individual differences in the way participants learn from feedback (Frank et al., 2005). However, the present results go beyond that of Frank and colleagues (2005) with two respects: First, the present data show learning biases not only in overall performance, but also in the learning effects, indicating that there are considerable individual differences in how much participants learn from positive or negative feedback.[15] Second, the present findings show that the learning biases are resistant against interference during the reversal phase, suggesting that these individual differences are a very robust phenomenon. In contrast to the predictions of the socioemotional selectivity theory (Carstensen, 2006, Mather & Carstensen, 2005) the findings on age differences in learning biases do not suggest that older adults tend to learn more from positive feedback. Rather, the present findings indicate that especially older negative learners tend to adopt a more conservative response strategy when feedback is ambiguous. The fact that overall more older adults were negative learners even points to the opposite view, indicating that older adults might have a tendency towards learning more from negative outcomes.

---

[15] In the study by Frank and colleagues (2005) participants learned to a select a certain stimulus of a pair of stimuli, based on feedback in the training phase. In a subsequent test phase participants were tested with novel combinations of the learned stimuli. Participants were defined as positive learners in the test phase if they decided to choose the rewarded stimuli from the training more frequently. In contrast, participants were defined as negative learners if they avoided the punished stimuli form the training phase more frequently. This could be termed an indirect measure of learning biases since learning is not measured during the training phase, but as a response bias during the test phase. Moreover, it could be generally argued that since learning itself is not measured it is only assumed that learning resulted in these biases.

Response-locked ERPs

The analysis of the response-locked ERPs revealed that younger adults were able to differentiate correct from incorrect responses in the learning, as well as the reversal phase. For older adults this was only the case for learning phase, whereas they were impaired in representing the correctness of the response when the stimulus response assignments were reversed. This finding is consistent with the accuracy data that showed that older adults performed much worse than younger adults and were actually only somewhat better than chance performance in the reversal blocks. Moreover, this indicates that the elderly were particularly impaired in representing the correctness of the response when the stimulus-response mappings were reversed. In this way, the present results support the findings from the first experiment that showed that the more participants are able to learn (the more valid the feedback) the larger the difference between the ERN to incorrect responses and the response-locked positivity to correct responses.

The analysis of the mean amplitude measures showed learning-related changes only in the response-locked positivity for correct responses. This supports the conclusions from the previous experiments and shows that learning-related changes are more pronounced in the response-locked positivity to correct trials compared to the ERN. As illustrated in Figure 32 similar to the learning effects in the accuracy data learning-related changes in the response-locked positivity were larger for younger compared to older adults.

Response - outcome relations. As illustrated in Figure 33 younger adults clearly differentiated correct from incorrect responses irrespectively of whether these responses were followed by ambiguous (neutral, *00) or unambiguous (positive, +50 or negative, -50) outcomes. This suggests that younger adults were able to represent that in the negative learning condition the neutral feedback is better than the alternative (negative) outcome, whereas in the positive learning condition the neutral feedback is worse than the alternative (positive) outcome. That is, they were able to build up a representation of the correctness of a response even if these responses were associated with ambiguous outcomes. In contrast, older adults showed a differentiation between correct and incorrect responses only if these responses led to unambiguous positive (+50) or negative (-50) outcomes. However, as can be seen in Figure 33 they did not differentiate at all between correct and incorrect responses that led to ambiguous neutral outcomes. This finding shows that older adults were impaired in building up a representation of the correctness of a response when the feedback that was associated with this response was ambiguous and had to be processed in relation to the alternative outcome.

Peak-to-peak analysis of the ERN. Similar to the results of Nieuwenhuis and colleagues (2002) and in contrast to the first experiment the peak-to-peak analysis revealed an age-related reduction of the ERN. However, as in the Nieuwenhuis et al. (2002) study, in the present data pronounced age differences in overall performance, as well as reduced learning effects were obtained for older compared to younger adults. Hence, in line with the interpretation offered in the discussion of the first experiment, age differences in the ERN might have been confounded with differences in performance levels between age groups in the present study.

Furthermore, the analysis showed a larger ERN for the learning compared to the reversal phase for younger, but not for older adults. Moreover, in younger adults the ERN increased with block half in the learning phase, but not the reversal phase, whereas no learning-related effect in the ERN was obtained for older adults. In line with the first experiment these findings suggest that when the ERN is captured most precisely using peak-to-peak measure there is evidence for a learning-related increase in the ERN, as it is suggested by the RL-theory (Holroyd & Coles, 2002). However, as in first experiment, in which learning-related changes were restricted to the condition in which feedback was always valid, in the present experiment no learning-related changes were found for the reversal phase. This is in line with the results in the accuracy data that show that participants were impaired in the reversal blocks compared to the learning blocks and that this effect was more pronounced for younger compared to older participants. Thus, the present findings point to the view that when there is interference on the level of the stimulus-response mappings participants are impaired in performance and in the ability to represent an incorrect response as reflected in the ERN.

However, it could be also argued that due to the reversal it took the participants longer to acquire the new stimulus-response mappings and hence learning-related effects in the ERN might occur later in time. However, this would suggest that there should be differences in the learning effects between the learning and the reversal phase. Since this is clearly not the case[16] the present data indicates that whether learning-related changes are found in the ERN might depend on overall accuracy level rather than the learning effects themselves. This view would be much more in line with the fact that for older adults no learning-related changes in the ERN occurred.

Learning biases in the ERN. As expected based on the findings by Frank and colleagues (2005), a larger ERN for negative compared to positive learners was obtained. However, this was only the case for younger adults in the learning phase, but not the reversal phase. This finding supports the results by Frank et al. (2005) and suggests that the behavioral learning biases in younger adults might indeed be reflected in the amplitude of the ERN. Yet, similar to the learning-related changes in the ERN the learning biases were only obtained for the learning phase, but not for the reversal phase. If the learning biases were a stable phenomenon that is resistant against interference, as it is suggested by the behavioral data, one would expect to find a larger ERN in the reversal phase as well. However, in the reversal phase the ERN was only somewhat larger for younger negative (M = -3.85 µV, SD = 0.59), compared to younger positive learners (M = -3.74 µV, SD = 0.69). Hence, it seems reasonable to assume that similar to the learning-related effects the learning biases might also depend on how good participants are in differentiating correct from incorrect responses.

---

[16] The interaction between block phase and block half was not significant F(1, 46) = 0.52, p = .47, indicating that the mean accuracy showed a similar increase with block half for the learning phase (younger adults: M = .13, SE = .01, older adults: M = .05, SE = .02) as well as for the reversal phase (younger adults: M = .13, SE = .01, older adults: M = .04, SE = .01).

Although older adults showed a similar response biases in overall accuracy as younger adults they did not show a significant difference for negative compared to positive learners in the ERN. On the first view this might be interpreted as a less pronounced bias towards learning from negative outcomes in older adults. Yet, in line with the argumentation presented above, it could also be argued that due to their general performance impairments older adults were not able to differentiate between the response types and hence no individual differences in the ERN were obtained.

Peak-to-peak analysis of the CRN. Similar to the ERN a larger CRN amplitude was found for younger negative compared to younger positive learner (see Figure 35). In line with previous findings on the functional significance of the CRN (Bartholow et al., 2005; Eppinger et al., 2007; Kray et al., 2005) the present findings suggests that negative learners are not only more sensitive to errors but also perceive more conflict when performing a correct response. Hence, negative learners might be generally more engaged in performance monitoring than positive learners, as reflected in larger CRN and ERN amplitudes. The idea that the CRN reflects post-response conflict processing is further supported by the fact that the CRN was increased in the reversal compared to the learning phase, indicating that participants perceived more response conflict when the previously learned mappings interfered with learning the new assignments. Moreover, the CRN decreased with learning, which suggests that the more participants learned the less conflict they perceived. This finding is consistent with the results in older adults in the first experiment, which showed that the CRN also decreased the more valid the feedback, indicating that the CRN was the smaller the more participants could learn.

Feedback-locked ERPs.

The analysis of the feedback-locked ERPs showed that the feedback-locked positivity was larger for negative (-50) compared to positive (+50) outcomes and for both types of outcomes compared to neutral outcomes. This pattern of results is in line with the findings of several ERP studies on valence processing (Ito, Larsen, Smith, & Cacioppo, 1998; Kisley, Wood, & Burrows, 2007; Wood & Kisley, 2005). In these studies it was found that participants are more sensitive to negative compared to positive stimuli and for both types of valence information compared to neutral stimuli. This negativity bias is reflected in a late positive potential (LPP). However, in contrast to recent ERP studies on the negativity effect in older adults (Kisley, Wood, & Burrows, 2007; Wood & Kisley, 2006) the present data does not suggest that the negativity bias is reduced in older compared to younger adults. In contrast, the present data suggest that older adults are less sensitive to both, positive and negative outcomes in relation to neutral outcomes. This less pronounced differentiation between neutral and positive and negative outcomes might have been one reason for the problems of older adults in disambiguating the neutral outcomes. However, when arguing against a reduction of the negativity bias in older adults it needs to be considered that there are several differences between the aforementioned studies and the present experiment. First, the LPP, which is typically used as a measure of valence processing in these studies, occurs much later (300 – 900) than the feedback-locked positivity (200 – 300 ms). Second, the stimuli that were used in these studies (pictures from the International Affective Picture System, IAPS) as well as the task (categorization tasks) differ considerably from the feedback stimuli and the learning task used in the present experiment. However, as could be observed in Figure 36 the differentiation between the different types of outcomes also occurs later in the time window of the LPP, suggesting that there might a considerable overlap between the feedback-locked positivity and the LPP.

Interestingly, in addition to the valence effects the feedback-locked data showed that in the learning phase participants did not differentiate between neutral outcomes. In contrast, in the reversal phase a larger feedback-locked positivity was obtained for positive compared to negative zero outcomes. This suggests that in the reversal phase participants were more engaged in processing positive compared to negative zero outcomes. Considering that positive zero outcomes are the ‘better’ outcomes in the negative learning condition, this finding is nicely in line with the findings in the accuracy data. The analysis of the accuracy data revealed that in contrast to the learning phase in the reversal phase participants performed better in the negative compared to the positive learning condition. Taken together, these findings may reflect a more conservative response strategy during the reversal phase in which interference from previously learned mappings is present.

Learning-related effects. The analysis of the learning-related effects in the feedback-locked positivity revealed that the amplitude of the feedback-locked positivity decreased with learning for neutral and positive outcomes, but not for negative outcomes. This suggests that in contrast to neutral and positive outcomes participants did not disengage from processing negative outcomes with learning. This is consistent with the feedback-locked data of the first experiment that also showed that learning-related effects were only obtained for positive but not for negative feedback. With respect to feedback-locked positivity neutral outcomes seem to be similarly processed as positive outcomes.

Learning biases. Moreover, also found evidence for learning biases in the feedback-locked positivity. As can be seen in Figure 36 negative learners showed a more pronounced difference between positive zero and negative (-50) feedback in the negative learning condition compared to positive learners. This might suggest that negative learners more engaged in differentiating the two types of outcomes in the condition in which they learned better.

Peak-to-peak analysis of the FRN. The analysis of the FRN revealed that in line with the results of the first experiments no effect of feedback type was obtained in the FRN in older adults. This suggests that the older adults did not differentiate between the feedback types. In contrast, in younger adults a significantly larger FRN for negative compared to positive outcomes was obtained, which is consistent with several previous findings and supports the view that the FRN might reflect the binary categorization of outcomes as favorable or unfavorable (see Hajcak et al., 2006; Holroyd & Coles, 2002; Nieuwenhuis et al., 2002; Yeung & Sanfey, 2004). However, in contrast to these findings the FRN was found to be even larger for neutral compared to positive and negative outcomes. This suggests that the ambiguous neutral outcomes were processed as being even worse than the negative outcomes. Yet, this result stands in contrast to the R-L theory (Holroyd & Coles, 2002), which would predict that that neutral outcomes should elicit a FRN that is midway between the FRNs to positive and negative outcomes (cf. Holroyd, Hajcak, & Larsen, 2006). Hence, the R-L theory in its original version does not seem to be able to account for these findings. In a recent study on the effects of neutral feedback on the FRN Holroyd and colleagues (2006) showed that across five experiments the FRN to neutral feedback was similar to that to negative feedback. Based on these findings they concluded that neutral and negative feedback are grouped together as events that indicate that task goals have not been satisfied and hence both lead to the generation of an FRN. In order to incorporate this finding into the RL-theory they suggest that whether an outcome is perceived as favorable or unfavorable depends on the subjective value that is attributed to this event rather than the objective value. This view is further supported by other findings that have shown that the FRN is context dependent that is, its amplitude depends on the value of the eliciting outcome relative to the range of possible outcomes rather than the objective value of the outcome (Holroyd, Larsen, & Cohen, 2004). With respect to the present findings such an interpretation would suggest that subjectively participants perceived the neutral feedback as being even worse than the negative feedback. In order reconcile the present results with this interpretation on needs to suggest that being worse in the current experiment does not only mean that participants won or lost a certain amount of money, but is related to the information content that the feedback conveys with respect to the task goal. For positive and negative feedback this means that the information is unambiguous and could be easily used for learning. In contrast, neutral feedback is ambiguous and needs to be processed in relation to the alternative outcome. Hence this feedback is unfavorable with respect to the fact that additional processing is necessary in order to determine its relative value. Thus, these findings suggest that the FRN reflects the activity of an evaluative system that detects whether events are favorable or unfavorable in reference to a certain task goal, not only with respect to the question whether an outcome is good or bad but also whether it is informative in relation to the task goal.

III. General Discussion

7 *Discussion*

The objective of the present thesis was to investigate age-related changes in error processing and learning across the lifespan by using an event-related potential (ERP-) approach. The focus of the three experiments presented in this thesis was on the question how internal and external error information is used for learning, how the role of errors for learning changes during the lifespan and how these changes are reflected in the ERP correlates of reinforcement learning. The first experiment examined the effects of aging on the ERP correlates of internal and external error processing (the error-related negativity, ERN and the feedback-related negativity, FRN) during learning (see also Nieuwenhuis et al., 2002). In this experiment a probabilistic learning task was used in which feedback validity was manipulated. The second experiment focused on developmental differences in the role of error processing for learning by comparing children to younger adults. A similar paradigm as in the first experiment was used, but it was adapted in order to meet the special requirements of children (see Methods Experiment 1). In the third experiment a slightly different approach was adopted. Based on the findings of the first experiment a learning paradigm was developed in order to investigate individual differences in the way participants learn form positive or negative feedback. The idea of this experiment was to test two competing hypotheses on how aging affects the learning from positive and negative feedback (see Frank et al., 2004; Carstensen, 2006).

The following general discussion is structured into three main sections in which the results of the three experiments will be integrated and discussed in the light of the recent literature. The first section will focus on the effects of life-span development on behavioral reinforcement learning, as well as on age differences in learning from positive and negative feedback. In the second section the role of internally generated error information for learning will be discussed. The aim of this section will be to integrate and discuss the present findings on life-span age differences in internal error processing and its ERP correlate, the ERN. One further focus of this section will be on the results of the third experiment, which point to age differences in response-outcome relations and individual differences in the ERN. Furthermore, learning-related changes in the CRN that occurred across the three experiments will be discussed. The third section will be on how the processing of external error information and its ERP correlate, the FRN changes across the lifespan. The purpose of the third section will be to discuss age-related asymmetries in feedback processing across the lifespan. Moreover, it will be discussed how the FRN changes with learning and how it is affected by expectancy violations and different types of feedback.

7.1 Age differences in reinforcement learning across the lifespan

The behavioral findings presented in this thesis point to the view that reinforcement learning is not generally impaired in children and older adults (see Figures 10 and 17). In contrast, the findings from Experiment 1 and 2 indicate that children and older adults are as well able to learn as younger adults if feedback is fully valid and if the learning task they have to perform is adapted to their requirements. Moreover, the analysis of the learning functions in Experiment 1 and 2 showed that across validity conditions children and older adults did not differ from younger adults with respect to the steepness of their learning rates. This stands in contrasts to recent findings, which suggested that older adults might be generally impaired in learning from their errors due to age-related deficits in the activity of the mesencephalic dopamine system (MDS) (Nieuwenhuis et al., 2002). Furthermore, these findings seem to question recent ideas that pointed to developmental changes in the systems underlying learning and error processing (Davies et al., 2004; Ramscar & Gitcho, 2007).

The major difference between the study by Nieuwenhuis et al. (2002) and the present work is that in the former study a fixed response deadline of 700 ms was used, whereas in the present experiments the response deadlines were adaptively adjusted depending on the number of time out trials. The reason for this adaptive procedure was that a similar deadline for all age groups, as in the Nieuwenhuis et al. (2002) study, would have produced a disproportionate time pressure on children and older adults and thereby impaired their ability to learn. Moreover, this time pressure could be expected to result in an increased number of time-outs, which lowers the number of trials in which they can learn from feedback. Yet, on the other hand, one might argue that the adaptive response deadline procedure has masked age differences in learning since it allows children and older adults to respond slower than younger adults[9]. However, it should be noted that slower and more variable reaction times in children and older adults are not a surprising phenomenon that per se points age differences in speed-accuracy trade-offs. In contrast, it has been shown that general slowing is one of the hallmarks of cognitive aging (see Birren & Fisher, 1995; Salthouse, 1996, 2000). Likewise, larger and especially more variable reaction times in children are a typical finding in developmental studies (see Leth-Steensen et al., 2000; Williams et al., 2005). Hence, the idea of the first two experiments was to account for these general age effects by individually adjusting the response deadlines. However, the first two experiments differed from that of Nieuwenhuis et al. (2002) not only with this respect, but also in the kind of feedback provided to the subjects. Nieuwenhuis et al. (2002) used rather ambiguous feedback stimuli (head of a lettuce and a carrot). In contrast, in the present study unambiguous feedback stimuli (German words for ‘correct’, printed in green and ‘incorrect’ printed in red) were used, which are easy to encode and process and might have helped children and older adults in learning.

However, the results of the first two experiments also provided evidence for age-related impairments in overall performance. As can be observed in Figures 10 and 17 children and older adults performed overall worse when invalid information occurred during learning. Moreover, in both age groups these impairments were most pronounced in the first bin of the learning blocks, but absent at the end of the learning blocks. This indicates that children and older adults are particularly sensitive to interference by invalid information at the beginning of learning, possibly because they are impaired in recruiting cognitive control mechanisms in order to protect learning from interference. Such a view would be supported by several findings on age differences in cognitive control, which have been associated with age-related changes in the prefrontal cortex (Bunge et al., 2002; Craik & Bialystok, 2006; West, 1996).

Given these findings on age-related changes in interference control during learning one prediction for the third experiment was that older adults should be particularly impaired in reversal learning, when the previously acquired mappings interfere with learning the new stimulus-response assignments. Interestingly, the opposite pattern was obtained. Younger adults had larger performance impairments when they had to switch from the learning to the reversal phase than older adults, indicating that they suffered more from interference than the elderly. Yet, this is probably due to the fact that in contrast to the first experiment older adults were impaired in learning compared to younger adults. Hence, they performed overall much worse than younger adults at the end of the learning phase and as a consequence perceived less conflict when the stimulus response assignments were reversed. The finding that older adults were impaired in learning in the third experiment is somewhat surprising, given the fact that feedback validity was increased compared to the first experiment and given that the number of stimuli that had to be learned was reduced from six to four stimuli.

The most likely reason for the impairments of older adults in the third experiment is that they had difficulties in disambiguating the neutral outcomes. That is, they were impaired in building up a representation of the correctness of a response when the feedback following this response was ambiguous and had to be processed in relation to the alternative feedback. This is supported by the fact that older positive learners tended to perform below chance level in the positive learning condition, suggesting that they chose more frequently for the response that leads to the neutral outcome, rather than choosing for the response which leads to the (better) positive outcome. Hence, the elderly seem to be less able to differentiate the responses in which the neutral feedback is better in relation to the alternative outcome (-50 in the negative learning condition) from those responses in which the neutral feedback is worse in relation to the alternative feedback (+50 in the positive learning condition). Support for this idea comes from the analysis of the response-locked ERPs, which showed that older adults did not differentiate between responses that were followed by ambiguous positive and negative zero feedback. In contrast, similar to younger adults the elderly clearly differentiated between responses that led to (unambiguous) positive and negative feedback (+50 and -50) (see Figure 33 and discussion of the response-locked ERPs). Taken together, these findings point to the view that older adults show pronounced impairments in learning when feedback is ambiguous and needs to be processed in relation to alternative outcomes. This indicates that older adults might be impaired in building up relational representations of feedback or reward value.

Age differences in learning biases. Apart from these general age-related changes in reinforcement learning the main objective of the third experiment was to investigate whether older adults differ from younger adults in the way they learn from positive and negative feedback. Based on the Go-NoGo model by Frank et al. (2004) and the neurocomputational account by Nieuwenhuis et al. (2002) one would expect that due to their deficits in the activity of the MDS older adults should tend to learn more form negative compared to positive feedback and hence could be termed negative learners. In contrast, according to the socioemotional selectivity theory (SST) by Carstensen and colleagues (Carstensen, 2006; Mather & Carstensen, 2004) older adults perceive their lifetime as more constrained than younger adults and hence are more engaged in emotion regulation and try to implement cognitive control mechanisms in order to enhance positive and diminish negative information. Thus, according to this model older adults could be expected to have a tendency towards learning more from positive rather than negative feedback and hence could be termed positive learners.

Overall the data of the third experiment revealed that a larger number of younger adults were positive learners, whereas more older adults were negative learners. At first glance this seems to support the Go-NoGo model (Frank et al., 2004) rather than SST (Carstensen, 2006), since it suggests that the majority of older adults are negative learners. However, it should be noted that older adults performed overall much worse than younger adults in the third experiment and due to this they may have adopted a more conservative response strategy than their younger counterparts. Being more conservative here means that they had a tendency towards deciding for those responses that lead to neutral feedback, irrespectively of whether the neutral feedback was better or worse with respect to its alternative. This view is supported by the fact that older negative learners performed below chance in the positive learning condition, especially at the beginning of learning (see Table 8). This indicates that at the beginning of learning they indeed decided more frequently for those responses that led to neutral rather than for those that led to positive feedback. Further support for this idea comes from the fact that in the reversal blocks younger and older adults performed better in the negative compared to the positive learning condition. This indicates that when they are uncertain about the stimulus-response mappings both age groups tend to be more conservative and decide for the responses that lead neutral feedback even if the alternative choice would have been better.

On the other hand, it could be argued that adopting a conservative response strategy in the present task means nothing else than trying to avoid the negative feedback even at the expense of performing much worse in the positive learning condition. From such a viewpoint it might be justified to assume that older adults tend to learn more from negative feedback than younger adults. However, rather than being driven by dopamine this response bias seems to be related to the fact that older adults are generally more cautious than younger adults. This view is supported by the fact that differences in the learning biases between younger and older adults were observed. In younger adults similar learning effects were obtained for positive as well as for negative learners. In contrast, older positive learners showed the expected pattern, whereas older negative learners showed small learning effects in both learning conditions. Yet, when considering the overall performance levels in older negative learners it appears that they perform below chance level in the positive learning condition in the first block half and end up at chance level in the second block half (see also Table 8). These findings again suggest that older negative learners adopted a more conservative response strategy, especially at the beginning of learning. Hence, the present behavioral findings do not point to the view that older adults have a positivity bias, that is, they do not seem to have a greater tendency towards learning more from positive compared to negative feedback. Yet, although there is evidence that the elderly are more engaged in avoiding negative feedback rather than in choosing for positive feedback, it seems questionable whether this response strategy is indeed associated with age-related changes in phasic dopaminergic activity. In contrast, it seems more plausible that this response bias is due to the fact the elderly generally tend to be more cautious.

7.2 Processing of internal error information

*Life-span developmental changes in the ER*N. One of the most important new findings of the present thesis is the fact that across three experiments the ERN turned out to be performance- rather than age-sensitive. In the first two experiments, in which performance levels were equated between age groups, neither for children nor for older adults the ERN was reduced compared to younger adults (see Figure 11 and 18). In contrast, in the third experiment, in which performance was severely impaired in older adults a reduction of the ERN was found in the elderly (see Figure 31). The findings in older adults are inconsistent with the predictions of a recent neurocomputational account (Nieuwenhuis et al., 2002), as well as with the results of several other studies on age differences in error processing (Band & Kok, 2000; Falkenstein et al., 2001; Mathewson et al., 2005; West, 2004). However, it must be noted that in most of these studies, including the Nieuwenhuis et al. (2002) study, age differences in the ERN were paralleled by age differences in accuracy. This is a crucial issue since when considering the basic ideas of the mismatch model (Bernstein et al., 1995; Scheffers & Coles, 2000) and the R-L theory (Holroyd & Coles, 2002), it appears that these models assume that the ERN depends on an intact internal representation of the correct response. This is because in these models the ERN is assumed to reflect the mismatch between the incorrect response and the representation of the correct response. Since in the present experiments the representation of the correct response has to be built up through learning, any impairments in learning would result in a weakened mismatch signal and hence in a reduced ERN. Hence, the worse participants perform the less they are able to represent the correct response and the smaller the mismatch signal, that is, the ERN. Thus, the present data indicates that it is not age per se, but differences in performance level (in the expectation on the correctness of the response) that drive the ERN (for a similar finding see Pietschmann, Endrass, & Kathmann, 2007).

If we have a look at the developmental data on error processing and the ERN a similar picture emerges. In most of the developmental studies the ERN was found to be reduced for children compared to younger adults (Davies et al., 2004; Ladouceur et al., 2004; Santesso et al., 2006; Wiersema et al., 2007). However, similar to the findings in older adults, in all of these studies reduced ERNs in children were paralleled by performance impairments when compared to adults. This supports the idea that in these studies developmental differences in the ERN might have been confounded with performance differences between age groups. Further evidence for this view comes from a recent study on developmental differences in error processing between adolescents and younger adults (Hogan et al., 2005). In this study age differences in the ERN were examined by manipulating task complexity. The authors found a reduced ERN for adolescents compared to adults only in the more complex task version in which performance was also impaired. In contrast, in the less complex task version, no significant age differences in the ERN or performance were obtained (for similar findings see Kim et al. 2007).

To summarize, there is ample evidence for the view that the ERN is performance- rather than age-sensitive. This questions the idea that children and older adults are generally impaired in error processing and points to the importance of equating accuracy levels between age groups in order to avoid confounding effects of performance differences when studying age differences in the ERN or any other ERP component.

Learning-related changes in the ERN. One of the central predictions of the R-L theory (Holroyd & Coles, 2002) is that the ERN should increase with learning. This is because the theory proposes that the ERN reflects a negative prediction that gets larger the more participants learn (see Holroyd & Coles, 2002, Holroyd et al., 2003; Nieuwenhuis et al., 2002). However, it must be noted that in these studies a difference wave approach was adopted to investigate learning-related changes in the ERN. Such an approach is ignorant with respect to how much either of the response types contributes to the observed effects.

In contrast, in the present thesis the response-locked data was analyzed separately for correct and incorrect responses. This was done since an inspection of the response-locked averages showed that learning-related changes were not restricted to the ERN but also occurred in a positivity to correct trials (see Figures 12, 19, and 32). The first experiment revealed that consistent with the R-L theory (Holroyd & Coles, 2002), the ERN increased with feedback validity, suggesting that the more participants are able to build up expectations on the correctness of their response the larger the ERN (see Figure 11). Moreover, the data in the first experiment revealed that when feedback was fully valid in the 100% validity condition, and the ERN was captured most precisely using peak-to-peak amplitude measures, there is also evidence that the ERN increases with learning (see Figure 12). However, no learning-related changes in the ERN were obtained in the 80% validity condition in which feedback was partially invalid and participants were impaired in overall performance. Similar to these results, in the third experiment it was found that the ERN increases with learning, however, only in the learning phase, but not in the reversal phase, in which participants performed overall worse. In contrast, the learning effects themselves were found to be similar for the two age groups across the validity conditions in the first experiment, as well as across the block phases in the second experiment.

Together, these findings indicate that whether or not learning-related changes are found in the ERN depends on the overall accuracy level rather than the learning effects themselves. Such a view implicates that the ERN does not reflect learning per se, but a performance monitoring process that comes into play when participants have already achieved a certain performance level and try to maintain it. (see Ridderinkhof et al., 2004; Rushworth, Walton, Kennerly, & Bannerman, 2004) This view is further supported by an investigation of the time course of the learning-related effects in the ERN in the first experiment. This analysis revealed that the differentiation between the 100% and the other validity conditions in the ERN occurs from the first to second bin of the learning blocks (see Figure 13). In contrast, the accuracy data showed a differentiation between the validity conditions already in the first bin. Hence, the increase of the ERN occurs later in time than the increase in accuracy, supporting the view that the ERN might reflect a performance monitoring process that signals the need for performance adjustments in order to maintain a certain performance level.

However, the response-locked data from Experiment 1 and 2 showed that not only the ERN, but also the response-locked positivity to correct trials increased with feedback validity (see Figures 11 and 18). Moreover, learning-related changes were much more pronounced in the response-locked positivity to correct trials compared to the ERN (see Figures 12 and 19). In contrast to the ERN, which increased with learning only if feedback was fully valid, the response-locked positivity showed a learning-related increase for the 80% condition as well. Hence, the learning-related effects in the response-locked positivity are much more consistent with the learning effects in the accuracy data. This is further supported by the fact that in the second experiment for children, as well as for younger adults the amplitude of the response-locked positivity correlated positively with overall accuracy. That is, better performance with learning is associated with a larger amplitude of the response-locked positivity (see Figure 21). Moreover, in the third experiment it was found that in contrast to the ERN, which only increased in the learning phase, the response-locked positivity increased in the learning, as well as in the reversal phase (see Figure 32). Again this is much more in line with the learning effects in the accuracy data, which did not differ between the learning and the reversal phase in the third experiment.

Thus, these findings show that learning-related changes are much more pronounced in the response-locked positivity on correct trials compared to the ERN on incorrect trials.

The learning-related increase in the response-locked positivity is nicely consistent with neurophysiological findings on reinforcement learning in monkeys (Mirenowicz & Schultz, 1996; Schultz et al., 1997). These findings show that at the beginning of learning phasic increases in the activity of the dopamine neurons are found for the reward. With learning this positive prediction error propagates back in time and is then elicited by the conditioned stimulus. Thus, the monkey is now able to predict the reward. In line with these findings, it seems reasonable to assume that the increase of the response-locked positivity with learning reflects the increasing ability of participants to predict reward (positive feedback) based on the knowledge they acquired through learning. Hence, the current data suggest that learning-related changes are reflected in the response-locked positivity to correct trials, as well as the ERN to incorrect trials. That is, the present findings are consistent with the idea that learning is driven by both, positive prediction errors when the outcome of an action is better than expected and negative prediction errors when the outcome of an action is worse than expected (see O’Doherty et al., 2004; Schultz, 2002; Seymour et al., 2004).

At first glance this interpretation seems inconsistent with the original version of the R-L theory (Holroyd & Coles, 2002), which focuses on the role of negative predictions errors and the ERN for learning. In order to integrate these findings with the R-L theory, one needs to suggest that a positive prediction error as reflected in phasic increases of mesencephalic dopaminergic activity inhibits the ACC, and by this leads to the generation of the response-locked positivity (see Holroyd et al., 2003; Holroyd, 2004). Similar to the original version of the R-L theory such a view presupposes several assumptions. At first, it suggests that a positive prediction error leads to an inhibition of ACC activity. This is probably difficult to show in humans, however, neurophysiological data from monkeys suggest that the ACC not only plays a role in error processing, but also in reward-based motor selection (Matsumoto et al., 2003; Shima & Tanji, 1998) and reward expectancy

(Shidara & Richmond, 2002). Thus, there is evidence that the ACC is implicated in reward processing, however, its exact role during reinforcement learning and the nature of its modulation by dopaminergic input remains to be established. A second assumption would be that the response-locked positivity should be generated in the ACC. Given the poor spatial resolution of the EEG data this question cannot be answered conclusively in the present study. However, the topographical maps in Figures 11 and 18 show that the reward-related variance in the difference wave is distributed fronto-centrally, which is generally in line with an involvement of the ACC in the generation of the ERN, as well as the response-locked positivity. Taken together, there is some evidence that the learning-related changes in the response-locked positivity might reflect modulations of ACC activity.

However, there is also an alternative explanation for this positivity that needs to be explored. According to this alternative account the component reflects the response-locked part of the P300 to the stimulus. Such a view would suggest that the learning-related effects in the response-locked positivity might reflect increasing decision confidence with learning (see Finnigan, Humphreys, Dennis, & Geffen, 2002; Cutmore & Muckert, 1998). This idea receives support by the fact that the positivity increases from frontal to parietal electrodes[17] as would be expected for the P300. However, this view would also suggest that the stimulus-evoked P300 should show similar learning-related changes as the response-locked positivity. That is, the P300 should increase with decision confidence in the two learning conditions but not in the 50% validity condition. Yet, an

---

[17] In the first experiment the response-locked positivity increased from anterior to posterior in the 100% validity condition for younger adults (Fz: M = 3.44 µV, SD = 5.44; Pz: M = 9.62 µV, SD = 5.31) and for older adults (Fz: M = 2.52 µV, SD = 3.34; Pz: M = 4.34 µV, SD = 3.79). This is in line with the view that the response-locked positivity reflects stimulus-evoked P300 activity. An inspection of the stimulus-locked averages at the electrode Pz showed that the stimulus-evoked P3 increased with learning for the 100% validity condition (p < .0002), (younger adults: Bin 1: M = 7.15 µV, SD = 4.58, Bin 4: M = 9.36 µV, SD = 5.12; older adults: Bin 1: M = 3.09 µV, SD = 4.02, Bin 4: M = 4.89 µV, SD = 3.9). However, the stimulus-evoked P300 increased for the 50% validity condition as well (p < .001), (younger adults: Bin 1: M = 5.71 µV, SD = 4.49, Bin 4: M = 7.09 µV, SD = 4.83; older adults: Bin 1: M = 2.64 µV, SD = 3.17, Bin 4: M = 3.48 µV, SD = 3.87). This result does not support the view that the response-locked positivity reflects stimulus-evoked P300 activity

analysis of the stimulus-locked averages[17] showed that in contrast to this prediction the P300 increased with time on task for all validity conditions. This does not support the P300 account since there is no reason why decision confidence should increase in the 50% condition.

To summarize, the present data points to the view that the response-locked positivity reflects response-related activity that is potentially driven by the ACC and reflects the increasing ability of participants to predict reward based on the information that has been acquired through learning. These findings provide an important extension to recent theoretical accounts (Holroyd & Coles, 2002; Nieuwenhuis et al., 2002) by showing that reward-related variance in the response-locked ERPs is driven by positive learning signals on correct trials rather than negative learning signals on incorrect trials.

Response-outcome relations in the third experiment. Most interestingly, the third experiment showed that older adults are impaired in representing the correctness of their responses if these responses lead to ambiguous neutral outcomes. As can be seen in Figure 33 younger adults clearly differentiated between correct and incorrect responses that led to positive (+50) and negative (-50) feedback, as well as between correct and incorrect responses that led positive and negative zero (*00) feedback. This suggests that younger adults are able to represent that in the negative learning condition the neutral feedback is better than the alternative negative feedback, whereas in the positive learning condition the neutral feedback is worse than the alternative positive feedback. That is, younger adults are able to build up a representation of the correctness of a response even if these responses are associated with ambiguous feedback that has to be processed in relation to the alternative feedback.

Older adults showed a differentiation between the ERPs to correct and incorrect responses only if these responses led to unambiguous (positive or negative) feedback. However, they did not differentiate at all between correct and incorrect responses that led to ambiguous neutral feedback. This finding shows that older adults are impaired in building up a representation of the correctness of a response when the feedback is ambiguous and has to be processed in relation to the alternative feedback. This impairment leads to pronounced performance decrements in the elderly, since learning in the third experiment critically depends on the ability to represent that in negative learning condition the neutral feedback is better than the alternative feedback, whereas in the positive learning condition it is worse than the alternative feedback.

There is considerable evidence for the view that the orbitofrontal cortex (OFC) plays a major role in building up relative representations of reward value (Schultz, Tremblay, & Hollermann, 2000, for reviews see Rolls, 2000; Wallis, 2007). Moreover, it has been shown that patients with lesions in the OFC are impaired in representing relations between actual and possible outcomes (Camille et al., 2007). This seems to be consistent with the pattern of results obtained in older adults in the third experiment, suggesting that the impairments of the elderly might be due to deficits in the function of the orbitofrontal cortex (see Resnick, Driscoll, & Lamar, in press).

Learning biases in the ERN. In line with data by Frank and colleagues (2005) the third experiment revealed a larger ERN for younger negative compared to younger positive learners, supporting the view that the ERN might be associated with individual differences in learning biases in younger adults. According to the Go-NoGo model (Frank, 2005), such a pattern of results would be due to reduced dopamine levels in negative learners, which result in a larger ERN and a bias towards learning more from negative feedback. However, similar to the learning-related changes the larger ERN for negative learners was only found for learning, but not for the reversal phase. One potential explanation for the absence of this effect in the reversal phase could be that after the reversal participants were less certain about the correctness of their response and hence perceived less mismatch and showed a smaller ERN. What appears from this pattern of results is that the error processing system, as reflected in the ERN, might be not directly implicated in learning the stimulus-response assignments, but in signaling the need for performance adjustments in order to maintain an already high performance level. This view would be in line with the findings from the first experiment that suggest that in contrast to the learning-related changes in the response-locked positivity the learning-related changes in the ERN occurred later in time. This suggests that the former component might reflect the initial learning of the stimulus-response assignments, whereas the latter is involved in signaling the need for performance adjustments if a mismatch between the representation of the correct response and the actual response occurs. According to such a view individual differences in learning biases might be most pronounced in those situations in which participants already have distinct representation of the correct response.

In contrast to younger adults, no significant learning biases in the ERN were obtained for older adults. Moreover, in the third experiment the ERN was found to be generally reduced for older compared to younger adults. This is inconsistent with the Go-NoGo model (Frank, 2005), which would suggest that especially older negative learners should show large ERN amplitudes, since they have reduced dopamine levels (Erixon-Lindroth et al., 2005; Wang et al., 1998; Volkow et al., 1996; Volkow et al., 1998). One interpretation for these results would be to assume that the absence of learning biases in the ERN in older adults might be due to the fact that they were impaired in representing the correctness of the response and showed generally reduced ERN amplitudes. The view that older adults had problems in differentiating correct and incorrect responses is supported by the results of the response-outcome analysis, which showed that older adults had representation of the correctness of the response only for those responses that led to unambiguous outcomes.

Taken together, the data of the third experiment supports the view that in younger adults the ERN is related to individual differences in learning biases (see Frank et al., 2005). Thus, the present findings point to the view that the larger ERN in negative learners reflects the fact that these participants are more engaged in performance monitoring and hence show more ACC activity than positive learners. In contrast, no such effects were obtained in older adults, suggesting that whether or not learning biases in the

ERN are obtained depends on how distinct a participants' representation of the correct response is.

Response conflict and learning biases in the CRN. The response-locked data in the first experiment showed an increased CRN component in the condition in which feedback was delivered randomly compared to the condition in which feedback was fully valid and could be used for learning (see Figure 11). This is consistent with results from the third experiment that showed that the CRN decreased with learning, suggesting that the CRN is enhanced when participants are uncertain about the outcome of their response. Moreover, the third experiment showed that that the CRN was larger for the reversal compared to the learning phase (see Figure 31), indicating that the component is related to conflict on the level of stimulus-response mappings. There is an ongoing debate on the question of what the CRN reflects (see Coles et al., 2001; Vidal et al., 2000, 2003). Yet, recent data suggests that the CRN might be generated in a similar network (involving the ACC) as the ERN and is related to post-response conflict processing (Bartholow et al., 2005; Eppinger et al., 2007; Kray et al., 2005). The data presented in this thesis further support this view and suggests that the CRN is related to response conflict that is elicited when participants are uncertain about the correctness of their response.

Interestingly, the results of the third experiment showed that the CRN was not only larger under conditions in which response conflict was enhanced, but was also increased for younger negative compared to younger positive learners. This finding stands in contrast to the results of the study by Frank and colleagues (2005), in which no learning biases in the CRN were found. This discrepancy in the results of the two studies may be related to the differences in the tasks that were applied. In contrast to the reinforcement learning paradigm used by Frank and colleagues (2005), in which participants had to decide between two stimulus pairs, in the present study only one stimulus was presented and participants had to decide for one of two responses. Moreover, in contrast to the Frank et al. (2005) study, in which three stimulus pairs were presented in each learning block in the present study six stimuli were presented per learning block. Together, these factors might have enhanced response conflict and hence increased the probability of detecting learning biases in the CRN.

To summarize, the present data suggests that learning biases are not only reflected in the ERN, but also in the CRN. This indicates that negative learners are not only more sensitive to errors, but also respond more strongly to conflict during learning. This points to the view, that they are generally more engaged in performance monitoring and show enhanced ACC activity when they are uncertain about the correctness of their response. In line with previous studies (Bartholow et al., 2005; Eppinger et al., 2007; Kray et al., 2005) the data of the first as well as third experiment point to the view that the CRN reflects response conflict processing and is enhanced if participants are uncertain about the outcome of their response.

7.3 Processing of external error information

Asymmetries in feedback processing across the lifespan. Experiments 1 and 2 showed an asymmetry in the processing of feedback information across the lifespan. Children had a larger FRN compared to younger adults, whereas both age groups showed a similar amplitude of the feedback-locked positivity (see Figure 22). In contrast, in older adults the FRN was strongly reduced compared to younger adults (see Figure 14). The findings in children suggest that they are more sensitive to negative feedback during learning than adults, whereas both age groups seem to be similarly affected by positive feedback. Since this experiment is one of the first developmental ERP-studies on feedback processing the existing literature is limited. However, there is some evidence from developmental studies on feedback processing using heart-rate measures that point to a similar asymmetry in feedback processing between children and adults (Crone et al., 2004; Somsen et al., 2000). A typical finding in studies using heart-rate measures is that in younger adults heart rate is slowed for negative compared to positive feedback when the feedback contains information that could be used for learning or for the adaptation of performance (Somsen et al., 2000). In contrast to younger adults, 8 - 10 year-old children showed heart rate slowing for informative as well as uninformative negative feedback

(Crone et al., 2004). This indicates that younger children may be more sensitive to negative feedback and less able to use external error information for learning. Moreover, older children (12-14 year-old), but not younger children (8-10 year-old) showed heart rate slowing following performance errors (Crone, Somsen, Zanolie, & van der Molen, 2006). Thus, younger children seem to be less able to internally represent incorrect responses. Taken together, these findings are consistent with the results of the second experiment suggesting that on the one hand children are more sensitive to negative feedback during learning, whereas on the other hand they seem to be impaired in representing the correctness of a response, especially when invalid information interferes with learning.

In contrast to children, older adults showed a strongly reduced FRN compared to younger adults. Hence, the present findings suggest that although older adults learned comparably to younger adults (in the first experiment) they showed reduced activity of the structures involved in the processing of negative feedback (presumably the ACC, but also the orbitofrontal cortex; see O'Doherty et al., 2001; Rolls, 2000). This result is somewhat surprising given the absence of age differences in the ERN in the first experiment and may point to a functional dissociation of both components (see Nieuwenhuis et al., 2005). Similar findings on an asymmetry in the processing of valence information in older adults have been obtained in research on episodic memory and decision making (Charles, Mather, & Carstensen, 2003; Mather & Johnson, 2000). These findings have been interpreted within the framework of the socio-emotional selectivity theory of aging, which proposes that the ratio between positive and negative affect improves through adulthood and leads to what is called a "positivity effect" (see Carstensen, 2006; Mather & Carstensen, 2005). The idea is that older adults focus more on emotion regulation and implement cognitive control mechanisms that enhance positive and diminish negative information. Interestingly, recent fMRI-findings from Larkin and colleagues (2007) using a gain and loss anticipation task support this view and suggest that older adults are less affected by potential losses than younger adults, whereas both age groups are equally excited by potential gains. Given the absence of age differences in the ERN in the first experiment, which suggests that error processing is not impaired per se in the elderly, the present data points to an age-related asymmetry in feedback processing. This is in line with the idea of a positivity effect in older adults (Carstensen, 2006; Mather & Carstensen, 2005).

However, the data from the third experiment does not entirely support this view. In fact, the behavioral findings point to the idea that if at all, older adults seem to be negative learners rather than positive learners, which is inconsistent with the proposed positivity effect in older adults. Yet, as shown above, this might be a consequence of the fact that older adults generally performed worse than younger adults in this experiment and hence tended to respond more conservatively. In line with this idea, the feedback-locked data of the third experiment showed that in contrast to younger adults in older adults the FRN was not sensitive to the type of feedback participants obtained. That is, they did not differentiate between positive, negative and neutral outcomes in the FRN. This seems to be consistent with the findings from the first experiment. Thus, the results of the third experiment are generally in line with the idea of an age-related asymmetry in feedback processing however, whether this asymmetry indeed reflects a positivity bias in older adults remains to be established in future studies. One point that needs to be considered in these studies is to avoid performance differences between age groups, as it has been done in the first two experiments of this thesis. One way to do so, by maintaining the general logic of the learning task applied in the third experiment, would be to disambiguate the neutral feedback by using color information in order to make it easier for older adults to assess the valence of the neutral feedback.

Taken together, the present data speaks for an asymmetry in feedback processing across the lifespan. Children seem to be more sensitive to negative feedback during learning than younger adults, whereas both age groups seem to be similarly affected by positive feedback. The opposite pattern was obtained for older adults who showed a strongly reduced FRN compared to younger adults, suggesting that they are less affected by negative feedback during learning.

A FRN to neutral feedback? One of the most surprising results was obtained in the analysis of the FRN in the third experiment. This analysis showed that in younger adults the largest FRN was obtained for the neutral feedback whereas the FRN was smaller for negative and smallest for positive feedback (see Figure 36). In contrast, for older adults no effect of feedback type was obtained. The fact that in younger adults the FRN was larger for neutral compared to negative feedback is inconsistent with the R-L theory (Holroyd & Coles, 2002), which would predict that the FRN to neutral feedback should be midway between the FRNs to negative and positive feedback (see Holroyd et al. 2006). Moreover, this finding stands in contrast to most of the other functional interpretations of the FRN, which suggest that the component is associated with the processing of negative feedback information (see Hajcak, Moser, Holroyd, & Simons, 2007; Miltner et al., 1996; Nieuwenhuis et al., 2004; Yeung & Sanfey, 2004).

However, a recent study by Holroyd et al. (2006) revealed similar results. In this study different paradigms, involving a time estimation task and several versions of trial and error learning tasks were used to investigate the effects of neutral feedback on the FRN. Across five experiments, the authors obtained similar FRNs to neutral and negative feedback, indicating that both types of feedback were processed similarly. Holroyd and colleagues (2006) concluded that neutral and negative feedback might be grouped together as events that indicate that task goals have not been satisfied and hence both lead to the generation of the FRN. However, in the present study the FRN was found to be even larger for neutral compared to negative feedback, which indicates that participants perceived the neutral feedback as even worse than the negative feedback. This makes sense if we assume that being worse in the present task might not only be related to the valence of the feedback, but to the information content that is conveyed by the feedback stimulus. This means that in contrast to positive and negative feedback the neutral feedback contains no direct information that could be used for learning. Rather, it needs to be processed in relation to the alternative feedback. Hence, neutral feedback could be conceived of as being even worse than negative feedback since it does not provide direct information for learning and needs to be processed in relation to the alternative feedback. Thus, in line with the suggestion by Holroyd et al. (2006) the present findings point to the view that FRN reflects the activity of an evaluative system that detects whether events are favorable or unfavorable in reference to a certain task goal. Yet, going beyond this, the present findings also show that depending on the task context neutral feedback could be perceived as being even worse than negative feedback. This is the case if the feedback does not provide direct information for learning and additional resources have to be recruited in order to build up a relational representation of feedback value.

The FRN and expectancy violations. According to the findings by Holroyd and Coles (2002) and Nieuwenhuis et al., (2002) it was expected that in the first two experiments the FRN should be affected by feedback validity, that is, it should be larger the more invalid the feedback. This is because the more participants expect to receive positive feedback the larger the mismatch (or prediction error) if they receive negative feedback (see Holroyd & Coles, 2002; Holroyd et al., 2003; Nieuwenhuis et al., 2002). However, in contrast to these predictions the FRN was not modulated by feedback validity (see Figures 14 and 22). Instead, the analysis revealed that a feedback-locked positivity to positive feedback increased the more invalid the feedback. This was further supported by a comparison of the 20% invalid trials with the 80% valid trials of the 80% validity condition in the first experiment [10]. For this comparison the R-L theory (Holroyd & Coles, 2002; Holroyd et al., 2004) would suggest that the FRN should be much larger for the 20% invalid trials since in that condition a strong expectation for positive feedback is violated. Indeed, the analysis showed that there is a marginally significant difference between the FRN for valid compared to invalid negative feedback (for similar findings see Holroyd & Coles, 2002; Nieuwenhuis et al., 2002). However, the effect was much more pronounced in the positivity for positive feedback, which was significantly larger for the 20% invalid trials compared to the 80% valid trials. Hence, the current data show that expectancy violations seem to primarily affect the processing of positive feedback rather than negative feedback[18]. This result is nicely consistent with recent findings by Potts and colleagues (2006) who used a passive reward prediction task and showed that a similar feedback-related positivity, the P2a, is elicited by unpredicted rewards. Potts and colleagues (2006) proposed that this positivity is generated by dopaminergic input to the medial frontal cortex if reward expectancy is violated. This interpretation is supported by the medial frontal topography of the difference wave in younger adults in the first experiment (see Figure 14). Moreover, the medial frontal topography of the difference wave is the more pronounced the more invalid the feedback, which indicates that the reward-related variance is fronto-centrally distributed (see also Figure 14). Thus, these findings suggest that the feedback-locked positivity to positive feedback trials rather than the FRN reflects errors in reward prediction. Moreover, the topography of the difference wave seems to be consistent with the idea that the component is generated by dopaminergic input to the medial frontal cortex.

However, as for the response-locked positivity, an alternative account to these effects would suggest that the feedback positivity reflects a P300-like modulation. According to such a view one would suggest that the increase of the feedback positivity the more invalid the feedback might reflect the amount of information that is extracted from the feedback stimulus (Donchin & Coles, 1988; Johnson, 1986). The argumentation would then be that the more participants are able to internally represent the correct response, the less they have to rely on the feedback and the smaller the P300. Moreover, in favor of the P300 account one might further suggest that since the P300 is sensitive to the stimulus probability (Donchin & Coles, 1988), learning-related effects in this study might have been obscured by probability effects. However, this has been recently addressed in a study by Gibson et al., (2006), who showed that the reward-related variance in the ERP difference wave for positive and negative feedback is fronto-centrally distributed across probabilities. This does not support the P300 account and favors the

---

[18] Again, it must be noted that similar effects might have been obtained in the studies by Holroyd & Coles, (2002) and Nieuwenhuis et al. (2002) if in these studies feedback-locked ERPs would have been analyzed separately for positive and negative feedback, rather than using a difference wave approach to study the FRN.

idea that the feedback-locked positivity reflects a reward-related modulation of activity in the medial frontal cortex (see Potts et al., 2006).

Moreover, the third experiment showed that the feedback-locked positivity not only varied with feedback validity, but was also affected by feedback valence. In line with previous ERP studies it was found that the feedback-locked positivity was larger for negative compared to positive feedback and for both types of feedback compared to neutral feedback. Similar patterns of results have been obtained in ERP studies on valence processing (Ito, et al., 1998; Kisley et al., 2007; Wood & Kisley, 2005). In these studies it was observed that participants are more sensitive to negative compared to positive stimuli and for both types of valence information compared to neutral stimuli. This negativity effect is reflected in a late positive potential (LPP). However, in contrast to recent ERP studies on the negativity effect in older adults (Kisley, et al., 2007; Wood & Kisley, 2006) the present data does not suggest that the negativity effect is reduced in older compared to younger adults. In contrast, the present data suggest that older adults are less sensitive to both, positive and negative outcomes in relation to neutral outcomes. This less pronounced differentiation between neutral and positive and negative outcomes might have been one reason for the problems of older adults in disambiguating the neutral outcomes. However, when arguing against a reduction of the negativity effect in older adults it needs to be considered that there are several differences between the aforementioned studies and the present experiment. First, the LPP, which is typically used as a measure of valence processing in these studies, occurs much later (300 – 900) than the feedback-locked positivity (200 – 300 ms). Second, the stimuli that were used in these studies (pictures from the International Affective Picture System, IAPS, Lang, Bradley and Cuthbert, 1998) as well as the task (categorization tasks) differ considerably from the feedback stimuli and the learning task used in the present experiment. This points to the view that it could be misleading to compare the valence effects in the feedback-locked positivity with those typically obtained in the LPP. However, as could be observed in Figure 36 the valence effects in the feedback-locked ERPs occurred not only in the time window of the feedback-locked positivity, but also later, in the time window of the LPP. Hence it seems reasonable to assume that the present findings reflect an overlap between the feedback-locked positivity and a later positive component similar to the LPP.

The FRN does not vary with learning. Across the three experiments there is no evidence for a reduction of the FRN with learning. In contrast, similar to the response-locked ERPs a feedback-locked positivity to positive feedback was observed that decreased the more participants learned (see Figures 15, 23, 37). This finding indicates that in contrast to the predictions of the R-L theory (Holroyd & Coles, 2002) learning-related changes are observed in the ERPs to positive, but not to negative feedback. However, it should be kept in mind that in the previous studies on the effects of learning on the FRN a difference wave approach was adopted (Holroyd & Coles, 2002; Nieuwenhuis et al., 2002). Hence, it may well be that learning-related changes on positive feedback trials have been overlooked in these studies. A comparison of the ERP learning functions and the behavioral learning functions (see Figures 10 and 16) in Experiment 1 that the more participants learn the smaller the feedback-locked positivity, indicating that they are more and more able to disengage from processing positive feedback during learning. In line with the findings of Potts et al., (2006), one interpretation of the learning effects in the feedback-locked positivity would be to assume that the component reflects a positive prediction error that decreases the more participants are able to internally represent the correctness of the response (see Holroyd et al., 2003; Holroyd, 2004). Recent data from Cohen, Elgar and Ranganath (2007) strongly support this view by showing that as reward expectation increases (and the positive prediction error decreases) during learning the feedback-locked positivity also decreases. Consistent with these ideas, the current findings point to the view that with learning participants rely less on the external feedback since they are increasingly able to internally predict the reward. Interestingly, children as well as older adults showed a less pronounced decrease of the feedback-locked positivity with learning. This indicates that they were less able to disengage from processing positive feedback during learning, probably because it takes them longer than younger adults to be able to internally predict the feedback.

Limitations of the present thesis. As most of the ERP and neuroimaging studies the experiments presented in this thesis suffer from limited sample sizes. Although the samples in the present thesis are large compared to other studies (Holroyd & Coles, 2002; Nieuwenhuis et al., 2002) larger samples would be desirable for several reasons. First, they would allow investigating interindividual differences in learning and error processing. How fruitful such an approach could be can be observed in the third experiment. Second, they would increase statistical power and hence the ability to detect smaller effects, which is important, especially with respect to the lower signal-to-noise ratio in children and older adults. Another important aspect that needs to be accomplished in future studies is to track developmental changes in learning and error processing during childhood and ageing, either by adopting a longitudinal approach or by including several groups of children and older adults with narrower age ranges.

Another problem that is inherent in the ERP approach adopted in the present study, is the fact that in order to obtain reliable ERN amplitudes a sufficient number of trials (typically more than 15 trials) have to be averaged. Yet, as could be observed in the learning functions in the first two experiments a great deal of learning takes place in the first two bins. Hence, it cannot be precluded that age differences, especially at the beginning of learning, might have been overlooked in the present experiments. Moreover, given the limited spatial resolution of the EEG data it would be highly desirable to apply functional imaging techniques (fMRI) in order to investigate which regions are implicated in learning. This would help to resolve the question whether the learning-related changes in the response-locked and the feedback-locked positivity are indeed associated with ACC activity.

One further limitation of the present thesis is that the theoretical model that underlies the present work assumes that age differences in learning and error processing are the result of age-related changes in function of the MDS. Although there is considerable evidence for this view, dopamine levels have not been directly manipulated in the present thesis. Hence, one idea for further research would either be to manipulate dopamine levels in older adults, or to separate groups of older adults based on their dopamine levels (e.g. by using genetic approaches) and to test for individual differences in learning and error processing.

8 Conclusion

The data presented in this thesis shows that reinforcement learning is not generally impaired in children and older adults. Rather, the present data points to the view that age-related impairments occur when feedback is invalid and interferes with learning, or when feedback is ambiguous and needs to be processed in relation to alternative feedback. This suggests that age differences in learning and error processing occur if participants have to recruit control mechanisms in order to resolve interference or to build up a relational representation of feedback value. Likewise, the results of the present experiments do not support the view that the ERN is reduced in children and older adults. In contrast, the present findings show that the ERN is performance- rather than age-sensitive. That is, when performance levels are equated between age groups, as in Experiment 1 and 2 similar ERN amplitudes are obtained in children, younger and older adults. In contrast, when older adults perform worse than younger adults, as in the third experiment there is also evidence for a reduction of the ERN. The implications of these findings are twofold: First, they point to the importance of equating performance levels when comparing the ERN (or any other ERP component) between age groups. Second, they question the idea that age differences in the ERN might be associated with age-related impairments in the function of the MDS (see Nieuwenhuis et al., 2002). This is not to say that dopamine does not play a role in learning or lifespan development. However, the present data points to the view that age differences are most pronounced when control processes have to be engaged during learning. This view is further supported by the findings from the third experiment, which suggest that older adults are impaired in learning and in differentiating correct from incorrect responses when the feedback is ambiguous and needs to be processed in relation the alternative outcome.

The feedback-locked ERP data presented in this thesis shows an age-related asymmetry in feedback processing across the lifespan. Children showed a larger FRN than younger adults, indicating that they are more sensitive to negative feedback. In contrast in older adults the FRN was strongly reduced, suggesting that they focus less on negative feedback during learning. Hence, in line with the socioemotional selectivity theory it could be suggested that older adults tend to learn more from positive compared to negative feedback. However, the results of the third experiment did not entirely support this view. In contrast, these findings suggest that if older adults are impaired in performance they tend to adopt a more conservative response strategy, which results in a bias towards learning more from negative feedback.

Apart from the age-related changes in error processing and learning, one important new finding of the present thesis is that across three experiments learning-related changes were much more pronounced on correct responses and positive feedback trials, compared to incorrect responses and negative feedback trials. These findings are inconsistent with the R-L theory (Holroyd & Coles, 2002), which focuses on the role of errors and negative feedback for learning. In contrast, these results are in line with a variety of neurophysiologial and neuroimaging studies that showed learning is driven by both, positive prediction errors when the outcome of an action is better than expected and negative prediction errors when the outcome of an action is worse than expected (see O'Doherty et al., 2004; Schultz, 2002; Seymour et al., 2004). Moreover, there is some evidence from other electrophysiological studies that showed similar results (Cohen et al., 2007; Potts et al, 2006). In line with these findings the present results point to the view that the response- and the feedback-locked positivity might reflect reward-related modulations of activity in the medial frontal cortex. In order to incorporate these findings into the R-L theory (Holroyd & Coles, 2002), one needs to suggest that positive learning signals from the MDS inhibit the neurons in the ACC and lead to the generation of a positivity (see Holroyd et al., 2003; Holroyd, 2004). One prediction that follows from such an assumption would be that the ACC activity should co vary with positive learning signals during reinforcement learning

In line with results from Frank and colleagues (2005), the data presented in thesis revealed that the ERN is indeed larger in younger negative compared to younger positive learners. In line with the Go-NoGo model (Frank, 2005), it seems reasonable to assume that the larger ERN in negative learners is due to individual differences in dopamine levels. However, similar to the learning-related changes the larger ERN for negative learners was only found for learning, but not for the reversal phase. One potential explanation for this would be to assume that in the reversal phase participants were less certain about the correctness of their response and thus generally smaller ERNs were elicited. Hence it might be argued that learning biases in the ERN are only obtained if participants have a distinct representation of the correct response. In line with this idea older adults showed impairments in representing the correctness of the response and no individual differences in learning biases in the ERN.

Surprisingly, in the third experiment the largest FRN was obtained for neutral feedback, indicating that participants might have perceived neutral feedback as even worse than negative feedback. This suggests that when feedback is ambiguous and contains no direct information for learning it is processed as being even worse than unambiguous negative feedback. Hence, these results are in line with recent findings by Holroyd et al. (2006) that suggested that neutral and negative feedback might be grouped together as events that indicate that task goals have not been satisfied and hence both lead to the generation of the FRN.

9 Appendix

Table 2

Accuracy in % correct	*Younger adults*			*Older adults*		
	Validity			**Validity**		
Bin	**100%**	**80%**	**50%**	**100%**	**80%**	**50%**
1	0.64 (0.07)	0.60 (0.07)	0.49 (0.04)	0.59 (0.10)	0.53 (0.08)	0.50 (0.03)
2	0.74 (0.11)	0.68 (0.10)	0.49 (0.05)	0.69 (0.13)	0.61 (0.10)	0.50 (0.05)
3	0.77 (0.09)	0.71 (0.11)	0.49 (0.04)	0.71 (0.15)	0.62 (0.13)	0.51 (0.04)
4	0.77 (0.10)	0.71 (0.11)	0.53 (0.05)	0.74 (0.14)	0.65 (0.15)	0.49 (0.04)

Table 3

Note: inv = inverse learning function, lin = linear learning function.

Age group	Validity	Accuracy				Response-locked positivity				ERN				Feedback-locked positivity			
		R^2		β		R^2		β		R^2		β		R^2		β	
		lin	inv	lin	inv	lin	inv	lin	inv	lin	inv	lin	inv	lin	Inv	lin	inv
Younger	100%	.61	.77	0.04	-0.18	.47	.54	0.90	-3.48	.41	.45	0.81	-3,47	.60	.62	-1.34	5.42
		(.27)	(.23)	(.02)	(.10)	(.32)	(.33)	(.80)	(3.2)	(.30)	(.30)	(1.4)	(5.3)	(.29)	(.33)	(0.7)	(3.2)
	80%	.48	.58	0.03	-0.15	.50	.57	0.83	-3.38	.28	.31	0.10	-0.76	.36	.42	-0.32	2.04
		(.32)	(.33)	(.03)	(.14)	(.33)	(.31)	(.84)	(3.7)	(.35)	(.32)	(1.0)	(3.5)	(.29)	(.31)	(1.1)	(4.5)
	50%	.24	.22	0.01	-0.02	.44	.41	0.08	-0.37	.38	.39	0.22	-0.64	.46	.40	0.08	-0.53
		(.31)	(.29)	(.02)	(.07)	(.31)	(.31)	(.79)	(2.7)	(.32)	(.32)	(0.9)	(3.0)	(.36)	(.31)	(1.0)	(3.8)
Older	100%	.69	.78	0.05	-0.20	.52	.51	0.78	-3.14	.35	.37	0.13	-0.97	.50	.50	-0.70	2.62
		(.31)	(.27)	(.03)	(.14)	(.28)	(.30)	(.80)	(3.0)	(.32)	(.30)	(0.7)	(2.3)	(.36)	(.34)	(0.8)	(3.4)
	80%	.61	.58	0.04	-0.15	.47	.40	0.43	-1.75	.39	.39	-0.05	-0.04	.47	.43	-0.41	1.82
		(.38)	(.35)	(.04)	(.15)	(.30)	(.36)	(.96)	(3.0)	(.33)	(.32)	(1.1)	(3.8)	(.31)	(.30)	(0.7)	(2.7)
	50%	.22	.16	0.00	0.00	.53	.52	0.41	-1.30	.34	.35	0.23	-0.73	.45	.43	0.03	-0.39
		(.25)	(.20)	(.02)	(.05)	(.34)	(.32)	(.86)	(3.3)	(.34)	(.33)	(0.8)	(2.9)	(.28)	(.34)	(0.9)	(3.7)

Table 5

Accuracy in % correct	***Younger adults***			***Children***		
	Validity			**Validity**		
Bin	**100%**	**80%**	**50%**	**100%**	**80%**	**50%**
1	0.64 (0.07)	0.60 (0.07)	0.49 (0.04)	0.61 (0.09)	0.52 (0.06)	0.51 (0.05)
2	0.74 (0.11)	0.68 (0.10)	0.49 (0.05)	0.67 (0.11)	0.57 (0.07)	0.49 (0.04)
3	0.76 (0.09)	0.70 (0.11)	0.50 (0.04)	0.69 (0.13)	0.58 (0.09)	0.48 (0.04)
4	0.77 (0.10)	0.70 (0.11)	0.53 (0.05)	0.73 (0.12)	0.63 (0.11)	0.50 (0.04)

Table 6

		Accuracy		**Response-locked positivity**		**Feedback-locked positivity**	
Age group	**Validity**	R^2	β	R^2	β	R^2	β
Adults	100%	.63 (.24)	.042 (.025)	.55 (.36)	1.00 (.800)	.60 (.33)	-1.34 (.726)
	80%	.50 (.31)	.035 (.034)	.58 (.31)	0.75 (.858)	.36 (.29)	-0.32 (1.11)
	50%	.26 (.31)	.009 (.017)	.36 (.30)	0.47 (.813)	.46 (.36)	0.08 (.998)
Children	100%	.52 (.36)	.037 (.031)	.40 (.28)	0.66 (1.28)	.30 (.26)	-0.20 (1.21)
	80%	.55 (.29)	.032 (.035)	.53 (31)	1.12 (1.46)	.29 (.26)	0.08 (1.28)
	50%	.32 (.32)	-.005 (.020)	.28 (.27)	-0.03 (1.58)	.44 (.33)	0.01 (1.69)

Table 8

Accuracy in % correct			***Younger adults***		***Older adults***	
Block phase	**Learning condition**	**Block half**	Positive learners	Negative Learners	Positive learners	Negative Learners
Learning	positive	1	.68 (.08)	.62 (.11)	.59 (.10)	.49 (.07)
		2	.81 (.08)	.73 (.12)	.66 (.10)	.54 (.10)
	negative	1	.63 (.07)	65 (.11)	.51 (.07)	.59 (.10)
		2	.74 (.09)	.81 (.08)	53 (.10)	.66 (.11)
Reversal	positive	1	.59 (.08)	.55 (.08)	.51 (.06)	.45 (.08)
		2	.74 (.11)	.65 (.12)	.58 (.09)	.49 (.09)
	negative	1	.61 (.03)	.60 (.07)	.52 (.06)	.61 (.08)
		2	.70 (.10)	.76 (.10)	.53 (.08)	.64 (.11)

10 Zusammenfassung

10.1 Einleitung

“Aus Fehlern wird man klug.” Dieses deutsche Sprichwort weist auf die Bedeutung hin, die Fehler für das Lernen haben können. Fehler signalisieren, dass ein Ziel nicht erreicht wurde, dass eine Erwartung nicht erfüllt wurde und dass Verhaltensanpassungen nötig sind um die eigenen Bedürfnisse zu befriedigen. In den letzten Jahren wurden verschiedene Modelle vorgeschlagen um die neuronalen Prozesse der Fehlerverarbeitung zu erklären (Botvinick et al., 2001; Coles et al., 1998; Holroyd & Coles, 2002; Holroyd et al., 2005; Yeung et al., 2004). Die meisten dieser Modelle basieren auf Befunden zur so genannten Fehlernegativierung („error-related negativity“, ERN (Gehring et al., 1993), oder “error negativity”, Ne (Falkenstein et al., 1995)). Die ERN ist eine Negativierung im ereigniskorrelierten Potential (EKP), die sich findet, wenn Probanden in Reaktionszeitaufgaben Fehler machen. Zahlreiche Befunde deuten darauf hin, dass die ERN im anterioren cingulären Kortex (ACC), also im medialen Teil des frontalen Kortex generiert wird (Carter et al., 1998; Holroyd et al., 2004; Miltner et al., 2003; Van Veen & Carter, 2002; Swick & Turken, 2002). Ein Grossteil der Modelle zur Fehlerverarbeitung geht davon aus, dass die ERN dann generiert wird, wenn es zu einem “mismatch”, also zu einer Nichtübereinstimmung der Repräsentation der richtigen und der tatsächlichen (falschen) Antwort kommt (Coles et al., 1998; Holroyd & Coles, 2002; Holroyd et al., 2005).

Befunde aus Untersuchungen zur strukturellen Entwicklung des Gehirns deuten darauf hin, dass insbesondere die frontalen Areale, die den Prozessen der Fehlerverarbeitung zugrunde liegen, bis in das frühe Erwachsenenalter hinein reifen (Gogtay et al., 2004, Sowell et al., 2003; Sowell et al., 2004). Konsistent mit diesen Befunden zeigten Studien zu entwicklungsbedingten Veränderungen in der Fehlernegativierung, dass die Amplitude der ERN bis ins frühe Erwachsenenalter zunimmt (Davies et al., 2004; Ladouceur et al., 2004; Santesso et al., 2006). Altersbedingte Unterschiede in der Fehlernegativierung finden sich aber auch im höheren

Erwachsenenalter. So zeigten altersvergleichende Studien, dass ältere Menschen Beeinträchtigungen in der Fehlerverarbeitung haben, was sich in einer reduzierten ERN Amplitude ausdrückt (Band & Kok, 2000; Falkenstein et al., 2001; Mathewson et al., 2005, Nieuwenhuis et al., 2002; Themanson et al., 2006; West, 2004).

Allerdings lernen wir nicht nur aus unseren Fehlern, sondern auch aus unseren richtigen Handlungen. Das heißt, wir lernen nicht nur durch Bestrafung, sondern auch durch Belohnung, nämlich dann, wenn eine Handlung zu einem angenehmen Effekt führt, wenn das Ergebnis einer Handlung besser als erwartet ist. Diese Tatsache wurde bereits 1911 von Edward Thorndike im “law of effect” explizit formuliert und stellt die Grundlage verschiedener Theorien des Verstärkungslernens dar (Holroyd & Coles, 2002; Pearce, 1987; Rescorla & Wagner, 1972; Schultz, 2007; Skinner, 1938; Sutton & Barto, 1981). In den letzten 15 Jahren sind diese grundlegenden Lernprozesse in den Fokus der neurowissenschaftlichen Forschung geraten. Dies ist vor allem das Verdienst der Gruppe um den Neurophysiologen Wolfram Schultz, die anhand von elektrophysiologischen Ableitungen bei Primaten zeigen konnte, dass der Neurotransmitter Dopamin eine zentrale Rolle für das Verstärkungslernen spielt (Übersichtsarbeiten siehe Schultz, 2000; Schultz, 2002; Schultz, 2007). Insbesondere konnten Schultz und Kollegen (1997) zeigen, dass dopaminerge Neurone im Mittelhirn positive und negative Lernsignale kodieren und diese an andere subkortikale Areale und insbesondere den frontalen Kortex projizieren. Diese Signale stellen die Grundlage des Lernens dar und lassen sich als so genannte Vorhersagefehler (siehe Rescorla & Wagner, 1972; Sutton & Barto, 1981) modellieren. Das heißt, die Dopaminsignale reflektieren die Abweichungen von Erwartungen, also dann, wenn ein Ereignis besser oder schlechter als erwartet ist. Ähnlich wie bei der Fehlerverarbeitung deuten verschiedene Befunde darauf hin, dass das Dopaminsystem und insbesondere die dopaminergen Projektionen in den frontalen Kortex stark von entwicklungsbedingten Veränderungen (Diamond, 1996; Diamond et al., 2004; Goldman-Rakic & Brown, 1982; Weickert et al., 2007), als auch von altersbedingten

Beeinträchtigungen betroffen sind (Erixon-Lindroth et al., 2005; Volkow et al., 1996, Volkow et al., 2000; Wang et al., 1998).

Das Ziel der vorliegenden Arbeit ist die Untersuchung von Altersunterschieden in der Fehlerverarbeitung und dem Verstärkungslernen über die Lebensspanne. Dabei liegt das Hauptaugenmerk dieser Arbeit auf der Frage, welche Bedeutung Fehler für das Verstärkungslernen haben und wie sich Fehlerverarbeitung und Lernen zwischen Kindern, jüngeren und älteren Erwachsenen unterscheiden. Die Grundlage dieser Arbeit ist die Verstärkungslerntheorie der ERN von Holroyd und Coles, (2002). Dieses Modell geht davon aus, dass die Fehlernegativierung (ERN) im anterioren cingulären Kortex (ACC) generiert wird, wenn ein negatives Lernsignal aus dem Dopaminsystem die Neurone im ACC disinhibiert. Das heißt, das Modell geht davon aus, dass die ERN einen negativen Vorhersagefehler reflektiert, also die Tatsache, dass ein Ereignis schlechter als erwartet ist. Interessanterweise führt nicht nur die Verarbeitung interner Fehlerinformation zur Generierung einer ERN, sondern auch externe Fehlerinformation. Es findet sich also auch dann eine ERN, wenn ein externer Stimulus signalisiert, dass ein Fehler begangen wurde. In diesem Fall spricht man von einer so genannten “Feedback-” ERN (FRN) (siehe Miltner et al., 1997). Die zentrale Vorhersage des Modells von Holroyd und Coles, (2002) ist, dass die ERN mit dem Lernen zunehmen sollte, während die FRN mit dem Lernen abnehmen sollte. Die ERN sollte mit dem Lernen zunehmen, weil der negative Vorhersagefehler mit zunehmendem Lernen größer wird. Der Grund dafür ist, dass mit zunehmendem Lernen eine Erwartung daran aufgebaut wird richtig zu antworten. Je größer diese Erwartung ist, umso größer ist auch die Erwartungsverletzung (der Vorhersagefehler), wenn es zu einer fehlerhaften Antwort kommt. Im Gegensatz dazu sollte die Amplitude der FRN mit zunehmendem Lernen abnehmen, weil die Probanden immer mehr in der Lage sind, selbst die Richtigkeit der Antwort vorherzusagen und weniger von der externen Information abhängig sind. Holroyd und Coles (2002) untersuchten die Vorhersagen ihres Modells anhand einer probabilistischen Lernaufgabe. In dieser Aufgabe wurden den Probanden Stimuli präsentiert und sie sollten mit einer von zwei Tasten auf diese Stimuli reagieren. Daraufhin erhielten sie eine Rückmeldung über die Richtigkeit der Antwort. Die Probanden wurden instruiert, die Stimulus-Reaktionsbeziehungen auf der Basis der Rückmeldung per Versuch und Irrtum zu erlernen. Die EKPs wurden auf die Reaktion und auf die Rückmeldung hin ausgewertet. Die Ergebnisse bestätigten das Modell und zeigten tatsächlich den erwarteten Anstieg der ERN und eine Reduktion der FRN mit dem Lernen.

In der Folge schlugen Nieuwenhuis und Kollegen (2002) eine Erweiterung dieses Modells vor um die Beeinträchtigungen älterer Menschen beim Lernen und in der Fehlerverarbeitung zu erklären. Basierend auf Befunden zu altersbedingten Beeinträchtigungen in der Fehlerverarbeitung und im Dopaminsystem geht das erweiterte Modell davon aus, dass ältere Menschen ein abgeschwächtes Dopaminsignal haben, das zu einer Reduktion der ERN und zu Beeinträchtigungen beim Lernen führt. Ähnlich wie Holroyd und Coles (2002) untersuchten Nieuwenhuis et al. (2002) ihren Ansatz anhand einer probabilistischen Lernaufgabe und untersuchten die ERN bei jüngeren und älteren Erwachsenen. Im Einklang mit dem Modell fanden sie sowohl eine reduzierte ERN als auch altersbedingte Beeinträchtigungen beim Lernen

10.2 Experimente

10.2.1 Experiment 1

Ziel des ersten Experimentes war es die Befunde zu altersbedingten Beeinträchtigungen in der Fehlerverarbeitung und beim Lernen zu erweitern und zwei wichtige Punkte zu klären, die in der Untersuchung von Nieuwenhuis et al. (2002) nicht oder unzureichend geklärt wurden. Bei dem ersten Punkt geht es um die Rolle von Performanz-Unterschieden zwischen Altersgruppen und deren Bedeutung für die Altersunterschiede in der ERN. In der Studie von Nieuwenhuis et al. (2002) waren Altersunterschiede in der ERN durch erhebliche Performanzunterscheide zwischen Altersgruppen konfundiert. Das ist bedeutsam, weil sowohl die Mismatchmodelle der ERN (Coles et al., 1998; Holroyd et al., 2005) als auch die Theorie von Holroyd und Coles (2002) davon ausgehen, dass die Amplitude der ERN von der Fähigkeit abhängt die richtige Antwort zu repräsentieren. Das heißt, je besser Probanden lernen, umso eindeutiger ist ihre Repräsentation der richtigen Antwort und umso größer ist der Mismatch und damit die ERN, wenn ein Fehler passiert. Ein zentrales Ziel des ersten Experiments war es die Performanz der älteren Probanden an die der jüngeren Probanden anzugleichen. Sollte die ERN performanz- und nicht altersabhängig sein, dann sollten sich keine Altersunterschiede in der ERN zwischen den Altersgruppen finden.

Der zweite Punkt, der in der Studie von Nieuwenhuis et al. (2002) nur unzureichend geklärt wurde, ist die Frage, wie sich die ERN über den Verlauf des Lernens hinweg bei jüngeren und älteren Probanden entwickelt. Nieuwenhuis et al. (2002) verglichen zwar verschiedene Lernbedingungen, betrachteten die ERN aber nicht über den Verlauf des Lernens hinweg. Hier stellt sich die Frage, ob sich tatsächlich altersbedingte Unterschiede im Lernverlauf in der ERN finden lassen, wenn die Performanzunterschiede zwischen den Altersgruppen ausgeglichen sind.

Um diese Fragestellungen zu untersuchen wurde eine probabilistische Lernaufgabe eingesetzt, in der die Validität der Rückmeldung manipuliert wurde. Ähnlich wie bei Holroyd und Coles (2002) und Nieuwenhuis et al. (2002) sollten die Probanden per Versuch und Irrtum, auf Basis der Rückmeldung, Stimulus-Reaktionsverknüpfungen erlernen. Die Validität der Rückmeldung wurde in drei Bedingungen (100% valide, 80% valide und 50% valide) manipuliert. In der 100% validen Bedingung war das Feedback immer valide und die Probanden konnten auf der Basis der Rückmeldung lernen. In der 80% valide Bedingung war die Rückmeldung in 80% der Durchgänge valide und in 20 % der Durchgänge invalide. Dementsprechend war das Lernen in dieser Bedingung beeinträchtigt, da die invalide Rückmeldung das Erlernen der Stimulus-Reaktionsverknüpungern stört. Die 50% valide Bedingung diente als Kontrollbedingung, in der kein Lernen möglich ist (siehe Abb. 8).

Um die Performanz zwischen den beiden Altersgruppen anzugleichen wurde ein Algorithmus in der Aufgabe implementiert, der adaptiv die Reaktionszeitgrenzen anpasst (siehe Methode Experiment 1). Die Idee hinter dieser adaptiven Anpassung ist es, die mit dem Altern einhergehenden Effekte genereller Verlangsamung zu kontrollieren (siehe Salthouse, 1996; 2000). Würde man die gleiche Reaktionszeitgrenze für beide Altersgruppen verwenden, wie es in der Nieuwenhuis et al. (2002) Studie getan wurde, dann würde dies zu einem überproportionalen Zeitdruck für die älteren Erwachsenen führen und damit ihre Möglichkeit zu lernen einschränken. Darüber hinaus führt ein solcher Zeitdruck zu einer vermehrten Anzahl von Reaktionszeitüberschreitungen bei den Älteren, was die Anzahl der Durchgänge reduziert, in denen sie auf Basis der Rückmeldung lernen können.

Die adaptive Anpassung der Reaktionszeitgrenzen sollte dazu führen, dass die Älteren besser lernen und zumindest in der 100% valide Bedingung eine ähnliche Performanz wie die Jüngeren erreichen sollten. Auf der Basis des Modells von Holroyd und Coles (2002) erwarteten wir, dass die ERN mit dem Lernen zunehmen sollte, während die FRN mit dem Lernen abnehmen sollte. Entsprechend des Ansatzes von Nieuwenhuis und Kollegen (2002) sollte sowohl die ERN als auch die FRN bei den Älteren reduziert sein und es sollten sich abgeschwächte lernbedingte Veränderungen in beiden Komponenten finden. Wenn die Altersunterschiede in der ERN in der Studie von Nieuwenhuis et al. (2002) jedoch durch Performanzunterschiede konfundiert waren, dann sollte eine Anpassung der Performanz zwischen den Altersgruppen zu vergleichbaren ERN Amplituden der Jüngeren und der Älteren führen.[19]

10.2.2 Experiment 2

Ziel des zweiten Experimentes war es entwicklungsbedingte Unterschiede in der Fehlerverarbeitung und dem Lernen zu untersuchen. Der Fokus lag dabei auf der Frage, welche Rolle die Verarbeitung interner und externer Fehlerinformation für das Lernen bei Kindern und jüngeren Erwachsenen spielt. Dazu wurde die Lernaufgabe aus dem ersten Experiment zur Verwendung mit Kindern angepasst (siehe Methode Experiment 2).

---

[19] Die Stichprobe in Experiment 1 bestand aus 18 älteren Erwachsenen (mittleres Alter: 69 Jahre, SD = 2.8) und 18 jüngeren Erwachsenen (mittleres Alter: 21 Jahre, SD = 1.8) (siehe Tabelle 1). Die Stichprobe in Experiment 2 bestand aus 17 Kindern (mittleres Alter: 21 Jahre, SD = 1.8) und 18 jüngeren Erwachsenen (mittleres Alter: 11 Jahre, SD = 0.8) (siehe Tabelle 4).

Ähnlich wie für die älteren Probanden erwarteten wir, dass die adaptive Anpassung der Reaktionszeitgrenzen dazu führen sollte, dass die Kinder zumindest in der 100% valide Bedingung eine ähnliche Performanz wie die Erwachsenen erreichen sollten. Auf Basis früherer Befunde (Davies et al., 2004; Ladouceur et al., 2004; Santesso et al., 2006) und der Vorhersage des Modells von Holroyd & Coles (2002) erwarteten wir, dass die Verarbeitung interner Fehlersignale bei Kindern im Verhältnis zu Erwachsenen beeinträchtigt sein sollte. Dies sollte sich in einer reduzierten ERN bei den Kindern zeigen. Allerdings deuten einige Befunde darauf hin, dass sich ähnlich wie bei den älteren Erwachsenen Reduktionen in der Amplitude der ERN nur dann finden, wenn sich auch Altersunterschiede in der Performanz zeigen (Hogan et al. 2005; Kim et al., 2007). Leider gibt es bis dato nach meinem Kenntnisstand keine altersvergleichenden Studien zur FRN bei Kindern. Studien zu Altersunterschieden in der Verarbeitung externer Fehlerinformation anhand von Herzratenmaßen (Crone et al., 2005, Crone et al., 2006) und eine Studie zur FRN bei Kindern mit Aufmerksamkeitsdefizitstörung (ADS) (van Meel et al., 2005) deuten aber darauf hin, dass Kinder im Vergleich zu Erwachsenen sensitiver für externe Fehlerinformation sind und sich mit dem Lernen weniger von der externer Rückmeldung lösen können. [1]

10.2.3 Experiment 3

Im dritten Experiment wurde ein etwas anderer Ansatz zur Untersuchung von Altersunterschieden in der Fehlerverarbeitung und dem Lernen gewählt. Das Ziel des dritten Experiments war es zu untersuchen, ob jüngere und ältere Erwachsenen sich darin unterscheiden, wie sie aus positiver oder negativer Rückmeldung lernen und ob sich diese individuellen Unterschiede in der Amplitude der ERN niederschlagen. Das dritte Experiment basiert auf Befunden aus dem ersten Experiment, die darauf hindeuteten, dass ältere Menschen die Tendenz haben weniger aus negativer als aus positiver Rückmeldung zu lernen. Dieser Befund ist konsistent mit der Sozioemotionalen Selektivitätstheorie von Carstensen und Kollegen (Carstensen, 2006; Mather & Carstensen, 2004), die davon ausgeht, dass ältere Menschen aufgrund der Tatsache, dass sie sich dem Ende ihres Lebens nähern, mehr Gewicht auf die Regulation ihrer Emotionen legen als jüngere Menschen. Infolge dessen setzen sie kognitive Kontrollmechanismen ein um positive Informationen zu verstärken und negative Informationen zu unterdrücken. Dieses Modell sagt also voraus, dass ältere Menschen im Vergleich zu jüngeren die Tendenz haben sollten eher aus positiver als aus negativer Rückmeldung zu lernen.

Im Gegensatz dazu lässt sich aus einem kürzlich von Frank und Kollegen (2004) vorgeschlagenen Modell eine konträre Hypothese entwickeln. Das Modell von Frank und Kollegen wurde auf der Basis von Untersuchungen bei Parkinsonpatienten entwickelt. In diesen Untersuchungen wurden Parkinsonpatienten in Abhängigkeit ihrer L-Dopa Medikation, also ihres Dopaminspiegels, anhand einer probabilistischen Lernaufgabe untersucht. In dieser Aufgabe werden Stimuluspaare präsentiert und die Probanden sollen lernen auf Basis der Rückmeldung einen der Stimuli auszuwählen. In einem darauf folgenden Test wurde anhand von Neukombinationen der Stimuli überprüft, ob die Probanden gelernt hatten bestimmte Stimuli zu vermeiden oder bestimmte Stimuli zu wählen, ob sie also eher aus negativer Rückmeldung oder eher aus positiver Rückmeldung lernen. Es zeigte sich, dass Parkinsonpatienten ohne Medikation, die also einen reduzierten Dopaminspiegel haben, eher aus negativer Rückmeldung lernen, während sie mit L-Dopa mehr aus positivem Feedback lernen. Das Modell von Frank und Kollegen (2004) geht davon aus, dass die Lerntendenz, ob also mehr aus positiver oder aus negativer Rückmeldung gelernt wird, vom Dopaminspiegel abhängt. Wenn der Dopaminspiegel niedrig ist, wie bei Parkinsonpatienten ohne Medikation, dann kommt es eher zur Generierung negativer Lernsignale, entsprechend lernen die Probanden eher aus negativer Rückmeldung. Im Gegensatz dazu kommt es bei einem hohen Dopaminspiegel eher zur Generierung positiver Lernsignale und die Probanden lernen eher aus positiver Rückmeldung. Interessanterweise finden sich diese Lerntendenzen nicht nur bei Parkinsonpatienten, sondern auch bei gesunden jungen Probanden (Frank et al., 2005). Darüber hinaus konnten Frank und Kollegen (2005) zeigen, dass bei jungen

Probanden die ERN als Funktion individueller Unterschiede in den Lerntendenzen variiert. Das heißt, Negativlerner zeigen eine größere ERN als Positivlerner. Etliche Befunde deuten darauf hin, dass der Dopaminspiegel mit zunehmendem Alter abnimmt (Erixon-Lindroth et al., 2005; Wang et al., 1998; Volkow et al., 1996, Volkow et al., 2000) und dass sich diese Veränderungen in der ERN niederschlagen (Nieuwenhuis et al., 2002). Auf der Basis dieser Befunde und dem Modell von Frank und Kollegen lässt sich die Vorhersage ableiten, dass ältere Menschen die Tendenz haben sollten mehr aus negativer Rückmeldung als aus positiver Rückmeldung zu lernen. Dies steht im Gegensatz zu dem Modell von Carstensen (2006), das vorhersagt, dass ältere Menschen Positivlerner sein sollten.

Um diese beiden konträren Hypothesen zu untersuchen, wurde ein Paradigma entwickelt, das es erlaubt Lerntendenzen bei jüngeren und älteren Probanden in zwei Lernbedingungen, einer Positivlernbedingung und einer Negativlernbedingung, zu untersuchen. Wie in den vorangegangenen Experimenten sollten die Probanden auf der Basis von Rückmeldung Stimulus-Reaktionsbeziehungen erlernen. In der Positivlernbedingung konnten die Probanden 50 Eurocent gewinnen, wenn sie die richtige Taste drückten, oder sie konnten eine neutrale Rückmeldung erhalten, wenn sie die falsche Taste drückten (siehe Abb. 26). In dieser Bedingung sollten die Probanden lernen die Taste zu wählen, die zur positiven Rückmeldung führt. In der Negativlernbedingung konnten die Probanden 50 Eurocent verlieren, wenn sie die falsche Reaktion abgaben, oder eine neutrale Rückmeldung erhalten, wenn sie die richtige Taste drückten. In dieser Bedingung sollten die Probanden lernen die Reaktion zu vermeiden, die zu einer negativen Rückmeldung führt. Um die Stabilität der Lerntendenzen zu untersuchen, wurde eine Umlernphase in den Lernblöcken implementiert. Jeder Lernblock beinhaltet also eine Lernphase, in der die Stimulus-Reaktionszuordnungen erlernt werden, und eine Umlernphase, in der die Stimulus-Reaktionszuordnungen umgelernt werden müssen. Dass heißt, Stimuli, die in der Lernphase mit der Positivlernbedingung assoziiert waren, sind in der Umlernphase mit der Negativlernbedingung assoziiert und umgekehrt (siehe

Abb. 27). Positivlerner werden in dieser Aufgabe als Probanden definiert, die eine bessere Performanz in der Positivlernbedingung im Vergleich zur Negativlernbedingung aufweisen. Negativlerner werden als Probanden definiert, die in der Negativlernbedingung eine bessere Performanz aufweisen als in der Positivlernbedingung.

Hinsichtlich der Lerntendenzen der Älteren gibt es zwei konkurrierende Hypothesen. Entsprechend des Modells von Carstensen (2006) würde man erwarten, dass die Älteren eher aus positiver Rückmeldung als aus negative Rückmeldung lernen sollten. Im Gegensatz dazu lässt sich aus dem Modell von Frank und Kollegen (2004) ableiten, dass ältere Menschen aufgrund ihrer Beeinträchtigungen im Dopaminhaushalt eher Negativlerner als Positivlerner sein sollten. Darüber hinaus ist aufgrund der Befunde von Frank et al. (2005) zu erwarten, dass Negativlerner eine größere ERN haben sollten als Positivlerner, wobei unklar ist, ob sich dieser Effekt auch bei älteren Probanden findet. Die Umlernphase sollte zu Interferenz auf der Ebene der Stimulus-Reaktionszuordnungen führen, weil die gelernten Verknüpfungen umgelernt werden müssen.[20]

10.3 Ergebnisse und Diskussion

Die Ergebnisse der ersten beiden Experimente zeigen, dass Kinder und ältere Probanden vergleichbare Lerneffekte zeigen wie jüngere Erwachsene, wenn die Lernaufgabe an ihre Anforderungen angepasst ist (siehe Abb. 10 und Abb. 17). Wenn also die zeitlichen Anforderungen der Aufgabe so angepasst werden, dass alle Altersgruppen genug Zeit zum Lernen haben und eine ähnliche Anzahl von Reaktionszeitüberschreitungen aufweisen. Die vorliegenden Befunde deuten also im Gegensatz zu der Studie von Nieuwenhuis et al. (2002) nicht darauf hin, dass ältere Probanden oder Kinder Beeinträchtigungen beim Verstärkungslernen haben. Allerdings fanden sich für beide Altersgruppen Beeinträchtigungen in der Performanz, wenn die Rückmeldung partiell invalide war und mit dem Lernen interferierte. Daraus lässt sich

---

[20] Die Stichprobe in Experiment 3 bestand aus 26 jüngeren Erwachsenen (mittleres Alter: 22 Jahre, SD = 2.3) und 24 älteren Erwachsenen (mittleres Alter: 70 Jahre, SD = 2.9) (siehe Tabelle 7).

folgern, dass Kinder und ältere Erwachsene dann beeinträchtigt sind, wenn sie Kontrollprozesse rekrutieren müssen um mit der interferierenden invaliden Information während des Lernens umzugehen (Bunge et al., 2002; Craik & Bialystok, 2006; West, 1998).

In Einklang mit unserer Hypothese zeigten die Analysen der reaktionsbezogenen EKPs der ersten beiden Experimente, dass die Angleichung der Performanz zwischen Kindern, jüngeren und älteren Erwachsenen zu vergleichbaren ERN Amplituden in den drei Altersgruppen führte (siehe Abb. 11 und Abb. 18). Dieses Ergebnis steht sowohl im Widerspruch zu den Vorhersagen des Modells von Nieuwenhuis und Kollegen (2002) als auch zu den Ergebnissen etlicher anderer Studien zu Altersunterschieden in der ERN (Davies et al., 2004; Ladouceur et al., 2004; Santesso et al., 2006; Band & Kok, 2000; Falkenstein et al., 2001; Mathewson et al., 2005, Nieuwenhuis et al., 2002; Themanson et al., 2006; West, 2004). Allerdings muss berücksichtigt werden, dass in diesen Studien Altersunterschiede in der ERN mit Performanzunterschieden zwischen den Altersgruppen konfundiert waren.

Aus diesen Befunden ergeben sich zweierlei Implikationen. Zum einen weisen diese Ergebnisse darauf hin, wie wichtig es ist, die Performanz zwischen Altersgruppen anzugleichen wenn man Altersunterschiede in der ERN untersuchen möchte. Zum anderen zeigen diese Befunde, dass die ERN performanz- und nicht alterssensitiv ist. Obwohl diese Ergebnisse im Widerspruch zu dem Ansatz von Nieuwenhuis et al. (2002) stehen, sind sie doch konsistent mit dem ursprünglichen Missmatch Model der ERN (Bernstein et al., 1995; Coles et al., 1998) Dieses Modell impliziert, dass die Amplitude der ERN davon abhängig ist, wie gut eine Versuchsperson in der Lage ist eine Repräsentation der richtigen Antwort aufzubauen. Je besser die Performanz einer Versuchsperson, umso besser die Repräsentation der richtigen Antwort und umso größer der „mismatch“, also die ERN, wenn eine falsche Reaktion erfolgt.

Darüber hinaus ergaben die ersten zwei Experimente zwei Befunde, die einen erheblichen Einfluss auf die Modelle von Holroyd und Coles (2002), sowie Nieuwenhuis et al. (2002) haben könnten. Zum einen zeigte die Analyse der reaktionsbezogenen EKPs, dass die lernbedingten Veränderungen in einer Positivierung für korrekte Reaktionen sehr viel stärker ausgeprägt waren als in der ERN für inkorrekte Reaktionen (siehe Abb. 12 und Abb. 19). Ein ähnlicher Befund ergab sich in den rückmeldungsbezogenen EKPs. Hier fand sich keine Reduktion der FRN mit dem Lernen. Im Gegensatz dazu fand sich eine erhebliche lernbezogene Reduktion einer Positivierung für positive Rückmeldungen (siehe Abb. 15 und Abb. 23). Zum Teil lassen sich die Diskrepanzen zwischen den Studien durch unterschiedliche Auswertungsansätze erklären. In den Studien von Holroyd and Coles (2002) und Nieuwenhuis et al. (2002) wurden die reaktions- und rückmeldungsbezogenen EKPs anhand von Differenzwellen ausgewertet. Im Gegensatz dazu wurden in den vorliegenden Experimenten die EKPs für korrekte und inkorrekte Antworten als auch positive und negative Rückmeldungen separat ausgewertet. Der Grund für dieses Vorgehen ist, dass die reaktions- als auch die rückmeldungsbezogenen EKPs lernbezogene Veränderungen hauptsächlich für korrekte Durchgänge zeigten. Es ist also durchaus anzunehmen, dass sich in den oben genannten Studien ähnliche Befunde ergeben hätten, wenn korrekte und inkorrekte Durchgänge separat analysiert worden wären. Die vorliegenden Befunde sind konsistent mit einer großen Anzahl von tierphysiologischen und bildgebenden Studien, die zeigen, dass Verstärkungslernen über negative und positive Verstärkungslernsignale aus dem Dopaminsystem getrieben wird (Mirenowicz & Schultz, 1996; Schultz et al., 1997; O Doherty et al., 2004; Seymore et al., 2004). Um diese Befunde mit dem Modell von Holroyd & Coles (2002) zu integrieren müsste das Modell dahingehend modifiziert werden, dass es davon ausgeht, dass positive Lernsignale die Neurone im ACC inhibieren und dies zur Generierung der reaktionsbezogenen Positivierung führt (siehe Holroyd et al., 2003; Holroyd, 2004).

Die Analyse der rückmeldungsbezogenen EKPs zeigte eine Asymmetrie in der Verarbeitung der Rückmeldung über die Lebensspanne. Ältere Erwachsene zeigten eine stark reduzierte FRN im Vergleich zu jüngeren Erwachsenen, was darauf hindeutet, dass die Älteren negative Rückmeldungen weniger stark verarbeiten als die Jüngeren (siehe

Abb. 14). Dies ist ein überraschender Befund, da die älteren Erwachsenen eine ähnliche Lernleistung und eine ähnlich große ERN wie die jüngeren Erwachsenen zeigen. In Übereinstimmung mit dem Modell von Carstensen (2006) deuten diese Befunde darauf hin, dass die älteren Erwachsenen während des Lernens mehr auf die positive als auf die negative Rückmeldung fokussieren. Im Gegensatz dazu zeigen Kinder eine größere FRN als junge Erwachsene (siehe Abb. 22). Das ist konsistent mit Studien zur Herzratenverlangsamung bei Kindern (Crone et al., 2004), also auch mit Befunden zur FRN bei Kindern mit ADS (van Meel et al., 2005) und deutet darauf hin, dass Kinder während des Lernens sensitiver für externe Fehlerinformation sind als jüngere Erwachsene. Die Verarbeitung von Rückmeldungsinformation ändert sich also im Verlaufe der Entwicklung über die Lebensspanne. Kinder scheinen sensitiver für negative Rückmeldung während des Lernens zu sein, während ältere Probanden negative Rückmeldungen weniger stark zu verarbeiten scheinen.

Die Verhaltensdaten des dritten Experiments zeigen erhebliche altersbedingte Unterschiede in der generellen Performanz, sowie altersbedingte Beeinträchtigungen beim Lernen (siehe Abb. 28 und Abb. 30). Dieses Ergebnis ist überraschend, da im ersten Experiment keine altersbedingten Unterschiede in der Performanz oder beim Lernen gefunden wurden. Eine Interpretation dieser Befunde, die insbesondere durch die EKP Daten unterstützt wird, ist, dass die älteren Erwachsenen Probleme beim Verarbeiten der ambigen neutralen Rückmeldung hatten. Das heißt, sie hatten Probleme eine Repräsentation darüber aufzubauen, wann die neutrale Rückmeldung besser und wann sie schlechter als die alternative Rückmeldung ist. Die reaktionsbezogenen EKPs zeigten, dass die jüngeren Probanden klar zwischen Reaktionen, die zu positiver (+50) Rückmeldung und Reaktionen die zu negativer (-50) Rückmeldung führen, differenzieren. Das heißt, sie zeigen eine ERN für die falschen als auch eine reaktionsbezogene Positivierung für die richtigen Antworten. Dasselbe Muster zeigt sich bei den richtigen und falschen Antworten, die zu neutraler Rückmeldung führen (siehe Abb. 33). Das heißt, die jüngeren Probanden sind fähig die Richtigkeit einer Antwort in Relation zur alternativen

Rückmeldung zu repräsentieren. Bei den Älteren findet sich ein solcher Effekt nur für den Vergleich von positiver und negativer Rückmeldung, während sich kein Unterschied in den EKPs für richtige und falsche Antworten findet, wenn diese zu ambiger neutraler Rückmeldung führen (siehe Abb. 33). Zusammengenommen deuten diese Daten darauf hin dass die älteren Erwachsenen insbesondere dann Beeinträchtigungen beim Lernen zeigen, wenn sie eine relationale Repräsentation der Bedeutung der Rückmeldung aufbauen müssen.

Konsistent mit unseren Vorhersagen zeigten sich Lerntendenzen bei jüngeren als auch bei älteren Erwachsenen (siehe Abb. 30). Insgesamt waren mehr ältere Erwachsene Negativlerner und mehr jüngere Erwachsene Positivlerner. Darüber hinaus hatten die älteren Negativlerner vor allem am Anfang des Lernens eine starke Tendenz das negative Feedback zu vermeiden. Einige Befunde des dritten Experiments deuten also darauf hin, dass ältere Probanden eine Tendenz zum Negativlernen haben, was konsistent mit dem Modell von Frank und Kollegen wäre. Die Performanzdaten des dritten Experiments sprechen also nicht dafür, dass ältere Erwachsene mehr aus positiver als aus negativer Rückmeldung lernen. Wie auf Basis der Befunde von Frank und Kollegen (2005) erwartet, zeigten junge Negativlerner eine größere ERN als junge Positivlerner (siehe Abb. 35). Allerdings zeigte sich kein solcher Effekt bei den älteren Negativlernern. Bei der Interpretation der Altersbefunde muss aber berücksichtigt werden, dass die ERN und die Performanz der Älteren generell stark reduziert war. Das heißt, es kann nicht ausgeschlossen werden, dass sich ein ähnliches Muster für die Älteren gefunden hätte, wenn diese eine ähnliche Performanz und eine vergleichbare ERN Amplitude wie die jüngeren Erwachsenen gehabt hätten, so wie es im ersten Experiment der Fall war.

Ein weiterer, unerwarteter Befund ergab sich bei der Analyse der rückmeldungsbezogenen EKPs. Hier zeigte sich die größte FRN für das neutrale und nicht für das negative Feedback (siehe Abb. 36). Dieser Befund ist konsistent mit Daten die kürzlich von Holroyd et al. (2006) publiziert wurden. Holroyd und Kollegen (2006) fanden über verschiedene Aufgaben hinweg vergleichbare ERNs für neutrale und negative

Rückmeldungen. Ihre Interpretation war, dass neutrale und negative Rückmeldungen ähnlich verarbeitet werden, weil sie beide anzeigen, dass ein Ziel nicht erreicht wurde. In diesem Experiment scheinen neutrale Rückmeldungen sogar noch negativer wahrgenommen zu werden als negative Rückmeldungen. Das könnte damit zu tun haben, dass die neutralen Rückmeldungen keine Information für das Lernen tragen und in Relation mit der alternativen Rückmeldung verarbeitet werden müssen. Eine Interpretation für diesen Effekt wäre also, dass die FRN sensitiv für den Informationsgehalt ist, den ein Stimulus in Bezug auf ein zu erreichendes Ziel vermittelt.

Zusammengenommen zeigen die Ergebnisse dieser Arbeit, dass Kinder und ältere Erwachsene keine generellen Beeinträchtigungen in der Fehlerverarbeitung und beim Verstärkungslernen haben. Allerdings zeigen sich altersbedingte Beeinträchtigungen dann, wenn invalide Information mit dem Lernen interferiert und wenn auf der Basis ambiger Information gelernt werden muss, wenn also eine relationale Repräsentation der Rückmeldung aufgebaut werden muss. Darüber hinaus zeigt die vorliegende Arbeit Hinweise auf eine Asymmetrie in der Verarbeitung von Rückmeldungsinformation über die Lebensspanne. Kinder scheinen sensitiver für externe Fehlerinformation zu sein als Erwachsene, während ältere Erwachsene externe Fehlerinformation weniger stark verarbeiten als jüngere Erwachsene. Dies könnte darauf hinweisen, dass Ältere eine Tendenz dazu haben positive, selbstwertsteigernde Information stärker und negative Information schwächer zu verarbeiten (Carstensen, 2006). Allerdings zeigte das dritte Experiment eher Evidenz für die Idee, dass die Älteren auf der Basis von negativer Rückmeldung lernen (siehe Frank et al., 2004).

11 *References*

Ahammer, I. M., & Goulet, L. R. (1969). Intrapair conceptual similarity and verbal reinforcement combinations as determinants of discrimination learning and retention in children. *Child Development*, 859-868.

Ahlskog, J. E. (2007). Beating a dead horse. *Neurology, 69*, 1701-1711.

Adleman, N. E., Menon, V., Blasey, C. M., White, C. D., Warsofsky, I. S., Glover, G. H., et al. (2002). A developmental fMRI study of the Stroop color-word task. *Neuroimage, 16*, 61-75.

Bäckman, L., Nyberg, L., Lindenberger, U., Li, S.-C., & Farde, L. (2006). The correlative triad among aging, dopamine, and cognition. *Neuroscience and Biobehavioral Reviews, 30*, 791-807.

Baltes, P. B., Staudinger, U., & Lindenberger, U. (1999). Lifespan psychology: Theory and application to intellectual functioning. *Annual Review of Psychology, 50*, 471-507.

Band, G. P., & Kok, A. (2000). Age effects on response monitoring in a mental rotation task. *Biological Psychology, 51*, 201-221.

Bartholow, B. D., Pearson, M. A., Dickter, C. L., Sher, K. J., Fabiani, M., & Gratton, G. (2005). Strategic control and medial frontal negativity: Beyond errors and response conflict. *Psychophysiology, 42*, 33-42.

Barto, A. G. (1995). Adaptive critics and the basal ganglia. In J. C. Houk, J. Davis & D. Beisler (Eds.), *Models of information processing in the basal ganglia* (pp. 215-232). Cambridge, MA: MIT Press.

Barto, A. G., & Sutton, R. S. (1997). Reinforcement learning in artificial intelligence. In J. W. Donahoe & V. P. Dorsel (Eds.), *Neural-network models of cognition* (Vol. 121, pp. 358-387). North-Holland.

Bernstein, P. S., Scheffers, M. K., & Coles, M. G. (1995). "Where did I go wrong?" A psychophysiological analysis of error detection. *Journal of Experimental Psychology: Human Perception and Performance, 21*, 1312-1322.

Botvinick, M., Braver, T. S., Barch, D. M., Carter, C. S., & Cohen, J. D. (2001). Conflict monitoring and cognitive control. *Psychological Review, 108*, 624-652.

Biederman, J., & Faraone, S. V. (2005). Attention-deficit hyperactivity disorder. *Lancet, 366*, 237-248.

Birren, J. E., & Fisher, L. M. (1995). Aging and speed of behavior: Possible consequences for psychological functioning. *Annual Review of Psychology, 46*, 329-353.

Braver, T. S., Barch, D. M., Keys, B. A., Carter, C. S., Cohen, J. D., Kaye, J. A., et al. (2001). Context processing in older adults: Evidence for a theory relating cognitive control to neurobiology in healthy aging. *Journal of Experimental Psychology: General, 130*, 746-763.

Braver, T. S., & Barch, D. M. (2002). A theory of cognitive control, aging cognition, and neuromodulation. *Neuroscience and Biobehavioral Reviews, 26*, 809-817.

Braver, T. S., Satpute, A. B., Rush, B. K., Racine, C. A., & Barch, D. M. (2005). Context processing and context maintenance in healthy aging and early stage of the alzheimer's type. *Psychology and Aging, 20*, 33-46.

Bunge, S., Dudukovic, N. M., Thomason, M. E., Vaidya, C. J., & Gabrieli, J. D. E. (2002). Immature frontal lobe contributions to cognitive control in children. Evidence from fMRI. *Neuron, 33*, 301-311.

Bush, G., Vogt, B. A., Holmes, J., Dale, A. M., Greve, D., Jenike, M. A., et al. (2002). Dorsal anterior cingulate cortex: A role in reward-based decision making. *Proceedings of the National Academy of Sciences, 99*, 507-512.

Camille, N., Coricelli, G., Sallet, J., Pradat-Diehl, P., Duhamel, J.-R., & Sirigu, A. (2007). The involvement of the orbitofrontal cortex in the experience of regret. *Science, 304*, 1167-1170.

Carstensen, L. L. (2006). The influence of a sense of time on human development. *Science, 312*, 1913-1915.

Carter, C. S., Braver, T. S., Barch, D. M., Botvinick, M. M., Noll, D., & Cohen, J. D. (1998). Anterior cingulate cortex, error detection, and the online monitoring of performance. *Science, 280*, 747-749.

Cepeda, N. J., Kramer, A. F., & Gonzalez de Sather, C. M. (2001). Changes in executive control over the life span: Examination of task switching performance. *Developmental Psychology, 37*, 715-730.

Charles, S. T., Mather, M., & Carstensen, L. L. (2003). Aging and emotional memory: The forgettable nature of negative images for older adults. *Journal of Experimental Psychology: General, 132*, 310-324.

Clarke, A. M., Viney, L. L., Waterhouse, I. K., & Lord, J. (1974). Instrumental learning in preschool children as a function of type of task, type of reward, and some organismic variables. *Journal of Experimental Child Psychology, 17*, 1-17.

Cohen, J. (1973). Eta-squared and partial eta-squared in fixed factor ANOVA designs. *Educational and Psychological Measurement, 33*, 107-112.

Cohen, J. D., Braver, T. S., & Brown, J. W. (202). Computational perspectives on dopamine function in prefrontal cortex. *Current Opinion in Neurobiology, 12*, 223-229.

Cohen, M. X., Elger, C. E., & Ranganath, C. (2007). Reward expectation modulates feedback-related negativity and eeg spectra. *Neuroimage, 35*, 968-978.

Coles, M. G., Scheffers, M. K., & Holroyd, C. B. (1998). Berger's dream? The error-related negativity and modern cognitive psychophysiology. In H. Witte, U. Zwiener, B. Schack & A. Döring (Eds.), *Quantitative and topological EEG and MEG analysis* (pp. 96-102). Jena-Erlangen: Druckhaus Mayer Verlag.

Coles, M. G. H., Scheffers, M. K., & Holroyd, C. B. (2001). Why is there an ERN/Ne on correct trials? Response representations, stimulus-related components, and the theory of error-processing. *Biological Psychology, 56*, 173-189.

Cunningham, M. G., Bhattacharyya, S., & Benes, F. M. (2002). Amygdalo-cortical sprouting continues into early adulthood: Implications for the development of normal and abnormal function during adolescence. *Journal of Comparative Neurology, 453*, 116-130.

Cutmore, T. R. H., & Muckert, T. D. (1998). Event-related potentials can reveal differences between two decision-making groups. *Biological Psychology, 47*, 159-179.

Craik, F. I. M., & Bialystok, E. (2006). Cognition through the lifespan: Mechanisms of change. *Trends in Cognitive Sciences, 10*, 131-138.

Crone, E., Jennings, J. R., & van der Molen, M. W. (2004). Developmental change in feedback processing as reflected by phasic heart rate changes. *Developmental Psychology, 40*, 1228-1238.

Crone, E. A., Somsen, R. J. M., Zanolie, K., & Van der Molen, M. W. (2006). A heart rate analysis of developmental change in feedback processing and rule shifting from childhood to early adulthood. *Journal of Experimental Child Psychology, 95*, 99-116.

Crone, E. A., & Van der Molen, M. W. (2007). Development of decision making in school-aged children and adolescents: Evidence from heart rate and skin conductance analysis. *Child Development, 78*, 1288-1301.

deBruijn, E. R. A., Hulstijn, W., Verkes, R. J., Ruigt, G. S. F., & Sabbe, B. G. C. (2004). Drug-induced stimulation and suppression of action monitoring in healthy volunteers. *Psychopharmacology, 177*, 151-160.

Donchin, E., & Coles, M. G. (1988). Is the P300 component a manifestation of context updating? *Behavioral and Brain Sciences, 11*, 357-427.

Davies, P. L., Segalowitz, S. J., & Gavin, W. J. (2004). Development of response-monitoring ERPs in 7-to 25-year-olds. *Developmental Neuropsychology, 25*, 355-376.

Diamond, A. (1996). Evidence for the importance of dopamine for prefrontal cortex functions early in life. *Philosphical Transactions of the Royal Society (London) Series B, 351*, 1483-1494.

Diamond, A., Briand, L., Fossella, J., & Gehlbach, L. (2004). Genetic and neurochemical modulation of prefrontal cognitive functions in children. *American Journal of psychiatry, 161*, 125-132.

Donchin, E., & Coles, M. G. (1988). Is the p300 component a manifestation of context updating? *Behavioral and Brain Sciences, 11*, 357-427.

Durstewitz, D., Seamans, J. K., & Sejnowski, T. J. (2000). Neurocomputational models of working memory. *Nature Neuroscience supplement, 3*, 1184-1191.

Eppinger, B., Kray, J., Mecklinger, A., & John, O. (2007). Age differences in task switching and response monitoring: Evidence from ERPs. *Biological Psychology, 75*, 52-67.

Erixon-Lindroth, N., Farde, L., Robins Wahlin, T.-B., Sovago, J., Halldin, C., & Bäckman, L. (2005). The role of the striatal dopamine transporter in cognitive aging. Psychiatry research: Neuroimaging, 138, 1-12.

Falkenstein, M., Hohnsbein, J., & Hoormann, J. (1990). Effects of errors in choice reaction tasks and the erp under focused and divided attention. In C. H. M. Brunia, A. W. K. Gaillard & A. Kok (Eds.), *Psychophysiological Brain Research* (pp. 192-195). Tilburg: Tilburg University Press.

Falkenstein, M., Hohnsbein, J., & Hoormann, J. (1995). Event-related potential correlates of errors in reaction tasks. *Electroencephalography and Clinical Neurophysiology, 44*, 287-296.

Falkenstein, M., Hielscher, H., Dziobek, I., Schwarzenau, P., Hoormann, J., Sundermann, B., et al. (2001). Action monitoring, error detection, and the basal ganglia: An ERP study. *Neuroreport, 12*, 157-161.

Falkenstein, M., Hoormann, J., & Hohnsbein, J. (2001). Changes of error-related ERPs with age. *Experimental Brain Research, 138*, 258-262.

Finnigan, S., Humphreys, M. S., Dennis, S., & Geffen, G. (2002). ERP 'old/new' effects: Memory strength and decisional factors. *Neuropsychologia, 40*, 2288-2304.

Frank, M. J. (2005). Dynamic dopamine modulation in the basal ganglia: A neurocomputational account of cognitive deficits in medicated and nonmediacted parkinsonism. *Journal of Cognitive Neuroscience, 17*, 51-72.

Frank, M. J., Seeberger, L. C., & O'Reilly. (2004). By carrot or by stick: Cognitive reinforcement learning in parkinsonism. *Science, 306*, 1940-1943.

Frank, M. J., Woroch, B. S., & Curran, T. (2005). Error-related negativity predicts reinforcement learning and conflict biases. *Neuron, 47*, 495-501.

Gehring, J. W., Goss, B., Coles, M. G., Meyer, D. E., & Donchin, E. (1993). A neural system for error detection and compensation. *Psychological Science, 4*, 385-390.

Gehring, J. W., & Knight, R. T. (2000). Prefrontal cingulate interactions in action monitoring. *Nature Neuroscience, 3*, 516-520.

Gehring, J. W., & Willoughby, A. R. (2002). The medial frontal cortex and the rapid processing of monetary gains and losses. *Science, 295*, 2279-2282.

Geisser, S., & Greenhouse, S. W. (1958). An extension of box’s results on the use of the F-distribution in multivariate analysis. *Annals of Mathematical Statistics, 29*, 885-891.

Gibson, J., Krigolson, O. E., & Holroyd, C. B. (2006). Sensitivity of the feedback error-related negativity to reward probability. *Psychophysiology, 43*, S41-S42.

Goldman-Rakic, P., & Brown, R. (1982). Postnatal development of monoamine content and synthesis in the cerebral cortex of rhesus monkeys. *Developmental Brain Research, 4*, 339-349.

Gogtay, N., Giedd, J. N., Lusk, L., Hayashi, K. M., Greenstein, D., Vaituzis, C., et al. (2004). Dynamic mapping of human cortical development during childhood through early adulthood. *Proceedings of the National Academy of Sciences, 101, 8174-8179*.

Gratton, G., Coles, M. G., & Donchin, E. (1983). A new method for off-line removal of ocular artifact. *Electroencephalography and Clinical Neurophysiology, 55*, 468-484.

Hajcak, G., Moser, J. S., Holroyd, C. B., & Simons, R. F. (2007). It's worse than you thought: The feedback negativity and violations of reward predicition in gambling tasks. *Psychophysiology, 44*, 905-912.

Hampton, A. N., & O'Doherty, J. (2007). Decoding the neural substrates of reward-related decision making with functional MRI. *Proceedings of the National Academy of Sciences, 104*, 1377-1382.

Hartman, M., Bolton, E., & Fehnel, S. E. (2001). Accounting for age differences in the Wisconsin Card Sorting Test: Decreased working memory, not inflexibility. *Psychology and Aging, 16*, 385-399.

Hazy, T. E., Frank, M. J., & O'Reilly. (in press). Towards an executive without an homunculus: Computational models of the prefrontal cortex/basal ganglia system. *Philosphical Transactions of the Royal Society (London) Series B*.

Heimer, L., & van Hoesen, G. W. (1979). The ventral striatum. In I. Divac (Ed.), *The Neostriatum* (pp. 24-35). New York: Plenum Press.

Hogan, A. M., Vargha-Khadem, F., Kirkham, F. J., & Baldeweg, T. (2005). Maturation of action monitoring from adolescence to adulthood: An ERP study. *Developmental Science, 8*, 525-534.

Holroyd, C. B., Dien, J., & Coles, M. G. (1998). Error-related scalp potentials elicited by hand and foot movements: Evidence for an output-independent error processing system. *Neuroscience Letters, 242*, 65-68.

Holroyd, C. B., Praamstra, P., Plat, E., & Coles, M. G. (2002). Spared error-related potentials in mild to moderate parkinsons's disease. *Neuropsychologia, 40*, 2116-2124.

Holroyd, C. B., & Coles, M. G. H. (2002). The neural basis of human error processing: Reinforcement learning, dopamine, and the error-related negativity. *Psychological Review, 109*, 679-709.

Holroyd, C. B., Nieuwenhuis, S., Yeung, N., & Cohen, J. D. (2003). Errors in reward prediction are reflected in the event-related potential. *Neuroreport, 14*, 2481-2484.

Holroyd, C. B. (2004). A note on the oddball N200 and the feedback ERN. In M. Ullsperger & M. Falkenstein (Eds.), *Errors, conflicts, and the brain. Current Opinions on Performance Monitoring* (pp. 211-218). Leipzig: MPI of Cognitive Neuroscience.

Holroyd, C. B., Nieuwenhuis, S., Yeung, N., Nystrom, L., Mars, R. B., Coles, M. G. H., et al. (2004). Dorsal anterior cingulate cortex shows fMRI response to internal and external error signals. *Nature Neuroscience, 7*, 1-2.

Holroyd, C. B., Hajcak, G., & Larsen, J. T. (2006). The good, the bad, and the neutral: Electrophysiological responses to feedback stimuli. *Brain Research, 1105*, 93-101.

Ito, T. A., Larsen, J. T., Smith, N. K., & Cacioppo, J. T. (1998). Negative information weighs more heavily on the brain: The negativity bias in evaluative categorizations. *Journal of Personality and Social Psychology, 75*, 887-900.

Ito, S., Stuphorn, V., J.W., B., & Schall, J. D. (2003). Performance monitoring by the anterior cingulate cortex during saccade countermanding. *Science, 302*, 120-122.

Johnson, R. J. (1986). A triarchic model of P300 amplitude. *Psychophysiology, 23*, 367-384.

Johnson, M. H., & Munakata, Y. (2005). Processes of change in brain and cognitive development. *Trends in Cognitive Sciences, 9*, 152-158.

Kim, E. Y., Iwaki, N., Imashioya, H., Uno, H., & Fujita, T. (2007). Error-related negativity in a visual go/no-go task: Children vs. Adults. Developmental Neuropsychology, 31, 181-191.

Kisely, M. A., Wood, S., & Burrows, C. L. (2007). Looking at the sunny side of life. Age-related change in an event-related potential measure of the negativity bias. *Psychological Science, 18*, 838-843.

Kray, J., Eber, J., & Lindenberger, U. (2004). Age differences in executive functioning across the life-span: The role of verbalization in task-preparation. *Acta Psychologica, 115*, 143-165.

Kray, J., Eppinger, B., & Mecklinger, A. (2005). Age differences in attentional control: An event-related potential approach. *Psychophysiology, 42*, 407-416.

Ladouceur, C. D., Dahl, R. E., & Carter, C. S. (2004). ERP correlates of action monitoring in adolescence. *Annals of the New York Academy of Sciences, 1021*, 329-336.

Lambe, E. K., Krimer, L. S., & Goldman-Rakic, P. (2000). Differential postnatal developemnt of catecholamine and serotonin inputs to identified neurons in prefrontal cortex of rhesus monkey. *Journal of Neuroscience, 20*, 8780-8787.

Lang, P. J., Bradley, M. M., & Cuthbert, B. N. (1998). Emotion and motivation: Measuring afective perception. *Journal of Clinical Neurophysiology, 15*, 397-408.

Larkin, G. R., Gibbs, S. E. B., Khanna, K., Nielsen, L., Carstensen, L. L., & Knutson, B. (2007). Anticipation of monetary gain but not loss in healthy older adults. *Nature Neuroscience, 10*, 787-791.

Lehrl, S. (1977). Mehrfachwahl-Wortschatz-Test B (MWT-B). [Multiple-Choice Knowledge Test-B (MWT-B)]. Erlangen, Germany: Straube.

Leth-Steensen, C., King Elbaz, Z., & Douglas, V. I. (2000). Mean response time, variability and skew in the responding of ADHD children: A response time distributional approach. *Acta Psychologica, 104*, 167-190.

Li, S.-C., Lindenberger, U., & Sikström, S. (2001). Aging cognition: From neuromodulation to representation. *Trends in Cognitive Sciences, 5*, 479-486.

Light, L. L., Chung, C., Pendergrass, R., & van Ocker, J. C. (2006). Effects of repetition and response deadline on item recognition in young and older adults. *Memory & Cognition, 34*, 335-343.

Liotti, M., Pliszka, S. R., Perez, R., Kothmann, D., & Woldorff, M. G. (2005). Abnormal brain activity related to performance monitoring and error detection in children with ADHD. *Cortex, 41*, 377-388.

MacDonald III A.W., Cohen, J. D., Stenger, V. A., & Carter, C. S. (2000). Dissociating the role of the dorsolateral prefrontal and anterior cingulate cortex in cognitive control. *Science, 288*, 1835-1838.

Mather, M., & Carstensen, L. L. (2005). Aging and motivated cognition: The positivity effect in attention and memory. *Trends in Cognitive Sciences, 9*, 496-502.

Mather, M., & Johnson, M. K. (2000). Choice-supportive source monitoring: Do our decisions seem better to us as we age? *Psychology and Aging, 15*, 596-606.

Mathewson, K. J., Dywan, J., & Segalowitz, S. J. (2005). Brain bases of error-related ERPs as influenced by age and task. *Biological Psychology, 70*, 88-104.

Matsumoto, K., Suzuki, W., & Tanaka, K. (2003). Neuronal correlates of goal-based motor selection in the prefrontal cortex. *Science, 301*(229-232).

Mell, T., Heekeren, H. R., Marschner, A., Wartenburger, I., Villringer, A., & Reischies, F. M. (2005). Effect of aging on stimulus-reward association learning. *Neuropsychologia, 43*, 554-563.

Milham, M. P., Banich, M. T., Claus, E. D., & Cohen, N. J. (2003). Practice-related effects demonstrate complementary roles of anterior cingulate and prefrontal cortices in attentional control. *Neuroimage, 18*, 483-493.

Mirenowicz, J., & Schultz, W. (1996). Preferential activation of midbrain dopamine neurons by appetitive rather than aversive stimuli. *Nature, 379*, 449-351.

Miller, E. K. (2000). The prefrontal cortex and cognitive control. *Nature Neuroscience Reviews, 1*, 59-65.

Miller, E. K., & Cohen, J. D. (2001). An integrative theory of prefrontal cortex function. *Annual Review of Neuroscience, 24*, 167-202.

Miltner, W. H., Braun, C. H., & Coles, M. G. H. (1997). Event-related brain potentials following incorrect feedback in a time-estimation task: Evidence for a 'generic' neural system for error detection. *Journal of Cognitive Neuroscience, 9*, 788-798.

Miltner, W. H., Lemke, U., Weiss, T., Holroyd, C. B., Scheffers, M. K., & Coles, M. G. H. (2003). Implementation of error processing in the human anterior cingulate cortex: A source analysis of the magnetic equivalent of the error-related negativity. *Biological Psychology, 64*, 157-166.

Montague, P. R., Hyman, S. E., & Cohen, J. D. (2004). Computational roles for dopamine in behavioral control. *Nature, 431*, 760-767.

Moran, J. D., & McCullers, J. C. (1979). Reward and number of choices in children's probability learning: An attempt to reconcile conflicting findings. *Journal of Experimental Child Psychology, 27*, 527-532.

Nieuwenhuis, S., Ridderinkhof, K. R., Blom, J., Band, G. P., & Kok, A. (2001). Error-related brain potentials are differentially related to awareness of response errors: Evidence from an antisaccade task. *Psychophysiology, 38*, 752-760.

Nieuwenhuis, S., Ridderinkhof, K. R., Talsma, D., Coles, M. G. H., Holroyd, C. B., Kok, A., et al. (2002). A computational account of altered error processing in older age: Dopamine and the error-related negativity. *Cognitive, Affective and Behavioral Neuroscience, 2*, 19-36.

Nieuwenhuis, S., Yeung, N., Holroyd, C. B., Schurger, A., & Cohen, J. D. (2004). Sensitivity of electrophysiological activity from medial frontal cortex to utilitarian and performance feedback. *Cerebral Cortex, 14*, 741-747.

Nieuwenhuis, S., Slagter, H. A., von Geusau, N. J. A., Heslenfeld, D., & Holroyd, C. B. (2005). Knowing good from bad: Differential activation of human cortical areas by positive and negative outcomes. *European Journal of Neuroscience, 21*(3161-3168).

Niv, Y., Daw, N. D., Joel, D., & Dayan, P. (2007). Tonic dopamine: Opportunity costs and the control of response vigor. *Psychopharmacology, 191*, 507-520.

O'Doherty, J., Kringelbach, M. L., Rolls, E. T., Hornak, J., & Andrews, C. (2001). Abstract reward and punishment representations in the human orbitofrontal cortex. *Nature Neuroscience, 4*, 95-102.

O'Doherty, J., Critchley, H., Deichmann, R., & Dolan, R. J. (2003). Dissociating valence of outcome from behavioral control in human orbital and ventral prefrontal cortices. *Journal of Neuroscience, 23*, 7931-7939.

O'Doherty, J. O., Dayan, P., Schultz, J., Deichmann, R., Friston, K., & Dolan, R. J. (2004). Dissociable roles of the ventral and dorsal striatum in instrumental conditioning. *Science, 304*, 452-454.

Offenbach, S. I. (1973). The effect of partially valid cues on children's discrimination learning. *Journal of Experimental Child Psychology, 16*, 91-97.

Overbeek, T. J. M., Nieuwenhuis, S., & Ridderinkhof, K. R. (2005). Dissociable components of error processing: On the functional significance of the Pe vis-à-vis the ERN/Ne. *Journal of Psychophysiology, 19*, 319-329.

Packard, M. G., & Knowlton, B. J. (2002). Learning and memory functions of the basal ganglia. *Annual Review of Neuroscience, 25*, 563-593.

Pailing, P. E., & Segalowitz, S. J. (2004). The effects of uncertainty in error monitoring on associated ERPs. *Brain and Cognition, 56*, 215-233.

Paus, T. (2001). Primate anterior cingulate cortex: Where motor control, drive and cognition interface. *Nature Neuroscience Reviews, 2*, 417-424.

Pavlov, I. (1898). Die Arbeit der Verdauungsdrüsen. Wiesbaden: J.F.Bergmann.

Perrin, P., Pernier, J., Bertrand, O., & Echallier, J. F. (1989). Spherical splines for scalp potential and current density mapping. *Electroencephalography and Clinical Neurophysiology, 72*, 184-187.

Piaget, J. (1969). The Mechanisms of Perception. London: Rutledge & Kegan Paul.

Picard, N., & Strick, P. L. (1996). Motor areas of the medial wall: A review of their location and functional activation. *Cerebral Cortex, 6*, 342-353.

Pietschmann, M., Endrass, T., & Kathmann, N. (2007). *Changes of performance monitoring in younger and older adults with learning.* Poster presented at the Annual meeting of the Society for Cognitive Neuroscience, New York.

Polich, J., & Criado, J. R. (2006). Neuropsychology and neuropharmacology of the P3a and P3b. *International Journal of Psychophysiology, 60*, 172-185.

Potts, G. F., Martin, L. E., Burton, P., & Montague, P. R. (2006). When things are better of worse than expected: The medial frontal cortex and the allocation of processing resources. *Journal of Cognitive Neuroscience, 18*, 1112-1119.

Ramscar, M., & Gitcho, N. (2007). Developmental change and the nature of learning in childhood. *Trends in Cognitive Sciences, 11*, 274-279.

Rescorla, R. A., & Wagner, A. R. (1972). A theory of pavlovian conditioning: Variations in the effectiveness of reinforcement and nonreinforcement. In A. H. Black & W. F. Prokasy (Eds.), *Classical conditioning II* (pp. 64-99): Appleton-Century-Croft.

Resnick, S. M., Lamar, M., & Driscoll, I. (in press). Vulnerability of the orbitofrontal cortex to age-associated structural and functional brain changes. *Annals of the New York Academy of Sciences*.

Rhodes, M. G. (2004). Age-related differences in perfromance on the wisconsin card sorting test: A meta-analytic review. *Psychology and Aging, 19*, 482-494.

Ridderinkhof, K. R., Span, M. M., & van der Molen, M. W. (2002). Perseverative behavior and adaptive control in older adults: Performance monitoring, rule induction, and set shifting. *Brain and Cognition, 49*, 382-401.

Rinkenauer, G., Osman, A., Ulrich, R., Müller-Gethmann, H., & Mattes, S. (2004). On the locus of speed-accuracy trade-off in reaction time: Inferences from the lateralized readiness potential. *Journal of Experimental Psychology: General, 133*, 261-282.

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.

Rolls, E. T. (2000). Memory systems in the brain. *Annual Review of Psychology, 51*, 599-630.

Rolls, E. T. (2000). The orbitofrontal cortex and reward. *Cerebral Cortex, 10*, 284-294.

Rubia, K., Smith, A. B., Woolley, J., Nosarti, C., Heyman, I., Taylor, E., et al. (2006). Progressive increase of frontostriatal brain activation from childhood to adulthood during event-related tasks of cognitive control. *Human Brain Mapping, 27*, 973-993.

Rueda, M. R., Rothbart, M. K., McCandliss, B. D., Saccomanno, L., & Posner, M. I. (2005). Training, maturation, and genetic influences on the development of executive attention. *Proceedings of the National Academy of Sciences, 11*, 14931-14936.

Rumelhart, D. & Norman, D. (1981). Analogical processes in learning. In J.R. Anderson (ed.), Cognitive Skills and their Acquisition. Hillsdale, NJ: Erlbaum.

Rushworth, M. F. S., Walton, M. E., Kennerly, S. W., & Bannerman, D. M. (2004). Action sets and decisions in the medial frontal cortex. *Trends in Cognitive Sciences, 8*, 410-417.

Salthouse, T. A. (1996). The processing-speed theory of adult age differences in cognition. *Psychological Review, 103*, 403-428.

Salthouse, T. A. (2000). Aging and measures of processing speed. *Biological Psychology, 54*, 35-54.

Samii, A., Nutt, J. G., & Ransom, B. R. (2004). Parkinson's disease. *Lancet, 363*, 1783-1793.

Santesso, D. L., Segalowitz, S. J., & Schmidt, L. A. (2006). Error-related electrocortical responses in 10-year-old children and young adults. *Developmental Science, 9*, 473-481.

Scheffers, M. K., & Coles, M. G. (2000). Performance monitoring in a confusing world: Error-related brain activity, judgments of response accuracy, and types of errors. *Journal of experimental psychology: Human perception and performance, 26*, 141-151.

Schultz, W., Dayan, P., & Montague, P. R. (1997). A neural substrate of prediction and reward. *Science, 275*, 1593-1599.

Schultz, W., Tremblay, L., & Hollerman, J. R. (2000). Reward processing in primate orbitofrontal cortex and basal ganglia. *Cerebral Cortex, 10*, 272-283.

Schultz, W. (2002). Getting formal with dopamine and reward. *Neuron, 36*, 241-263.

Schultz, W. (2006). Behavioral theories and the neurophysiology of reward. *Annual Review of Psychology, 57*, 87-115.

Schultz, W. (2007). Behavioral dopamine signals. *Trends in Neurosciences, 30*, 203-210.

Schultz, W. (2007). Multiple dopamine functions at different time courses. *Annual Review of Neuroscience, 30*, 259-288.

Seymore, B., O'Doherty, J., Dayan, P., Koltzenburg, M., Jones, A. K., Dolan, R. J., et al. (2004). Temporal difference models describe higher-order learning in humans. *Nature, 429*, 664-667.

Shidara, M., & Richmond, B. J. (2002). Anterior cingulate: Single neuronal signals related to degree of reward expectancy. *Science, 296*, 1709-1711.

Shima, K., & Tanji, J. (1998). Role for cingulate motor area cells in voluntary movement selection based on reward. *Science, 282*, 1335-1338.

Skinner, B. F. (1938). The behavior of organisms: An experimental analysis. New York: Appleton-Century-Crofts.

Smith, E. E., & Jonides, J. (1999). Storage and executive processes in the frontal lobes. *Science, 283*, 1657-1661.

Snodgrass, J. G., & Vanderwart, M. (1980). A standardized set of 260 pictures: Norms for name agreement, image agreement, familiarity, and visual complexity. *Journal of Experimental Psychology: Human Learning & Memory, 6*, 174-215.

Sowell, E. R., Thompson, P. M., Leonard, C. M., Welcome, S. E., Kan, E., & Toga, A. W. (2004). Longitudinal mapping of cortical thickness and brain growth in normal children. *Journal of Neuroscience, 24*, 8223-8231.

Somsen, R. J. M., van der Molen, M. W., Jennings, J. R., & van Beek, B. (2000). Wisconsin card sorting in adolescents: Analysis of performance, response times and heart rate. *Acta Psychologica, 104*, 227-257.

Sutton, R. S., & Barto, A. G. (1990). Time-derivate models of pavlovian reinforcement. In M. Gabriel & J. Moore (Eds.), *Learning and compuational neuroscience: Foundations of adaptive networks* (pp. 497-537): MIT Press.

Swick, D., & Turken, A. U. (2002). Dissociation between conflict detection and error monitoring in the human anterior cingulate cortex. *Proceedings of the National Academy of Sciences, 99*, 16354-16359.

Themanson, J. R., Hillman, C. H., & Curtin, J. J. (2006). Age and physical activity influences on action monitoring during task switching. *Neurobiology of Aging, 27*, 1335-1345.

Thorndike, E. L. (1911). *Animal intelligence*. New York: Macmillan.

Ungless, M. (2004). Dopamine: The salient issue. *Trends in Neurosciences, 27*, 702-706.

van Meel, C. S., Oosterlaan, J., Heslenfeld, D. J., & Seargent, J., A. (2005). Telling good from bad news: ADHD differentially affects the processing of positive and negative feedback during guessing. *Neuropsychologia, 43*, 1946-1954.

Van Veen, V., Cohen, J. D., Botvinick, M. M., Stenger, V. A., & Carter, C. S. (2001). Anterior cingulate cortex, conflict monitoring, and levels of processing. *Neuroimage, 14*, 1302-1308.

van Veen, V., & Carter, C. S. (2002). The timing of action-monitoring processes in the anterior cingulate cortex. *Journal of Cognitive Neuroscience, 14*, 593-602.

Verhaeghen, P., & Salthouse, T. A. (1997). Meta-analyses of age-cognition relations in adulthood: Estimates of linear and nonlinear age effects and structural models. *Psychological Bulletin, 122*, 231-249.

Vidal, F., Burle, B., Bonnet, M., Grapperon, J., & Hasbroucq, T. (2003). Error negativity on correct trials: A reexamination of available data. *Biological Psychology, 64*, 265-282.

Vidal, F., Hasbroucq, T., Grapperon, J., & Bonnet, M. (2000). Is the 'error negativity' specific to errors? *Biological Psychology, 51*, 109-128.

Volkow, N. D., Gur, R. C., Wang, G.-J., Fowler, J. S., Mpberg, P. J., Ding, Y.-S., et al. (1998). Association between decline in brain dopamine activity with age and cognitive and motor impairment in healthy individuals. *American Journal of Psychiatry, 155*, 344-349.

Volkow, N. D., Logan, J., Fowler, J. S., Wang, G.-J., Gur, R. C., Wong, C., et al. (2000). Association between age-related decline in brain dopamine activity and impairment in frontal and cingulate metabolism. *American Journal of Psychiatry, 157*, 75-80.

Vidal, F., Burle, B., Bonnet, M., Grapperon, J., & Hasbroucq, T. (2003). Error negativity on correct trials: A reexamination of available data. *Biological Psychology, 64*, 265-282.

Vidal, F., Hasbroucq, T., Grapperon, J., & Bonnet, M. (2000). Is the 'error negativity' specific to errors? *Biological Psychology, 51*, 109-128.

Waelti, P., Dickinson, A., & Schultz, W. (2001). Dopamine responses comply with basic assumptions of learning theory. *Nature, 412*, 43-48.

Wallis, J. D. (2007). Oorbitofrontal cortex and its contribution to decision making. *Annual Review of Neuroscience, 30*, 31-56.

Wang, Y., Chan, G. L. Y., Holden, J. E., Dobko, T., Mak, E., Schulzer, M., et al. (1998). Age-dependent decline of dopamine d1 receptors in human brain: A PET study. *Synapse, 30*, 56-61.

Watson, J. B. (1913). Psychology as the behaviorist views it. *Psychological Review, 20*, 158-177.

Wechsler, W. (1982). *Handanweisung zum Hamburg-Wechsler-Intelligenztest für Erwachsene (HAWIE). [Manual for the Hamburg-Wechsler Intelligence Test for Adults]*. Bern, Switzerland: Huber.

Weickert, C. S., Webster, M. J., Gondipalli, P., Rothmond, D., Fatula, R. J., Herman, M. M., et al. (2007). Postantal alterations in dopaminergic markers in the human prefrontal cortex. *Neuroscience, 144*, 1109-1119.

West, R. L. (1996). An application of prefrontal cortex functioning theory to cognitive aging. *Psychological Bulletin, 120*, 272-292.

West, R. (2004). The effects of aging on controlled attention and conflict processing in the Stroop task. *Journal of Cognitive Neuroscience, 16*, 103-113.

Wiersema, J. R., van der Meere, J. J., & Roeyers, H. (2007). Developmental changes in error monitoring: An event-related potential study. *Neuropsychologia, 45*, 1649-1657.

Williams, B. R., Hultsch, D. F., Strauss, E. H., Hunter, M. A., & Tannock, R. (2005). Inconsistency in reaction time across the life span. *Neuropsychology, 19*, 88-96.

Williams, Z. M., Bush, G., Rauch, S. L., Cosgrove, G. R., & Eskander, E. N. (2004). Human anterior cingulate neurons and the integration of monetary reward with motor responses. *Nature Neuroscience, 7*, 1370-1375.

Wise, R. A. (2004). Dopamine, learning and motivation. *Nature Reviews Neuroscience, 5*, 1-11.

Wise, R. A., Spindler, J., DeWit, H., & Gerber, G. J. (1978). Neuroleptic induced "anhedonia" hypothesis in rats: Pimozide blocks reward quality of food. *Science, 201*, 655-656.

Witte, K. L., & Grossman, E. E. (1971). The effects of reward and punishment upon children's attention, motivation, and discrimination learning. *Child Development, 42*, 537-542.

Wood, S., & Kisely, M. A. (2006). The negativity bias is eliminated in older adults: Age-related reduction in event-related brain potentials asscoiated with evaluative categorization. *Psychology and Aging, 21*, 815-820.

Yeung, N., Botvinick, M. M., & Cohen, J. D. (2004). The neural basis of error detection: Conflict monitoring and the error-related negativity. *Psychological Review, 111*, 931-959.

Yeung, N., & Sanfey, A. G. (2004). Independent coding of reward magnitude and valence in the human brain. *Journal of Neuroscience, 24*, 6258-6264.

Zirnheld, P. J., Carroll, C. A., Kieffaber, P. D., O'Donnell, B. F., Shekhar, A., & Hetrick, W. P. (2004). Haloperidol impairs learning and error-related negativity in humans. *Journal of Cognitive Neuroscience, 16*, 1098-1112.

Abbreviations

ACC	anterior cingulate gyrus
ANOVA	analysis of variance
CCZ	caudal cingulate zone
EEG	electroencephalogram
ERP	event-related potential
ERN	error-related negativtity
fMRI	functional magnetic resonance imaging
FRN	feedback-related negativity
MDS	mesencephalic dopamine system
MEG	magnetoencephalography
ms	milliseconds
Ne	error negativity
OFC	orbitofrontal cortex
PFC	prefrontal cortex
pre-SMA	pre-supplementary motor area
RCZ	rostral cingulate zone
RT	reaction time
SST	socioemotional selectivity theory

Annotation

Some of the data reported in the present thesis are also included in the following manuscripts:

Eppinger, B., Kray, J., Mock, B., & Mecklinger, A. (in press). Better or worse than expected? Aging, learning, and the ERN. *Neuropsychologia*.

Eppinger, B., Mock, B., & Kray, J. (under review). Insights into the development of reinforcement learning: Evidence from ERPs. *Developmental Psychology*.

Danksagung

Mein Dank gilt Prof. Dr. J. Kray und Prof. Dr. Axel Mecklinger, die mich vorbildlich betreut und mir viele Freiheiten gelassen haben. Ohne dieses Vertrauen wäre diese Arbeit nicht möglich gewesen.

Des Weiteren möchte ich mich bei den studentischen Hilfskräften Barbara Mock, Philippe Gainsbourger und Michael Herbert bedanken, die einen großen Anteil an dieser Arbeit haben. Darüber hinaus gilt mein Dank meinen (ehemaligen) Kollegen, Julia Karbach, Denise Veith, Christian Doeller, und Oliver John, mit denen ich eine schöne Zeit verbracht habe und die mir stets geholfen und mich unterstützt haben.

Meinen Eltern und Lena danke ich für die Unterstützung in den letzten Monaten. Letztlich gilt mein größter Dank Dir, Geertje.