THE IMPACT OF PREDICTION ERRORS ON PERCEPTION AND LEARNING:

A SYSTEMS APPROACH

Esther De Loof

Promotor: Prof. Dr. Tom Verguts
Copromotor: Dr. Filip Van Opstal

UNIVERSITEIT
GENT

FACULTEIT PSYCHOLOGIE EN
PEDAGOGISCHE WETENSCHAPPEN

THE IMPACT OF PREDICTION ERRORS ON PERCEPTION AND LEARNING:

A SYSTEMS APPROACH

Esther De Loof

Promotor: Prof. Dr. Tom Verguts
Copromotor: Dr. Filip Van Opstal

Proefschrift ingediend tot het behalen van de academische graad
van Doctor in de Psychologie

2016

The best way to predict the future is to invent it.

Alan Kay

CONTENTS

DATA STORAGE FACT SHEETS   **189**

ACKNOWLEDGEMENTS

Looking back at the past four years, I realize we have come a long way. The initial goal of demonstrating the beneficial effect of reward prediction errors on learning proved a hard nut to crack. Our research took a meandering path as the initial plan of performing an fMRI study morphed into an eye tracking study, which eventually led to the time-frequency EEG study reported in this dissertation. In the meantime, we also continued to explore the effect of prediction errors on visual awareness, building upon a line of research started during my master's thesis. Eventually, we succeeded to reach the double goal set at the start of the project even though it proved to be a tough hill to climb. Luckily, I wasn't alone on this journey and my deepest gratitude goes out to everyone who guided and inspired me along the way.

My first thoughts of gratitude go to my wonderful supervisor and the leader of our team, **Tom Verguts**. Choosing to do a PhD under your guidance is definitely among the best decisions of my life. You taught us to look beyond the hype of the day and to critically asses any theory on its explanatory and predictive merits. The elegance of a theory, the clarity of an idea, the insights gained from discovering a simple underlying mechanism in a complex pattern of findings; it all becomes that more exciting when reflecting on our work with you. I'm still impressed by the width and depth of your theoretical knowledge, which you combine with great technical insight and skills. On top of all that, I learned a great deal from the exemplary way in which you tackle challenges and the admirable care you have for the people you work with. You truly are a beautiful person and it

has been an honor to work with you.

A profound word of thanks also goes to my co-supervisor **Filip Van Opstal**. Dear Filip, we have come a long way since we first started working together in 2010. From building our first CFS setup to publishing our first paper, you have helped me in taking those first steps and have continued to actively guide me during my PhD. You're the kind of person who gets goose bumps from a great idea and your bold ideas and incredible enthusiasm make you a pleasure to work with. I greatly admire your wit and resilience in life. Undoubtedly the best is yet to come for you and I wish you all the best in your future adventures and thrilling projects (just be carefull with those bees).

The complementary combination of Tom and Filip was further supported by the members of my guidance committee. I would like to thank **Florian Waszak**, **Massimo Silvetti** and **Daniele Marinazzo** for supporting the projects with their thoughtful comments during our annual meetings. They provided me with excellent feedback to reevaluate the road taken so far and to adjust our future course.

A special mention goes out to **Clio Janssens**, who has been my office mate and partner in crime for four years now. I couldn't have wished for a better friend to share my thoughts and ideas with on a daily basis. Clio, you are definitely wise beyond your years, calm and patient under any circumstances but always with a great sense of humor. I can count on you for fun and laughter every day and for putting everything in perspective. You also taught me a lot about EEG, provided me with excellent input when designing experiments and were of great help in interpreting the data and planning out the analyses. I wish we could have collaborated more on

projects because I'm convinced we form a great team. I'll greatly miss you if our paths diverge after we finish our PhDs. All the best of luck in your future endeavors and see you at our next under the radar Game of Thrones in depth quantum analysis!

Next in line is my PhD godmother (and soon-to-be proud mother of two), **Eliana Vassena**. Dear Eliana, thank you so much for all your care and good advice over the past years. You really went out of your way to guide Clio and me in our research trajectory. I'm very grateful for your helpful comments during lab meetings and your witty stories during random encounters. The department won't be the same when you leave and you'll be sorely missed by all of us.

If Eliana is my godmother, than surely **Massimo Silvetti** has been my PhD godfather. Massimo, you have not only been a mentor in science but also in life. I remember you quickly understood what makes me tick and your counseling has been of great value to me. Thank you for all your kind attention and care over the years, I'm definitely greatly indebted to you.

I'm very grateful to all members at our department for creating a wonderful working environment. It was a pleasure to work with you, hear about your research at the lunch seminars or over drinks, and getting to know you personally. **Elger Abrahamse**, life just isn't what it used to be without you in the office next to ours. Thank you for always being there to share ideas and for your advice throughout the years. **Wim Fias** and **Muhammet Ikbal Sahan**, thanks for exchanging your ideas about our research and science in general. **James Deraeve**, I greatly enjoyed our late-night trips to the cinema and trailer analyses, lots more to come! Thanks to **Marc Brysbaert** for his time and wise guidance. **Michael Stevens** also

CHAPTER 1
INTRODUCTION

THE ADAPTIVE RELEVANCE OF PREDICTION ERRORS

Throughout our lives we are shaped by our experiences and by the decisions we made in the past. Our prior experiences color how we view the world around us and direct where our curiosity will lead us next. Yet, we live our everyday life without realizing how our view of the world is shaped by the entirety of our accumulated experiences that keeps evolving over time. This is not only a semi-philosophical way of looking at our everyday lives; according to the *predictive coding* framework it is also a good description of how the brain works (Friston & Kiebel, 2009; Rao & Ballard, 1999; Summerfield et al., 2006). Over time, we build up an understanding of the world that subsequently guides our perception and our actions.

Each moment, the brain is tasked with processing a rich stream of (sensory) information and determining our next actions. Accomplishing this daunting task requires a large investment of energy. In order to function as efficiently as possible, the brain has therefore been hypothesized to construct an internal model of the world based on our previous experiences. This internal model allows us to infer what the future will look like. Next, the resulting predictions about the nearby future are cascaded down the processing hierarchy in the brain, e.g. from higher processing stages to the primary sensory cortex. There, the top-down predictions explain away most of the incoming (sensory) input. What remains to be processed is mainly limited to the unpredicted events. Thus, the brain can proactively anticipate

its (sensory) input before its actual occurrence, simplifying the processing requirements.

The information that is not accounted for by the top-down predictions then enters the bottom-up processing stream. These bottom-up prediction errors (PEs) not only inform higher processing stages about the outside world, they also function as a teaching signal. Indeed, PEs signal a learning opportunity as the internal model was unable to predict the actual events and should be adjusted. Thus, by repeated adjustments based on the PE feedback, the brain learns and thereby optimizes its processing efficiency.

Taken together, the predictive coding framework entails that on the short term top-down predictions steer perception. In addition, the bottom-up PEs serve as a teaching signal and hence result in learning on the long term. In the current dissertation we will explore both of these phenomena: the short term effect of PEs on perception and the long term effect of PEs on learning. However, before moving on to the specific research questions we first discuss our general research approach.

Arguably, although the predictive coding framework mainly endeavors to characterize the fundamental principles and organization of the brain, it comes at the cost of limited specificity. In particular, it attempts to cover a wide variety of processes, encompassing the role of PEs on the scale of a single cortical column, as well as on the scale of a network that spans the entire cortex. Therefore, although the broad framework has been sketched, more specific mechanisms are needed to fill in how PEs guide perception and learning.

Therefore, we aimed to explore more deeply how PEs determine perception and learning by using a systems approach (hence the subtitle of

this dissertation). In this systems approach, we drew upon formal models of how neural circuits perform a number of computations. In the first part of the dissertation, a formal model is applied to investigate how a perceptual decision can be influenced by prior information. In the second part, we formally quantified RPEs during a declarative learning task to probe their effect on learning. Below we will introduce both formal approaches separately and conclude with an overview of the dissertation outline.

THE SHORT TERM EFFECT OF PES ON PERCEPTION

We start by turning our attention toward the short term effect of PEs on perception. More specifically, we will investigate the effect of PEs on vision. Visual processing lends itself ideally for exploring the implications of the predictive coding framework. First, the visual processing areas and pathways belong to the most extensively studied structures in the brain. Also, the strong hierarchical organization of the visual processing stream is ideally fit to test the feedforward and feedback mechanisms described by the predictive coding framework. It therefore comes as no surprise that visual processing has been among the first empirical research lines established to test the effect of PEs on perception (Enns & Lleras, 2008; Rao & Ballard, 1999; Summerfield et al., 2006).

Exposure to prior information has long been known to influence subsequent perception. For example in priming studies, information that was presented subliminally (either through masking or by presenting it for only a few milliseconds) has been demonstrated to influence the reaction to subsequent stimuli (for a review, see Neely, 1991). Matching or related

stimuli are typically processed faster and more accurately. However, it is often difficult to differentiate between the impact of prior expectations on perception itself and the impact these prior expectations have on subsequent processing stages such as decision making and response execution.

Interestingly, research on visual awareness has offered a number of paradigms that are specifically designed to focus on the earliest stages of visual processing. One of these paradigms of particular interest for examining the role of PEs on perception, is binocular rivalry (Clark, 2013; Dayan, 1998; Hohwy, Roepstorff, & Friston, 2008). As the name suggests, in this paradigm rivalry is created between both eyes by presenting each eye with a different picture. For example, a face could be shown to the right eye, whereas a house could be shown to the left eye. Instead of blending both pictures into one transparently overlaid image, perception alternates between the two pictures. Importantly, this bistable visual experience could be the result of the predictive mechanisms that try to explain away the conflicting input. When for instance the house is the dominant percept, a top-down prediction about the house is projected along the processing hierarchy toward the earliest processing stages. There, the top-down predictions explain away the perception of the house presented to the left eye. However, the face picture presented to the right eye will give rise to a bottom-up PE signal, accumulating over time. Once these accumulated PEs can no longer be ignored, the predictive mechanism switches to the face percept in an attempt to minimize PEs. Thus, the face becomes the dominant percept which is predicted in a top-down fashion, starting another cycle in the alternating perception sequence.

Taken together, these examples illustrate how research on visual awareness has already sparked a theoretical and empirical interest in the role of PEs in perception (Hohwy, 2012; Rauss, Schwartz, & Pourtois, 2011; Seth, Suzuki, & Critchley, 2011). We will build upon this research in the current dissertation, but before we introduce the current research question we will first discuss the role of PEs in learning.

Probing the effect of PEs on perception through the drift diffusion model

In order to determine how PEs influence perception on the short term, we used a formal model from the decision making literature: the drift diffusion model (DDM; Bogacz, 2007; Gold & Shadlen, 2007; Ratcliff & McKoon, 2008; Ratcliff & Rouder, 1998; Ratcliff, 1985). This model is typically applied to a two-alternative forced choice task in which a (visually) presented stimulus must be assigned to one of two categories. Participants are requested to respond as fast and as accurate as possible. This instruction creates a speed-accuracy trade-off: either a decision is made quickly and thus with higher error likelihood or the participant alternatively chooses to increase accuracy but at the cost of a less speeded response.

Based on behavioral choices and the response times of these choices, the DDM disentangles how speed and accuracy are weighed against each other. To illustrate the DDM in analogy to our previous example, suppose that the picture of a house or a face is presented and that participants are asked to categorize the picture accordingly. From the onset of the stimulus presentation (e.g., a house), evidence in favor of each option accumulates over time until a decision threshold is reached. When the accumulated evidence reaches the decision threshold for the detection of a house first, the

participant will correctly report seeing a house. Alternatively, if the evidence accumulation accidently reaches the threshold associated with the face first, the picture will incorrectly be categorized as a face.

The DDM contains a number of parameters that detail how the speed-accuracy trade-off can be solved in different ways. Generally speaking, a (perceptual) decision can be reached faster by either lowering the decision threshold or by increasing the accumulation speed (i.e., the processing efficiency). Each of these adjustments has a distinct influence on the response times for the correct and incorrect categorizations. Thus, the model enables us to infer the parameter settings based on merely the reaction times and behavioral choices.

In previous research, the DDM has been succesfully to differentiate under what circumstances perception is influenced by either altered processing efficiency or threshold setting (Mulder, Wagenmakers, Ratcliff, Boekel, & Forstmann, 2012; Summerfield & de Lange, 2014; Summerfield & Egner, 2009). Thus, we will use the DDM to probe whether PEs influence visual awareness by altering the threshold setting or by changing the processing efficiency.

THE LONG TERM EFFECT OF PES ON LEARNING

A second component of this dissertation is the long term effect of PEs on learning. Based on our internal model of the world, the brain learns by comparing the rewarding value of the outcome it predicted to the actual reward feedback provided by the environment. This reward feedback is not only evaluated on its predictability (i.e., the size of the reward prediction

error; RPE), but also on its valence (i.e., is the outcome more positive or negative than expected). In that sense, any experience can be categorized as being more rewarding (positive RPE) or less rewarding than expected (negative RPE). These positive and negative RPEs subsequently serve as a teaching signal to optimize the accuracy of future predictions. Hence, research on learning has not just focused on PEs, but more specifically on RPEs (Rescorla & Wagner, 1972; Sutton & Barto, 1998).

RPEs are known to be signaled by the dopaminergic reward system, encompassing midbrain structures such as the substantia nigra and the ventral tegmental area (VTA; Bayer & Glimcher, 2005; Schultz, Dayan, & Montague, 1997). Dopaminergic cells in these structures respond to the difference between the expected and received reward. Each time an event proves to be more rewarding than anticipated a phasic dopamine burst is elicited, lasting about 200 to 500 milliseconds. These dopaminergic bursts evoked by RPEs subsequently trigger the activation of a dopaminergic pathway with projection ranging across a wide variety of brain regions such as the hippocampus, the frontal cortex and the anterior cingulate cortex (ACC).

The notion that any event can be valued as rewarding or unrewarding is supported by neuroscientific evidence and computational models. Of particular importance in this regard is the *reward value and prediction model* (RVPM) put forward by Silvetti, Seurinck, and Verguts (2011). Centered on the involvement of the ACC in various functions such as error detection and value estimation, the RVPM model states that any event (e.g., any stimulus or action) has its own predicted value (Sutton & Barto, 1998). This reward value evolves over time as the value of an event is updated each time this

event occurs. Again, the difference between the predicted and actual value of the event (i.e., the RPE) functions as a teaching signal and results in an updated value associated with that event.

RPEs influence behavior on the long term by shaping stable adaptations to the environment through various forms of learning (for a historical overview, see Squire, 2004). A long research tradition has focused on the influence of RPEs on procedural learning. Typically, these experiments reveal how RPEs gradually shape the acquisition of for instance stimulus-response associations. In this type of trial-by-trial learning, associations are formed by integrating across all previous encounters with the stimulus material and extracting which response is on average preferable. The associative strength between the stimulus and the response is weakened when the outcome is less beneficial than expected (a negative RPE) and is enhanced when the outcome is better than expected (a positive RPE).

Aside from procedural learning, much of our knowledge is acquired through declarative learning. Both types of learning differ in a number of ways. First, declarative learning mainly involves the acquisition of stimulus-stimulus associations. Thus, the type of information being learned typically involves facts and events. Second, declarative learning is usually very fast as knowledge can be acquired based on a single encounter. Lastly, whereas procedural knowledge is often difficult to describe verbally, declarative knowledge can be probed through explicit recall or reporting. Also, this allows the acquired information to be describing in relation to other knowledge.

Interestingly, although the influence of RPEs on procedural learning is well established, the empirical evidence for the impact of RPEs on

declarative learning remains remarkably absent in the literature. Therefore, in the second part of this dissertation we set out to test whether RPEs indeed serve as a teaching signal in declarative learning.

Probing the effect of PEs on learning through the neoHebbian learning account

In order to examine the long term effect of RPEs on declarative learning, we combined the well-established role of RPEs in procedural learning with a recent framework on declarative learning. To start, we drew upon the reinforcement learning literature to quantify RPEs as the difference between the obtained reward and the expected reward (Bush & Mosteller, 1951a, 1951b; Rescorla & Wagner, 1972). This parametric quantification of RPEs is traditionally used as the teaching signal in procedural learning. Next, we combined these quantified RPEs with the recent neoHebbian learning framework put forward by Lisman, Grace, & Duzel (2011), detailing how RPEs can influence declarative learning.

According to this neoHebbian account (Lisman et al., 2011), a large amount of information is temporarily stored by the hippocampus during initial memory encoding. However, not all of this information will be consolidated in long-term memory through long-term potentiation (LTP). Only those memory traces that were tagged as important during initial encoding will be engrained in long-term memory, a process that is known as the *tagging and capture model* (Frey & Morris, 1998). At the level of the individual synapses, a connection can be tagged as important by a concurrent dopamine burst. These tags have little influence on the strength of the synaptic connection on the short term. Instead, the tags trigger the synthesis of plasticity-related proteins through late LTP, thus strengthening the

synaptic connection on the long term. Notably, the tagging and capture model implies that the effect of dopamine bursts on memory performance will become more evident in a delayed memory test (e.g., after a week) compared to an immediate test (e.g., after twenty minutes).

The neoHebbian learning framework thus asserts that dopamine plays a crucial role in the acquisition of declarative knowledge on the long term. Events that necessitate learning such as novelty, salience and reward give rise to dopamine-dependent tagging and thus determine what information will be stored through LTP and what will be forgotten (Lisman et al., 2011). As described previously, RPEs elicit such dopamine bursts in the substantia nigra and VTA, triggering a dopaminergic pathway with projections to the hippocampal area. Thus, RPEs are hypothesized to tag the information that is being encoded in the hippocampus at that time, resulting in enhanced declarative learning.

In sum, we will calculate RPEs based on the reinforcement learning quantification and test whether this teaching signal predicts long term declarative learning, as described by the neoHebbian learning framework.

OUTLINE OF THE DISSERTATION

In **chapter 2**, we will use the DDM to investigate how PEs influence perception. More specifically, the brief presentation of a house or face picture will be cued with either an icon that correctly predicts the identity of the picture (congruent trials), an icon that predicts the alternative picture (incongruent or PE trials) or an icon that gives no prior information (neutral trials). In contrast to previous studies, participants were not asked to

categorize the stimulus based on its identity (i.e., identification task); instead they indicated whether the picture was presented above or below fixation (i.e., individuation task). Thus, although the cue was informative about the identity of the upcoming picture, it was not predictive for the correct response in the individuation task. This allowed us to exclude a bias toward a specific motor response. Applying the DDM we compared the parameter estimates for the congruent, incongruent and neutral trials in order to test whether PEs (incongruent trials) influenced visual processing by modulating the processing efficiency parameter or the threshold setting parameter.

In **chapter 3**, we turn our attention toward the effect of RPEs on declarative learning. During a declarative learning task, we parametrically manipulated the RPEs that participants experienced during the encoding of word pairs. Next, we tested declarative memory in a recognition task that was performed either immediately (i.e., after a brief filler task) or after a one-day delay. In our analysis, we probed how positive and negative RPEs influence declarative memory on the short term (immediate test) and on the long term (delayed test). In addition, we excluded alternative interpretations of the results such as the time-on-task.

In **chapter 4**, we built upon the behavioral paradigm used in chapter 3 to further explore whether RPEs have a direct influence on declarative learning or whether this influence is mediated by attentional modulations. In order to do so, the behavioral task was adapted to an electroencephalography (EEG) design suited for a time-frequency analysis. This allowed us to measure the oscillatory power in several frequency bands. Furthermore, we extracted these power estimates on a trial-by-trial basis. This approach allowed us to verify whether our quantification of the RPEs matches how

participants experience these teaching signals, by testing whether RPEs were reflected in the oscillatory power estimates during reward feedback. Next, we probed whether these RPE signatures directly predicted the enhancement of declarative learning, or whether they improved learning indirectly (e.g., through increased attention during encoding).

Finally, in the **general discussion**, we will evaluate the implications of our findings for the literature on perception and learning. Also, we will discuss future perspectives on how our approach could further our current understanding of the influence of PEs on perception and learning.

REFERENCES

Bayer, H. M., & Glimcher, P. W. (2005). Midbrain dopamine neurons encode a quantitative reward prediction error signal. *Neuron*, *47*(1), 129–141. doi:10.1016/j.neuron.2005.05.020

Bogacz, R. (2007). Optimal decision-making theories: Linking neurobiology with behaviour. *Trends in Cognitive Sciences*, *11*(3), 118–125. doi:10.1016/j.tics.2006.12.006

Bush, R., & Mosteller, F. (1951a). A mathematical model for simple learning. *Psychological Review*, *58*(5), 313–123.

Bush, R., & Mosteller, F. (1951b). A model for stimulus generalization and discrimination. *Psychological Review*, *58*(6), 413–423.

Clark, A. (2013). Whatever next? Predictive brains, situated agents, and the future of cognitive science. *The Behavioral and Brain Sciences*, *36*(3), 181–204. doi:10.1017/S0140525X12000477

Dayan, P. (1998). A hierarchical model of binocular rivalry. *Neural Computation, 10*(5), 1119–1135.

Enns, J. T., & Lleras, A. (2008). What's next? New evidence for prediction in human vision. *Trends in Cognitive Sciences*, *12*(9), 327–333. doi:10.1016/j.tics.2008.06.001

Frey, U., & Morris, R. G. M. (1998). Synaptic tagging: Implications for late maintenance of hippocampal long-term potentiation. *Trends in Neurosciences*, *21*(5), 181–188. doi:10.1016/S0166-2236(97)01189-2

Friston, K., & Kiebel, S. (2009). Predictive coding under the free-energy principle. *Philosophical Transactions of the Royal Society of London. Series B, Biological Sciences*, *364*(1521), 1211–1221. doi:10.1098/rstb.2008.0300

Gold, J. I., & Shadlen, M. N. (2007). The neural basis of decision making. *Annual Review of Neuroscience*, *30*, 535–574. doi:10.1146/annurev.neuro.29.051605.113038

Hohwy, J. (2012). Attention and conscious perception in the hypothesis testing brain. *Frontiers in Psychology*, *3*(96). doi:10.3389/fpsyg.2012.00096

Hohwy, J., Roepstorff, A., & Friston, K. (2008). Predictive coding explains binocular rivalry: An epistemological review. *Cognition*, *108*(3), 687–701. doi:10.1016/j.cognition.2008.05.010

Lisman, J., Grace, A. A., & Duzel, E. (2011). A neoHebbian framework for episodic memory; role of dopamine-dependent late LTP. *Trends in Neurosciences*, *34*(10), 536–547. doi:10.1016/j.tins.2011.07.006

Mulder, M. J., Wagenmakers, E.-J., Ratcliff, R., Boekel, W., & Forstmann, B. U. (2012). Bias in the brain: A diffusion model analysis of prior probability and potential payoff. *The Journal of Neuroscience : The Official Journal of the Society for Neuroscience*, *32*(7), 2335–2343. doi:10.1523/JNEUROSCI.4156-11.2012

Neely, J. (1991). Semantic priming effects in visual word recognition: A selective review of current findings and theories. *Basic Processes in Reading: Visual Word Recognition*, *11*, 264–336.

Rao, R. P., & Ballard, D. H. (1999). Predictive coding in the visual cortex: A functional interpretation of some extra-classical receptive-field effects. *Nature Neuroscience*, *2*(1), 79–87. doi:10.1038/458

Ratcliff, R. (1985). Theoretical interpretations of the speed and accuracy of positive and negative responses. *Psychological Review*, *92*(2), 212–225. doi:10.1037//0033-295X.92.2.212

Ratcliff, R., & McKoon, G. (2008). The diffusion decision model: Theory and data for two-choice decision tasks. *Neural Computation*, *20*(4), 873–922. doi:10.1162/neco.2008.12-06-420

Ratcliff, R., & Rouder, J. N. (1998). Modeling response times for two-choice decisions. *Psychological Science*, *9*(5), 347–356. doi:10.1111/1467-9280.00067

Rauss, K., Schwartz, S., & Pourtois, G. (2011). Top-down effects on early visual processing in humans: A predictive coding framework. *Neuroscience & Biobehavioral Reviews*, *35*(5), 1237–1253. doi:10.1016/j.neubiorev.2010.12.011

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

Schultz, W., Dayan, P., & Montague, R. (1997). A neural substrate of prediction and reward. *Science (New York, N.Y.)*, *275*(5306), 1593–1599. doi:10.1126/science.275.5306.1593

Seth, A. K., Suzuki, K., & Critchley, H. D. (2011). An interoceptive predictive coding model of conscious presence. *Frontiers in Psychology*, *2*, 395. doi:10.3389/fpsyg.2011.00395

Silvetti, M., Seurinck, R., & Verguts, T. (2011). Value and prediction error in medial frontal cortex: Integrating the single-unit and systems levels of analysis. *Frontiers in Human Neuroscience*, *5*(75). doi:10.3389/fnhum.2011.00075

Squire, L. R. (2004). Memory systems of the brain: A brief history and current perspective. *Neurobiology of Learning and Memory*, *82*(3), 171–177. doi:10.1016/j.nlm.2004.06.005

Summerfield, C., & de Lange, F. P. (2014). Expectation in perceptual decision making: Neural and computational mechanisms. *Nature Reviews Neuroscience*, *15*(11), 745–756. doi:10.1038/nrn3838

Summerfield, C., & Egner, T. (2009). Expectation (and attention) in visual cognition. *Trends in Cognitive Sciences*, *13*(9), 403–409. doi:10.1016/j.tics.2009.06.003

Summerfield, C., Egner, T., Greene, M., Koechlin, E., Mangels, J., & Hirsch, J. (2006). Predictive codes for forthcoming perception in the frontal cortex. *Science (New York, N.Y.)*, *314*(5803), 1311–1314. doi:10.1126/science.1132028

Sutton, R. S., & Barto, A. G. (1998). *Reinforcement Learning: An Introduction*. Cambridge, MA: MIT Press.

CHAPTER 2
PREDICTIVE INFORMATION SPEEDS UP VISUAL AWARENESS IN AN INDIVIDUATION TASK BY MODULATING THRESHOLD SETTING, NOT PROCESSING EFFICIENCY[1]

Theories on visual awareness claim that predicted stimuli reach awareness faster than unpredicted ones. In the current study, we disentangle whether prior information about the upcoming stimulus affects visual awareness of stimulus location (i.e., individuation*) by modulating processing efficiency or threshold setting. Analogous research on stimulus* identification *revealed that prior information modulates threshold setting. However, as identification and individuation are two functionally and neurally distinct processes, the mechanisms underlying identification cannot simply be extrapolated directly to individuation. The goal of this study was therefore to investigate how individuation is influenced by prior information about the upcoming stimulus. To do so, a drift diffusion model was fitted to estimate the processing efficiency and threshold setting for predicted versus unpredicted stimuli in a cued individuation paradigm. Participants were asked to locate a picture, following a cue that was congruent, incongruent or neutral with respect to the picture's identity. Pictures were individuated faster in the congruent and neutral condition compared to the incongruent condition. In the diffusion model analysis, the processing efficiency was not significantly different across conditions. However, the threshold setting was significantly higher following an incongruent cue compared to both congruent and neutral cues. Our results indicate that predictive information about the upcoming stimulus influences visual awareness by shifting the threshold for individuation rather than by enhancing processing efficiency.*

INTRODUCTION

Driving through an unfamiliar city, looking for the colleagues you promised to pick up, you might face a challenging visual perception task. Luckily, having some prior (i.e., predictive) knowledge about what your colleagues look like will facilitate becoming aware of them. Indeed, several consciousness theories have proposed mechanisms by which prior information modulates visual awareness. For example, according to Clark (2013) prior information is one of the key aspects to determine which stimuli reach visual awareness and at what speed. While the effect of prior information on visual perception has already been investigated extensively in paradigms that require stimulus *identification*, it remains unclear how it influences the distinct visual process of stimulus *individuation* (i.e., the spatial tagging of an object in a visual scene). The current study investigates whether and how prior information influences visual awareness in individuation.

The influence of prior information on visual perception has typically been investigated in paradigms that require the identification of visual input. In these experiments, participants need to categorize a degraded or masked stimulus (e.g., distinguishing a face from a house picture masked by noise). Predicted stimuli are typically perceived faster and more accurately. In order to gain more insight into the modulations of identification by prior information, formal models such as the drift diffusion model (DDM) and signal detection theory (SDT) have been used to disentangle the underlying mechanisms (for a review, see Mulder, van Maanen, & Forstmann, 2014; Summerfield & de Lange, 2014). In the DDM (see Figure 1; Ratcliff & Rouder, 1998), evidence is accumulated

at a certain rate (i.e., drift rate) from a starting point toward an upper or lower criterion bound. The distance between the upper and lower bounds is called boundary separation. The total response time is the sum of this evidence accumulation time plus any cognitive processes preceding or following the decision process (i.e., non-decision time). Critically, the parameters of the decision process (e.g., drift rate, boundary separation and starting point) can be mapped onto distinct mechanisms by which expectations can influence the accumulation process. First, prior information can improve visual *processing efficiency*. This is reflected in increased drift rate (see Figure 1a). Second, prior information can reduce the required amount of accumulated information. This is reflected by the distance between the starting point and decision boundaries henceforth referred to as *threshold setting* (see Figure 1b). Importantly, threshold setting encompasses both starting point placement and boundary separation, although only the latter is relevant in the current experimental paradigm (see below).

a. Processing efficiency
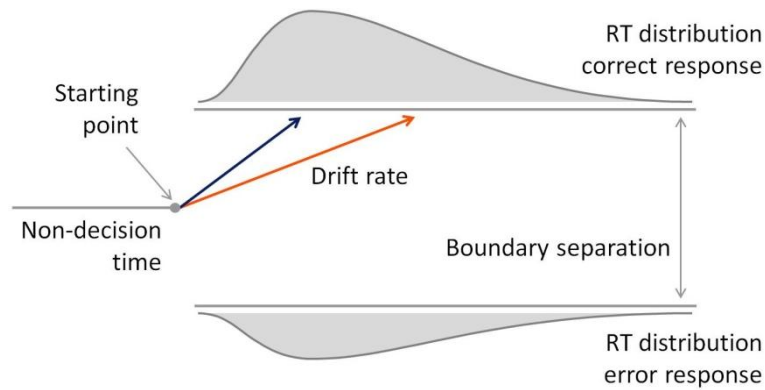
b. Threshold setting
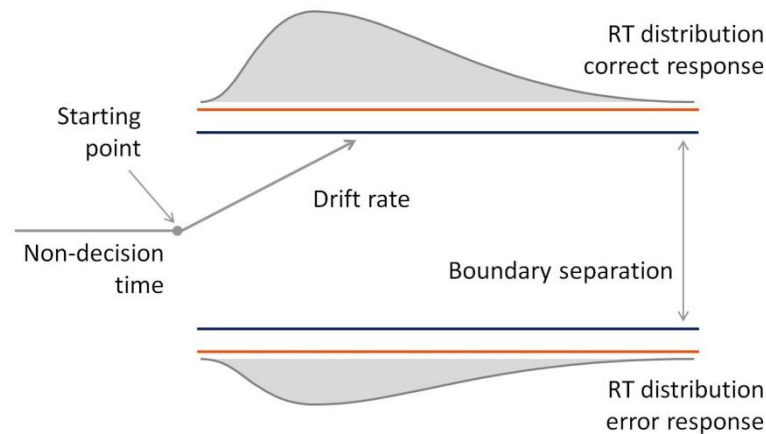
Figure 1: Prior information can influence visual perception by modulating processing efficiency (panel a) or threshold setting (panel b), respectively mapped onto the DDM parameters drift rate and boundary separation. The DDM is depicted including the non-decision time and starting point parameter. Hypothetical reaction time (RT) distributions for the correct and error responses are plotted at the corresponding upper and lower boundary. Increased processing efficiency and more lenient boundaries are indicated in blue, while decreased processing efficiency and more conservative boundary settings are depicted in orange.

Using these and related formal models, a number of studies have investigated how prior information influences stimulus identification. By manipulating the predictability of a shape in a shape discrimination task, Domenech and Dreher (2010) found using the LATER model (Reddi & Carpenter, 2000) that prior information influences threshold setting rather than processing efficiency. A cue predicting the movement direction in a random-dot motion paradigm influenced threshold setting but not processing efficiency (using the DDM: Mulder, Wagenmakers, Ratcliff, Boekel, & Forstmann, 2012; using a linear ballistic accumulator model: Forstmann, Brown, Dutilh, Neumann, & Wagenmakers, 2010). Using the DDM, Dunovan, Tremel, and Wheeler (2014) found that the identification of a house or face masked by noise was influenced by a house or face cue through the modulation of threshold setting. Interestingly, this modulation increased with the reliability of the cue (50, 70 or 90% accuracy) establishing a clear causal link between prior information and threshold setting. By contrast, using SDT Lupyan and Ward (2013) showed that cueing the word 'circle' or 'square' in a shape discrimination paradigm modulated processing efficiency (i.e., d') but not threshold setting (i.e., response criterion).

The results from these identification paradigms suggest that prior information influences visual awareness by modulating threshold setting. However, visual awareness studies usually require participants to report whether any item was perceived, irrespective of its identity (Baars, 1994; Overgaard & Sandberg, 2012; Sandberg, Timmermans, Overgaard, & Cleeremans, 2010; Sergent & Dehaene, 2004; Tononi & Koch, 2008). Interestingly, participants can often report where something was seen without knowing what was presented (Ramsøy & Overgaard, 2004). Similarly, to corroborate awareness reports, participants are commonly asked to report the location of a stimulus

(i.e., individuation) rather than its identity (e.g., Yang & Blake, 2012). Therefore, to investigate how prior information influences visual awareness, it is critical to probe its effect on stimulus individuation. According to the individuation-identification theory (Leslie, Xu, Tremoulet, & Scholl, 1998), the number of objects in a scene (i.e., individuation) and object identity are determined in two separate processes. This idea resonates with theories claiming that spatial information plays a unique role in visual processing, separate from the identification process (Sagi & Julesz, 1984). This notion is also supported in object file theory (Kahneman, Treisman, & Gibbs, 1992), where an object file is created based on spatial and temporal information, while its content is determined separately. As the individuation and identification stage are functionally and neurally different (Xu, 2009), prior information may influence perception via different mechanisms in these two visual processes.

Indirect evidence for distinct mechanisms underlying stimulus identification and individuation comes from the spatial attention literature. First, while object-based attention (crucial for identification) is associated with the ventral processing stream, location-based attention (crucial for individuation) depends on the dorsal processing stream (Arrington, Carr, Mayer, & Rao, 2000; Chen, 2009; Chou, Yeh, & Chen, 2014). Second, in stark contrast to the modulation of threshold setting presented above, prior information about the location of the upcoming stimulus has been argued to enhance stimulus identification by increasing processing efficiency (Anton-Erxleben, Abrams, & Carrasco, 2010; Smith, Ratcliff, & Wolfgang, 2004; however, for an alternative interpretation see Schneider, 2011). So in similar identification paradigms, location cueing boosts processing efficiency while identity cueing modulates threshold setting. It could be argued that − as locating a stimulus is crucial to individuation − location cueing boosted processing efficiency in the

individuation process and not in the identification process. However, this interpretation of the results remains to be tested as an identification task was used rather than an individuation task.

To investigate how prior information affects visual awareness of stimulus location in an individuation paradigm, we developed a cued masking task analogous to the identification studies described above. The picture of a house or a face (i.e., the target) was briefly presented above or below fixation, followed by a masking stimulus. Prior to the target presentation, participants were presented with a house or face cue that predicted the target identity with 80% accuracy, or with a cue that provided no prior information (a question mark). This manipulation generated three trial types: congruent, incongruent and neutral trials. Participants responded as fast and as accurate as possible to the location and not to the identity of the target picture by pressing an upper or lower response button. The visibility of the target picture was tailored to the individuation threshold of the individual participants in a staircase procedure. Furthermore, as the identity cues informed on target identity but were orthogonal to the target location that participants responded to, no motor response could be primed by the cue. Therefore, the starting point was restricted to be half the boundary separation. A DDM was fitted to compare drift rate and boundary separation estimates across the three trial types (i.e., congruent, incongruent and neutral trials), revealing how prior information influences individuation by modulating the processing efficiency (drift rate) or the threshold setting (boundary separation).

METHOD

Participants

Twenty Ghent University students were paid 20 euro for taking part in the current experiment combined with another experiment. The order of the experiments was counterbalanced across subjects and spread across two days. The experiment order and the results of the other experiment did not interact with the current experiment and will not be discussed further. The experiment lasted approximately one hour. All participants (5 male; on average 19 years old with a range of 18 to 25) had normal or corrected to normal vision. Prior to the experiment they gave their informed consent in accordance with the Code of Ethics of the World Medical Association (Declaration of Helsinki) and received a debriefing form afterwards.

Stimuli and material

The stimulus set consisted of ten pictures of Caucasian faces from the Face Database of the Park Aging Mind Laboratory (5 males, age raging from 19 to 79; Minear & Park, 2004) and ten pictures of houses taken from the Scene Understanding Database from the Princeton Vision Group (Xiao, Hays, Ehinger, Oliva, & Torralba, 2010). The face, house and question mark cue were taken from the website of The Noun Project (www.thenounproject.com; Person designed by Alex Fuller, House designed by OCHA Visual Information Unit, Question designed by Vicons Design). Scrambled versions of the pictures were constructed by dividing the picture in a 7×7 grid and randomizing the location of its 49 cells. All pictures and cues were luminance scaled to the average

luminance of all stimuli (mean HSV luminance of 0.62; mean Michelson contrast of 0.97) to avoid additional luminance-based variation in RTs.

The experiment was run on a DELL Latitude E6430 laptop running Windows 7 Professional and an external DELL E2213 screen with a 1680 by 1050 resolution. The refresh rate of the screen was set to 60 Hz. The stimulus presentation was programmed in MATLAB 2013a (Mathworks Inc.) with a Psychtoolbox extension (Brainard, 1997; Pelli, 1997). Answers were registered through a Cedrus RB-730 response box enhanced with four time-accurate push buttons (Cedrus Corporation, San Pedro, California). Participants were seated at approximately 60 cm from the screen. The display extended over a 41° × 28° visual angle black background.

Procedure and design

Each trial started with the presentation of the cue (house, face or question mark; 3.8° × 3.8° visual angle) for two seconds (see Figure 2a). Next, a full and a scrambled version of a picture (the targets; 10.5° x 10.5° visual angle) were presented for 33 ms at opposite sides of the cue (above and below; note that the locations of the cue and pictures show no overlap). Another scrambled version of the same picture (the mask) was subsequently presented at both sides of the cue for 33 ms. Next, an empty display was presented until a response was given. Participants were instructed to respond as fast as possible by pressing the button that matched the location of the full picture. Their right index and middle finger were positioned in corresponding positions on the response box. The trial ended with the feedback message 'correct' or 'wrong' presented for one second.

a. Trial design

b. Experimental design

Figure 2: Trial design (panel a) and experimental design (panel b) of the experiment. At the start of the trial a cue (the icon of a face, house or question mark) is presented for two seconds. Next a full picture and a scrambled version are shown as targets (33 ms), and subsequently masked by another scrambled version (33 ms). An empty display is presented until a response is given, followed by accuracy feedback (1000 ms). The three possible cues can be followed by a picture of a house or a face. The combinations of the cues and target pictures results in three trial types: neutral, congruent and incongruent trials. The number of trials is indicated for each cell of the design.

The experiment started with a practice phase of 60 trials, consisting of 20 neutral and 40 congruent trials. The remaining 600 trials consisted of 200 neutral trials and 400 trials with an informative cue (see Figure 2b). Of the 400 informative cue trials, 80% were congruent trials (320 trials) and 20% were incongruent trials (80 trials). There were an equal number of house and face pictures in the congruent, incongruent and neutral trials. The incongruent trials were randomly dispersed across the experiment with the restriction that two incongruent trials were always separated by at least two congruent or neutral trials. The position of the full picture on the screen (top versus bottom) was randomized across all trial types (congruent, incongruent and neutral trials) and picture types (house versus face). Participants were informed that the neutral cues had no predictive value, and that house and face cues would correctly predict the picture type in 80% of the trials.

Accuracy staircase procedure

In order to acquire sufficient error trials for the DDM analysis, the visibility of the target pictures was varied in a staircase procedure such that participants localized the full picture incorrectly in 30% of the trials. The visibility was manipulated on a trial-to-trial basis by varying the luminance of the pictures from 0% (not visible) to 100% (fully visible). If errors remained above 30% when the target pictures were at maximum luminance, the luminance of the masks would subsequently be lowered from 100% to 0%. The luminance of the target pictures and masks started at 50% and 100% respectively. The practice phase allowed the staircase procedure to reach a stable plateau at 30% errors. The staircase procedure was further applied throughout the experiment to sustain the 30% error rate (Busch & VanRullen, 2010; Gilzenrat, Nieuwenhuis, Jepma, & Cohen, 2010).

A separate staircase was applied to the house and face picture trials. On each trial the average accuracy was calculated for the preceding ten house trials or ten face trials. When less than seven out of ten houses (faces) were individuated correctly, the luminance on the house (face) trials was increased with one percent. The luminance decreased with one percent when more than seven out of ten individuations were correct. To avoid unnecessary fluctuations in the staircase procedure, the less frequent incongruent trials were not taken into account for calculating this average accuracy. However, the luminance was adjusted on all trials, irrespective of the trial type (congruent, incongruent and neutral trials).

Drift diffusion model

The DDM parameters were estimated using the DMAT toolbox (Vandekerckhove & Tuerlinckx, 2007) running on MATLAB 2013a. The DMAT toolbox allows the estimation of seven parameters (see Figure 1) on a participant level: drift rate, boundary separation, starting point, non-decision time, drift rate variability, starting point variability, and non-decision time variability. The two parameters of interest (drift rate and boundary separation) were estimated separately for the three trial types (congruent, incongruent and neutral trials). In contrast to earlier work (Dunovan et al., 2014; Mulder et al., 2012) the cue only informed participants about the identity of the upcoming stimulus. It did not inform on the to-be-reported location. Therefore, the starting point was restricted to half the boundary separation estimates as the cue (or indeed any other information) is unable to bias the starting point toward one or the other boundary when judging the stimulus location. All other parameters (the non-decision time and the three variability parameters) were estimated but not allowed to vary across trial types. The estimation method was set to

multinomial likelihood estimation based on percentiles with four percentile bins separated at the 20[th], 50[th] and 80[th] percentile. The model was fitted for each participant individually, resulting in one drift rate and boundary separation estimate per trial type per participant.

RESULTS

Descriptive statistics

The staircase procedure was successfully applied in all 20 participants. The average number of incorrectly individuated full pictures per participant ranged from 30% to 39% for the houses (mean = 33%, sd = 2.1%) and from 28% to 34% for the faces (mean = 29%, sd = 1.3%).

Trials on which the interval between target pictures and masks exceeded 33 ms due to software slowing were excluded (0.12% of the trials removed because interval lasted for 48 ms). RTs in the individuation task were subjected to a lower cutoff of 200 ms and an upper cutoff of 4000 ms (0.09% and 0.06% of the trials removed respectively). The average RTs per participant ranged from 458 ms to 1266 ms (mean average = 695 ms). The remaining data were entered into the DDM and accuracy analysis; the RT analysis was performed separately for the correct and error trials.

Accuracy and RT analysis

Accuracies were entered into a generalized linear mixed effects model with a random intercept across participants and the trial type as a fixed effects predictor. There was no significant main effect of trial type, $\chi^2(2, N = 20) = 1.03$, $p = 0.60$ (see Figure 3a).

The RTs for the correct and error trials were entered into a separate linear mixed effects model with a random intercept across participants and the trial type as a fixed effects predictor. For the correct RTs, trial type was a significant predictor (average RT for correct congruent, incongruent and neutral trials was 670 ms, 724 ms and 676 ms respectively; $\chi^2(2, N = 20) = 42.3$, $p < 0.001$; see Figure 3b). Follow-up tests revealed that the RTs on the incongruent trials were significantly slower compared to the congruent and neutral trials (respectively $\chi^2(1, N = 20) = 40.9$, $p < 0.001$, and $\chi^2(1, N = 20) = 27.5$, $p < 0.001$). There was no significant difference in RTs between the congruent and neutral trials, $\chi^2(1, N = 20) = 1.10$, $p = 0.29$.

The trial type also significantly predicted the RTs on the error trials (average RT for error congruent, incongruent and neutral trials was 733 ms, 764 ms and 717 ms respectively; $\chi^2(2, N = 20) = 11.0$, $p = 0.0040$; see Figure 3c). Again, the RTs on the incongruent trials were higher compared to the congruent and neutral trials (respectively $\chi^2(1, N = 20) = 5.22$, $p = 0.022$, and $\chi^2(1, N = 20) = 11.3$, $p < 0.001$), with no significant difference between congruent and neutral trials, $\chi^2(1, N = 20) = 2.51$, $p = 0.11$.

a. Accuracy: all trials
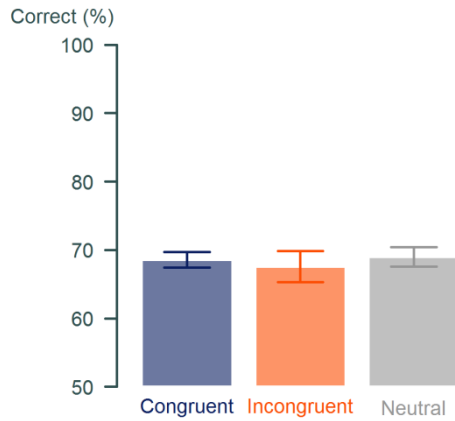
b. RT: correct trials
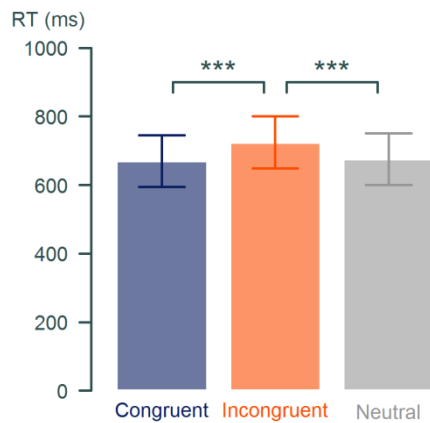
c. RT: error trials
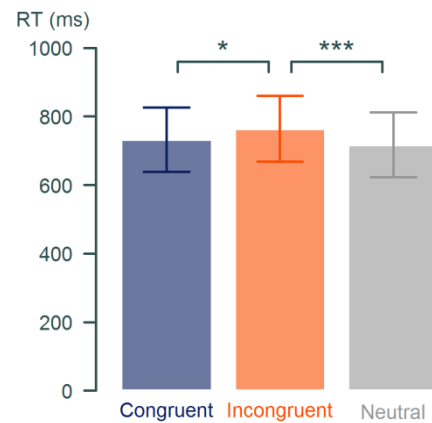
Figure 3: Accuracy (panel a) and RT analysis (panels b and c). The average accuracies and RTs are plotted with their 95% confidence intervals for all three trial types: congruent (blue), incongruent (orange) and neutral (gray) trials. The analysis revealed no significant effect of trial type on accuracy ($p = 0.60$). By contrast, RTs were significantly predicted by trial type on the correct trials (panel b; $p < 0.001$) and the error trials (panel c; $p = 0.0040$). Follow-up tests indicated that the RTs were higher on incongruent trials compared to neutral and congruent trials (* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$).

DDM analysis

The drift rate and boundary separation estimates were entered into separate linear mixed effects models with a random intercept across participants and the trial type as a fixed effects predictor. There was no significant difference between the drift rate estimates across trial types, $\chi^2(2, N = 20) = 3.08$, $p = 0.21$ (see Figure 4a). Conversely, the boundary separation estimates were predicted significantly by the trial type, $\chi^2(2, N = 20) = 27.2$, $p < 0.001$ (see Figure 4b). Follow-up tests revealed an increased boundary separation in incongruent trials compared to neutral trials, $\chi^2(1, N = 20) = 12.14$, $p < 0.001$, and congruent trials, $\chi^2(1, N = 20) = 18.79$, $p < 0.001$. The boundary separation was not significantly different for the neutral and congruent trials, $\chi^2(1, N = 20) = 0.068$, $p = 0.79$.

a. Drift rate comparison

b. Boundary separation comparison
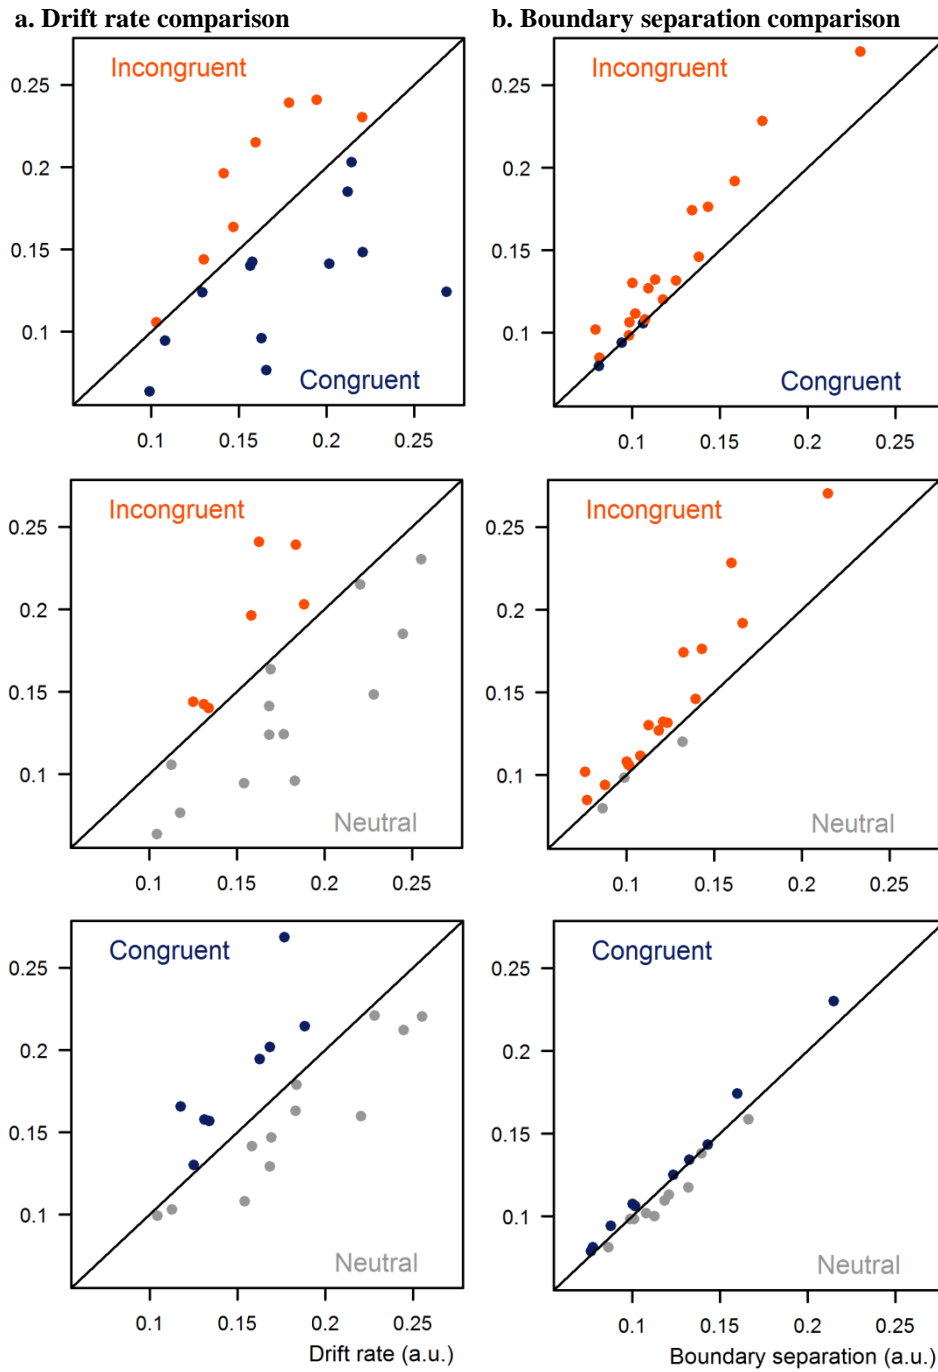
Figure 4: The drift rate estimates (panel a) and boundary separation estimates (panel b) are compared across the three trial types (congruent (blue), incongruent (orange) and neutral (gray) trials), with each point representing a participant. The analysis revealed no significant difference between the drift rate estimates ($p = 0.21$) but a significantly higher boundary separation for incongruent trials compared to congruent and neutral trials (all $p < 0.01$).

Four validation tests were performed to assess the overall quality of our DDM implementation. First, a quantile probability plot (see Figure 5) demonstrates a good fit between the observed data and the data simulated by our model, with only in the highest quantile a (typical) overestimation of the RTs (Leite & Ratcliff, 2011). Second, a bottom-up model building approach was used to test whether the model fit of a null model with no condition-specific parameter estimates (model 1; M1) would significantly benefit from adding a condition-specific estimate for the drift rate (M2), boundary separation (M3) or non-decision time parameter (M4). For each participant separately, all four models were fitted and the deviances of models M2 to M4 were subtracted from the deviance of the null model M1, resulting in a chi-square value with the difference between the number of estimated parameters as the degrees of freedom. A chi-square test across participants revealed that the model fit of M1 was significantly improved by adding condition-specific boundary separation (M3) and non-decision time estimates (M4), but not by adding a condition-specific drift rate (M2; see Table 1). Next, we tested whether model M3 could be significantly improved by adding a condition-specific drift rate (M5; the model reported in detail above) or non-decision time estimate (M6). Statistical tests revealed no significantly improved model fits (see Table 1). Importantly, in model M6 the non-decision time estimates were not significantly different across conditions, $\chi^2(2, N = 20) = 0.17$, $p = 0.92$, while there was a significant

difference between boundary separation estimates, $\chi^2(2, N = 20) = 15.4$, $p <$ 0.001, with a higher boundary separation for the incongruent trials compared to congruent trials, $\chi^2(2, N = 20) = 7.81$, $p = 0.0052$, and neutral trials, $\chi^2(2, N = 20) = 8.86$, $p = 0.0029$, but no difference between congruent and neutral trials, $\chi^2(2, N = 20) = 0.20$, $p = 0.66$. Thus, the bottom-up model building approach confirmed the results from the main DDM implementation: Only the boundary separation varies significantly across conditions. As a third validation test, the main model was fitted separately for the trials with house and face pictures to control for stimulus-specific effects or artifacts caused by the separate staircase procedure for house and face trials. Both models confirmed our main conclusion. Indeed, boundary separation estimates were higher for the incongruent trials compared to the neutral and congruent trials (all $p < 0.001$), while the trial type failed to significantly predict drift rate estimates. The fourth and final validation test considered that the estimation could be biased by the unbalanced design (320, 80 and 200 trials in the congruent, incongruent and neutral condition respectively). In a bootstrapping approach, an equal number of trials per picture type and trial type were randomly selected to fit the model and this procedure was repeated fifty times. The drift rate and boundary separation estimates per participant and trial type were subsequently entered in the linear mixed effects model with a random intercept across participants and the trial type as a fixed effects predictor. Confirming the results from the main model, there was no significant difference between the drift rate estimates, $\chi^2(2, N = 20) = 2.47$, $p = 0.29$, while the boundary separation estimates were significantly different, $\chi^2(2, N = 20) = 15.95$, $p < 0.001$, with higher estimates in the incongruent trials compared to the neutral trials, $\chi^2(1, N = 20) = 10.7$, $p = 0.0011$, and congruent trials, $\chi^2(1, N = 20) = 15.4$, $p < 0.001$. In sum, our validation tests demonstrate the quality of the model, and exclude a distortion of

the results due to the specific parameter restrictions, stimulus material, the staircase procedure or the unbalanced design.

Table 1: Model comparisons

M1: null	↔ M2: drift rate	χ^2 (40, $N = 20$) = 35.7, $p = 0.66$
	↔ M3: boundary	χ^2 (40, $N = 20$) = 82.9, $p < 0.001$
	↔ M4: non-decision time	χ^2 (40, $N = 20$) = 80.1, $p < 0.001$
M3: boundary	↔ M5: boundary + drift rate	χ^2 (40, $N = 20$) = 41.8, $p = 0.39$
	↔ M6: boundary + non-decision time	χ^2 (40, $N = 20$) = 42.8, $p = 0.35$

Figure 5: The quantile probability plot for the DDM analysis is plotted for the observed (full lines) and simulated data (dashed lines). The average RTs across participants (y-axis) are plotted separately for error/correct trials in the incongruent (blue), congruent (orange) and neutral condition (gray), divided in 5 quantiles.

DISCUSSION

In a cued individuation paradigm we investigated how visual awareness of location is affected by prior information. The results show that prior information had no impact on the individuation accuracy, but it was highly predictive for the RTs: Stimuli following incongruent cues were individuated slower compared to stimuli following congruent or neutral cues. A drift diffusion analysis revealed that prior information modulates visual awareness by shifting threshold setting (implemented as boundary separation) and not by the modulation of processing efficiency (i.e., drift rate). Thus, when a stimulus is unexpected more evidence needs to be accumulated before the threshold for individuation is reached.

The current study is to our knowledge the first to investigate whether and how prior information modulates performance in an individuation task, the typical measure for visual awareness (Baars, 1994; Sergent & Dehaene, 2004; for implementations see e.g. research using the continuous flash suppression technique based on Tsuchiya & Koch, 2005). This is highly relevant as prior information is hypothesized to be one of the key elements that determine which input enters awareness and how fast (Enns & Lleras, 2008; Hohwy, Roepstorff, & Friston, 2008; Kouider, de Gardelle, Sackur, & Dupoux, 2010). The idea that prior information plays a pivotal role in the transition from unconscious processing to conscious perception has inspired a number of models on sensory awareness and visual awareness more specifically (Grossberg, 1999; King & Dehaene, 2014; Seth, Suzuki, & Critchley, 2011; Thilakarathne, 2015). In addition, it has spurred multiple lines of empirical research. For example, our

actions and their sensory effects − whether visual, auditory or proprioceptive − are highly predictable. Research in this domain has produced valuable insights on how the predictability of proprioceptive and visual input produced by our own actions alters awareness, especially when the input is ambiguous (Desantis, Hughes, & Waszak, 2012; Moore & Haggard, 2008; Salomon, Lim, Herbelin, Hesselmann, & Blanke, 2013; Stenner et al., 2014). In the current study we add to this literature by demonstrating by what mechanism prior information influences individuation, namely threshold setting.

The modulation of threshold setting in the current individuation task parallels the mechanism by which prior information about the upcoming stimulus influences stimulus identification (Domenech & Dreher, 2010; Dunovan et al., 2014; Forstmann et al., 2010; Mulder et al., 2012). Importantly, the modulation of threshold setting by identity cueing is measured differently when fitting a DDM to identification or individuation paradigms. Depending on the paradigm, identity cues can alter the required amount of evidence accumulation by influencing either the boundary separation or the starting point parameter (both contributing to the threshold setting; see our discussion on Figure 1). In the identification paradigm, the identity cue is directly relevant for (correlated with) the response options, leading to response priming. In that case, the boundary separation remains fixed and the starting point can be positioned closer to the boundary associated with the predicted identity response (Dunovan et al., 2014; Mulder et al., 2012). Conversely, in the current identification paradigm the cue is unrelated to the response options, causing no response priming. In this case, the starting point cannot be biased toward one boundary or the other. Instead, the required amount of evidence can be raised by increasing the boundary separation symmetrically. Thus, unlike in identification paradigms, we can exclude the possibility that the observed difference in

threshold setting reflects response priming. However, note that by excluding response priming we do not claim that the identity cue had no influence on the response mechanisms involved in the decision process (Schneider, 2011); threshold setting is part of the response mechanism.

As discussed in the introduction, the question how prior information influences visual awareness has also been tackled by using location cueing, however still within identification paradigms (Anton-Erxleben et al., 2010; Smith et al., 2004). Information at cued spatial locations is identified faster and this effect is ascribed to the overlap between the frontoparietal network supporting visual awareness and the parietal orienting system (for a review see Chica & Bartolomeo, 2012). As individuation and location cueing are both embedded in the dorsal processing stream, we could have investigated how identification is influenced by location cueing rather than identity cueing. However, as consciousness theories mainly focus on how visual awareness is affected by prior information about the identity of an upcoming stimulus rather than prior information about its location (Clark, 2013), we applied identity cueing to the individuation task. It would be interesting to apply a formal model (e.g., DDM) to disentangle whether location cueing in an individuation paradigm would boost processing efficiency as in the identification paradigms (Anton-Erxleben et al., 2010; Smith et al., 2004) or threshold setting as in the current individuation paradigm. Using SDT, Chica et al. (2011) already demonstrated that location cueing mainly influences processing efficiency (i.e., d'). However, also some effects on threshold setting (i.e., response criterion) were observed. Using DDM, future research could expound upon these results by clarifying under what conditions location cueing can influence threshold setting in individuation paradigms.

Now that we have identified threshold setting as crucial in individuation, we can speculate about its neural basis. Whereas drift rate is associated with processing efficiency mechanisms such as neural gain and the tuning of response curves (Isaacson & Scanziani, 2011; Kok, Jehee, & de Lange, 2012; Liu, Larsson, & Carrasco, 2007; Martinez-Trujillo & Treue, 2004), shifts in threshold setting have been linked to altered baseline activation in regions coding for the predicted stimulus feature (Giesbrecht, Weissman, Woldorff, & Mangun, 2006; Langner et al., 2011; Macaluso, Eimer, Frith, & Driver, 2003). In the current study, face and house cues might trigger increased baseline activity in respectively the fusiform face area (FFA) and the parahippocampal place area (PPA). We propose that this boosted activity in relevant cortical areas will facilitate resonance with the parietal individuation areas (for a similar influence of stimulus-driven ventral processes on dorsal processes, see Macaluso & Doricchi, 2013). Correspondingly, Summerfield and colleagues demonstrated how increased baseline activation in FFA can cause a house to be misperceived as a face, activating the frontoparietal network for awareness (Summerfield, Egner, Mangels, & Hirsch, 2006). Functionally, this would lead to a lower threshold in the individuation process. Similarly, in our analysis lower thresholds were observed for the congruent trials compared to the incongruent trials, which would lead to faster but more error-prone stimulus individuation on congruent trials. Although no significant difference in error rate was found in the current paradigm higher error rates on congruent trials are possible. The latter could occur because the (task-irrelevant) target identity is quickly processed in the congruent condition, with the (task-relevant) target location processing lagging behind, potentially resulting in premature and incorrect localization responses. The exact neural basis of such a process remains to be studied. Interestingly, our study and those outlined above lay out

the empirical restrictions necessary for constructing more detailed neural models on how prior information influences visual awareness.

REFERENCES

Anton-Erxleben, K., Abrams, J., & Carrasco, M. (2010). Evaluating comparative and equality judgments in contrast perception: attention alters appearance. *Journal of Vision*, *10*(11), 6. doi:10.1167/10.11.6

Arrington, C. M., Carr, T. H., Mayer, A. R., & Rao, S. M. (2000). Neural mechanisms of visual attention: object-based selection of a region in space. *Journal of Cognitive Neuroscience*, *12*(Supplement 2), 106–117. doi:10.1162/089892900563975

Baars, B. (1994). A global workspace theory of conscious experience. *Consciousness In Philosophy And Cognitive Neuroscience*, 149–171.

Brainard, D. H. (1997). The psychophysics toolbox. *Spatial Vision*, *10*(4), 433–436. doi:10.1163/156856897X00357

Busch, N. A., & VanRullen, R. (2010). Spontaneous EEG oscillations reveal periodic sampling of visual attention. *Proceedings of the National Academy of Sciences of the United States of America*, *107*(37), 16048–53. doi:10.1073/pnas.1004801107

Chen, Z. (2009). Not all features are created equal: Processing asymmetries between location and object features. *Vision Research*, *49*(11), 1481–1491. doi:10.1016/j.visres.2009.03.008

Chica, A. B., & Bartolomeo, P. (2012). Attentional routes to conscious perception. *Frontiers in Psychology*, *3*, 1. doi:10.3389/fpsyg.2012.00001

Chica, A. B., Lasaponara, S., Chanes, L., Valero-Cabré, A., Doricchi, F., Lupiáñez, J., & Bartolomeo, P. (2011). Spatial attention and conscious perception: the role of endogenous and exogenous orienting. *Attention, Perception & Psychophysics*, *73*(4), 1065–1081. doi:10.3758/s13414-010-0082-6

Chou, W.-L., Yeh, S.-L., & Chen, C.-C. (2014). Distinct mechanisms subserve location- and object-based visual attention. *Frontiers in Psychology*, *5*, 456. doi:10.3389/fpsyg.2014.00456

Clark, A. (2013). Whatever next? Predictive brains, situated agents, and the future of cognitive science. *The Behavioral and Brain Sciences*, *36*(3), 181–204. doi:10.1017/S0140525X12000477

Desantis, A., Hughes, G., & Waszak, F. (2012). Intentional binding is driven by the mere presence of an action and not by motor prediction. *PLoS ONE*, *7*(1), e29557. doi:10.1371/journal.pone.0029557

Domenech, P., & Dreher, J.-C. (2010). Decision threshold modulation in the human brain. *The Journal of Neuroscience : The Official Journal of the Society for Neuroscience*, *30*(43), 14305–14317. doi:10.1523/JNEUROSCI.2371-10.2010

Dunovan, K., Tremel, J., & Wheeler, M. (2014). Prior probability and feature predictability interactively bias perceptual decisions. *Neuropsychologia*, *61*, 210–221. doi:10.1016/j.neuropsychologia.2014.06.024.Prior

Enns, J. T., & Lleras, A. (2008). What's next? New evidence for prediction in human vision. *Trends in Cognitive Sciences*, *12*(9), 327–333. doi:10.1016/j.tics.2008.06.001

Forstmann, B. U., Brown, S., Dutilh, G., Neumann, J., & Wagenmakers, E.-J. (2010). The neural substrate of prior information in perceptual decision making: a model-based analysis. *Frontiers in Human Neuroscience*, *4*, 40. doi:10.3389/fnhum.2010.00040

Giesbrecht, B., Weissman, D. H., Woldorff, M. G., & Mangun, G. R. (2006). Pre-target activity in visual cortex predicts behavioral performance on spatial and feature attention tasks. *Brain Research*, *1080*(1), 63–72. doi:10.1016/j.brainres.2005.09.068

Gilzenrat, M. S., Nieuwenhuis, S., Jepma, M., & Cohen, J. D. (2010). Pupil diameter tracks changes in control state predicted by the adaptive gain theory of locus coeruleus function. *Cognitive, Affective & Behavioral Neuroscience*, *10*(2), 252–69. doi:10.3758/CABN.10.2.252

Grossberg, S. (1999). The link between brain learning, attention, and consciousness. *Consciousness and Cognition*, *8*(1), 1–44. doi:10.1006/ccog.1998.0372

Hohwy, J., Roepstorff, A., & Friston, K. (2008). Predictive coding explains binocular rivalry: an epistemological review. *Cognition*, *108*(3), 687–701. doi:10.1016/j.cognition.2008.05.010

Isaacson, J. S., & Scanziani, M. (2011). How inhibition shapes cortical activity. *Neuron*, *72*(2), 231–243. doi:10.1016/j.neuron.2011.09.027

Kahneman, D., Treisman, A., & Gibbs, B. J. (1992). The reviewing of object files: object-specific integration of information. *Cognitive Psychology*, *24*(2), 175–219. doi:10.1016/0010-0285(92)90007-O

King, J.-R., & Dehaene, S. (2014). A model of subjective report and objective discrimination as categorical decisions in a vast representational space. *Philosophical Transactions of the Royal Society of London. Series B, Biological Sciences*, *369*(1641), 20130204. doi:10.1098/rstb.2013.0204

Kok, P., Jehee, J. F. M., & de Lange, F. P. (2012). Less is more: expectation sharpens representations in the primary visual cortex. *Neuron*, *75*(2), 265–270. doi:10.1016/j.neuron.2012.04.034

Kouider, S., de Gardelle, V., Sackur, J., & Dupoux, E. (2010). How rich is consciousness? The partial awareness hypothesis. *Trends in Cognitive Sciences*, *14*(7), 301–307. doi:10.1016/j.tics.2010.04.006

Langner, R., Kellermann, T., Boers, F., Sturm, W., Willmes, K., & Eickhoff, S. B. (2011). Modality-specific perceptual expectations selectively modulate baseline activity in auditory, somatosensory, and visual cortices. *Cerebral Cortex (New York, N.Y. : 1991)*, *21*(12), 2850–2862. doi:10.1093/cercor/bhr083

Leite, F. P., & Ratcliff, R. (2011). What cognitive processes drive response biases? A diffusion model analysis. *JUDGMENT AND DECISION MAKING*, *6*(7), 651–687.

Leslie, A. M., Xu, F., Tremoulet, P. D., & Scholl, B. J. (1998). Indexing and the object concept: developing "what" and "where" systems. *Trends in Cognitive Sciences*, *2*(1), 10–18. doi:10.1016/S1364-6613(97)01113-3

Liu, T., Larsson, J., & Carrasco, M. (2007). Feature-based attention modulates orientation-selective responses in human visual cortex. *Neuron*, *55*(2), 313–323. doi:10.1016/j.neuron.2007.06.030

Lupyan, G., & Ward, E. J. (2013). Language can boost otherwise unseen objects into visual awareness. *Proceedings of the National Academy of Sciences of the United States of America*, *110*(35), 14196–14201. doi:10.1073/pnas.1303312110

Macaluso, E., & Doricchi, F. (2013). Attention and predictions: control of spatial attention beyond the endogenous-exogenous dichotomy. *Frontiers in Human Neuroscience*, *7*, 685. doi:10.3389/fnhum.2013.00685

Macaluso, E., Eimer, M., Frith, C. D., & Driver, J. (2003). Preparatory states in crossmodal spatial attention: spatial specificity and possible control mechanisms. *Experimental Brain Research*, *149*(1), 62–74. doi:10.1007/s00221-002-1335-y

Martinez-Trujillo, J. C., & Treue, S. (2004). Feature-based attention increases the selectivity of population responses in primate visual cortex. *Current Biology*, *14*(9), 744–751. doi:10.1016/j.cub.2004.04.028

Minear, M., & Park, D. C. (2004). A lifespan database of adult facial stimuli. *Behavior Research Methods, Instruments, & Computers*, *36*(4), 630–633. doi:10.3758/BF03206543

Moore, J., & Haggard, P. (2008). Awareness of action: inference and prediction. *Consciousness and Cognition*, *17*(1), 136–144. doi:10.1016/j.concog.2006.12.004

Mulder, M. J., van Maanen, L., & Forstmann, B. U. (2014). Perceptual decision neurosciences - a model-based review. *Neuroscience*, *277*, 872–884. doi:10.1016/j.neuroscience.2014.07.031

Mulder, M. J., Wagenmakers, E.-J., Ratcliff, R., Boekel, W., & Forstmann, B. U. (2012). Bias in the brain: a diffusion model analysis of prior probability and potential payoff. *The Journal of Neuroscience : The Official Journal of the Society for Neuroscience*, *32*(7), 2335–2343. doi:10.1523/JNEUROSCI.4156-11.2012

Overgaard, M., & Sandberg, K. (2012). Kinds of access: different methods for report reveal different kinds of metacognitive access. *Philosophical Transactions of the Royal Society of London. Series B, Biological Sciences*, *367*(1594), 1287–1296. doi:10.1098/rstb.2011.0425

Pelli, D. G. (1997). The VideoToolbox software for visual psychophysics: transforming numbers into movies. *Spatial Vision*, *10*(4), 437–442. doi:10.1163/156856897X00366

Ramsøy, T. Z., & Overgaard, M. (2004). Introspection and subliminal perception. *Phenomenology and the Cognitive Sciences*, *3*(1), 1–23. doi:10.1023/B:PHEN.0000041900.30172.e8

Ratcliff, R., & Rouder, J. N. (1998). Modeling response times for two-choice decisions. *Psychological Science*, *9*(5), 347–356. doi:10.1111/1467-9280.00067

Reddi, B., & Carpenter, R. (2000). The influence of urgency on decision time. *Nature Neuroscience*, *3*(8), 827–830.

Sagi, D., & Julesz, B. (1984). Detection versus discrimination of visual orientation. *Perception*, *13*(5), 619–628. doi:10.1068/p130619

Salomon, R., Lim, M., Herbelin, B., Hesselmann, G., & Blanke, O. (2013). Posing for awareness: proprioception modulates access to visual consciousness in a continuous flash suppression task. *Journal of Vision*, *13*(7), 2. doi:10.1167/13.7.2

Sandberg, K., Timmermans, B., Overgaard, M., & Cleeremans, A. (2010). Measuring consciousness: is one measure better than the other? *Consciousness and Cognition*, *19*(4), 1069–1078. doi:10.1016/j.concog.2009.12.013

Schneider, K. A. (2011). Attention alters decision criteria but not appearance: a reanalysis of Anton-Erxleben, Abrams, and Carrasco (2010). *Journal of Vision*, *11*(13), 7. doi:10.1167/11.13.7

Sergent, C., & Dehaene, S. (2004). Is consciousness a gradual phenomenon? Evidence for an all-or-none bifurcation during the attentional blink. *Psychological Science*, *15*(11), 720–728. doi:10.1111/j.0956-7976.2004.00748.x

Seth, A. K., Suzuki, K., & Critchley, H. D. (2011). An interoceptive predictive coding model of conscious presence. *Frontiers in Psychology*, *2*, 395. doi:10.3389/fpsyg.2011.00395

Smith, P. L., Ratcliff, R., & Wolfgang, B. J. (2004). Attention orienting and the time course of perceptual decisions: response time distributions with masked and unmasked displays. *Vision Research*, *44*(12), 1297–1320. doi:10.1016/j.visres.2004.01.002

Stenner, M.-P., Bauer, M., Machts, J., Heinze, H.-J., Haggard, P., & Dolan, R. J. (2014). Re-construction of action awareness depends on an internal model of action-outcome timing. *Consciousness and Cognition*, *25*, 11–16. doi:10.1016/j.concog.2014.01.007

Summerfield, C., & de Lange, F. P. (2014). Expectation in perceptual decision making: neural and computational mechanisms. *Nature Reviews Neuroscience*, *15*(11), 745–756. doi:10.1038/nrn3838

Summerfield, C., Egner, T., Mangels, J., & Hirsch, J. (2006). Mistaking a house for a face: neural correlates of misperception in healthy humans. *Cerebral Cortex (New York, N.Y. : 1991)*, *16*(4), 500–508. doi:10.1093/cercor/bhi129

Thilakarathne, D. J. (2015). Modelling of situation awareness with perception, attention, and prior and retrospective awareness. *Biologically Inspired Cognitive Architectures*, *12*, 77–104. doi:10.1016/j.bica.2015.04.010

Tononi, G., & Koch, C. (2008). The neural correlates of consciousness: an update. *Annals of the New York Academy of Sciences*, *1124*, 239–261. doi:10.1196/annals.1440.004

Tsuchiya, N., & Koch, C. (2005). Continuous flash suppression reduces negative afterimages. *Nature Neuroscience*, *8*(8), 1096–1101. doi:10.1038/nn1500

Vandekerckhove, J., & Tuerlinckx, F. (2007). Fitting the ratcliff diffusion model to experimental data. *Psychonomic Bulletin & Review*, *14*(6), 1011–1026. doi:10.3758/BF03193087

Xiao, J., Hays, J., Ehinger, K. A., Oliva, A., & Torralba, A. (2010). SUN database: large-scale scene recognition from abbey to zoo. In *2010 IEEE Computer Society Conference on Computer Vision and Pattern Recognition* (pp. 3485–3492). IEEE. doi:10.1109/CVPR.2010.5539970

Xu, Y. (2009). Distinctive neural mechanisms supporting visual object individuation and identification. *Journal of Cognitive Neuroscience*, *21*(3), 511–518. doi:10.1162/jocn.2008.21024

Yang, E., & Blake, R. (2012). Deconstructing continuous flash suppression. *Journal of Vision*, *12*(3), 1–14. doi:10.1167/12.3.8

CHAPTER 3
SIGNED REWARD PREDICTION ERRORS DRIVE DECLARATIVE LEARNING [1]

Reward prediction errors (RPEs) are thought to drive learning. This has been firmly established in procedural learning paradigms (i.e., classical and operant conditioning). However, empirical evidence on whether RPEs drive declarative learning – a quintessentially human form of learning – remains surprisingly absent. In this study, we used a declarative learning paradigm in which RPEs were coupled to the acquisition of Dutch-Swahili word pairs. The memory performance for these word pairs was subsequently tested in a recognition test, either immediate or after a one-day delay. The results demonstrate a causal effect of signed RPEs on declarative learning, with larger RPEs leading to better recognition on the immediate test and even stronger benefits on the delayed test. In addition, we demonstrate that classic declarative memory mechanisms such as time-on-task fail to explain the recognition performance. Importantly, these results offer a powerful reinterpretation of the testing effect, with key implications for education.

---

[1] De Loof, E., Naert, L., Van Opstal, F., & Verguts, T. (submitted). Signed reward prediction errors drive declarative learning.

INTRODUCTION

Declarative and procedural learning are key assets of the human brain. Ever since Thorndike (Thorndike, 1932), it has been thought that reward is crucial for both forms of learning. Additionally, inspired by the phenomenon of blocking (Kamin, 1969), Rescorla and Wagner proposed and modeled the concept that reward prediction is crucial for learning, and that learning occurs mainly for unexpected reward outcomes (i.e., reward prediction errors, RPEs; Rescorla & Wagner, 1972). Their classic model foreshadowed many decades of work to come in the conditioning literature (Mackintosh, 1975; Pearce & Hall, 1980). A recent surge of interest in this concept results from the remarkable synergy between dopaminergic recordings in the mammal brainstem (i.e., the neural signature of RPEs; Schultz, Dayan, & Montague, 1997) and the temporal-difference RPE model (Montague, Dayan, & Sejnowski, 1996; Sutton & Barto, 1998). Similar views on the role of RPEs in learning were developed in other prominent theoretical frameworks (e.g., predictive coding; Friston & Kiebel, 2009; neoHebbian account; Lisman, Grace, & Duzel, 2011). In ensuing empirical research, the effect of RPEs has been amply demonstrated in procedural learning paradigms such as classical and operant conditioning (Fiorillo, Tobler, & Schultz, 2003; Pagnoni, Zink, Montague, & Berns, 2002; Sevenster, Beckers, & Kindt, 2013; Steinberg et al., 2013). However, in these procedural learning paradigms RPEs gradually shape the acquisition of stimulus-response contingencies over multiple encounters. This is distinct from the typically human ability to learn (verbal, stimulus-stimulus) information through a single encounter by declarative learning. This type of learning is, however, costly and subject to interference. The

brain must therefore decide what information to prioritize for storage. How it achieves this remains a critical gap in the literature.

We investigated whether RPEs prioritize information during declarative learning. While RPEs have been demonstrated to enhance procedural learning, the current lack of evidence on the link between RPEs and declarative learning is a crucial gap in the literature as declarative learning is a quintessentially human form of learning that is important in everyday life (e.g., in education). Nevertheless, findings from procedural learning provide clear predictions on the role of RPEs in declarative learning. According to the neoHebbian learning framework (Lisman et al., 2011), dopamine bursts generated by the ventral tegmental area (VTA) and projected to the hippocampus amplify long term potentiation (LTP). As a consequence, phasic dopamine bursts during learning result in better memory, especially after a delay including sleep (O'Neill, Pleydell-Bouverie, Dupret, & Csicsvari, 2010). Rodent research has demonstrated that dopamine bursts enhance learning of spatial information, even through a single encounter (Bethus, Tse, & Morris, 2010). Given that dopamine is thought to implement RPEs (Cohen, Haesler, Vong, Lowell, & Uchida, 2012; Eshel et al., 2015; Montague et al., 1996), it can be expected that RPEs enhance declarative learning. To date, however, there is no direct empirical evidence for a beneficial role of RPEs in declarative learning.

In the current study we set out to experimentally manipulate RPEs in declarative learning. We administered Dutch-Swahili word pairs to participants by presenting them with a Dutch word accompanied by one, two or four possible Swahili translations to choose from (acquisition phase). By manipulating the number of available options, we manipulated the reward probability and hence the reward prediction (error). In this way, during feedback, positive and negative

RPEs of known and various sizes were coupled to the valid Dutch-Swahili word pairs (see Figure 1); allowing us to empirically test for the first time whether RPEs drive declarative learning. We subsequently probed recognition in an immediate or one-day delayed test (recognition phase). Forty participants performed the experiment (twenty in each group) after giving informed consent.

To start, we tested the predictions from the classic time-on-task account. According to this account, longer deliberation on a particular Dutch-Swahili word pair would improve its retention. Next, we distinguished two possible RPE effects. A first possibility is that signed RPEs (SRPEs) ("better than expected" signals) determine learning. This account predicts that positive RPEs (i.e., receiving a higher reward than expected) improve learning while negative RPEs (i.e., receiving a lower reward than expected) abate learning. This would be consistent for example with how SRPEs are used to train Actors in Actor-Critic models, for example using the delta-rule (Rescorla & Wagner, 1972) or the temporal-difference model (Holroyd & Coles, 2002; Sutton & Barto, 1998). A second possibility is that unsigned RPEs (URPEs) ("different than expected" signals) drive learning, with large (both positive and negative) RPEs bolstering learning. This would reflect the established role of surprise in learning (Bryden, Johnson, Tobia, Kashtelyan, & Roesch, 2011; Hayden, Heilbronner, Pearson, & Platt, 2011). Hence, by testing whether negative RPEs enhance or abate learning we empirically disentangle two theoretical accounts on how RPEs can drive declarative learning.
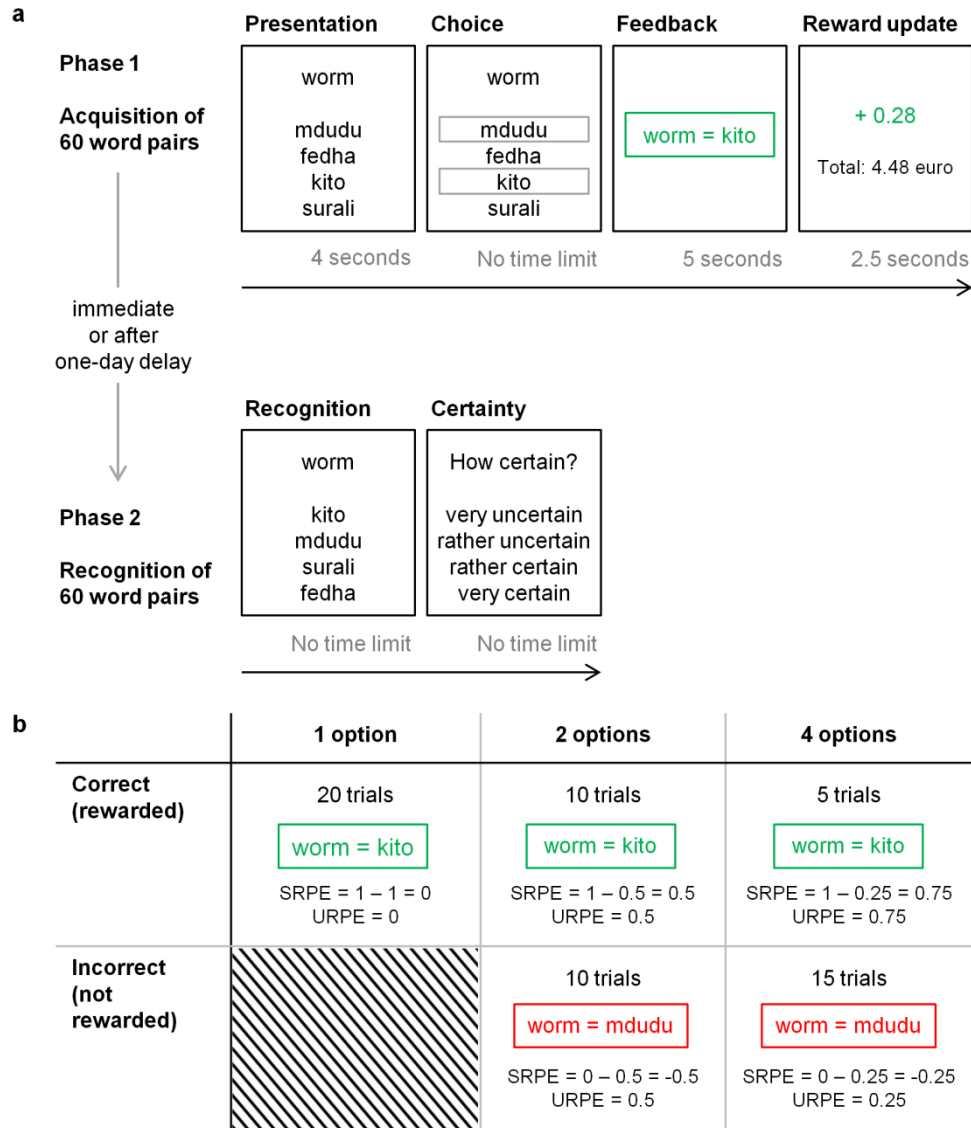
Figure 1: Experiment overview (a) and experimental design (b). (a) Participants chose between one, two or four Swahili translations in the acquisition phase; the two-option condition with correct choice is illustrated. Recognition and certainty were probed immediately or after a one-day delay. (b) The 3 (number of options) × 2 (accuracy of choice) experimental design, including number of trials and associated signed and

unsigned RPE (SRPE and URPE). SRPE was calculated by subtracting probability of reward from actual reward; URPE is the absolute value of SRPE. The feedback is illustrated assuming that the participant chose 'kito' as the translation for 'worm'.

METHODS

Participants

Forty participants (all university students; 8 male) enrolled in the study and were rewarded € 10 for their participation. All participants were naive to the purpose of the experiment, had no prior knowledge of Swahili and had not previously taken part in any experiment involving Swahili words. Half of the participants were randomly assigned to perform the recognition test immediately after the acquisition task and the other half performed the recognition test one day later. One gift voucher of € 20 was awarded to the participant with the best performance on the immediate recognition test and a second voucher was given to the participant with the best performance on the recognition test one day later. All participants signed an informed consent before the start of the experiment and were debriefed afterwards.

Material

The experiment was run on an Asus 1215N netbook running Eprime software (Schneider, Eschman, & Zuccolotto, 2012). For the declarative learning task, 60 Dutch and 240 Swahili words were selected (Table 1).

Table 1: Stimulus material

Swahili words (240)

adhabu	chupi	jeraha	kioo	maisha	msitu	nyundo	surali
adui	daima	jibini	kisiwa	maji	msumari	nyundu	takatak
afya	dakika	jikoni	kisu	mali	mtawa	nzuri	tamasha
aibu	daraja	jiwe	kitanda	mamba	mtirka	ofisi	tanuri
akili	dari	jokofu	kitande	mapafu	mundamo	osha	tembo
alizeti	dizeli	jua	kiti	mashua	mungu	panya	trekta
amani	duka	jumatu	kito	matumai	mvringo	petye	tumbili
asili	elfu	juuya	kitovu	matumbo	mvua	picha	tumbo
baadaye	farasi	kaburi	kofia	maua	mvuke	pombe	twai
bafuni	fedha	kahawa	kovuli	mazishi	mwanake	punda	uadui
bahari	filimbi	kalamu	kuacha	mbolea	mwanga	punguza	uchorai
baharia	funzi	kamba	kuandika	mbuzi	mwezi	pwani	ufagio
baiski	furaha	kamwe	kubale	mbwa	mzungu	rafiki	ugomvi
bandari	garisi	kartasi	kubwa	mchanga	nanga	rangi	uhuru
barua	geza	katika	kudhibi	mchawi	nchi	rombus	ukame
basi	godoro	kawaida	kuhesa	mchuzi	ndaniya	sabuni	ukweli
bega	goti	kazi	kujenga	mdudu	ndege	sahani	umasijo
bendi	gundi	kelele	kukimba	mechezo	ndevu	samaki	uongo
bilaska	guruwe	kemia	kumba	mekno	ndizi	sayari	usiku
bloke	haki	kengele	kumbuka	mfuko	ndogo	seesaw	uyoga
buli	hamsi	kesho	kununa	mgonjwa	ndoora	sehemu	viatu
bunifu	hasira	kiatu	kunywa	miaka	ndugu	seri	wakala
bustani	hatua	kichwa	kupanda	mkasi	neyemba	shimoni	washia
chaki	hazini	kidole	kusanya	mkate	ngazi	shule	welder
chombo	hofu	kifua	kushoto	mkoba	ngono	simu	wengine
choori	ijayo	kihozi	kusikiza	mkuu	ngozi	singizi	wimbo
chubani	imani	kijiko	kuzama	mlango	nopya	soko	wingi
chuki	ishara	kikapu	kweli	moyo	nyange	starehe	wingu
chuma	ishiri	kimysa	leso	mpishi	nyeusi	stork	yatima
chupa	jansa	kinywa	mageho	mraba	nyota	sufuria	zeituni

Dutch words (60)

agent	bord	ezel	kaas	mest	rijst	stoel	wolk
anker	brief	fiets	kassa	nacht	schat	stoom	wonde
appel	bril	goud	knie	neus	sjaal	stuur	worm
bezem	broek	graf	laken	olijf	slaap	touw	zomer
bier	brood	hamer	lamp	oven	slang	trein	
bloem	doos	haven	lepel	paard	slot	tuin	
boer	eend	hond	lijm	poort	stier	verf	
boot	emmer	hoofd	melk	regen	stift	water	

Procedure

At the start of the experiment, participants were informed about the four parts of the study: the familiarization task, the acquisition task, the filler task and the recognition test (see below for a detailed description of each part). Participants knew they could earn at least € 8 but possibly more than € 10 during the acquisition task and could receive an additional gift voucher of € 20 if they had the best recognition performance. The gift voucher was shown to the participants at the start of the experiment and again at the start of the recognition test. Participants took part in the experiment in pairs to increase their involvement in the acquisition task and recognition test.

Familiarization task. In order to familiarize the participants with the stimuli used in the experiment, all Dutch and Swahili words were presented in random order at the start of the experiment. Each word appeared at the center of the screen for two seconds. Participants were instructed to read the words in silence and push a response button when a Dutch word appeared. The familiarization task lasted about nine minutes.

Acquisition task. At the start of the acquisition task, participants were informed that they were about to learn 60 Dutch-Swahili word pairs. During this task they would be able to gain at least € 8 and possibly more than € 10. In addition they were reminded of the recognition test that would follow the experiment and of the additional gift voucher of € 20 for the participant with the best recognition performance.

At the start of each trial, one Dutch word was presented at the top of the screen with four Swahili words below (Figure 1a). All words remained on screen for four seconds and participants were instructed to read the options. Next, a

frame appeared around the possible Swahili translations for the Dutch word. In the one-option condition only one Swahili word was framed, thus immediately indicating the correct Swahili translation. In the two-option condition a frame appeared around two Swahili words so participants had a 50% chance of choosing the correct translation. Finally, in the four-option condition all four Swahili words were framed and the participant thus had a 25% chance of choosing the correct Swahili translation. Four response buttons were assigned to the four word positions and participants responded with the index and middle finger of their left and right hand. There was no time constraint on the decision but participants were encouraged to follow their first impression.

After the participants chose a Swahili translation, feedback on the correct translation was given. The Dutch word, an equation sign and the correct Swahili word appeared at the center of the screen. If the chosen Swahili translation was correct, a green frame was presented around the Dutch word and the chosen Swahili word, while participants heard the sound of money tumbling in a cup (three seconds). Alternatively, if the chosen Swahili translation was incorrect, a red frame appeared around the Dutch word and one of the other possible Swahili word options, while an error buzz was played (three seconds). The words remained on the screen for five seconds and participants were instructed to use this time to learn the word pair by heart for the recognition test. Finally, the trial ended with a 2.5 seconds presentation of the total amount of reward collected thus far. Participants won € 0.28 on correct trials; no money was added on error trials.

Filler task. A brief filler task was inserted to reduce recency effects in the immediate recognition test. In order to keep both versions of the experiment as similar as possible, the filler task was also presented to participants who would

perform the recognition test one day later. Participants categorized numbers between 1 and 9 (excluding 5) as being smaller or larger than 5 (left and right button presses respectively). A total of 400 numbers were presented and the filler task took about four minutes.

Recognition test. At the start of the recognition test, participants were reminded about the additional gift voucher of € 20 for the participant with the best recognition performance. The display layout for the recognition test was similar to that of the acquisition task. The Dutch word appeared at the top of the screen with the same four Swahili words below. However, the order of the four Swahili words was randomized and participants were warned about this change. As soon as the words appeared, participants could choose between the four Swahili words by using the same four response buttons as in the acquisition task. No time constraints were imposed on their answer. After a Swahili word was chosen, participants indicated how certain they were about their answer: 'very uncertain', 'rather uncertain', 'rather certain' or 'very certain' (measured on a scale from 1 'very uncertain' to 4 'very certain'). No feedback was given about the accuracy of their Swahili translation.

Design

Unbeknownst to the participants, the accuracy of the chosen translations in the acquisition task was determined in advance. More specifically, a fixed number of trials was assigned to each cell of the design and trials were predetermined to have one, two or four valid Swahili options and to be correct or incorrect (Figure 1b). Note that by predetermining whether a translation would be correct or not, participants did not necessarily learn the correct Swahili translations of the Dutch words. For example, if a trial had been determined to be a two-option trial with a correct answer, the participants would be rewarded

irrespective of their choice and this chosen word would be the translation they had to memorize for the recognition test. Moreover, for each Dutch word four randomly drawn Swahili words were presented, often not including the actual translation. Participants were informed about this manipulation afterwards.

The SRPEs were calculated by subtracting the reward probability (i.e., 1, 0.5 and 0.25 probability of a correct choice in the one-, two- and four-option condition, respectively) from the actual received reward (i.e., 1 reward on correct trials and 0 reward on incorrect trials). Thus a unique SRPE was calculated for each cell in the design, ranging from -0.50 to 0.75 (see Figure 1b for a full overview). The URPEs were calculated by taking the absolute value of the SRPEs. Note that this merely reverses the sign of the RPEs for the unrewarded word pairs.

RESULTS

Three participants were removed from the dataset because of technical failures during the experiment. Further analyses are performed on the data of the remaining 19 participants in the immediate recognition test group and 18 in the delayed recognition test group. All participants performed the acquisition task according to the instructions, choosing a valid option in 99.5% of the trials (trials with invalid choices were excluded from further analyses). As participants were free to deliberate on their answer in the recognition test without any time restriction, all further analyses will focus exclusively on recognition accuracies and certainty ratings.

The effect of reward and number of options

First, we probed how the reward and the number of options in the acquisition task influenced the recognition rate. Recognition accuracies were entered into a generalized linear mixed effects model with a random intercept across participants and the test delay, the number of options and the reward as fixed effects predictors. Recognition accuracy was significantly higher in the immediate test than in the delayed test, $\chi^2(1, N = 37) = 15.7$, $p < 0.001$. Recognition accuracies ranged from 40% to 90% (M = 67.4%, SD = 14.4%) for the immediate test group and from 27% to 73% (M = 50.7%, SD = 11.6%) for the delayed test group. As Figure 2 reveals, there was a significant main effect of reward, $\chi^2(1, N = 37) = 24.5$, $p < 0.001$, with rewarded choices being remembered more accurately. In addition, the recognition rate significantly increased with an increasing number of options, $\chi^2(1, N = 37) = 36.8$, $p < 0.001$.
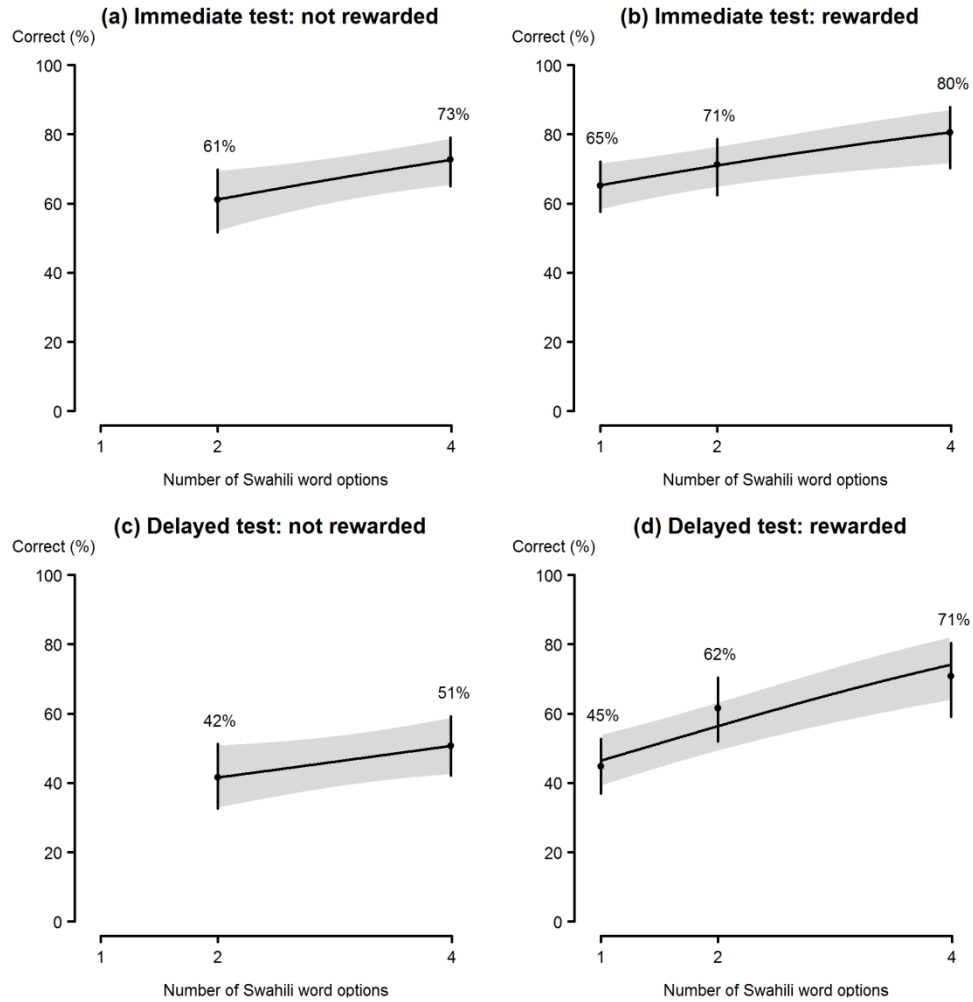
Figure 2: For the unrewarded (panel a) and rewarded word pairs (panel b) in the immediate test group and their equivalent in the delayed test (panels c and d), the recognition rate (y-axis) is plotted as a function of the number of options (x-axis). Note that in the one-option condition the chosen translation was always rewarded. Recognition performance was higher on the immediate test compared to the delayed test and a generalized linear mixed effects model revealed a significant increase in recognition rate with an increasing number of options and a better performance for rewarded word pairs

(black regression line with gray 95% confidence band). For each number of options and depending on the reward and delay, the average RT and 95% confidence interval was estimated and superimposed.

Classic declarative learning mechanisms: time-on-task

We verified whether variations in time-on-task could account for the positive effect of number of options on recognition. As participants were allowed an unlimited amount of time to choose between the Swahili word options during the acquisition task, variations in time-on-task could result in better recognition. However, the time-on-task account failed to explain the positive effect of the number of options as revealed by the following two additional tests.

First, we tested whether longer deliberation on individual trials would lead to better recognition. To approximate the time devoted to each option (word) on a particular trial, we divided the time-on-task by the number of options. The resulting time-on-task per option (time-on-word) revealed that each option was examined longer when less options were available (the mean time-on-word on the one-, two- or four-option trials was 2880 ms, 1826 ms and 1169 ms, respectively). Next, we tested whether increased time-on-word would lead to better recognition. Recognition accuracies were entered into a generalized linear mixed effects model with a random intercept across participants and the test delay and time-on-word as fixed effects predictors. Counter to the predictions from the time-on-task account, there was no significant influence of time-on-word on recognition, $\chi^2(1, N = 37) = 1.48$, $p = 0.22$. Follow-up tests for one-, two- or four-option trials separately confirmed that recognition was not significantly influenced by the time-on-word (one-option trials, $\chi^2(1, N = 37) = 0.096$, $p = 0.76$; two-option trials, $\chi^2(1, N = 37) = 0.026$, $p = 0.87$; four-option trials, $\chi^2(1, N = 37) = 2.52$, $p = 0.11$). The result of the one-option trials is particularly

interesting as participant could already start learning the word pair during the deliberation time. Still, even in the one-option condition longer deliberation on the valid Dutch-Swahili word pair failed to result in better declarative learning.

Second, we tested whether participants who deliberated longer recognized more word pairs. The recognition accuracies were entered into a generalized linear mixed effects model with a random intercept across participants and the test delay and average time-on-task per participant as fixed effects predictors. The results show that recognition was not significantly influenced by the average time-on-task across participants, $\chi^2(1, N = 37) = 0.0022$, $p = 0.96$. The same analysis was repeated separately for trials with one, two or four options. None of these analyses revealed an influence of the average time-on-task across participants on the recognition accuracies (one-option trials, $\chi^2(1, N = 37) = 0.35$, $p = 0.55$; two-option trials, $\chi^2(1, N = 37) = 0.57$, $p = 0.45$; four-option trials, $\chi^2(1, N = 37) = 1.69$, $p = 0.19$). Again, it is interesting to note that even in the one-option trials, people who examined word pairs longer had no increased recognition accuracy. In sum, declarative learning was not accurately predicted by the time-on-task account, as allotting more time to the word pairs had no impact on subsequent recognition.

The effect of positive and negative RPEs

To further investigate the effect of RPEs on declarative learning, we disentangled the role of positive and negative RPEs by probing the interaction between the number of options and reward. Recognition accuracies were entered into a generalized linear mixed effects model with a random intercept across participants and the test delay, the number of options and the reward as fixed effects predictors. Contrary to the URPE account, but consistent with the SRPE account, there was no significant interaction between the number of options and

reward, $\chi^2(1, N = 37) = 1.42$, $p = 0.23$. There were also no interactions with the test delay (all $p > 0.20$).

The same pattern of results were obtained when the certainty ratings were entered into a linear mixed effects model with a random intercept across participants and the test delay, recognition accuracy, the number of options and the reward as fixed effects predictors. There was a main effect of the number of options, $\chi^2(1, N = 37) = 27.7$, $p < 0.001$, a main effect of reward, $\chi^2(1, N = 37) = 8.38$, $p = 0.0038$, and no significant interaction between the number of options and reward, $\chi^2(1, N = 37) = 1.29$, $p = 0.26$, or any other significant interactions involving the number of options or the reward (all $p > 0.10$).

So far, the data failed to support the URPE account as there was no significant interaction between the number of options and the reward. Moreover, the positive effect of the number of Swahili word options on recognition rates in the unrewarded trials (Figure 2a and 2c) supports the SRPE account and refutes the URPE account. To pit the SRPE and URPE accounts against one another more directly, the recognition accuracies of the unrewarded trials were entered into a generalized linear mixed effects model with a random intercept across participants and the test delay and the number of options as fixed effects predictors. In line with the positive effect suggested in Figure 2a and 2c, there was a significant effect of the number of options, $\chi^2(1, N = 37) = 9.49$, $p = 0.0021$, with higher recognition rates for trials associated with more Swahili word options. There was also a significant main effect of test delay, $\chi^2(1, N = 37) = 18.09$, $p < 0.001$, but there was no significant interaction between the test delay and the number of options, $\chi^2(1, N = 37) = 0.25$, $p = 0.61$.

SRPEs as a linear predictor of declarative learning

To formally probe whether declarative memory performance increases linearly with SRPEs, recognition accuracies were entered into a generalized linear mixed effects model with a random intercept across participants and the test delay and SRPEs as fixed effects predictors. As Figure 3a and 3b illustrates, recognition improved significantly with increasing SRPEs, $\chi^2(1, N = 37) = 27.4$, $p < 0.001$. The interaction between SRPEs and delay was not significant, $\chi^2(1, N = 37) = 2.52$, $p = 0.11$, and Figure 3a and 3b clearly illustrates a significant effect of SRPEs in the immediate test group, $\chi^2(1, N = 19) = 6.20$, $p = 0.013$, and an even stronger effect in the delayed test group, $\chi^2(1, N = 18) = 23.6$, $p < 0.001$. Note that the recognition in the delayed test is especially lower for negative RPEs but remains high for positive RPEs.

The linear effect of SRPEs on declarative learning was further validated by the certainty ratings (Figure 3c and 3d). The certainty ratings were entered into a linear mixed effects model with a random intercept across participants and the test delay, recognition accuracy and SRPEs as fixed effects predictors. In line with the effect of SRPEs on recognition, there was a significant main effect of SRPEs on the certainty ratings, $\chi^2(1, N = 37) = 9.49$, $p = 0.0021$, with higher SRPEs leading to higher certainty ratings. While there was no significant interaction between the effect of SRPEs and test delay, $\chi^2(1, N = 37) = 0.039$, $p = 0.84$, there was a significant interaction between SRPEs and recognition accuracy, $\chi^2(1, N = 37) = 4.56$, $p = 0.033$, and a marginally significant three-way interaction, $\chi^2(1, N = 37) = 3.25$, $p = 0.071$. Follow-up tests revealed that SRPEs had no significant effect on certainty ratings for the false recognitions (neither in the immediate test group, $\chi^2(1, N = 19) = 1.76$, $p = 0.18$, nor the delayed test group, $\chi^2(1, N = 18) = 2.021$, $p = 0.16$), but did significantly predict certainty

ratings for the correctly recognized word pairs in the immediate, $\chi^2(1, N = 19) = 4.24$, $p = 0.039$, and delayed test group, $\chi^2(1, N = 18) = 7.27$, $p = 0.0070$. The fact that the SRPEs only influence certainty ratings for the correct recognitions and not for false alarms further corroborates our finding that SRPEs indeed drive declarative learning.
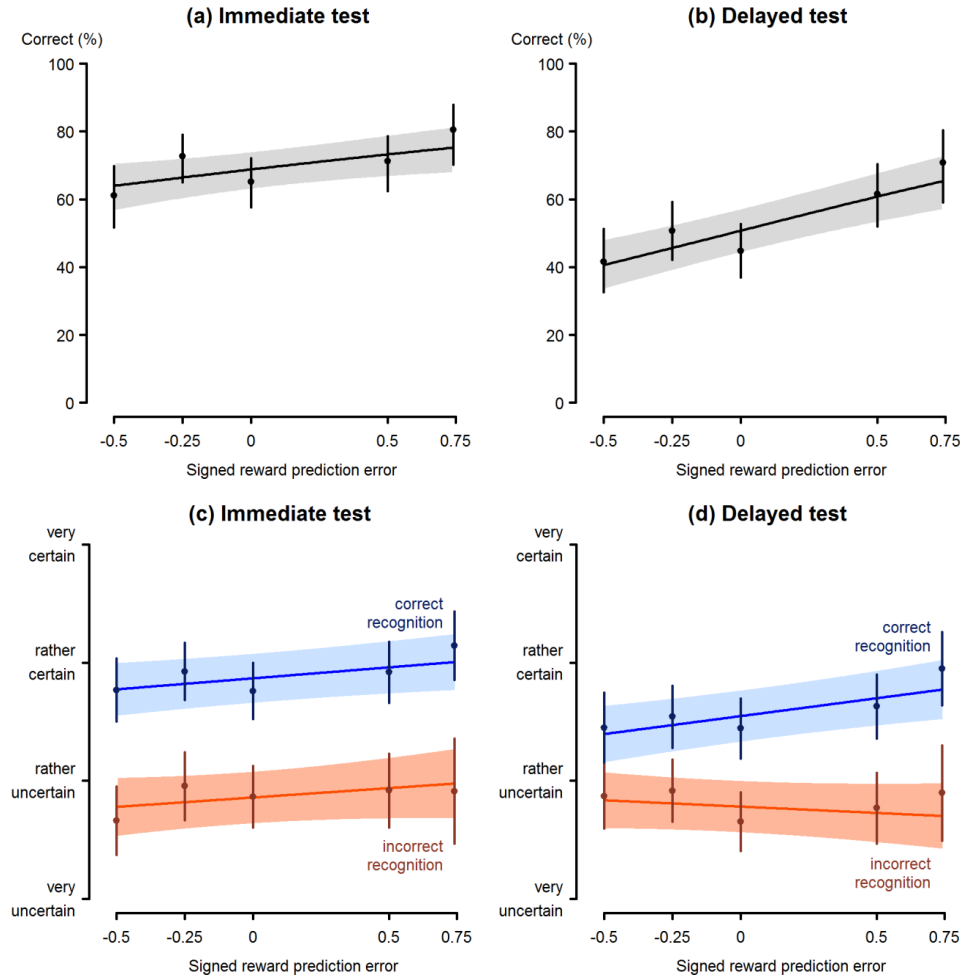
Figure 3: For the immediate and delayed test, the recognition (panel a and b; y-axis) and certainty ratings (panel c and d; y-axis) are plotted as a function of the SRPEs (x-axis). For the word pairs associated with each SRPE, the average recognition and certainty ratings and their 95% confidence intervals were superimposed. (a and b) Recognition increased significantly with SRPEs (black regression line with gray 95% confidence band). (c and d) Certainty ratings for the false recognitions (depicted in orange) were not significantly different across SRPEs, while higher SRPEs led to significantly higher certainty ratings for correctly recognized word pairs (depicted in blue).

DISCUSSION

Our results provide the first empirical demonstration that RPEs prioritize declarative learning, in both immediate and delayed recognition. Positive and negative RPEs of known and various sizes were generated by manipulating the number of options available during word pair acquisition. Our analysis revealed that larger and more positive RPEs lead to better recognition and increased certainty of the word pairs, as evident in both an immediate and delayed recognition test. Thus, while the importance of URPEs ("different than expected" signals; Bryden et al., 2011; Hayden et al., 2011) has been shown in procedural learning paradigms (Pearce & Hall, 1980; Sevenster et al., 2013; Steinberg et al., 2013) our analysis indicates that declarative learning is driven by SRPEs ("better than expected" signals).

These results further our understanding of how motivational cues determine which information is prioritized for encoding in memory. As discussed previously, the neoHebbian learning framework (Lisman et al., 2011) describes how phasic dopamine bursts enhance hippocampal LTP by strengthened consolidation. Consolidation is thought to occur especially during sleep, which may explain why our effects were stronger after a night of sleep. Critically, these dopamine bursts can be caused by a variety of motivational cues such as RPEs, novelty and salience. Previous research has indeed demonstrated the effect of reward anticipation in declarative learning (Adcock, Thangavel, Whitfield-Gabrieli, Knutson, & Gabrieli, 2006; Wittmann et al., 2005), the effect of exposure to novel environments (Fenker et al., 2008) and the effect of salient (emotional) stimuli more generally (Anderson, Wais, & Gabrieli, 2006).

Critically, we provide a first empirical validation of the effect of RPEs on declarative memory.

In addition to an empirical validation of the effect of RPEs on declarative memory, our results also offer an alternative explanation for established learning strategies such as learning by testing (i.e., the testing effect; Gates, 1917). In a seminal publication, Karpicke and Roediger empirically manipulated the amount of study and test trials allotted to Swahili-English word pairs (Karpicke & Roediger, 2008). In a follow-up test one week later, the authors found that additional study trials during the acquisition session had no strong beneficial effect on retention. Conversely, recall was strongly enhanced by additional test trials during acquisition. Although this testing effect has consistently been observed to drive declarative learning and holds major educational implications (Howard-Jones, 2014), its origin has remained unclear. Based on our current findings, we argue that the testing effect could be ascribed to RPEs. Specifically, a test crucially involves making a prediction about what the correct answer might be and about the probability of success. These active predictions and their entailing RPEs may drive declarative learning (even in the absence of external feedback; Aarts, Houwer, & Pourtois, 2012; Schouppe et al., 2015). An interesting case in point is a study in which participants learned cue-target word pairs with a strong or weak semantic association (Carpenter, 2009). Whereas restudying the material equally improved the retention of strongly and weakly associated word pairs, repeated testing improved recall of weakly associated work pairs more compared to strongly associated words. Moreover, by the final test the recall for the weak semantic associations surpassed that of the strong semantic associations. Although counter-intuitive at first glance, these findings follow naturally from the beneficial effect of RPEs on declarative learning. More broadly, the natural occurrence of RPEs during learning might be why testing,

elaborative interrogation and self-explanation outperform other active learning strategies such as summarizing, keyword mnemonics and imagery (Dunlosky, Rawson, Marsh, Nathan, & Willingham, 2013).

In sum, we demonstrate for the first time that SRPEs drive declarative learning, closing the gap between research on reward learning and declarative memory. Our results are in line with the neoHebbian learning framework and suggest new avenues to improve learning in both informal and educational settings.

REFERENCES

Aarts, K., Houwer, J. De, & Pourtois, G. (2012). Evidence for the automatic evaluation of self-generated actions. *Cognition*, *124*(2), 117–127. doi:10.1016/j.cognition.2012.05.009

Adcock, R. A., Thangavel, A., Whitfield-Gabrieli, S., Knutson, B., & Gabrieli, J. D. E. (2006). Reward-motivated learning: Mesolimbic activation precedes memory formation. *Neuron*, *50*(3), 507–517. doi:10.1016/j.neuron.2006.03.036

Anderson, A. K., Wais, P. E., & Gabrieli, J. D. E. (2006). Emotion enhances remembrance of neutral events past. *Proceedings of the National Academy of Sciences of the United States of America*, *103*(5), 1599–1604. doi:10.1073/pnas.0506308103

Bethus, I., Tse, D., & Morris, R. G. M. (2010). Dopamine and memory: Modulation of the persistence of memory for novel hippocampal NMDA receptor-dependent paired associates. *The Journal of Neuroscience : The Official Journal of the Society for Neuroscience*, *30*(5), 1610–1618. doi:10.1523/JNEUROSCI.2721-09.2010

Bryden, D. W., Johnson, E. E., Tobia, S. C., Kashtelyan, V., & Roesch, M. R. (2011). Attention for learning signals in anterior cingulate cortex. *The Journal of Neuroscience : The Official Journal of the Society for Neuroscience*, *31*(50), 18266–18274. doi:10.1523/JNEUROSCI.4715-11.2011

Carpenter, S. K. (2009). Cue strength as a moderator of the testing effect: the benefits of elaborative retrieval. *Journal of Experimental Psychology. Learning, Memory, and Cognition*, *35*(6), 1563–1569. doi:10.1037/a0017021

Cohen, J. Y., Haesler, S., Vong, L., Lowell, B. B., & Uchida, N. (2012, February 2). Neuron-type-specific signals for reward and punishment in the ventral tegmental area. *Nature*. doi:10.1038/nature10754

Dunlosky, J., Rawson, K. A., Marsh, E. J., Nathan, M. J., & Willingham, D. T. (2013). Improving students' learning with effective learning techniques: Promising directions from cognitive and educational psychology. *Psychological Science in the Public Interest*, *14*(1), 4–58. doi:10.1177/1529100612453266

Eshel, N., Bukwich, M., Rao, V., Hemmelder, V., Tian, J., & Uchida, N. (2015). Arithmetic and local circuitry underlying dopamine prediction errors. *Nature*, *525*(7568), 243–246. doi:10.1038/nature14855

Fenker, D. B., Frey, J. U., Schuetze, H., Heipertz, D., Heinze, H.-J., & Duzel, E. (2008). Novel scenes improve recollection and recall of words. *Journal of Cognitive Neuroscience*, *20*(7), 1250–1265. doi:10.1162/jocn.2008.20086

Fiorillo, C. D., Tobler, P. N., & Schultz, W. (2003). Discrete coding of reward probability and uncertainty by dopamine neurons. *Science (New York, N.Y.)*, *299*(5614), 1898–1902. doi:10.1126/science.1077349

Friston, K., & Kiebel, S. (2009). Predictive coding under the free-energy principle. *Philosophical Transactions of the Royal Society of London. Series B, Biological Sciences*, *364*(1521), 1211–1221. doi:10.1098/rstb.2008.0300

Gates, A. I. (1917). Recitation as a factor in memorization. *Archives of Psychology*, *6*(40), 112–121.

Hayden, B. Y., Heilbronner, S. R., Pearson, J. M., & Platt, M. L. (2011). Surprise signals in anterior cingulate cortex: Neuronal encoding of unsigned reward prediction errors driving adjustment in behavior. *The Journal of Neuroscience : The Official Journal of the Society for Neuroscience*, *31*(11), 4178–4187. doi:10.1523/JNEUROSCI.4652-10.2011

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*(4), 679–709. doi:10.1037//0033-295X.109.4.679

Howard-Jones, P. A. (2014). Neuroscience and education: myths and messages. *Nature Reviews. Neuroscience*, *15*(12), 817–824. doi:10.1038/nrn3817

Kamin, L. J. (1969). Predictability, surprise, attention, and conditioning. *Punishment and Aversive Behavior*, 279–296.

Karpicke, J. D., & Roediger, H. L. (2008). The critical importance of retrieval for learning. *Science (New York, N.Y.)*, *319*(5865), 966–968. doi:10.1126/science.1152408

Lisman, J., Grace, A. A., & Duzel, E. (2011). A neoHebbian framework for episodic memory; role of dopamine-dependent late LTP. *Trends in Neurosciences*, *34*(10), 536–547. doi:10.1016/j.tins.2011.07.006

Mackintosh, N. J. (1975). A theory of attention: Variations in the associability of stimuli with reinforcement. *Psychological Review*, *82*(4), 276–298.

Montague, P. R., Dayan, P., & Sejnowski, T. J. (1996). A framework for mesencephalic dopamine systems based on predictive Hebbian learning. *The Journal of Neuroscience : The Official Journal of the Society for Neuroscience*, *16*(5), 1936–1947.

O'Neill, J., Pleydell-Bouverie, B., Dupret, D., & Csicsvari, J. (2010). Play it again: reactivation of waking experience and memory. *Trends in Neurosciences*, *33*(5), 220–229. doi:doi:10.1016/j.tins.2010.01.006

Pagnoni, G., Zink, C. F., Montague, P. R., & Berns, G. S. (2002). Activity in human ventral striatum locked to errors of reward prediction. *Nature Neuroscience*, *5*(2), 97–98. doi:10.1038/nn802

Pearce, J. M., & Hall, G. (1980). A model for Pavlovian learning: Variations in the effectiveness of conditioned but not of unconditioned stimuli. *Psychological Review*, *87*, 532–552. doi:10.1037//0033-295X.87.6.532

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

Schneider, W., Eschman, A., & Zuccolotto, A. (2012). *E-Prime User's Guide.* Pittsburgh: Psychology Software Tools, Inc.

Schouppe, N., Braem, S., De Houwer, J., Silvetti, M., Verguts, T., Ridderinkhof, K. R., & Notebaert, W. (2015). No pain, no gain: the affective valence of congruency conditions changes following a successful response. *Cognitive, Affective, & Behavioral Neuroscience*, *15*(1), 251–261. doi:10.3758/s13415-014-0318-3

Schultz, W., Dayan, P., & Montague, R. (1997). A neural substrate of prediction and reward. *Science (New York, N.Y.)*, *275*(5306), 1593–1599. doi:10.1126/science.275.5306.1593

Sevenster, D., Beckers, T., & Kindt, M. (2013). Prediction error governs pharmacologically induced amnesia for learned fear. *Science (New York, N.Y.)*, *339*(6121), 830–833. doi:10.1126/science.1231357

Steinberg, E. E., Keiflin, R., Boivin, J. R., Witten, I. B., Deisseroth, K., & Janak, P. H. (2013). A causal link between prediction errors, dopamine neurons and learning. *Nature Neuroscience*, *16*(7), 966–973. doi:10.1038/nn.3413

Sutton, R. S., & Barto, A. G. (1998). *Reinforcement Learning: An Introduction*. Cambridge, MA: MIT Press.

Thorndike, E. (1932). *The fundamentals of learning*. New York: Teachers College Press.

Wittmann, B. C., Schott, B. H., Guderian, S., Frey, J. U., Heinze, H.-J., & Düzel, E. (2005). Reward-related FMRI activation of dopaminergic midbrain is associated with enhanced hippocampus-dependent long-term memory formation. *Neuron*, *45*(3), 459–467. doi:10.1016/j.neuron.2005.01.010

CHAPTER 4
THE MODULATION OF EEG OSCILLATIONS BY REWARD PREDICTION ERRORS DRIVES DECLARATIVE LEARNING [1]

The effect of reward prediction errors (RPEs) on learning has been well established, mainly in procedural learning paradigms. It is generally assumed that RPEs also drive declarative learning, a typically human form of learning widely used in everyday life (e.g., in education). However, little empirical evidence has validated this claim. In a previous study (chapter 3) we therefore provided the first empirical demonstration of enhanced declarative learning due to RPEs. By parametrically manipulating the RPEs experienced during reward feedback we were able to boost the acquisition of Dutch-Swahili word pairs. In the current time-frequency EEG study we verify whether participants experience RPEs during reward feedback and whether declarative learning is indeed driven by these RPEs, rather than by related acquisition processes. The behavioral results offer a full replication of the previous study: Word pairs associated with a large, positive RPE are recognized with higher accuracy and certainty. Additionally, our results confirm that RPEs modulate EEG oscillations during reward feedback, confirming the experience of a "better-than-expected" signal. Moreover, the alpha suppression during reward feedback was predictive for the recognition performance in the delayed recognition test. In sum, our results confirm that the neural mechanisms activated by RPEs during procedural learning are also activated during declarative learning. In addition, the activation of this reward mechanism predicted declarative recognition in a follow-up test.

INTRODUCTION

Reward is one of the most important arbiters to determine what information should be prioritized for learning. Research in reward learning (Rescorla & Wagner, 1972), computational modeling (Montague, Dayan, & Sejnowski, 1996) and single-cell recoding (Schultz, Dayan, & Montague, 1997), has amply demonstrated that learning is not predominantly regulated by obtained rewards but mainly by the difference between the obtained and expected reward (the reward prediction error, RPE). However, this research has been largely restricted to procedural learning, probably because of its long tradition in animal research (Squire, 2004). In these procedural learning paradigms RPEs gradually shape the acquisition of stimulus-response contingencies over multiple encounters. This is distinct from the typically human ability to learn (verbal, stimulus-stimulus) information through a single encounter by declarative learning. Although it is often suggested that RPEs also boost declarative learning (den Ouden, Friston, Daw, McIntosh, & Stephan, 2009; Hyman, Malenka, & Nestler, 2006), this claim had not been verified in empirical research on humans. Interestingly, a number of well-established effects in declarative learning hint at an involvement of RPEs. One of these effects is the testing effect (Gates, 1917; Karpicke & Roediger, 2008) which entails that learned information is quickly forgotten after mere repeated practice but remains high when one is being tested on the acquired knowledge. From a reward learning perspective, this testing effect could be interpreted as the effect of RPEs, as especially during testing predictions may be generated and evaluated.

To bridge this gap, in a previous declarative learning experiment (chapter 3) we parametrically manipulated RPEs during the acquisition of 60 Dutch-Swahili word pairs. By letting participants choose among one, two or four possible Swahili translations for a Dutch word, we varied the reward probability and hence the ensuing reward prediction (error). In this way, during feedback, RPEs of manipulated (and known) intensity were coupled to the Dutch-Swahili word pairs, allowing us to empirically test for the first time whether RPEs drive declarative learning. By probing word pair recognition in an immediate and delayed test we demonstrated that declarative learning was indeed boosted by RPEs. More specifically, word pairs that were accompanied by a large RPE resulted in better recognition in a follow-up recognition test, both in terms of accuracy and certainty ratings.

In addition to providing a first experimental demonstration that RPEs indeed play a crucial role in declarative learning, we also excluded a number of alternative interpretations. To start, the alternative time-on-task interpretation (Hebb, 1949) was excluded by demonstrating that longer deliberation (e.g., caused by choosing among one, two or four options) failed to foster better recognition. We also formally differentiated between the influence of negative and positive RPEs on recognition. A first theoretical possibility we envisioned, was that only the size of the RPE matters and not its valence. These unsigned RPEs (URPEs) can be interpreted as "different-than-expected" signals and their influence on learning parallels the role of surprise in learning (Bryden, Johnson, Tobia, Kashtelyan, & Roesch, 2011; Hayden, Heilbronner, Pearson, & Platt, 2011). In a second possible scenario, large positive RPEs would enhance learning while large negative RPEs hinder learning. These signed RPEs (SRPEs) can thus be interpreted as "better-than-expected" signals, reflecting how RPEs are used to train Actors

in Actor-Critic models (Sutton & Barto, 1998), for example using the temporal-difference model (Holroyd & Coles, 2002; Montague et al., 1996).

The results of our previous study (chapter 3) clearly support the role of RPEs as "better-than-expected" signals, as recall was driven by SRPEs. Indeed, recognition accuracy and certainty ratings were mainly enhanced by large positive RPEs but not by large negative RPEs. These results were overall in line with the neoHebbian learning framework (Lisman, Grace, & Duzel, 2011) which details the way in which SRPEs can modulate memory formation. According to this framework, SRPEs cause the ventral tegmental area (VTA) to trigger a dopaminergic pathway with projections to the hippocampus. Memory traces in the hippocampus that were active when the dopaminergic pathway was activated are consolidated more strongly through long term potentiation (LTP) during sleep. This LTP account fits nicely with the data from our previous study, where the effect of SRPEs on learning was even more pronounced after a one-day delay. It is also consistent with the literature on the testing effect (Karpicke & Blunt, 2011; Karpicke & Roediger, 2008), where the advantages of being tested become more apparent in the recognition test as time progresses (e.g., after a delay of one week compared to an immediate test).

However, to confirm that the experience of a SRPE is the actual driver of improved declarative learning, we need to demonstrate that participants are actually experiencing a SRPE; and second we need to demonstrate that these experienced SRPEs indeed predict improved recognition. To do so, we adapted our declarative learning paradigm (chapter 3) to an EEG design suited for time-frequency analysis. This EEG experiment not only allowed us to replicate the behavioral findings from our previous study, it also

offered a validation of the proposed neural underpinnings of the effect of SRPEs on declarative learning by detailing more thoroughly how reward feedback influences word pair acquisition. First, we hypothesized that SRPEs would modulate EEG oscillations during reward feedback, reflecting the experience of a "better-than-expected" signal. In line with previous research on the modulation of the oscillatory power due to RPEs, we expected to see the influence of RPEs reflected in the theta (4-8 Hz; Cohen, Elger, & Ranganath, 2007), alpha (8-12 Hz; Oya et al., 2005) or high-beta frequency band (20-30 Hz; HajiHosseini & Holroyd, 2015a). Second, we tested whether the oscillatory power during acquisition could predict the accuracy and certainty ratings during the recognition test. Importantly, compared to the previous study we differentiated the feedback during the acquisition task into three separate components: reward anticipation, reward feedback (correct versus incorrect) and word pair encoding. This allowed us to pinpoint in which phase SRPEs were reflected by oscillatory power modulations and in which phase these power modulations were predictive for improved performance in the recognition test. Based on the reinforcement learning framework, we hypothesized that SRPEs influence learning during reward feedback. Alternatively, the SRPEs could influence declarative learning in a more indirect fashion, e.g. through adaptive processes during word pair encoding (Ridderinkhof, Van Den Wildenberg, Segalowitz, & Carter, 2004).

METHOD

Participants

Forty-one Dutch speaking, right-handed participants were recruited for the study. All participants had normal color vision, were naive to the purpose of the experiment, had no prior knowledge of Swahili and had not previously been enrolled in any experiment involving Swahili words. A random selection of twenty participants (15 female, 25.6 years on average) performed the recognition test immediately after the acquisition task and the other twenty-one participants (18 female, 25.5 years on average) performed the recognition test one day later. The participants were rewarded € 25 for their participation. An additional gift voucher of € 20 was awarded to the participant with the best performance on the recognition test. All participants signed an informed consent at the start of the experiment and received a debriefing afterwards.

Material

For the declarative learning task, the same 60 Dutch and 240 Swahili words were used as in the previous study (Table 1). The experiment was run on a Dell Optiplex 9010 mini-tower running Eprime software (Schneider, Eschman, & Zuccolotto, 2012). Button presses were registered through a Cedrus RB-730 response box enhanced with four time-accurate push buttons (Cedrus Corporation, San Pedro, California).

Table 1: Stimulus material

Swahili words (240)

adhabu	chupi	jeraha	kioo	maisha	msitu	nyundo	surali
adui	daima	jibini	kisiwa	maji	msumari	nyundu	takatak
afya	dakika	jikoni	kisu	mali	mtawa	nzuri	tamasha
aibu	daraja	jiwe	kitanda	mamba	mtirka	ofisi	tanuri
akili	dari	jokofu	kitande	mapafu	mundamo	osha	tembo
alizeti	dizeli	jua	kiti	mashua	mungu	panya	trekta
amani	duka	jumatu	kito	matumai	mvringo	petye	tumbili
asili	elfu	juuya	kitovu	matumbo	mvua	picha	tumbo
baadaye	farasi	kaburi	kofia	maua	mvuke	pombe	twai
bafuni	fedha	kahawa	kovuli	mazishi	mwanake	punda	uadui
bahari	filimbi	kalamu	kuacha	mbolea	mwanga	punguza	uchorai
baharia	funzi	kamba	kuandika	mbuzi	mwezi	pwani	ufagio
baiski	furaha	kamwe	kubale	mbwa	mzungu	rafiki	ugomvi
bandari	garisi	kartasi	kubwa	mchanga	nanga	rangi	uhuru
barua	geza	katika	kudhibi	mchawi	nchi	rombus	ukame
basi	godoro	kawaida	kuhesa	mchuzi	ndaniya	sabuni	ukweli
bega	goti	kazi	kujenga	mdudu	ndege	sahani	umasijo
bendi	gundi	kelele	kukimba	mechezo	ndevu	samaki	uongo
bilaska	guruwe	kemia	kumba	mekno	ndizi	sayari	usiku
bloke	haki	kengele	kumbuka	mfuko	ndogo	seesaw	uyoga
buli	hamsi	kesho	kununa	mgonjwa	ndoora	sehemu	viatu
bunifu	hasira	kiatu	kunywa	miaka	ndugu	seri	wakala
bustani	hatua	kichwa	kupanda	mkasi	neyemba	shimoni	washia
chaki	hazini	kidole	kusanya	mkate	ngazi	shule	welder
chombo	hofu	kifua	kushoto	mkoba	ngono	simu	wengine
choori	ijayo	kihozi	kusikiza	mkuu	ngozi	singizi	wimbo
chubani	imani	kijiko	kuzama	mlango	nopya	soko	wingi
chuki	ishara	kikapu	kweli	moyo	nyange	starehe	wingu
chuma	ishiri	kimysa	leso	mpishi	nyeusi	stork	yatima
chupa	jansa	kinywa	mageho	mraba	nyota	sufuria	zeituni

Dutch words (60)

agent	bord	ezel	kaas	mest	rijst	stoel	wolk
anker	brief	fiets	kassa	nacht	schat	stoom	wonde
appel	bril	goud	knie	neus	sjaal	stuur	worm
bezem	broek	graf	laken	olijf	slaap	touw	zomer
bier	brood	hamer	lamp	oven	slang	trein	
bloem	doos	haven	lepel	paard	slot	tuin	
boer	eend	hond	lijm	poort	stier	verf	
boot	emmer	hoofd	melk	regen	stift	water	

Electrophysiological data were recorded using a Biosemi ActiveTwo system (Biosemi, Amsterdam, Netherlands) with 64 Ag-AgCl electrodes arranged in the standard international 10-20 electrode mapping, with a posterior CMS-DRL electrode pair. Two reference electrodes were positioned at the left and right mastoids. Eye movements were registered with a pair of electrodes above and below the left eye and two additional electrodes at the outer canti of both eyes. Because of the planned time-frequency analysis, EEG signals were recorded at a 1024 Hz sampling rate.

Behavioral task

Analogous to the previous study (chapter 3), the experiment consisted of four parts: the familiarization task, the acquisition task, the filler task and the recognition test (either immediate or after a one-day delay). Due to the planned time-frequency analysis, the performance feedback in the acquisition task and recognition test was spaced out more widely in time (see below). Apart from these adaptations to the timing of the experiment to accommodate the analysis requirements, the current experiment provides an integral replication of the previous study.

At the start of the experiment, participants were informed they could earn at least € 20 but possibly more than € 25 during the acquisition task. Furthermore, a gift voucher was presented that would be awarded to the participant with the best performance in the recognition test. One gift voucher of € 20 was awarded to the participant with the best performance on the immediate recognition test and a second voucher was given to the participant with the best performance on the recognition test one day later. The voucher was presented for a second time at the start of the recognition test to remind participants of this additional incentive.

Familiarization task. All Dutch and Swahili words were presented in random order at the start of the experiment. Each word was displayed for two seconds at the center of the screen and participants were asked to read the words in silence. To ensure task performance, a button press was required for each Dutch word. The familiarization task lasted about ten minutes.

Acquisition task. During the acquisition task participants would learn 60 Dutch-Swahili word pairs while gaining at least € 20 and possibly more than € 25. To encourage learning, participants were reminded of the additional gift voucher of € 20 that would be awarded to the participant with the best performance in the recognition test following the acquisition task.

The trial started with the presentation of a fixation cue for two seconds. Next, one Dutch word and four Swahili words appeared on the screen for four seconds, allowing participants to read through the options (Figure 1a). Subsequently, one, two or four Swahili words were framed and participants could choose among the framed options to guess the Swahili translation matching the Dutch word. Depending on the number of available options, participants had a chance of respectively 100%, 50% or 25% of guessing the correct translation. Participants were encouraged to follow their first impression, although no response deadline was imposed. Responses were entered with the index and middle finger of the left and right hand, each placed on one of four response buttons assigned to the four Swahili word positions.

When a Swahili translation was chosen, an ellipsis (…) was presented for three seconds followed by the reward feedback ('correct' or 'error') presented for three seconds. The ellipsis was included in the design in order

to separate the reward feedback from the EEG response evoked by the preceding button press. Next, the Dutch word and its correct Swahili translation appeared for five seconds, to allow participants sufficient time to learn the word pair for the recognition test. The reward feedback and the word pair were presented in green on correct trials and in red on error trials. Lastly, participants were rewarded with € 0.70 on correct trials and € 0 on error trials, presented along with the cumulative monetary reward for 2.5 seconds.

Prior to the onset of the acquisition task, participants were informed that EEG recordings would be made and were asked to assume a comfortable position and to avoid unnecessary movements.

Filler task. Recency effects in the recognition test were avoided by inserting a filler task: the categorization of 400 numbers between 1 and 9 (excluding 5) as being smaller or larger than 5 (left and right button presses respectively). Participants who performed the delayed recognition test one day later were also presented with the filler task in order to keep both versions of the experiment as similar as possible. The filler task took about four minutes.

Recognition test. The recognition test was performed either immediately or after a one-day delay. In the group performing the immediate test, EEG was recorded during the recognition test and participants were asked to assume a comfortable position and avoid unnecessary movements.

At the start of the recognition test, participants were reminded about the additional gift voucher of € 20. The display layout in the recognition test paralleled that in the acquisition task. First, a fixation cross was presented for two seconds. Next, the Dutch word appeared on the screen, accompanied

by the same four Swahili words as in the acquisition task but in a randomized order. The Dutch word turned red after four seconds, indicating to participants that they could start making a choice by using the same four response buttons as in the acquisition task. Participants also rated their certainty: 'very uncertain', 'rather uncertain', 'rather certain' or 'very certain' (measured on a scale from 1 'very uncertain' to 4 'very certain'). No response deadline was imposed on the recognition test and certainty rating. Next, an ellipsis was presented for three seconds, followed by a blank screen for half a second in order to separate the EEG response evoked by the button press from that of the ensuing accuracy feedback (the ellipsis was also presented to the participants who performed the recognition test after a one-day delay although no EEG recordings were made in this group). The trials ended with the presentation of the accuracy feedback ('correct' or 'error') for three seconds.

a   <u>Acquisition task trial</u>

	Presentation	Choice	Reward anticipation	Reward feedback	Word pair encoding	Reward update
	worm mdudu fedha kito surali	worm mdudu fedha kito surali	...	Correct!	worm = kito	+ 0.70 Total: 11.2 euro
	4 seconds	No time limit	3 seconds	3 seconds	5 seconds	2.5 seconds

<u>Recognition test trial</u>

immediate or after one-day delay	Recognition	Certainty	Anticipation	Blank	Feedback
	worm kito mdudu surali fedha	How certain? very uncertain rather uncertain rather certain very certain	...		Correct!
	No time limit	No time limit	3 seconds	0.5 seconds	3 seconds

b

	1 option	2 options	4 options
Correct (rewarded)	20 trials worm = kito SRPE = 1 – 1 = 0 URPE = 0	10 trials worm = kito SRPE = 1 – 0.5 = 0.5 URPE = 0.5	5 trials worm = kito SRPE = 1 – 0.25 = 0.75 URPE = 0.75
Incorrect (not rewarded)		10 trials worm = mdudu SRPE = 0 – 0.5 = -0.5 URPE = 0.5	15 trials worm = mdudu SRPE = 0 – 0.25 = -0.25 URPE = 0.25

<u>Figure 1</u>: The trial design for the acquisition task and recognition test (panel a) and the experimental design of the acquisition task (panel b) are demonstrated. (a) A trial in the two-option condition is illustrated. Here, the participant correctly guessed the correct Swahili translation 'kito' for the Dutch word 'worm' in the acquisition task and accurately recognized the correct translation in the recognition test. (b) The 3 (number of options) x 2 (reward feedback in the acquisition task) experimental design is illustrated, replicating the design from our previous study (chapter 3). The number of trials and the associated signed and unsigned RPE (SRPE and URPE) are indicated per cell. The RPE was calculated by subtracting the reward probability from the obtained reward.

Design

The 60 Dutch-Swahili word pairs were arranged in the same 3 (number of options) x 2 (reward feedback in the acquisition task) experimental design as in the previous study (chapter 3). Because a fixed number of trials was assigned to each cell of the design, we not only predetermined the number of options (one, two or four) on each trial of the acquisition task but also whether the chosen Swahili translation would be considered correct or incorrect. Thus, participants would learn a random pairing of Dutch and Swahili words instead of the actual translation. This not only made sure we had a fixed number of trials in each cell of the design (Figure 1b); it also excluded any linguistic regularities in Dutch-Swahili word pairs that could influence learning. Participants were informed about this manipulation at the end of the recognition test.

The SRPEs were calculated by subtracting the reward probability (i.e., 1, 0.5 and 0.25 probability of a correct choice in the one-, two- and four-option condition, respectively) from the obtained reward (i.e., 1 reward on correct trials and 0 reward on incorrect trials). Thus a unique SRPE ranging from -0.50 to 0.75 was calculated for each cell in the design (see Figure 1b for a full overview). The URPEs were calculated by taking the absolute value of the SRPEs, thus reversing the sign of the RPEs for the unrewarded word pairs.

Data analysis

Behavioral data. In the familiarization task, accuracies ranged from 99% to 100% (M = 99.6%, SD = 0.50%). Accuracies in the filler task ranged from 91% to 100% (M = 97.7%, SD = 1.84%). Participants chose a valid

option in 99.6% of the trials in the acquisition task, indicating that all participants performed the experiment according to the instructions and only 0.4% of the trials had to be removed.

Unless mentioned otherwise, all statistical analyses on the behavioral data were performed within the linear mixed effects models framework. A linear mixed effects model was applied for a continuous dependent variable (e.g., certainty ratings in the recognition test) and a generalized linear mixed effects model was applied for binary dependent variables (e.g., recognition accuracy). Each model contained a random intercept across participants and centered predictors (e.g., number of options, obtained reward and SRPEs during the acquisition task). All analyses were run in R (R Core Team, 2014).

EEG preprocessing. Although EEG recordings were made during the recognition test in the group who performed the test immediately, the current analysis is restricted to the EEG recordings made during the acquisition task (recorded in all participants, irrespective of whether they performed the recognition test immediately of after a one-day delay). Therefore, from here on we will only discuss the processing of the EEG data recorded in the acquisition task.

The data were preprocessed in MATLAB (MATLAB R2013a, The MathWorks, Inc., Natick, Massachusetts, United State) by applying a custom EEGLAB preprocessing pipeline to the data of all participants (Delorme & Makeig, 2004). Because of the planned time-frequency analysis, the 1024 Hz data were not downsampled. The data were re-referenced offline to the average of the mastoid electrodes and visually inspected to remove stretches of data with excessive noise (e.g., large movement artifacts). Next, an

independent component analysis (ICA) was performed in order to remove the eyeblink artifacts. In preparation of the ICA procedure, the data were filtered with a 0.5-30 Hz band-pass Butterworth filter and the resulting filtered data were visually inspected for filtering artifacts. The filtered data were then entered into the EEGLAB ICA procedure with standard settings. Afterwards, the resulting ICA weights were added to the unfiltered (but re-referenced and cleaned) EEG data. The ICA components representing eyeblinks and lateral eye movements were selected by taking into account the topographic map of the components as well as a comparison between the time course of the components and the time course of the frontal EEG channels and eye movement electrodes. After the removal of these components from the data, the resulting data were again visually inspected. Next, electrodes were interpolated when necessary: eleven participants required the interpolation of one electrode and one participant required the interpolation of two electrodes. Electrode P2 was particularly noisy and made up for ten of the thirteen electrode interpolations. The cleaned data were subsequently filtered with a 60 Hz low-pass filter. Finally, the data were epoched time-locked to the response in the acquisition task (i.e., the choice between one, two or four available Swahili translations). The epoch extended over 2000 ms before and 13000 ms after the response (including the reward anticipation phase, the reward feedback and the word pair encoding).

Time-frequency analysis. The procedure for the time-frequency analysis was based on the code provided in chapter 16 of Cohen (2014). In order to extract the oscillatory power, a Hanning taper was first applied to the epoched EEG data and the tapered data were then subjected to a Fourier transformation. Because of the intended single-trial analysis on the power

estimates a single Hanning taper was used (Cohen, 2014; Haegens et al., 2011). The tapering and a fast Fourier transform were performed in a sliding time window of 600 ms which was advanced in steps of 100 ms between -1500 and 12500 ms relative to the choice of the Swahili translation. The oscillatory power was extracted in 18 frequency bands, spaced linearly between 1.67 and 30 Hz in steps of 1.67 Hz. Because we were mainly interested in the brain response to the reward feedback and during the word pair encoding (both during the acquisition task), a baseline correction of 300 ms before the onset of the reward feedback was applied to the power estimates. Lastly, the baseline-corrected data underwent a decibel conversion. This time-frequency analysis procedure was performed separately for each participant and each electrode.

The resulting baseline-corrected and decibel-converted oscillatory power estimates were coupled to the behavioral data on a trial-by-trial basis. In analogy to the behavioral analysis, trials were removed when an invalid Swahili translation was chosen in the acquisition task. Combined with the removal of EEG data during preprocessing this resulted in a removal of on average 1.63% of the power estimates per participant (ranging between 0% and 5%).

Clustering procedure. The goal of the subsequent clustering analyses on the power estimates was twofold. First, we wanted to test whether oscillatory power in several frequency bands could be predicted by parameters from the experimental design (Figure 1b), such as the number of options, the obtained reward and the SRPE. Second, we tested whether the power estimates could predict learning outcomes in the recognition test such as recognition accuracy and certainty ratings. For both types of analyses, the

numerous comparisons available across all power estimates necessitated a multiple comparisons correction. For example, when performing these analyses in the reward feedback phase a comparison would be made in each cell of the frequency(18) x time(31) x channels(64) tensor of power estimates, resulting in 35712 elements (we will further refer to the cells of this tensor as *voxels*). The same was true for the analyses on the frequency(18) x time(51) x channels(64) tensor of power estimates in the word pair encoding phase. To tackle this problem, a non-parametric clustering procedure based on Maris and Oostenveld (2007) was tailored to our paradigm. The clustering procedure will be described for the power estimates from the reward feedback phase, but the same procedure was applied to the power estimates from the word pair encoding phase.

For each participant, the power estimates in each voxel of the frequency(18) x time(31) x channels(64) tensor was standardized across all 60 trials. By standardizing across trials, any differences in average power estimates across frequencies, time and channels were removed within a participant. In addition, differences in average power between participants were also removed by the standardization. In this way we simultaneously accounted for the difference in oscillatory power across frequencies (e.g., lower frequencies typically have higher power estimates) and for different noise levels across time, channels and participants. The standardized power estimates of all participants and trials (41 participants x 60 trials = 2460 trials in total) were collected in a frequency(18) x time(31) x channels(64) x trials(2460) tensor, further referred to as the *TF-tensor*.

Next, five *regressor tensors* were constructed, each containing either one of the three predictive parameters from the experimental design (number of options, obtained reward or SRPE) or one of the two learning outcomes in the recognition test (recognition accuracy and certainty ratings). These regressor tensors had the same dimensions as the TF-tensor, with the regressor value of a particular trial repeated across all frequencies, time points and electrodes corresponding to that trial. In each of the five regressor tensors, the regressor values were standardized for each participant separately. In order to determine the predictive relation between the TF-tensor and a particular regressor tensor, both tensors were multiplied in a voxel-by-voxel way. The resulting frequency(18) x time(31) x channels(64) x trials(2460) tensor of trial-level statistics was subsequently summed across all 2460 trials, resulting in a frequency(30) x time(31) x channels(64) tensor with the voxel-level statistics, further referred to as the *VLS-tensor*.

Next a clustering analysis was performed on the VLS-tensor. In this clustering analysis, VLS-voxels with extremely high or low voxel-level statistics were clustered based on their proximity in frequency, time and space (for the spatial domain, the Euclidean distance between the electrodes was calculated). A separate clustering analysis was performed on the VLS-voxels with the 0.5% highest positive statistics and an additional clustering was performed on the VLS-voxels with the 0.5% lowest negative statistics. In the clustering procedure, the selected VLS-voxels were considered neighbors when their electrode position was within 4 cm of one another (see Figure 2a) and when they were within one frequency step (1.67 Hz) and one time step (100 ms; see Figure 2b). When two selected VLS-voxels were in each other's neighborhood in the frequency, time and spatial domain, they were grouped in the same cluster. All resulting clusters were ordered based

on the number of voxels in the cluster. For each of the clusters a cluster-level statistic was computer by multiplying the number of voxels in the cluster with the largest voxel-level statistic within that cluster (the highest positive statistic for the positive clusters and the lowest negative statistic for the negative clusters).

a. Proximity in space   **b. Proximity in frequency and time**
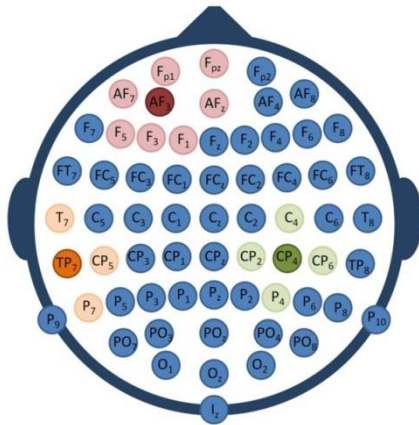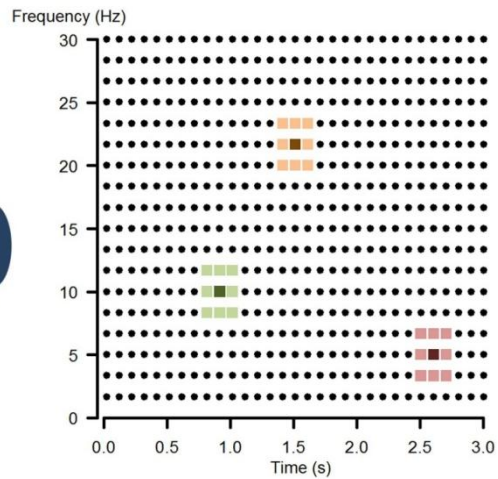
Figure 2: For three hypothetical data points in the clustering analysis (depicted in dark green, red and orange), their potential neighbors in (a) space as well as (b) frequency and time are depicted (in light green, red and orange, respectively). (a) Electrodes within 4 cm from the hypothetical data points were considered neighbors. Because of the unequal distance between electrodes, the number of spatial neighbors varies across hypothetical data points. (b) In addition, neighbors were required to fall within one step in the frequency and time domain. Each hypothetical data point thus has eight neighbors in the time-frequency domain. Note that although plotted separately, simultaneous proximity in frequency, time and space was required for being neighbors in the clustering analysis.

The resulting clusters were inspected in the frequency-time-space domain to confirm the quality of the clustering procedure. An exploratory analysis with less stringent criteria for the selection of the voxels based on their voxel-level statistic (e.g., selecting the voxels with the 1% highest or lowest statistics) or for the determination of the neighborhood (e.g., a wider frequency and time range, as well as a larger distance between neighboring electrodes) confirmed that the resulting clusters remained stable.

Non-parametric permutation test. In order to determine the significance of the observed positive and negative clusters, a non-parametric permutation test was performed on the observed cluster-level statistics. In the permutation procedure, a randomized version of the regressor tensor was first constructed by randomizing the regressor values for each participant. More specifically, the regressor values of all 60 trials for each participant were randomized separately before they were entered into the regression tensor by repeating them across frequencies, time points and channels. This randomized regressor tensor was then multiplied with the empirical TF-tensor and the resulting (randomized) VLS-tensor were calculated as before. This VLS-tensor was again clustered, resulting in a cluster-level statistic for each of the random positive and negative clusters. This randomization procedure was repeated 1000 times. The resulting random cluster-level statics from the randomization procedure were then used to determine the significance of the observed cluster-level statistics from the actual data. To do so, each observed cluster was matched to a comparable random cluster in each of the 1000 iterations in the randomization procedure. The comparable random cluster had to match the observed cluster in valence and in relative size (e.g., the second largest positive observed cluster was matched to the second largest positive randomized cluster within each iteration of the

randomization procedure). Thus, the cluster-level statistic from each observed cluster could be compared to the randomization distribution of the 1000 matching random cluster-level statistics. A significance threshold of 5% was applied to the resulting non-parametric p-values.

The significant clusters were visualized in a time-frequency plot and a topographical plot (see Figure 5). For the time-frequency plot, the voxel-level statistics pertaining to each significant cluster were summarized across channels for each time-frequency combination. The resulting statistics were log transformed before plotting. For the topographical plot, a separate plot was made per significant cluster. For each channel, the voxel-level statistics were summarized across all significant time-frequency combinations. Again, the resulting statistics were log transformed before plotting. Thus, a topographical plot was constructed for each significant cluster, with a white asterisk indicating the channels represented in the cluster.

RESULTS

Behavioral results

Recognition performance on immediate and delayed test. Paralleling the results from the previous study (chapter 3), recognition accuracy was significantly higher in the immediate test than in the delayed test, $\chi^2(1, N = 41) = 18.6$, $p < 0.001$. Recognition accuracies ranged from 37% to 100% for the immediate test group (M = 73.3%, SD = 16.7%) and from 32% to 81% for the delayed test group (M = 51.5%, SD = 14.2%).

Also, participants were significantly more certain of correctly recognized translations compared to false recognitions, $\chi^2(1, N = 41) = 471$,

$p < 0.001$. Although there was no significant main effect of test delay on the certainty ratings, $\chi^2(1, N = 41) = 1.46$, $p = 0.23$, there was a significant interaction between test delay and recognition accuracy, $\chi^2(1, N = 41) = 26.1$, $p < 0.001$. Follow-up tests revealed that whereas certainty about the false recognitions was not significantly different for the immediate and delayed test group, $\chi^2(1, N = 41) = 1.36$, $p = 0.24$, certainty about correctly recognized translations was significantly higher in the immediate test group compared to the delayed test group, $\chi^2(1, N = 41) = 4.69$, $p = 0.030$. In sum, the pattern of the recognition and certainty ratings across the immediate and delayed recognition test provides a full replication of the findings from our previous study.

The effect of positive and negative RPEs. Next, we tested whether recognition performance was predicted by the number of options and the obtained reward in the acquisition task. Also, to test whether the recognition performance was again consistent with the SRPE account we compared the effect of positive RPEs (in rewarded trials) to that of negative RPEs (in unrewarded trials; see Figure 1b).

Replicating our previous findings, the recognition rate significantly increased with an increasing number of options, $\chi^2(1, N = 41) = 24.0$, $p < 0.001$. Also, there was a significant main effect of obtained reward, $\chi^2(1, N = 41) = 18.9$, $p < 0.001$, with rewarded choices being remembered more accurately. There were no interactions with the test delay (all $p > 0.30$). Of importance for the comparison between positive and negative RPEs, there was a marginally significant interaction between the number of options and the obtained reward, $\chi^2(1, N = 41) = 3.44$, $p = 0.064$. Although the effect of the number of options was thus remarkably different for rewarded and

unrewarded word pairs, there was still a marginally significant positive effect of the number of options in the unrewarded trials, $\chi^2(1, N = 41) = 3.21$, $p = 0.073$, as suggested in Figure 3a and 3c. In sum, these results again support the SRPE account and refute the URPE account for recognition in declarative learning, confirming our previous findings.

A similar pattern of results was obtained for the certainty ratings. There was again a main effect of the number of options, $\chi^2(1, N = 41) = 7.85$, $p = 0.0051$, and an additional significant interaction between the number of options and the test delay, $\chi^2(1, N = 41) = 5.10$, $p = 0.024$. Separate follow-up analyses on the data of the immediate and delayed test revealed that in the delayed test certainty ratings increased significantly with an increasing number of options, $\chi^2(1, N = 41) = 14.9$, $p < 0.001$. However, in contrast to our previous study there was no significant effect in the certainty ratings of the immediate test, $\chi^2(1, N = 41) = 0.14$, $p = 0.71$. Next, there was no replication of the main effect of reward, $\chi^2(1, N = 41) = 0.10$, $p = 0.75$, but there was a significant interaction between reward and the test delay, $\chi^2(1, N = 41) = 6.72$, $p = 0.0095$. Follow-up tests indicated that certainty ratings on the immediate test were higher for rewarded word pairs though the effect was marginally significant, $\chi^2(1, N = 41) = 3.19$, $p = 0.074$, and this effect disappeared in the delayed test, $\chi^2(1, N = 41) = 2.52$, $p = 0.11$. Interestingly, there was again no significant interaction between the number of options and reward, $\chi^2(1, N = 41) = 1.46$, $p = 0.23$, confirming that the certainty ratings were in line with the SRPE account. All other interaction remained insignificant (all $p > 0.10$).

(a) Immediate test: not rewarded

(b) Immediate test: rewarded

(c) Delayed test: not rewarded

(d) Delayed test: rewarded
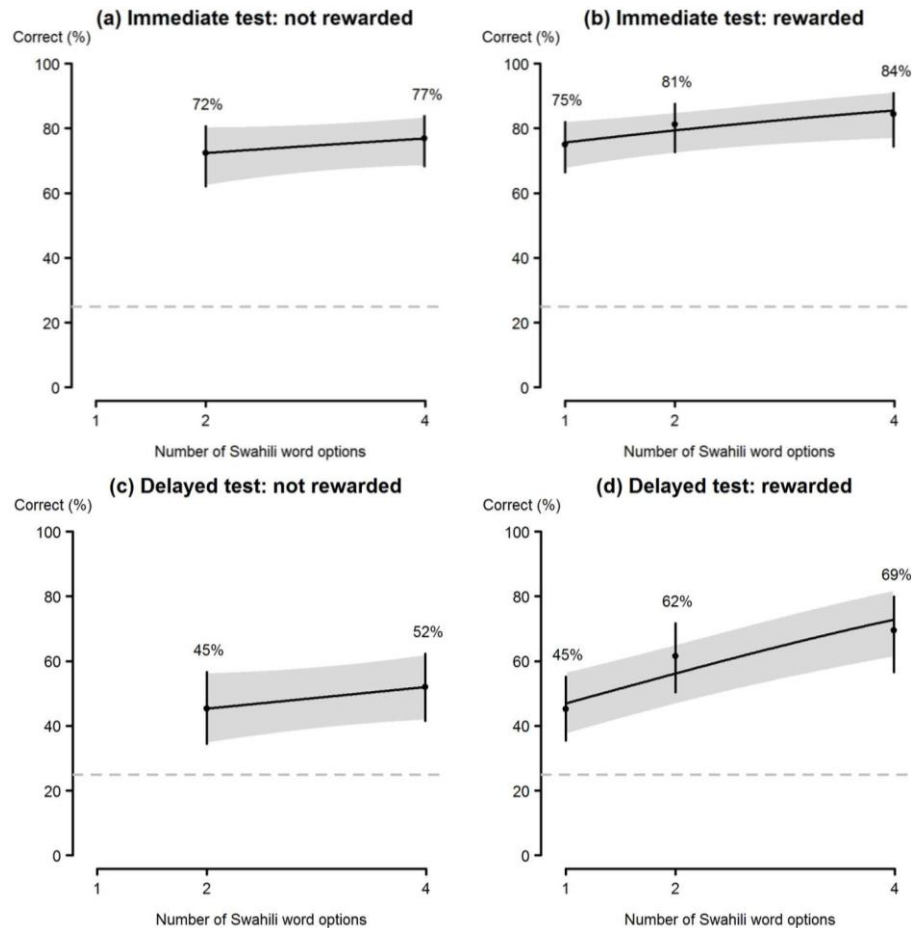
<u>Figure 3</u>: The recognition rate (y-axis) is plotted as a function of the number of options (x-axis) for the unrewarded (panel a) and rewarded word pairs (panel b) in the immediate test group and their equivalent in the delayed test (panels c and d). Note that in the one-option condition the chosen translation was always rewarded. Replicating our previous study, recognition increased significantly with an increasing number of options and reward (black regression line with gray 95% confidence band). Across the number of options, reward and delay, the average RT and 95% confidence interval was estimated and superimposed. Performance at chance level is indicated by the gray dashed line at 25% accuracy.

In sum, the analysis of the recognition accuracy fully replicates the findings from our previous study and supports the SRPE account for declarative learning. The same was true for the certainty ratings, although the differences between the immediate and delayed recognition test were more pronounced compared to the previous study.

SRPEs as a linear predictor of declarative learning. In analogy to our previous study, we formally probed whether declarative memory performance increases linearly with SRPEs. Indeed, as Figure 4a and 4b illustrate, recognition improved linearly with increasing SRPEs. This relation was highly significant, $\chi^2(1, N = 41) = 24.5$, $p < 0.001$. There was no significant interaction between SRPEs and delay, $\chi^2(1, N = 41) = 0.68$, $p = 0.41$.

Concerning the effect of SRPEs on certainty ratings, the pattern of results found in the previous study was again more consistently replicated in the delayed test group compared to the immediate test group (Figure 4c and 4d). As a consequence, there was no significant main effect of SRPEs on the certainty ratings, $\chi^2(1, N = 41) = 0.023$, $p = 0.88$, but there was a significant interaction between the effect of SRPEs and delay, $\chi^2(1, N = 41) = 7.01$, $p = 0.0081$. Because in our previous study SRPEs predicted certainty ratings only for the correctly recognized word pairs in the immediate and delayed test, we performed a follow-up test to compare this effect across both test delays. These follow-up tests revealed that SRPEs had no significant effect on certainty ratings for the correctly recognized word pairs in the immediate test, $\chi^2(1, N = 20) = 1.15$, $p = 0.28$, but resulted in significantly higher certainty ratings for correct recognitions in the delayed test, $\chi^2(1, N = 21) = 3.90$, $p = 0.048$. Thus, the linear effect of SRPEs on certainty ratings largely

replicates the pattern of findings from our previous study, mainly in the delayed test.

So far, the behavioral data offer a replication of the previous study (chapter 3) and confirm that SRPEs predict declarative learning, with more pronounced effects on recognition after a one-day delay. To further validate these results, we will subsequently verify whether participants indeed experience a SRPE during the acquisition task. If such "better-than-expected" signals are experienced, these SRPEs should be reflected by a modulation of the oscillatory power estimates (Cohen et al., 2007; HajiHosseini & Holroyd, 2015b; Oya et al., 2005). In addition, these power modulations during acquisition should predict the enhancement of the recognition performance. We will now turn to the analysis of the EEG data recorded during the acquisition task in order to test both predictions; and to pinpoint whether they only hold during the reward feedback phase or also extend into the word pair encoding phase (see Figure 1a).
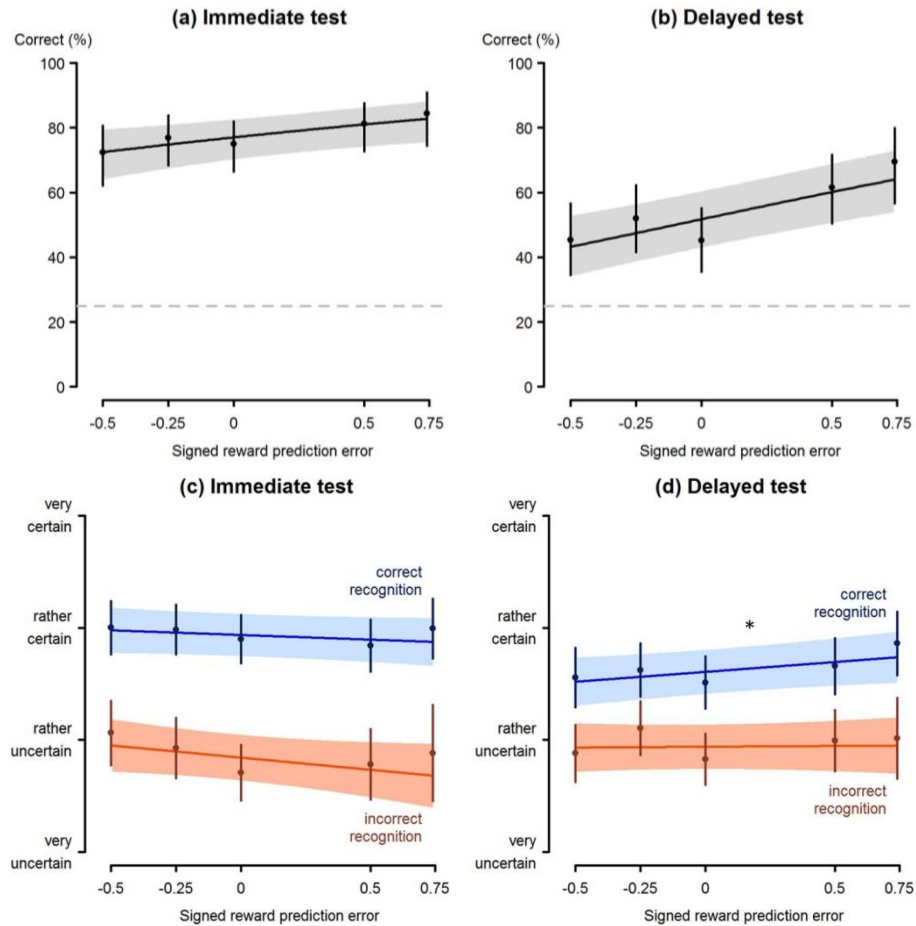
Figure 4: The recognition (panel a and b; y-axis) and certainty ratings (panel c and d; y-axis) are plotted as a function of the SRPEs (x-axis) for the immediate and delayed test group. For the word pairs associated with each SRPE, the average recognition and certainty ratings and their 95% confidence intervals were estimated separately and superimposed. (a and b) Recognition increased significantly with higher SRPEs (black regression line with gray 95% confidence band). Performance at chance level is indicated by the gray dashed line at 25% accuracy. (c and d) SRPEs only predicted certainty ratings for correctly recognized word pairs (depicted in blue) in the delayed test (indicated with an asterisk), but not in the immediate test.

Oscillatory power modulations during acquisition reflect SRPEs

To start, we examined the reflection of SRPEs in oscillatory power during reward feedback and word pair encoding. In analogy to the behavioral data analysis, we will first examine the separate modulating effects of the number of options and the obtained reward before proceeding to modulation by SRPEs.

Oscillatory power modulated by number of options and obtained reward. During reward feedback, oscillatory power was modulated by the number of options and obtained reward in a number of partially overlapping clusters (see Figure 5a and 5b) as revealed by the cluster-based non-parametric permutation test. A similar pair of partially overlapping clusters was found during word pair encoding (see Figure 6a and 6b). Because of the similarity of the clusters found in both phases of the acquisition task, they will jointly be discussed below.

First, immediately after reward feedback onset a pair of overlapping clusters in the beta frequency band (12-30 Hz) showed increased power when fewer options were available ($p = 0.000$; see Figure 5a) and when positive reward feedback was received ($p = 0.000$; see Figure 5b). Of all the voxels in the beta cluster predicted by the number of options, 43.4% was represented in the beta cluster predicted by obtained reward. The reverse comparison revealed an overlap of 42.1%. These clusters extended over roughly 500 ms and were more pronounced over posterior electrodes although both had a broad topography. The combination of increased power in the beta band due to fewer options on the one hand and positive reward feedback on the other hand suggests that this beta band does not represent an SRPE. Indeed, SRPEs should be represented by either a simultaneous

positive effect of the number of options and reward, or alternatively a simultaneous negative effect of both predictors. Therefore, this initial beta cluster is not a likely representation of the experience of a SRPE (nor a URPE). A similar pair of overlapping clusters in the alpha/low-beta range (8-20 Hz) was found during word pair encoding. Like in the reward feedback phase, power estimates increased when fewer options were available ($p = 0.001$; see Figure 6a) and when positive reward feedback was received ($p = 0.000$; see Figure 6b). Thus, these clusters are highly reminiscent of the beta clusters observed during reward feedback (Figure 5a and 5b), mimicking their topography and their relation to the number of options and the obtained reward. Hence, the alpha/low-beta clusters found during word pair encoding are likewise no probable reflection of the experienced SRPEs.

Second, the reward feedback evoked a positive cluster in the theta frequency band (4-8 Hz) with higher power as the number of options increased ($p = 0.000$; see Figure 5a). An overlapping but smaller negative cluster showed increased power in the theta band when error feedback was provided ($p = 0.000$; see Figure 5b). The small negative cluster showed an overlap of 88.7% with the large positive cluster. Both clusters peaked approximately 300 ms after feedback onset. The positive relation between the number of options and power estimates in the theta band was followed by a secondary cluster that demonstrated a reversed, negative relation ($p = 0.008$; see Figure 5a). Taken together, activation in the theta cluster during reward feedback is mainly modulated by the number of options and not by the obtained reward. No such modulations in the theta frequency range were found during word pair encoding (see Figure 6).

Third, a cluster in the low-beta range (12-20 Hz) showed increased power estimates during error feedback ($p = 0.000$; see Figure 5b). The cluster extended from 500 ms to 1500 ms after feedback onset. No clear topography could be distilled from the channels represented in the cluster. Importantly, this cluster was uniquely related to the obtained reward as it showed no relation to the number of options during reward feedback (see Figure 5a). In addition, no equivalent cluster was found during word pair encoding (see Figure 6).

In sum, the number of options and the obtained reward predicted oscillatory power modulations mainly during reward feedback and to a lesser extent during word pair encoding. During word pair encoding, the number of options and obtained reward only modulated oscillatory power in the alpha/low-beta band, though not likely reflecting the experience of a SRPE. During reward feedback, a similar modulation pattern was found in the beta band. This beta cluster was followed by a cluster in the theta band that was mainly modulated by the number of options, and a third modulation in the low-beta band specifically caused by the obtained reward. Thus, the separate effects of the number of options on the one hand and the obtained reward on the other hand are both reflected in the power modulations during the acquisition task. Taking these separate effects into account, we will now test whether SRPEs are likewise reflected during reward feedback and word pair encoding.

Figure 5: Oscillatory power during reward feedback is predicted by (a) the number of options and (b) the obtained reward. The significant positive (red) and negative (blue) clusters are plotted in the time-frequency domain, accompanied by their topographic representations (see p.109 for details).

a   Oscillatory power during word pair encoding predicted by number of options
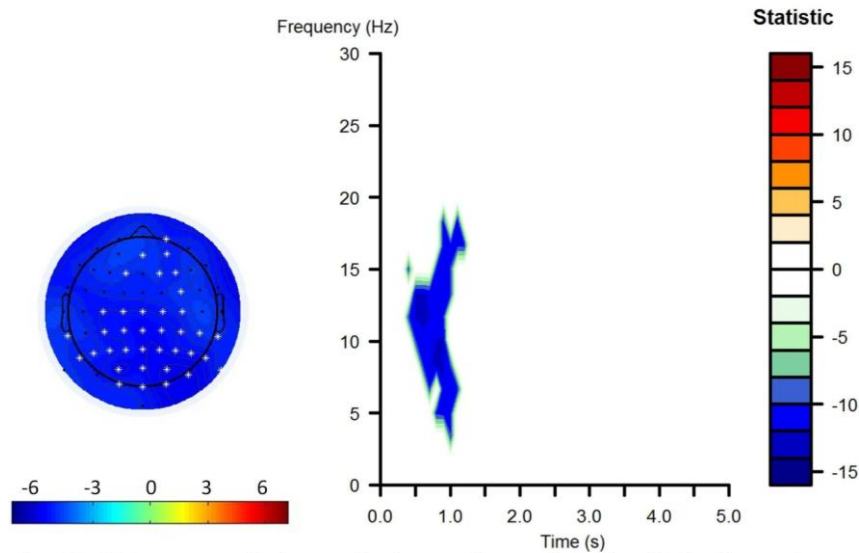
b   Oscillatory power during word pair encoding predicted by obtained reward
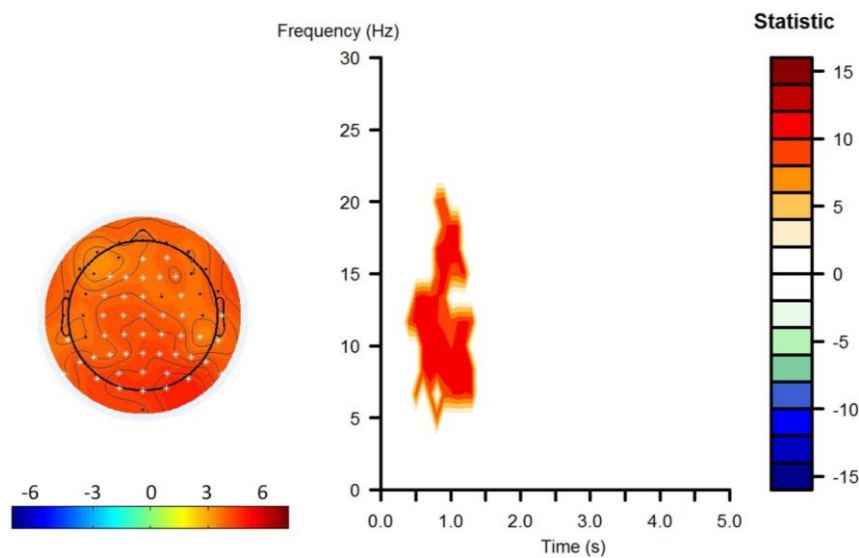
Figure 6: Oscillatory power during word pair encoding is predicted by (a) the number of options and (b) the obtained reward. No further relations between oscillatory power and SRPEs or recognition performance were found in this phase. (For details on the graphical representation, see p. 109)

Oscillatory power modulated by SRPEs. The cluster-based non-parametric permutation test on the relation between SRPEs and oscillatory power revealed no significant clusters during word pair encoding. Conversely, three significant clusters were found during reward feedback (see Figure 7) demonstrating that SRPEs were uniquely experienced when reward feedback was provided and not during the subsequent encoding phase.

To start, the reward feedback evoked two clusters that highly resemble the theta and low-beta clusters described in the previous analysis (see above). A positive relation was found between SRPEs and power estimates in the theta band ($p = 0.000$; see Figure 7), again peaking at approximately 300 ms. The relation in this theta cluster is likely based on the increased power estimates due to an increasing number of options, demonstrated in the analysis described above. Indeed, 52.5% of the voxels in the current theta cluster were also represented in the theta cluster in Figure 5a. Likewise, the increased low-beta power due to error feedback demonstrated in the previous analysis likely forms the basis for the negative relation between SRPEs and the current cluster of low-beta estimates ($p = 0.000$; see Figure 7), also peaking between 500 ms and 1000 ms after feedback onset. Half of the voxels in the current low-beta cluster were also present in the beta cluster predicted by the obtained reward (Figure 5b).

Thus, although power oscillations in the theta range and low-beta range were both predicted by SRPEs these relations seem to rely on different underlying processes. To further demonstrate this point, we calculated the average power within each cluster on a trial-by-trial basis and attempted to predict activation in the low-beta band based on the activation in the theta

band by applying a linear mixed effects model with a random intercept across participants. This analysis clearly demonstrated the absence of a predictive relation between the activation in both clusters, $\chi^2(1, N = 41) = 0.24$, $p = 0.63$, further suggesting two different mechanisms.

In contrast, there was a cluster in the high-beta range (20-30 Hz) showing a unique relation with SRPEs not previously found in the analyses with the number of options and obtained reward separately. The cluster also showed no overlap with any of the other clusters in the previous analyses. Oscillatory power in this high-beta cluster increased with increasing SRPEs ($p = 0.008$; see Figure 7), peaking around 500 ms post feedback onset. Its topography was relatively pronounced and showed a frontal mapping. In contrast to the absence of a predictive relation between activation in the theta and low-beta cluster, activation in the high-beta cluster was significantly predicted by the activation in the preceding theta cluster, $\chi^2(1, N = 41) = 7.31$, $p = 0.0068$, and in its turn predicted the power estimates in the subsequent low-beta cluster, $\chi^2(1, N = 41) = 12.22$, $p < 0.001$.

In sum, the high-beta cluster offers the main signature of the experienced SRPEs, not mainly accounted for by the separate effect of either the number of options or the obtained reward. Thus the clustering analysis confirms that participants experienced "better-than-expected" signals during the acquisition task, further supporting our claim that SRPEs drive declarative learning. This relation was uniquely found during reward feedback and did not extend into the word pair encoding phase.
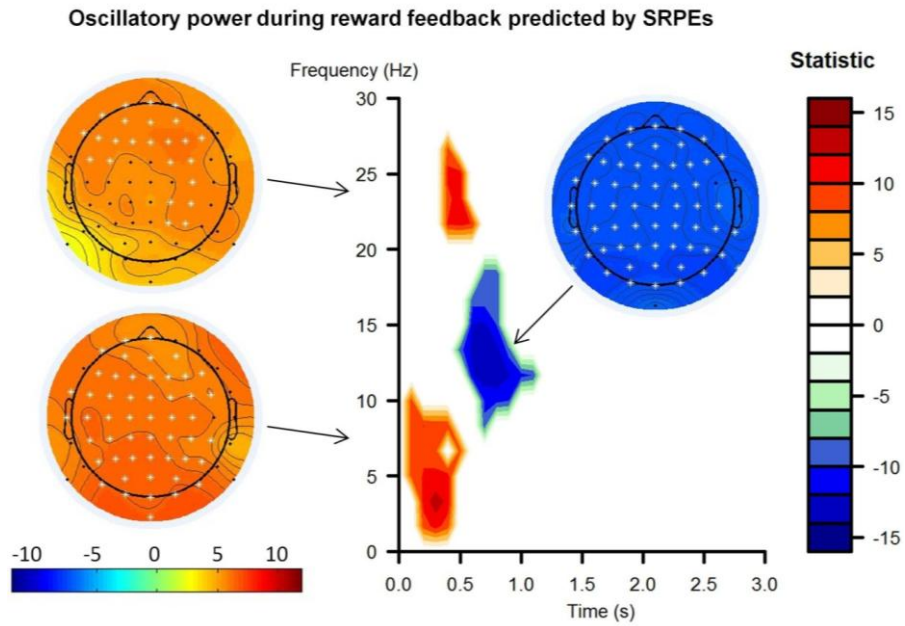
Figure 7: The oscillatory power during reward feedback reveals a positive relation with SRPEs in the theta range (peaking at 300 ms post feedback onset) and a negative relation with SRPEs in the low-beta range (peaking between 500 ms and 1000 ms post feedback onset). In addition, increasing SRPEs cause increased power in the high-beta range, peaking at 500 ms post feedback onset. (For details on the graphical representation, see p. 109)

Recognition performance predicted by oscillatory power

So far we have demonstrated that SRPEs are indeed reflected in the modulation of oscillatory power during reward feedback. Next, we probed whether power modulations during the acquisition task could predict recognition accuracy in the subsequent recognition test. The cluster-based non-parametric permutation test revealed no significant clusters, indicating that variations in oscillatory power during acquisition could not predict subsequent recognition accuracy. This was true in the reward feedback phase and during word pair encoding.

However, the recognition accuracy might not yield the most sensitive performance measure because of its dichotomous nature and because of its likely contamination by correct guesses (i.e., a 25% chance of guessing the correct translation). Therefore, in line with previous studies that demonstrated a link between electrophysiological activity during acquisition and subsequent confidence ratings (Johnson, 1995; Otten & Rugg, 2001; Sommer, Komoss, & Schweinberger, 1997), we also tested whether oscillatory power during reward feedback could predict the certainty ratings for the correctly recognized word pairs. We thus probed whether the pattern in the behavioral analysis (i.e., a significant effect of SRPEs on the certainty rating for the correctly recognized word pairs in the delayed test group; see Figure 4d) would be reflected by the power estimates. Indeed, in the delayed test group there was a large significant cluster in the alpha band (8-12 Hz) that predicted certainty ratings for correctly recognized word pairs ($p = 0.000$; see Figure 8). More specifically, suppressed power in the alpha band during reward feedback resulted in higher certainty ratings. This alpha cluster extended from 500 ms to 1700 ms post feedback onset and had a

midline topography. The predictive relation between oscillatory power and certainty ratings was specific for the correctly recognized word pairs in the delayed test group ($p > 0.26$ for all other comparisons involving the immediate test group and/or the false recognitions), mirroring the specificity of the relation between SRPEs and certainty ratings in the behavioral analysis.

Next, we probed the relation between the alpha band cluster predictive for recognition performance and the clusters related to the experience of SRPEs (see above). Confirming the overlap suggested on Figure 7 and Figure 8, 15.4% of the voxels in the current alpha cluster were also represented in the low-beta cluster predicted by SRPEs. In addition, activity in the alpha cluster was predicted by activation in the low-beta cluster (excluding overlapping voxels), $\chi^2(1, N = 20) = 376$, $p < 0.001$, and activity in the upper-beta cluster, $\chi^2(1, N = 20) = 5.45$, $p = 0.020$, but not by that in the theta cluster, $\chi^2(1, N = 20) = 0.0054$, $p = 0.94$. Thus, these results confirm that the current alpha cluster predictive for memory performance (see Figure 8) and the clusters related to the experience of SRPEs (see Figure 7) show a partial overlap and are functionally related.

Finally, we probed whether the certainty ratings were uniquely predicted by power modulations during reward feedback or whether they were also influenced by activity in the subsequent word pair encoding phase. Therefore, we repeated the same tests outlined above on the power estimates during word pair encoding, but no significant clusters were found (all p > 0.09).

In sum, power estimates during reward feedback are predictive for certainty ratings, but only when declarative memory is probed after a one-

day delay. The alpha band cluster predictive for these certainty ratings overlaps partially with the low-beta cluster (Figure 7) and its activity is predicted by both this low-beta cluster and the high-beta cluster that reflect the experience of SRPEs. Thus, there is a clear link between power oscillations during reward feedback and both the experience of SRPEs and the resulting enhancement of declarative learning.
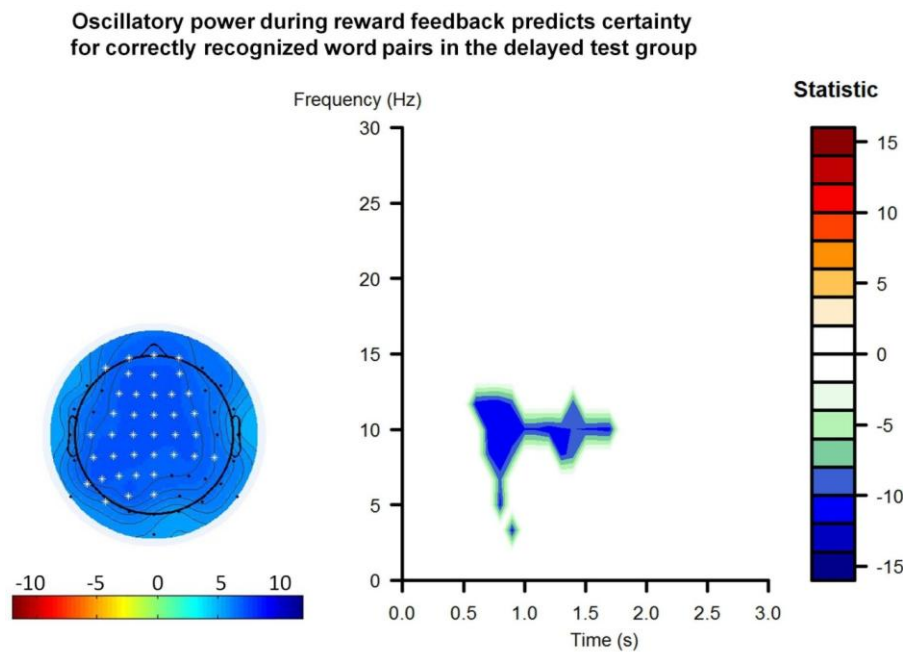
Figure 8: The certainty rating for correctly recognized word pairs in the delayed test group are predicted by oscillatory power during reward feedback. Lower alpha suppression (peaking from 500 ms to 1700 ms post feedback onset) was predictive for higher certainty ratings. (For details on the graphical representation, see p. 109)

DISCUSSION

The current experiment offers a full replication of the behavioral results from the previous study (chapter 3) and confirms that SRPEs drive declarative learning. Word pairs associated with large, positive RPEs were recognized more often and with higher certainty. Also, this effect grew stronger after a one-day delay as predicted by the neoHebbian learning account (Lisman et al., 2011). As a further validation of the beneficial effect of SRPEs on declarative learning, we showed that participants actually experience these "better-than-expected" signals by uncovering their modulating effect on the oscillatory power in the theta, high-beta and low-beta frequency band. Furthermore, a partially overlapping and functionally related cluster of activation in the alpha band was predictive for the recognition performance after a one-day delay, mirroring the pattern of findings in the behavioral analysis. Thus, the results from the EEG analysis confirm the experience of SRPEs and elucidate how these teaching signals influence declarative learning. The EEG analysis also revealed that both effects were uniquely found during reward feedback and did not extend into the ensuing word pair encoding phase.

Overall, our results nicely tie in with established findings on the role of RPEs in procedural learning and bridge the gap with the literature on declarative learning. First, our findings demonstrate that the same reward mechanism involved in procedural learning is also activated by SRPEs during declarative learning. Second, the current experiment makes a timely contribution to a recent line of research on how neural activity prior to

encoding can influence declarative learning. Each of these points will be further unpacked below.

RPEs modulate oscillatory power and influence both procedural and declarative learning

The current results uncover the involvement of the classical RPE mechanisms typically found in procedural learning. In line with previous research we found activation in the theta, high-beta and low-beta band. This pattern of results echoes how RPEs are traditionally reflected in power oscillations during procedural learning, where they serve as a teaching signal.

First, SRPEs were reflected by a power increase in the theta band (4-8 Hz) during reward feedback, even if this relation was mainly driven by the number of options. Activation in the theta band has been associated with event-related potentials such as the feedback related negativity (FRN), as well as other components such as the error related negativity (ERN) and the correct related negativity (CRN). All of these components signal an evaluation of the (reward) feedback based on prior knowledge and current predictions (Cavanagh, Cohen, & Allen, 2009; Gehring, Liu, Orr, & Carp, 2012; Luu, Tucker, & Makeig, 2004; Oliveira, McDonald, & Goodman, 2007; Yordanova, Falkenstein, Hohnsbein, & Kolev, 2004). The source of these theta oscillations has been ascribed to the medial prefrontal cortex and the anterior cingulate cortex (Cavanagh, Figueroa, Cohen, & Frank, 2011; Cohen, Elger, & Fell, 2009; Holroyd & Coles, 2002; Luu et al., 2004) where they have been suggested to drive several adaptive processes such as cognitive control (Botvinick, Braver, Barch, Carter, & Cohen, 2001; Cavanagh & Frank, 2014; Ridderinkhof, Ullsperger, Crone, & Nieuwenhuis,

2004). Interestingly, because the hippocampal area is also dominated by theta oscillations, learning has been hinted at as one of these adaptive processes triggered by RPEs (Cavanagh & Frank, 2014). Indeed, based on the reinforcement learning account for the FRN (Cavanagh, Frank, Klein, & Allen, 2010; Holroyd & Coles, 2002; Holroyd, Pakzad-Vaezi, & Krigolson, 2008; Silvetti, Seurinck, & Verguts, 2011), the theta oscillations have been hypothesized to function as a teaching signal and to reflect increased activity in the mesolimbic dopamine system. This complex interaction between the dopaminergic system and the theta oscillations in the prefrontal area and hippocampus (Fujisawa & Buzsáki, 2011) distinctively relates to the neoHebbian learning framework put forward in the introduction (Lisman et al., 2011).

Second, SRPEs were also reflected in the high- and low-beta frequency band during reward feedback. Although (U)RPEs have predominantly been related to activation in the theta band, activation in the beta band has been implicated in the thalamic network that regulates memory formation based on prefrontal teaching signals (Ketz, Jensen, O'Reilly, & O'Reilly, 2015). According to this thalamic coordination account, the theta oscillations regulate episodic-like (declarative) memory, while the beta oscillations are predictive for familiarity-based recognition. This dissociation between the role of oscillatory activity in the theta and beta band might inform why in the current experiment oscillatory power in the beta band – but critically not the theta band – was predictive for the recognition performance. However, further research is necessary to elucidate whether both frequency bands have a distinctive influence on recall and recognition (Ketz et al., 2015; Merkow, Burke, Stein, & Kahana, 2014). Interestingly, our study not only revealed the predicted negative relation

between SRPEs and power in the low-beta band (12-20 Hz; Ketz et al., 2015; Oya et al., 2005), it also revealed a positive relation with power in the high-beta band (20-30 Hz). RPEs have indeed been shown to evoke a power increase in the high-beta band (HajiHosseini & Holroyd, 2015a). These power modulations were source-localized to the dorsolateral prefrontal cortex, in keeping with the frontal topography of the high-beta cluster in the current experiment. In their procedural learning experiment, HajiHosseini and Holroyd (2015b) also probed whether activity in the high-beta band varied with reward valence and reward probability. Across separate blocks, reward probabilities were manipulated through a staircase procedure that determined the width of a response window resulting in either a 75% chance of positive feedback or a 75% chance of error feedback. However, counter to the current results, only the reward valence but not the distinct reward probabilities had a detectable influence on the oscillatory power estimates. It is unclear why the high-beta activity in the current experiment does reflect the reward probability, but we might speculate that the manipulation of the reward probability was more salient and task-relevant in the current study (i.e., varied on a trial-by-trial basis by choosing among one, two or four Swahili words).

On the whole, the current results established that the parametrically manipulated SRPEs that drive declarative learning indeed activate the same dopaminergic reward system that has been demonstrated to drive procedural learning. Next, we will turn to the implications of the current findings for the literature on declarative learning.

Oscillatory power prior to encoding is predictive for declarative learning

Recently, there has been a surge of interest in the contribution that reward mechanisms make to declarative learning (Hidi, 2016; Howard-Jones & Jay, 2016; Miendlarzewska, Bavelier, & Schwartz, 2016), with the goal of merging the gap between neuroscientific research on reinforcement learning and its application in everyday settings such as in education. This revived interest offers a fresh perspective on declarative learning as historically most research in this field has focused on neural activity during retrieval instead of during encoding (Wilding & Ranganath, 2011). In addition, studies that focused on encoding have largely been limited to the *difference due to memory effect* (Paller, 1990) and the *subsequent memory effect* (Paller & Wagner, 2002). In these paradigms, spontaneous variations in neural activity during encoding are retrospectively contrasted between subsequently recalled and forgotten items. Using electrophysiological recordings, oscillations in the theta frequency band have been shown to predict subsequent memory performance (Kleberg, Kitajo, Kawasaki, & Yamaguchi, 2014) and these oscillations have been ascribed to a network involving the prefrontal cortex and the hippocampal area (Cohen et al., 2015; Hsieh & Ranganath, 2014; Lisman & Jensen, 2013; Nyhus & Curran, 2010). A number of intracranial EEG studies have recently validated this claim by demonstrating how prefrontal and hippocampal theta oscillations – along with subsequent modulations in the alpha band – predicted successful memory encoding (Fell et al., 2011; Greenberg, Burke, Haque, Kahana, & Zaghloul, 2015; Haque, Wittig, Damera, Inati, & Zaghloul, 2015; Merkow et al., 2014; Sweeney-Reed et al., 2016).

Overall, these results are in line with the involvement of theta oscillations in procedural learning (see above) and the neoHebbian learning account (Lisman et al., 2011). However, because these studies are based on spontaneous neural activity it is hard to exclude alternative interpretations such as attentional processes and resource allocation (Wilding & Ranganath, 2011). In order to test whether oscillatory power variations are causally related to subsequent memory improvement, a number of studies have tried to directly manipulate oscillatory power prior to encoding. To start, converging evidence from fMRI (Bollinger, Rubens, Zanto, & Gazzaley, 2010; Galli, Bauch, & Gruber, 2011; Gruber, Gelman, & Ranganath, 2014; Loh et al., 2015) MEG (Garrido, Barnes, Kumaran, Maguire, & Dolan, 2015) and behavioral experiments (Liu, Grady, & Moscovitch, 2016; Otten, Quayle, Akram, Ditewig, & Rugg, 2006; Oyarzún, Packard, de Diego-Balaguer, & Fuentemilla, 2016) support the claim that prior knowledge and reward expectations are the actual driver of improved declarative learning. This was further confirmed in an EEG study, showing that reward expectations boost declarative learning through anticipatory theta activity over frontal regions (Gruber, Watrous, Ekstrom, Ranganath, & Otten, 2013). Interestingly, Gruber and colleagues only found this effect during the anticipation of large positive rewards, in accordance with the findings from the current study. It is likely that reward anticipation and SRPEs evoke similar power bursts in the theta band and increased dopaminergic activity that subsequently causes enhanced recognition performance.

On an important note, the current study only provided support for a direct relation between activity in the alpha cluster and subsequent recognition performance. Therefore, we need to take into account that the observed clusters in the theta, high-beta and low-beta band were only

indirectly related to memory performance through their functional relation with activity in the alpha cluster. This sparks the intriguing question whether SRPEs have a direct impact on declarative learning or whether they influence encoding indirectly by triggering other adaptive processes. Indeed, not only activation in the theta band but also activation in the beta and alpha band has been shown to improve declarative learning. For example, alpha oscillations have been associated with enhanced semantic encoding (Klimesch, 1999) and with the gating of relevant items and suppression of irrelevant information (Ketz et al., 2015; Park et al., 2014). In the current study, the increased alpha activity on unrewarded trials (i.e., during negative RPEs) might thus reflect the suppression of the incorrectly guessed translation in anticipation of the correct translation. Next to the alpha activity, spontaneous activation in the beta band has also been shown to be predictive for successful encoding, both in a surprise recall test (Salari & Rose, 2016) as well as during intentional encoding (Schneider & Rose, 2016). Still, the pattern of findings is highly variable across studies, especially when the analyses are based on spontaneous power fluctuations and various behavioral tasks. Also, the inconsistent pattern of findings across studies might be caused by insufficient statistical power (i.e., due to the low number of trials in declarative learning paradigms) or because of the variety of memory test used (e.g., planned or surprise recall, recognition accuracy and recognition certainty). By the parametric manipulation of RPEs, the current study provides a useful paradigm to help disentangle the involvement of several adaptive mechanisms that enhance declarative recall and recognition.

CONCLUSION

In the current experiment we replicated the beneficial effect of SRPEs on declarative learning found in our previous study (chapter 3). During reward feedback, SRPEs were shown to modulate oscillatory power in the theta, high-beta and low-beta band; SRPEs thus triggered the same reward mechanisms also activated by RPEs in procedural learning paradigms. Also during reward feedback, activity in the alpha band was predictive for the recognition performance after a one-day delay. Together, these findings provide the first empirical demonstration of how SRPEs can enhance declarative memory and make a timely contribution to the application of insights from reinforcement learning to the field of declarative learning.

REFERENCES

Bollinger, J., Rubens, M. T., Zanto, T. P., & Gazzaley, A. (2010). Expectation-driven changes in cortical functional connectivity influence working memory and long-term memory performance. *Journal of Neuroscience*, *30*(43), 14399–14410. doi:10.1523/JNEUROSCI.1547-10.2010

Botvinick, M. M., Braver, T. S., Barch, D. M., Carter, C. S., & Cohen, J. D. (2001). Conflict monitoring and cognitive control. *Psychological Review*, *108*(3), 624–652. doi:10.1037//0033-295X.108.3.624

Bryden, D. W., Johnson, E. E., Tobia, S. C., Kashtelyan, V., & Roesch, M. R. (2011). Attention for learning signals in anterior cingulate cortex. *The Journal of Neuroscience : The Official Journal of the Society for Neuroscience*, *31*(50), 18266–18274. doi:10.1523/JNEUROSCI.4715-11.2011

Cavanagh, J. F., Cohen, M. X., & Allen, J. J. B. (2009). Prelude to and resolution of an error: EEG phase synchrony reveals cognitive control dynamics during action monitoring. *Journal of Neuroscience*, *29*(1), 98–105. doi:10.1523/JNEUROSCI.4137-08.2009

Cavanagh, J. F., Figueroa, C. M., Cohen, M. X., & Frank, M. J. (2011). Frontal theta reflects uncertainty and unexpectedness during exploration and exploitation. *Cerebral Cortex (New York, N.Y. : 1991)*, *22*(11), 2575–2586. doi:10.1093/cercor/bhr332

Cavanagh, J. F., & Frank, M. J. (2014). Frontal theta as a mechanism for cognitive control. *Trends in Cognitive Sciences*, *18*(8), 414–421. doi:10.1016/j.tics.2014.04.012

Cavanagh, J. F., Frank, M. J., Klein, T. J., & Allen, J. J. B. (2010). Frontal theta links prediction errors to behavioral adaptation in reinforcement learning. *NeuroImage*, *49*(4), 3198–3209. doi:10.1016/j.neuroimage.2009.11.080

Cohen, M. X. (2014). *Analyzing Neural Time Series Data: Theory and Practice*. Cambridge, MA: MIT Press.

Cohen, M. X., Elger, C. E., & Fell, J. (2009). Oscillatory activity and phase-amplitude coupling in the human medial frontal cortex during decision making. *Journal of Cognitive Neuroscience*, *21*(2), 390–402. doi:10.1162/jocn.2008.21020

Cohen, M. X., Elger, C. E., & Ranganath, C. (2007). Reward expectation modulates feedback-related negativity and EEG spectra. *NeuroImage*, *35*(2), 968–978. doi:10.1016/j.neuroimage.2006.11.056

Cohen, N., Pell, L., Edelson, M. G., Ben-Yakov, A., Pine, A., & Dudai, Y. (2015). Peri-encoding predictors of memory encoding and consolidation. *Neuroscience & Biobehavioral Reviews*, *50*, 128–142. doi:10.1016/j.neubiorev.2014.11.002

Delorme, A., & Makeig, S. (2004). EEGLAB: an open source toolbox for analysis of single-trial EEG dynamics including independent component analysis. *Journal of Neuroscience Methods*, *134*(1), 9–21. doi:10.1016/j.jneumeth.2003.10.009

den Ouden, H. E. M., Friston, K. J., Daw, N. D., McIntosh, A. R., & Stephan, K. E. (2009). A dual role for prediction error in associative learning. *Cerebral Cortex (New York, N.Y. : 1991)*, *19*(5), 1175–1185. doi:10.1093/cercor/bhn161

Fell, J., Ludowig, E., Staresina, B. P., Wagner, T., Kranz, T., Elger, C. E., & Axmacher, N. (2011). Medial temporal theta/alpha power enhancement precedes successful memory encoding: Evidence based on intracranial EEG. *Journal of Neuroscience*, *31*(14), 5392–5397. doi:10.1523/JNEUROSCI.3668-10.2011

Fujisawa, S., & Buzsáki, G. (2011). A 4 Hz oscillation adaptively synchronizes prefrontal, VTA, and hippocampal activities. *Neuron*, *72*(1), 153–165. doi:10.1016/j.neuron.2011.08.018

Galli, G., Bauch, E. M., & Gruber, M. J. (2011). When anticipation aids long-term memory: What cognitive and neural processes are involved? *Journal of Neuroscience*, *31*(12), 4355–4356. doi:10.1523/JNEUROSCI.6369-10.2011

Garrido, M. I., Barnes, G. R., Kumaran, D., Maguire, E. A., & Dolan, R. J. (2015). Ventromedial prefrontal cortex drives hippocampal theta oscillations induced by mismatch computations. *NeuroImage*, *120*, 362–370. doi:10.1016/j.neuroimage.2015.07.016

Gates, A. I. (1917). Recitation as a factor in memorization. *Archives of Psychology*, *6*(40), 112–121.

Gehring, W., Liu, Y., Orr, J., & Carp, J. (2012). The error-related negativity (ERN/Ne). In S. J. Luck & E. S. Kappenman (Eds.), *The Oxford handbook of event-related potential components* (pp. 231–291). Oxford: Oxford University Press.

Greenberg, J. A., Burke, J. F., Haque, R., Kahana, M. J., & Zaghloul, K. A. (2015). Decreases in theta and increases in high frequency activity underlie associative memory encoding. *NeuroImage*, *114*, 257–263. doi:10.1016/j.neuroimage.2015.03.077

Gruber, M. J., Gelman, B. D., & Ranganath, C. (2014). States of curiosity modulate hippocampus-dependent learning via the dopaminergic circuit. *Neuron*, *84*(2), 486–496. doi:10.1016/j.neuron.2014.08.060

Gruber, M. J., Watrous, A. J., Ekstrom, A. D., Ranganath, C., & Otten, L. J. (2013). Expected reward modulates encoding-related theta activity before an event. *NeuroImage*, *64*, 68–74. doi:10.1016/j.neuroimage.2012.07.064

Haegens, S., Nácher, V., Hernández, A., Luna, R., Jensen, O., & Romo, R. (2011). Beta oscillations in the monkey sensorimotor network reflect somatosensory decision making. *Proceedings of the National Academy of Sciences of the United States of America*, *108*(26), 10708–10713. doi:10.1073/pnas.1107297108

HajiHosseini, A., & Holroyd, C. B. (2015a). Reward feedback stimuli elicit high-beta EEG oscillations in human dorsolateral prefrontal cortex. *Scientific Reports*, *5*, 13021. doi:10.1038/srep13021

HajiHosseini, A., & Holroyd, C. B. (2015b). Sensitivity of frontal beta oscillations to reward valence but not probability. *Neuroscience Letters*, *602*, 99–103. doi:10.1016/j.neulet.2015.06.054

Haque, R. U., Wittig, J. H., Damera, S. R., Inati, S. K., & Zaghloul, K. A. (2015). Cortical low-frequency power and progressive phase synchrony precede successful memory encoding. *Journal of Neuroscience*, *35*(40), 13577–13586. doi:10.1523/JNEUROSCI.0687-15.2015

Hayden, B. Y., Heilbronner, S. R., Pearson, J. M., & Platt, M. L. (2011). Surprise signals in anterior cingulate cortex: Neuronal encoding of unsigned reward prediction errors driving adjustment in behavior. *The Journal of Neuroscience : The Official Journal of the Society for Neuroscience*, *31*(11), 4178–4187. doi:10.1523/JNEUROSCI.4652-10.2011

Hebb, D. O. (1949). *The Organization of Behavior*. New York: Wiley & Sons.

Hidi, S. (2016). Revisiting the role of rewards in motivation and learning: Implications of neuroscientific research. *Educational Psychology Review*, *28*(1), 61–93. doi:10.1007/s10648-015-9307-5

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*(4), 679–709. doi:10.1037//0033-295X.109.4.679

Holroyd, C. B., Pakzad-Vaezi, K. L., & Krigolson, O. E. (2008). The feedback correct-related positivity: Sensitivity of the event-related brain potential to unexpected positive feedback. *Psychophysiology*, *45*(5), 688–697. doi:10.1111/j.1469-8986.2008.00668.x

Howard-Jones, P. A., & Jay, T. (2016). Reward, learning and games. *Current Opinion in Behavioral Sciences*, *10*, 65–72. doi:10.1016/j.cobeha.2016.04.015

Hsieh, L.-T., & Ranganath, C. (2014). Frontal midline theta oscillations during working memory maintenance and episodic encoding and retrieval. *NeuroImage*, *85*, 721–729. doi:10.1016/j.neuroimage.2013.08.003

Hyman, S. E., Malenka, R. C., & Nestler, E. J. (2006). Neural mechanisms of addiction: The role of reward-related learning and memory. *Annual Review of Neuroscience*, *29*, 565–598. doi:10.1146/annurev.neuro.29.051605.113009

Johnson, R. (1995). Event-related potential insights into the neurobiology of memory systems. *Handbook of Neuropsychology*, *10*, 135.

Karpicke, J. D., & Blunt, J. R. (2011). Retrieval practice produces more learning than elaborative studying with concept mapping. *Science (New York, N.Y.)*, *331*(6018), 772–775. doi:10.1126/science.1199327

Karpicke, J. D., & Roediger, H. L. (2008). The critical importance of retrieval for learning. *Science (New York, N.Y.)*, *319*(5865), 966–968. doi:10.1126/science.1152408

Ketz, N. a., Jensen, O., O'Reilly, R. C., & O'Reilly, R. C. (2015). Thalamic pathways underlying prefrontal cortex-medial temporal lobe oscillatory interactions. *Trends in Neurosciences*, *38*(1), 3–12. doi:10.1016/j.tins.2014.09.007

Kleberg, F. I., Kitajo, K., Kawasaki, M., & Yamaguchi, Y. (2014). Ongoing theta oscillations predict encoding of subjective memory type. *Neuroscience Research*, *83*, 69–80. doi:10.1016/j.neures.2014.02.010

Klimesch, W. (1999). EEG alpha and theta oscillations reflect cognitive and memory performance: A review and analysis. *Brain Research. Brain Research Reviews*, *29*(2-3), 169–195.

Lisman, J., Grace, A. A., & Duzel, E. (2011). A neoHebbian framework for episodic memory; role of dopamine-dependent late LTP. *Trends in Neurosciences*, *34*(10), 536–547. doi:10.1016/j.tins.2011.07.006

Lisman, J., & Jensen, O. (2013). The theta-gamma neural code. *Neuron*, *77*(6), 1002–1016. doi:10.1016/j.neuron.2013.03.007

Liu, Z.-X., Grady, C., & Moscovitch, M. (2016). Effects of prior-knowledge on brain activation and connectivity during associative memory encoding. *Cerebral Cortex (New York, N.Y. : 1991)*, bhw047. doi:10.1093/cercor/bhw047

Loh, E., Kumaran, D., Koster, R., Berron, D., Dolan, R., & Duzel, E. (2015). Context-specific activation of hippocampus and SN/VTA by reward is related to enhanced long-term memory for embedded objects. *Neurobiology of Learning and Memory*. doi:10.1016/j.nlm.2015.11.018

Luu, P., Tucker, D. M., & Makeig, S. (2004). Frontal midline theta and the error-related negativity: Neurophysiological mechanisms of action regulation. *Clinical Neurophysiology : Official Journal of the International Federation of Clinical Neurophysiology*, *115*(8), 1821–1835. doi:10.1016/j.clinph.2004.03.031

Maris, E., & Oostenveld, R. (2007). Nonparametric statistical testing of EEG- and MEG-data. *Journal of Neuroscience Methods*, *164*(1), 177–190. doi:10.1016/j.jneumeth.2007.03.024

Merkow, M. B., Burke, J. F., Stein, J. M., & Kahana, M. J. (2014). Prestimulus theta in the human hippocampus predicts subsequent recognition but not recall. *Hippocampus*, *24*(12), 1562–1569. doi:10.1002/hipo.22335

Miendlarzewska, E. A., Bavelier, D., & Schwartz, S. (2016). Influence of reward motivation on human declarative memory. *Neuroscience & Biobehavioral Reviews*, *61*, 156–176. doi:10.1016/j.neubiorev.2015.11.015

Montague, P. R., Dayan, P., & Sejnowski, T. J. (1996). A framework for mesencephalic dopamine systems based on predictive Hebbian learning. *The Journal of Neuroscience : The Official Journal of the Society for Neuroscience*, *16*(5), 1936–1947.

Nyhus, E., & Curran, T. (2010). Functional role of gamma and theta oscillations in episodic memory. *Neuroscience & Biobehavioral Reviews*, *34*(7), 1023–1035. doi:10.1016/j.neubiorev.2009.12.014

Oliveira, F. T. P., McDonald, J. J., & Goodman, D. (2007). Performance monitoring in the anterior cingulate is not all error related: Expectancy deviation and the representation of action-outcome associations. *Journal of Cognitive Neuroscience*, *19*(12), 1994–2004. doi:10.1162/jocn.2007.19.12.1994

Otten, L. J., Quayle, A. H., Akram, S., Ditewig, T. A., & Rugg, M. D. (2006). Brain activity before an event predicts later recollection. *Nature Neuroscience*, *9*(4), 489–491. doi:10.1038/nn1663

Otten, L. J., & Rugg, M. D. (2001). Electrophysiological correlates of memory encoding are task-dependent. *Cognitive Brain Research*, *12*(1), 11–18. doi:10.1016/S0926-6410(01)00015-5

Oya, H., Adolphs, R., Kawasaki, H., Bechara, A., Damasio, A., & Howard, M. A. (2005). Electrophysiological correlates of reward prediction error recorded in the human prefrontal cortex. *Proceedings of the National Academy of Sciences of the United States of America*, *102*(23), 8351–8356. doi:10.1073/pnas.0500899102

Oyarzún, J. P., Packard, P. A., de Diego-Balaguer, R., & Fuentemilla, L. (2016). Motivated encoding selectively promotes memory for future inconsequential semantically-related events. *Neurobiology of Learning and Memory*, *133*, 1–6. doi:10.1016/j.nlm.2016.05.005

Paller, K. A. (1990). Recall and stem-completion priming have different electrophysiological correlates and are modified differentially by directed forgetting. *Journal of Experimental Psychology. Learning, Memory, and Cognition*, *16*(6), 1021.

Paller, K. A., & Wagner, A. D. (2002). Observing the transformation of experience into memory. *Trends in Cognitive Sciences*, *6*(2), 93–102. doi:10.1016/S1364-6613(00)01845-3

Park, H., Lee, D. S., Kang, E., Kang, H., Hahm, J., Kim, J. S., … Jensen, O. (2014). Blocking of irrelevant memories by posterior alpha activity boosts memory encoding. *Human Brain Mapping*, *35*(8), 3972–3987. doi:10.1002/hbm.22452

R Core Team. (2014). R: A language and environment for statistical computing. *R Foundation for Statistical Computing, Vienna, Austria.*

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

Ridderinkhof, K. R., Ullsperger, M., Crone, E. A., & Nieuwenhuis, S. (2004). The role of the medial frontal cortex in cognitive control. *Science (New York, N.Y.)*, *306*(5695), 443–447. doi:10.1126/science.1100301

Ridderinkhof, K. R., Van Den Wildenberg, W. P. M., Segalowitz, S. J., & Carter, C. S. (2004). Neurocognitive mechanisms of cognitive control: The role of prefrontal cortex in action selection, response inhibition, performance monitoring, and reward-based learning. *Brain and Cognition*, *56*(2), 129–140. doi:10.1016/j.bandc.2004.09.016

Salari, N., & Rose, M. (2016). Dissociation of the functional relevance of different pre-stimulus oscillatory activity for memory formation. *NeuroImage*, *125*, 1013–1021. doi:10.1016/j.neuroimage.2015.10.037

Schneider, S. L., & Rose, M. (2016). Intention to encode boosts memory-related pre-stimulus EEG beta power. *NeuroImage*, *125*, 978–987. doi:10.1016/j.neuroimage.2015.11.024

Schneider, W., Eschman, A., & Zuccolotto, A. (2012). *E-Prime User's Guide.* Pittsburgh: Psychology Software Tools, Inc.

Schultz, W., Dayan, P., & Montague, R. (1997). A neural substrate of prediction and reward. *Science (New York, N.Y.)*, *275*(5306), 1593–1599. doi:10.1126/science.275.5306.1593

Silvetti, M., Seurinck, R., & Verguts, T. (2011). Value and prediction error in medial frontal cortex: integrating the single-unit and systems levels of analysis. *Frontiers in Human Neuroscience*, *5*, 75. Retrieved from https://biblio.ugent.be/publication/1989166/file/6764574

Sommer, W., Komoss, E., & Schweinberger, S. R. (1997). Differential localization of brain systems subserving memory for names and faces in normal subjects with event-related potentials. *Electroencephalography and Clinical Neurophysiology*, *102*(3), 192–199. doi:10.1016/S0013-4694(96)95577-0

Squire, L. R. (2004). Memory systems of the brain: A brief history and current perspective. *Neurobiology of Learning and Memory*, *82*(3), 171–177. doi:10.1016/j.nlm.2004.06.005

Sutton, R. S., & Barto, A. G. (1998). *Reinforcement Learning: An Introduction*. Cambridge, MA: MIT Press.

Sweeney-Reed, C. M., Zaehle, T., Voges, J., Schmitt, F. C., Buentjen, L., Kopitzki, K., … Rugg, M. D. (2016). Pre-stimulus thalamic theta power predicts human memory formation. *NeuroImage*. doi:10.1016/j.neuroimage.2016.05.042

Wilding, E., & Ranganath, C. (2011). Electrophysiological correlates of episodic memory processes. In S. J. Luck & E. S. Kappenman (Eds.), *The Oxford handbook of event-related potential components* (pp. 373–396). Oxford: Oxford University Press.

Yordanova, J., Falkenstein, M., Hohnsbein, J., & Kolev, V. (2004). Parallel systems of error processing in the brain. *NeuroImage*, *22*(2), 590–602. doi:10.1016/j.neuroimage.2004.01.040

CHAPTER 5
GENERAL DISCUSSION

In this dissertation we have set out to investigate how prediction errors (PEs) influence perception and learning. Indeed, the prediction of upcoming events offers a double adaptive advantage. First, predictions facilitate perception on the short term, for instance by prioritizing what information enters our visual awareness. They thus determine the speed and accuracy with which matching and mismatching information is detected. Second, PEs serve as a teaching signal, allowing us to perfect our predictive model of the world on the long term. More specifically, reward prediction errors (RPEs) prioritize what information will be remembered and what information will likely be forgotten.

Overall, the goal of this dissertation was to apply a systems approach to get a deeper insight into the mechanisms by which PEs influence perception and learning. In order to do so we drew upon formal models from perceptual decision making studies and the reinforcement learning literature, respectively. First, we investigated by what mechanism visual awareness is influenced by PEs. By applying a drift diffusion model (DDM; Ratcliff, 1985) to a cued individuation task, we demonstrated that PEs influence visual awareness by modulating the threshold setting (but importantly not the processing efficiency). Second, we verified whether RPEs not only drive procedural learning but also enhance declarative learning. In a declarative learning paradigm, we therefore manipulated the size of positive and negative RPEs experienced during word pair encoding. The RPEs were quantified based on the reinforcement learning models typically used in the

literature on procedural learning. In addition, we probed memory performance in an immediate and delayed recognition test in order to verify whether the effect of RPEs on declarative learning became more evident over time as predicted by the neoHebbian learning framework (Lisman, Grace, & Duzel, 2011). Overall, the results of the behavioral and EEG experiment demonstrate that increasingly positive RPEs lead to enhanced declarative learning. In line with the neoHebbian account, this effect was more pronounced in the delayed test compared to the immediate test. Attesting the validity of our findings, the results from the time-frequency analysis further showed that experiencing these RPEs evokes the same neural response as typically observed for RPE during procedural learning. Also, oscillatory power during reward feedback was predictive for the recognition performance.

Below, we will unpack the implications of the current results for the literature on visual awareness and declarative learning. Also, we will discuss future perspectives on how our approach can further deepen our understanding of the underlying cognitive mechanisms.

THE SHORT TERM EFFECT OF PRIOR INFORMATION ON VISUAL AWARENESS

In **chapter 2** we investigated the short term effect of PEs on perception. More specifically, we aimed to clarify how PEs influence stimulus individuation, a process involving the spatial tagging of an object irrespective of its identity (Leslie, Xu, Tremoulet, & Scholl, 1998; Xu, 2009). In an individuation task, participants were therefore asked to locate a picture above or below fixation. At the start of each trial, an icon appeared

which was congruent, incongruent or neutral with respect to the identity of the ensuing picture (i.e., a house or a face). Thus, participants were provided with prior information about the identity of the upcoming stimulus, without creating a bias toward the correct response on the individuation task (i.e., locating the picture above or below fixation). The results indicated that PEs (i.e., caused by incongruent icons) significantly delayed stimulus individuation but had no influence on individuation accuracy.

In order to clarify by what mechanism PEs influence visual awareness, we applied a drift diffusion model (DDM; Ratcliff & Rouder, 1998; Ratcliff, 1985) to the individuation task. This allowed us to disentangle whether PEs delayed individuation by altering either the speed of the evidence accumulation (i.e., the processing efficiency) or by modulating the amount of evidence accumulation needed before the location of a stimulus is detected (i.e., the threshold setting). The results revealed that PEs influence individuation by increasing the threshold setting. Importantly, the PE manipulation had no influence on the processing efficiency.

Next we will discuss how our results feed into the ongoing debate in the consciousness literature on the role of prior information in early (visual) awareness. We will also provide some practical pointers on the use of formal models such as the DDM in future research on the mechanisms by which prior information influences visual awareness.

Using the DDM to investigate visual awareness

With this experiment we provide the first formal test on how prior information influences individuation by differentiating between its effect on the threshold setting and the processing efficiency. Although diffusion

models have already been deployed to investigate the effect of prior information in identification tasks (for a review, see Mulder, van Maanen, & Forstmann, 2014; Summerfield & de Lange, 2014), this was the first study to apply a DDM to an individuation task. As discussed in chapter 2, this is an important step forward as individuation is a typical measure for visual awareness in the consciousness literature (Baars, 1994; Sergent & Dehaene, 2004; Tsuchiya & Koch, 2005).

Using the DDM to investigate the role of prior information in visual awareness makes a timely contribution to the consciousness literature. Indeed, there is currently a great interest in the role that prior information plays in consciousness. Traditionally, conscious awareness has been characterized as a limited capacity system, accommodating only a select set of items (Baars, 1994; Dehaene & Naccache, 2001). In order to reach consciousness, concurrent items are hypothesized to compete for access to consciousness. This competition model has brought about a number of attempts to determine why certain stimuli will receive prioritized access to conscious perception while others are excluded.

In one such line of research, we previously started to explore how the overlap between consciousness and working memory influences the threshold for visual awareness. In a first study, we demonstrated how increasing working memory load increases the threshold for visual awareness (De Loof, Verguts, Fias, & Van Opstal, 2013). Next, we also probed how different types of working memory load had distinct effects on the detectability of masked stimuli (De Loof, Poppe, Cleeremans, Gevers, & Van Opstal, 2015). Other studies have also demonstrated similar interactions between working memory and visual awareness (Gayet, Paffen, & Van der

Stigchel, 2013; Lupyan & Ward, 2013; Pan, Lin, Zhao, & Soto, 2014; Pinto, van Gaal, de Lange, Lamme, & Seth, 2015). Interestingly, some of these studies have probed the influence of the contents of working memory on conscious perception. Using a wide variety of stimuli and tasks, these studies indicated that stimuli that match the content of working memory typically reach awareness faster than mismatching items. These results are in line with the current findings and indicate that prior information on the upcoming stimulus facilitates the detection of related information.

However, there are a number of studies that suggest that unexpected information reaches visual awareness faster, such as pictures depicting incongruent scenes (e.g., a woman placing a chess board in the oven; Mudrik, Breska, Lamy, & Deouell, 2011) or sentences containing semantic errors (Sklar et al., 2012). This has given rise to an ongoing debate in the literature on the extent to which prior information can influence visual awareness or whether these top-down effects are the result of confounds in the experimental design (Gayet, Van der Stigchel, & Paffen, 2014). Recently, Firestone and Scholl (2015) have given a detailed overview of the pitfalls that can occur when testing the effect of top-down expectations on visual awareness, such as the creation of a response bias or the unclear distinction between perception and judgment.

With the DDM we offer a fresh approach to this ongoing debate. First, in addition to determining whether prior information will facilitate or delay visual awareness for matching and mismatching stimuli, it could be useful to map these effects onto their underlying mechanisms to differentiate between seemingly contradictory findings. Second, using diffusion models also provides a direct link with the established literature on the effect of prior

information on more general visual processing (Summerfield & de Lange, 2014; Summerfield & Egner, 2009). Third, because of the link between the model parameters and their neural substrate, augmenting future studies with a formal framework such as the DDM might add to the call for a deeper understanding of the influence of prior information on visual awareness (O'Callaghan, Kveraga, Shine, Adams, & Bar, 2016) and the first steps recently taking in that direction (e.g. using MEG; Aru, Rutiku, Wibral, & Singer, 2016).

Practical pointers for the use of diffusion models in research on visual awareness

Below we will briefly discuss some practical insights gleaned from applying the DDM to classical paradigms used to measure visual awareness.

A necessary adaptation to many classical visual awareness paradigms (e.g., continuous flash suppression and backward masking) will likely be to increase the number of error trials. That is to say, one of the potential drawbacks of the DDM is its reliance on a separate reaction time (RT) distribution for the correct trials as well as the error trials. This necessitates a sufficient number of error trials in order to reliably fit their RT distribution. In an earlier version of the experiment reported in chapter 2, we were unable to reliably fit the DDM due to a low number of error trials. To test the reliability of the parameter estimates, a randomization test was performed by pairing the RTs to the accuracy data in random order. By performing this randomization test on the parameter estimates, we found that the difference between the estimates across the conditions was a mere byproduct of the low number of error trials. More specifically, the low error rate had a particularly distorting effect on the parameter estimates for the incongruent condition

because of the low total number of incongruent trials in the design (i.e., only 20% of the trials). Thus, because the low error rate distorted the parameter estimates more for the incongruent condition compared to the congruent condition, a spurious difference between conditions was created.

A first possible solution to overcome a modest or insufficient number of error trials is to use a hierarchical form of the DDM. The hierarchical DDM (HDDM; Wiecki, Sofer, & Frank, 2013) allows to estimate the various parameters of the model across all participants instead of on an individual basis, as is the case in the traditional DDM. This reduces the detrimental effect of a low number of error trials in a subset of participants. Still, in line with the validations test performed in chapter 2, a thorough test of the model assumptions and possible confounds remains imperative.

In cases where the number of errors is extremely low (as was the case in our earlier version of the experiment), the hierarchical model might still not suffice to solve the problem. As a second solution we therefore added the staircase procedure to the masked priming paradigm (see page 41 for a description). This staircase procedure guaranteed a sufficiently high error rate for each participant. Applying the randomization test to the data reported in chapter 2 indeed demonstrated that the number of errors was sufficiently high in order to yield reliable parameter estimates.

A similar attempt to increase the number of error trials was made when applying the staircase procedure to the continuous flash suppression paradigm (CFS; Tsuchiya & Koch, 2005). During CFS, participants are presented with a continuous stream of colorful flashes to the one eye and a stationary image to the other eye. This causes a sustained suppression of the stationary image by the colorfull flashes, with the stationary image only

breaking through this suppression after several seconds or even minutes. By reducing the visibility of the stationary image (i.e., by increasing its transparacy) we tried to evoke more detection errors. However, after multiple attempts we failed to warrant a sufficiently high level of error trials. Probably the attempts failed because the CFS paradigm lacks a clear speed-accuracy trade-off. That is, participants are asked to detect the stimulus as fast as possible but there is no response deadline. Also, because the RT distribution is characterized by a high variance and many extremely high RTs, there is no clear sense of an internal response deadline. Additional adaptations to the paradigm will thus be necessary before a diffusion model can be fitted to the CFS paradigm and related paradigms that lack a clear speed-accuracy trade-off. Overall, this case illustrates that any visual awareness paradigm will have to be tailored to the requirements of the applied diffusion model before it can be used to differentiate by what mechanisms prior information influences visual awareness.

THE LONG TERM EFFECT OF "BETTER-THAN-EXPECTED" SIGNALS ON DECLARATIVE LEARNING

After establishing how PEs influence visual awareness, we next probed how PEs guide behavior on the long term through their impact on learning. In the introduction of this dissertation, we discussed how the anticipation of reward and the ensuing reward prediction errors (RPEs; Rescorla & Wagner, 1972; Sutton & Barto, 1998) play a crucial role in learning. Ample empirical evidence confirms that RPEs are signaled through dopamine release in the substantia nigra and ventral tegmental area, triggering a dopaminergic pathway with wide projections to cortical and

subcortical structures (Bayer & Glimcher, 2005; Eshel et al., 2015; Montague, Dayan, & Sejnowski, 1996; Schultz, Dayan, & Montague, 1997). In addition, a wide range of models have detailed the way in which the reward value of upcoming events is anticipated and evaluated (e.g. the reward value and prediction model; Silvetti, Seurinck, & Verguts, 2011) and how RPEs function as a teaching signal during procedural learning (Holroyd & Coles, 2002; Rescorla & Wagner, 1972; Sutton & Barto, 1998).

Interestingly, although the effect of RPEs has been firmly established and validated in research on procedural learning, empirical evidence for the impact of RPEs on declarative learning remains surprisingly absent (Squire, 2004). Importantly, a similar effect of RPEs on declarative learning could be expected based on the role of dopamine bursts in hippocampal long-term potentiation (LTP). According to the neoHebbian learning account (Lisman et al., 2011), dopamine bursts tag information that is concurrently encoded in the hippocampus, causing enhanced late LTP for these tagged memory traces. Thus, because RPEs trigger activation in a dopaminergic pathway with projections to the hippocampus, they likely enhance declarative learning. Also, because dopamine is hypothesized to enhance late LTP, the impact of RPEs on declarative memory should become more apparent on a delayed test (e.g. after a delay of a day or a week) compared to an immediate test (e.g., after twenty minutes).

To test this hypothesis and fill this gap between the literature on procedural and declarative learning, we constructed a declarative learning paradigm in which RPEs were coupled to the acquisition of Dutch-Swahili word pairs. Based on the reinforcement learning literature, we quantified the RPEs as the difference between the obtained reward and the reward

expectancy (Bush & Mosteller, 1951a, 1951b; Rescorla & Wagner, 1972; Sutton & Barto, 1998). The reward expectancy was manipulated by letting participants guess among one, two or four Swahili words for each Dutch word. In addition, the obtained reward was manipulated by rewarding the correct guesses and giving no reward for incorrect guesses. Thus, participants experienced RPEs of various quantifiable sizes before the correct word pair was presented. Participants were encouraged to memorize these word pairs for a recognition test that they performed either immediately or after a one-day delay.

The results reported in **chapter 3** indeed demonstrate that word pairs that were coupled to large, positive RPEs during encoding were subsequently recognized with significantly higher accuracy and certainty. Thus, declarative learning was boosted by "better-than-expected" teaching signals (i.e., signed reward prediction errors; SRPEs). The experiment provided an internal replication of the results, as the SRPEs predicted memory performance in the immediate test group as well as the delayed test group. Interestingly, the effect of SRPEs on recognition performance was even stronger after a one-day delay, although this between-subjects comparison failed to reach significance. In addition, we ruled out alternative explanations such as the time-on-task account (Hebb, 1949): longer deliberation on the possible translations had no effect on the subsequent memory performance.

Thus, the current experiment provides the first empirical demonstration of how RPEs influence declarative learning in humans. In order to further validate our findings, we performed a follow-up EEG study (reported in chapter 4). Apart from providing a replication of our previous

results (chapter 3), the first aim of this study was to test whether our quantification of SRPEs could actually account for the neural response to the reward feedback. If so, oscillatory power during reward feedback should be significantly modulated by the SRPEs. Second, we also tested whether SRPEs had a direct effect on declarative learning or whether they influenced learning indirectly (e.g., by increasing attention during encoding; Ridderinkhof, Van Den Wildenberg, Segalowitz, & Carter, 2004). Therefore, we tested whether SRPEs predicted oscillatory power only during reward feedback or also during word pair encoding. In addition, we also probed whether oscillatory power modulations during reward feedback and word pair encoding were predictive for the subsequent recognition performance.

The results of the follow-up EEG study presented in **chapter 4** confirmed the findings from our previous study. SRPEs significantly predicted declarative learning and this effect was again more pronounced when recognition was probed after a one-day delay. The time-frequency analysis on the EEG data additionally allowed us to validate the neural underpinning we proposed for our findings. First, we confirmed that SRPEs significantly predict the oscillatory power in the theta, high-beta and low-beta frequency band. Importantly, this modulating effect of SRPEs on power estimates was only found during reward feedback and did not extend into the word pair encoding phase. These results suggest that participants actually experience a "better-than-expected" signal during reward feedback. Also, the SRPEs evoked by our reward manipulation activated RPE mechanisms similar to those found during procedural learning (Cavanagh, Cohen, & Allen, 2009; Gehring, Liu, Orr, & Carp, 2012; Luu, Tucker, & Makeig, 2004; Oliveira, McDonald, & Goodman, 2007; Yordanova, Falkenstein, Hohnsbein, & Kolev, 2004). Second, we probed whether SRPEs had a direct

influence on declarative memory or whether their influence was more indirect, for instance through increased attention during word pair encoding. The results showed that recognition performance was not significantly predicted by oscillatory power during word pair encoding. However, power oscillations in the alpha band during reward feedback did predict the certainty of the subsequently recognized word pairs, but only when the recognition test was performed after a one-day delay.

Taken together, the results summarized above provide the first empirical evidence demonstrating that SRPEs reliably boost declarative learning. In line with the neoHebbian learning account (Lisman et al., 2011), the effect was more prominent on a delayed test. Also, we confirmed that participants actually experience SRPEs during reward feedback. Our results further suggest that SRPEs have a direct effect on declarative learning, and provide no evidence for an indirect effect through increased attention during word pair encoding. These results make a timely contribution to the literature on declarative learning in light of the recent debate on the application of reward mechanisms in educational settings (Hidi, 2016; Howard-Jones & Jay, 2016; Miendlarzewska, Bavelier, & Schwartz, 2016). Next, we will review what questions remain unanswered and how our current paradigm might provide a useful tool for future research.

Reliability and generalizability of the current paradigm

The reliability and replicability of our research findings deserves special attention as many studies on declarative learning suffer from low statistical power. That is, sample sizes are usually very low as a participant can only learn that many stimuli in a traditional one-hour experiment. Thus, studies often report on the results of paradigms that consist of as few as 48

learned items (Carpenter, 2009). This makes research in declarative learning significantly more challenging compared to procedural learning paradigms. To counter this potential pitfall, a number of initiatives are being taken to conduct large scale studies in educational settings and through (online) educational learning games (for a discussion, see Howard-Jones & Jay, 2016). However, these large scale studies are rare as they require a coordinated, multidisciplinary approach. In the current paradigm (with 60 trials per participant), we alternatively support the reliability of our findings by performing multiple replications. More precisely, SRPEs were a significant predictor for the recognition accuracies in the immediate and delayed group in both the behavioral and EEG experiment. Thus, the results were reliably replicated across four separate test groups, each consisting of approximately twenty participants.

However, the current EEG study might still suffer from low statistical power when relating the trial-by-trial power estimates with the recognition performance. Although both the oscillatory power and the recognition accuracy were significantly predicted by the SRPEs, there was no direct relation between trial-by-trial power and the recognition accuracy. The absence of this relation has many possible causes. First, the relation between the SRPEs and the power estimates might be easier to detect, as the SRPEs were precisely calculated. In contrast, the recognition accuracies only provide a dichotomized measure of the memory performance, and are likely contaminated with correct guesses as participants had a 25% chance of selecting the correct answer in the recognition test. Second, the SRPEs are directly and immediately evoked by the reward feedback, whereas the oscillatory power fluctuations related to enhanced encoding cannot be time-locked as accurately. Third, the trial-by-trial oscillatory power measured at

the scalp likely results in rather noisy estimates. In a paradigm with only 60 trials per participant, the combination of these factors might have resulted in a type II error. By comparison, 160 and 240 stimuli were used in two comparable studies (i.e., studies with a similar number of subjects that also related oscillatory power during encoding with the subsequency recognition accuracy; Gruber, Watrous, Ekstrom, Ranganath, & Otten, 2013; Kleberg, Kitajo, Kawasaki, & Yamaguchi, 2014). In future experiments we should therefore take care to increase the statistical power by increasing the number of participants, the number of trials or the sensitivity of the memory test.

Concerning the generalizability of the current findings, we considered whether the pattern of results might be tightly linked to the memorization instructions and the recognition test. First, one might wonder whether RPEs have an equal influence on intentional and incidental encoding. In the current paradigm, participants were clearly instructed to memorize the word pairs for the recognition test so the encoding was intentional. However, Lisman and colleagues (2011) point out that reward anticipation has been demonstrated to enhance learning during both intentional and incidental memorization (see also Stark & Okado, 2003). In line with these findings, we might argue that our results could be generalized to incidental declarative learning, especially as we found no evidence for a modulating effect of SRPEs on the neural response during word pair encoding.

Second, although we clearly demonstrated the effect of SRPEs on recognition accuracy and certainty ratings, it is unclear whether the results would generalize to other tests such as free recall. Previous studies have already demonstrated distinct effects on recall and recognition (Ketz, Jensen, O'Reilly, & O'Reilly, 2015; Merkow, Burke, Stein, & Kahana, 2014;

Wilding & Ranganath, 2011). On a practical note, given the low statistical power of many declarative learning paradigms, it is quintessential to choose the outcome variable carefully. For instance, we based our paradigm on the testing effect (Gates, 1917; Karpicke & Roediger, 2008), which has been proven to result in a large memory enhancement. In the original study (Karpicke & Roediger, 2008), memory performance increased from approximately 35% accuracy to 80% accuracy, an effect size rarely observed in psychology. Testing the generalizability of the current findings to more stringent memory tests such as free recall, might therefore first require increased statistical power.

Characteristics of the evoked RPEs

A central question in constructing the current paradigm was how to evoke PEs that remained stable throughout the experiment. Indeed, it is hard to create a surprising environment that does not become predictable after a few trials. For example, pairing the presentation of the word pair with the infrequent presence or absence of a concurrent sound might lose its effectiveness after a few trials. With the current paradigm we managed to create stable RPEs, likely because the predictions remained task relevant throughout the experiment. Still, the specific construction of RPEs in the current paradigm highlights a number of characteristics that deserve further attention.

First, in the current experiment, positive and negative RPEs of various sizes were generated by manipulating the obtained reward and the a priori chance of receiving a reward (i.e., by varying the number of response options). By doing so, we also needed to control for the separate effect of the number of response options and the obtained reward, as well as the

interaction between both variables. Especially in the time-frequency analysis, the separate effect of the number of options and obtained reward was clearly detectable in the oscillatory power estimates during reward feedback. Alternative manipulations of the RPEs during declarative learning are needed to fully differentiate between the separate effects of the SRPEs, the number of trials and the obtained reward found in the current study.

Second, our results demonstrated that increasingly negative RPEs have an increasingly negative effect on declarative learning. That is, memory performance was worse in the unrewarded two-option condition compared to the unrewarded four-option condition. This effect was significant in the behavioral experiment (see page 78) as well as the EEG experiment (see page 111). However, this effect should be interpreted with care as participants still performed better than chance on the recognition test (chance performance located at 25% accuracy). Overall, it would be interesting to focus more specifically on the effect of negative RPEs in subsequent studies, as also negative outcomes have been shown to elicit dopamine bursts (Bromberg-Martin, Matsumoto, & Hikosaka, 2010), which might reverse the currently observed effect.

Third, in the discussion of chapter 3 we speculate that the testing effect is largely driven by the RPEs created during the recall tests. We argued that actively engaging in learning by making (reward) predictions might also explain why testing, elaborative interrogation and self-explanation outperform other active learning strategies such as summarizing, keyword mnemonics and imagery (Dunlosky, Rawson, Marsh, Nathan, & Willingham, 2013). However, it mains to be tested empirically whether the

level of agency during prediction and the specificity of the generated predictions play a crucial role in whether RPEs enhance declarative learning.

Lastly, in order to claim that the testing effect is largely driven by SRPEs, we need to address the importance of explicit (monetary) feedback. To start, the monetary reward was mainly provided to encourage participants to actively engage in the decision between the one, two or four Swahili options. After all, the participants had no prior knowledge that could guide them in their choice, so we gave them an additional incentive for deliberating on their guess. However, we don't assume that a monetary incentive is necessary to evoke the SRPEs. More importantly, in the original study on the testing effect, no feedback was provided during testing. This raises the question whether SRPEs offer a valid explanation for the testing effect: How can SRPEs explain the testing effect when no feedback was provided during testing? However, research on the functions of the anterior cingulate cortex (ACC) indicate that both internal and external feedback result in similar teaching signals represented in the dorsal ACC (Holroyd et al., 2004). Thus, a similar internal evaluation of the test performance might have driven the original testing effect even in the absence of external feedback (see also Aarts, Houwer, & Pourtois, 2012; Carpenter, 2009; Schouppe et al., 2015).

The temporal dynamics of the effect of SRPEs on learning

Considering the temporal dynamics of the reported result, we will first consider the importance of a close temporal relation between the reward feedback and the presentation of the to-be-encoded items. It remains to be tested whether the influence of dopamine bursts on declarative information is limited to the stimuli that triggered the dopamine burst, or whether it

extends to events in the temporal proximity of the dopamine triggering event. Based on their review of the literature, Lisman and colleagues (2011) suggest that there is a time window (a *penumbra*) surrounding the dopamine burst and that all stimuli that fall within this penumbra benefit from the enhanced LPT. This penumbra is suggested to potentially span several minutes for events of large adaptive importance. This emphasis on the temporal relation between the reward feedback and the to-be-encoded material, may help elucidate why the SRPEs had a more pronounced effect on declarative learning in the study reported in chapter 3 compared to the study reported in chapter 4. Indeed, whereas the reward feedback and the to-be-remembered word pairs were presented simultaneously in the behavioral study (see page 67), we separated the reward feedback from the presentation of the correct word pair in the EEG experiment reported in chapter 4 (see page 100). By delaying the presentation of the word pair with three seconds, we might have reduced the impact of the SRPE on declarative learning. Clearly, a follow-up study is needed to further elucidate this point.

Next, we evaluate our results in light of the temporal dynamics of the late LTP hypothesized to implement the effect of SRPEs on declarative learning. In line with the neoHebbian learning account, the combined results of chapter 3 and 4 indicate that the effect of SRPEs on declarative learning was more pronounced after a one-day delay. Importantly, one might wonder whether it is odd that we also found a significant effect of SRPEs on declarative learning in the immediate test. For example, in their original demonstration of the testing effect, Karpicke and Roediger (2008) found that participant in the testing and studying condition showed no significant difference in their recognition performance at the end of the learning task. However, the absence of any difference might have been caused by the way

in which recognition performance was measured in this particular study (by calculating the accumulated recognition accuracy over multiple tests) and a ceiling effect (nearly 100% accumulated accuracy). Also, a number of studies have demonstrated that there can be immediate effects on declarative learning (Carpenter, 2009; Gruber et al., 2013; Kleberg et al., 2014).

Finally, the driving effect of RPEs on declarative learning could also be used to refine reconsolidation through distributed rehearsal. Each time a stimulus is reencountered after some delay, the memory trace becomes plastic for a moment, opening a window of opportunity to enhance the strength of the memory trace (Lee, 2009; McGaugh, 1966). In line with the current results, a number of studies have demonstrated that PEs are needed to foster reconsolidation during stimulus rehearsal (Jarome, Ferrara, Kwapis, & Helmstetter, 2015; Sevenster, Beckers, & Kindt, 2012). Applying this reasoning to the testing effect (Karpicke & Roediger, 2008) might offer a more profound explanation of why bulk rehearsal is an inefficient way of learning (Dunlosky et al., 2013). That is, during bulk rehearsal the stimulus material quickly loses its novelty or surprising value, resulting in increasingly small dopamine bursts. Contrastingly, spreading out rehearsal in time allows the material to regain its relative novelty and cause larger dopamine bursts.

Dopamine and declarative learning: beyond the RPE account

Dopaminergic activity is not solely triggered by RPEs. There is a wide variety of other sources of dopamine bursts that might be of interest (Lisman et al., 2011). First, novelty is another classical trigger of dopamine bursts. In addition, novelty can be triggered by a specific stimulus at a specific time, but novelty might also be a characteristic of the environment, encompassing

a wider variety of stimuli over a longer time. For example, exploring a novel environment boosts memory for subsequently presented items, an effect that has been ascribed to the tonic increase of dopamine levels induced by the exploration (Li, Cullen, Anwyl, & Rowan, 2003). Second, a volatile environment in which the reward contingencies change over time has also been demonstrated to boost dopamine release and enhance procedural learning. Typically, a volatile environment leads to a higher learning rate, as demonstrated through modeling and neurophysiological research (Behrens, Woolrich, Walton, & Rushworth, 2007; Rushworth & Behrens, 2008; Silvetti et al., 2011). It would be interesting to test whether novelty and volatility can likewise be demonstrated to enhance declarative learning.

Finally, although in the EEG study we try to confirm that SRPEs drive declarative learning through the activation of the dopaminergic reward system, a causal relation between the proposed neural mechanisms and the learning outcome can only be validated through interventional studies. Pharmacological interventions (e.g. by administering levodopa) can point out whether blocking/boosting the dopaminergic input indeed cancels/facilitates the effect of SRPEs on declarative learning. In addition, intracranial EEG (iEEG) data recorded in the hippocampal area during learning might further our understanding of the impact of SRPEs during declarative learning. Similar iEEG studies have demonstrated that hippocampal oscillatory activity was predictive for successful memory encoding (Fell et al., 2011; Greenberg, Burke, Haque, Kahana, & Zaghloul, 2015; Haque, Wittig, Damera, Inati, & Zaghloul, 2015; Merkow et al., 2014; Sweeney-Reed et al., 2016).

CONCLUSION

The overview presented above demonstrates how the current results add to our understanding of the role that PEs play in perception and learning. We have demonstrated the usefulness of a formal approach such as the DDM to disentangle how prior information influences visual awareness by differentiating between its impact on threshold setting and processing efficiency. Building upon this paradigm, future research could make a useful contribution to the current debate about the influence of prior knowledge and experiences on early visual processing. Also, based on combined insights from the literature on reinforcement learning and the neoHebbian learning framework, we provided the first empirical demonstration for the beneficial impact of SRPEs ("better-than-expected" signals) on declarative learning. The reliability of our paradigm and findings was validated in a follow-up EEG study that provided a full replication of our previous results. In addition, a time-frequency analysis on the EEG data confirmed that participants indeed experience a SRPE during reward feedback and that power modulations during reward feedback were predictive for the recognition performance. Thus, our declarative learning paradigm offers a dependable tool to further explore the interplay between the dopaminergic reward system and declarative learning.

REFERENCES

Aarts, K., Houwer, J. De, & Pourtois, G. (2012). Evidence for the automatic evaluation of self-generated actions. *Cognition*, *124*(2), 117–127. doi:10.1016/j.cognition.2012.05.009

Aru, J., Rutiku, R., Wibral, M., & Singer, W. (2016). Early effects of previous experience on conscious perception. *Neuroscience of Consciousness*. doi:http://dx.doi.org/10.1093/nc/niw004

Baars, B. (1994). A global workspace theory of conscious experience. *Consciousness In Philosophy And Cognitive Neuroscience*, 149–171.

Bayer, H. M., & Glimcher, P. W. (2005). Midbrain dopamine neurons encode a quantitative reward prediction error signal. *Neuron*, *47*(1), 129–141. doi:10.1016/j.neuron.2005.05.020

Behrens, T. E. J., Woolrich, M. W., Walton, M. E., & Rushworth, M. F. S. (2007). Learning the value of information in an uncertain world. *Nature Neuroscience*, *10*(9), 1214–1221. doi:10.1038/nn1954

Bogacz, R. (2007). Optimal decision-making theories: Linking neurobiology with behaviour. *Trends in Cognitive Sciences*, *11*(3), 118–125. doi:10.1016/j.tics.2006.12.006

Bromberg-Martin, E. S., Matsumoto, M., & Hikosaka, O. (2010). Dopamine in motivational control: Rewarding, aversive, and alerting. *Neuron*, *68*(5), 815–834. doi:10.1016/j.neuron.2010.11.022

Bush, R., & Mosteller, F. (1951a). A mathematical model for simple learning. *Psychological Review*, *58*(5), 313–123.

Bush, R., & Mosteller, F. (1951b). A model for stimulus generalization and discrimination. *Psychological Review*, *58*(6), 413–423.

Carpenter, S. K. (2009). Cue strength as a moderator of the testing effect: The benefits of elaborative retrieval. *Journal of Experimental Psychology. Learning, Memory, and Cognition*, *35*(6), 1563–1569. doi:10.1037/a0017021

Cavanagh, J. F., Cohen, M. X., & Allen, J. J. B. (2009). Prelude to and resolution of an error: EEG phase synchrony reveals cognitive control dynamics during action monitoring. *Journal of Neuroscience*, *29*(1), 98–105. doi:10.1523/JNEUROSCI.4137-08.2009

De Loof, E., Poppe, L., Cleeremans, A., Gevers, W., & Van Opstal, F. (2015). Different effects of executive and visuospatial working memory on visual consciousness. *Attention, Perception, & Psychophysics*, *77*(8), 2523–2528. doi:10.3758/s13414-015-1000-8

De Loof, E., Verguts, T., Fias, W., & Van Opstal, F. (2013). Opposite effects of working memory on subjective visibility and priming. *Journal of Experimental Psychology. Learning, Memory, and Cognition*, *39*(6), 1959–1965. doi:10.1037/a0033093

Dehaene, S., & Naccache, L. (2001). Towards a cognitive neuroscience of consciousness: Basic evidence and a workspace framework. *Cognition*, *79*(1-2), 1–37. doi:10.1016/S0010-0277(00)00123-2

Dunlosky, J., Rawson, K. A., Marsh, E. J., Nathan, M. J., & Willingham, D. T. (2013). Improving students' learning with effective learning techniques: Promising directions from cognitive and educational psychology. *Psychological Science in the Public Interest*, *14*(1), 4–58. doi:10.1177/1529100612453266

Eshel, N., Bukwich, M., Rao, V., Hemmelder, V., Tian, J., & Uchida, N. (2015). Arithmetic and local circuitry underlying dopamine prediction errors. *Nature*, *525*(7568), 243–246. doi:10.1038/nature14855

Fell, J., Ludowig, E., Staresina, B. P., Wagner, T., Kranz, T., Elger, C. E., & Axmacher, N. (2011). Medial temporal theta/alpha power enhancement precedes successful memory encoding: Evidence based on intracranial EEG. *Journal of Neuroscience*, *31*(14), 5392–5397. doi:10.1523/JNEUROSCI.3668-10.2011

Firestone, C., & Scholl, B. B. (2015). Cognition does not affect perception: Evaluating the evidence for ' top-down ' effects. *Behavioral and Brain Sciences*, 1–72.

Gates, A. I. (1917). Recitation as a factor in memorization. *Archives of Psychology*, *6*(40), 112–121.

Gayet, S., Paffen, C. L. E., & Van der Stigchel, S. (2013). Information matching the content of visual working memory is prioritized for conscious access. *Psychological Science*, *24*(12), 2472–2480. doi:10.1177/0956797613495882

Gayet, S., Van der Stigchel, S., & Paffen, C. L. E. (2014). Breaking continuous flash suppression: Competing for consciousness on the pre-semantic battlefield. *Frontiers in Psychology*, *5*, 1–10. doi:10.3389/fpsyg.2014.00460

Gehring, W., Liu, Y., Orr, J., & Carp, J. (2012). The error-related negativity (ERN/Ne). In S. J. Luck & E. S. Kappenman (Eds.), *The Oxford handbook of event-related potential components* (pp. 231–291). Oxford: Oxford University Press.

Gold, J. I., & Shadlen, M. N. (2007). The neural basis of decision making. *Annual Review of Neuroscience*, *30*, 535–574. doi:10.1146/annurev.neuro.29.051605.113038

Greenberg, J. A., Burke, J. F., Haque, R., Kahana, M. J., & Zaghloul, K. A. (2015). Decreases in theta and increases in high frequency activity underlie associative memory encoding. *NeuroImage*, *114*, 257–263. doi:10.1016/j.neuroimage.2015.03.077

Gruber, M. J., Watrous, A. J., Ekstrom, A. D., Ranganath, C., & Otten, L. J. (2013). Expected reward modulates encoding-related theta activity before an event. *NeuroImage*, *64*, 68–74. doi:10.1016/j.neuroimage.2012.07.064

Haque, R. U., Wittig, J. H., Damera, S. R., Inati, S. K., & Zaghloul, K. A. (2015). Cortical low-frequency power and progressive phase synchrony precede successful memory encoding. *Journal of Neuroscience*, *35*(40), 13577–13586. doi:10.1523/JNEUROSCI.0687-15.2015

Hebb, D. O. (1949). *The Organization of Behavior.* New York: Wiley & Sons.

Hidi, S. (2016). Revisiting the role of rewards in motivation and learning: Implications of neuroscientific research. *Educational Psychology Review*, *28*(1), 61–93. doi:10.1007/s10648-015-9307-5

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*(4), 679–709. doi:10.1037//0033-295X.109.4.679

Holroyd, C. B., Nieuwenhuis, S., Yeung, N., Nystrom, L., Mars, R. B., Coles, M. G. H., & Cohen, J. D. (2004). Dorsal anterior cingulate cortex shows fMRI response to internal and external error signals. *Nature Neuroscience*, *7*(5), 497–498. doi:10.1038/nn1238

Howard-Jones, P. A., & Jay, T. (2016). Reward, learning and games. *Current Opinion in Behavioral Sciences*, *10*, 65–72. doi:10.1016/j.cobeha.2016.04.015

Jarome, T. J., Ferrara, N. C., Kwapis, J. L., & Helmstetter, F. J. (2015). Contextual information drives the reconsolidation-dependent updating of retrieved fear memories. *Neuropsychopharmacology*, *40*(13), 3044–3052. doi:10.1038/npp.2015.161

Karpicke, J. D., & Roediger, H. L. (2008). The critical importance of retrieval for learning. *Science (New York, N.Y.)*, *319*(5865), 966–968. doi:10.1126/science.1152408

Ketz, N. a., Jensen, O., O'Reilly, R. C., & O'Reilly, R. C. (2015). Thalamic pathways underlying prefrontal cortex-medial temporal lobe oscillatory interactions. *Trends in Neurosciences*, *38*(1), 3–12. doi:10.1016/j.tins.2014.09.007

Kleberg, F. I., Kitajo, K., Kawasaki, M., & Yamaguchi, Y. (2014). Ongoing theta oscillations predict encoding of subjective memory type. *Neuroscience Research*, *83*, 69–80. doi:10.1016/j.neures.2014.02.010

Lee, J. L. C. (2009). Reconsolidation: maintaining memory relevance. *Trends in Neurosciences*, *32*(8), 413–420. doi:10.1016/j.tins.2009.05.002

Leslie, A. M., Xu, F., Tremoulet, P. D., & Scholl, B. J. (1998). Indexing and the object concept: Developing "what" and "where" systems. *Trends in Cognitive Sciences*, *2*(1), 10–18. doi:10.1016/S1364-6613(97)01113-3

Li, S., Cullen, W. K., Anwyl, R., & Rowan, M. J. (2003). Dopamine-dependent facilitation of LTP induction in hippocampal CA1 by exposure to spatial novelty. *Nature Neuroscience*, *6*(5), 526. doi:10.1038/nn1049

Lisman, J., Grace, A. A., & Duzel, E. (2011). A neoHebbian framework for episodic memory; role of dopamine-dependent late LTP. *Trends in Neurosciences*, *34*(10), 536–547. doi:10.1016/j.tins.2011.07.006

Lupyan, G., & Ward, E. J. (2013). Language can boost otherwise unseen objects into visual awareness. *Proceedings of the National Academy of Sciences of the United States of America*, *110*(35), 14196–14201. doi:10.1073/pnas.1303312110

Luu, P., Tucker, D. M., & Makeig, S. (2004). Frontal midline theta and the error-related negativity: Neurophysiological mechanisms of action regulation. *Clinical Neurophysiology : Official Journal of the International Federation of Clinical Neurophysiology*, *115*(8), 1821–1835. doi:10.1016/j.clinph.2004.03.031

McGaugh, J. L. (1966). Time-dependent processes in memory storage. *Science (New York, N.Y.)*, *153*(3742), 1351–1358. doi:10.1126/science.153.3742.1351

Merkow, M. B., Burke, J. F., Stein, J. M., & Kahana, M. J. (2014). Prestimulus theta in the human hippocampus predicts subsequent recognition but not recall. *Hippocampus*, *24*(12), 1562–1569. doi:10.1002/hipo.22335

Miendlarzewska, E. A., Bavelier, D., & Schwartz, S. (2016). Influence of reward motivation on human declarative memory. *Neuroscience & Biobehavioral Reviews*, *61*, 156–176. doi:10.1016/j.neubiorev.2015.11.015

Montague, P. R., Dayan, P., & Sejnowski, T. J. (1996). A framework for mesencephalic dopamine systems based on predictive Hebbian learning. *The Journal of Neuroscience : The Official Journal of the Society for Neuroscience*, *16*(5), 1936–1947.

Mudrik, L., Breska, A., Lamy, D., & Deouell, L. Y. (2011). Integration without awareness: expanding the limits of unconscious processing. *Psychological Science*, *22*(6), 764–770. doi:10.1177/0956797611408736

Mulder, M. J., van Maanen, L., & Forstmann, B. U. (2014). Perceptual decision neurosciences - a model-based review. *Neuroscience*, *277*, 872–884. doi:10.1016/j.neuroscience.2014.07.031

O'Callaghan, C., Kveraga, K., Shine, J. M., Adams, R. B., & Bar, M. (2016). Predictions penetrate perception: Converging insights from brain, behaviour and disorder. *Consciousness and Cognition*. doi:10.1016/j.concog.2016.05.003

Oliveira, F. T. P., McDonald, J. J., & Goodman, D. (2007). Performance monitoring in the anterior cingulate is not all error related: Expectancy deviation and the representation of action-outcome associations. *Journal of Cognitive Neuroscience*, *19*(12), 1994–2004. doi:10.1162/jocn.2007.19.12.1994

Pan, Y., Lin, B., Zhao, Y., & Soto, D. (2014). Working memory biasing of visual perception without awareness. *Attention, Perception & Psychophysics*, *76*(7), 2051–2062. doi:10.3758/s13414-013-0566-2

Pinto, Y., van Gaal, S., de Lange, F. P., Lamme, V. A. F., & Seth, A. K. (2015). Expectations accelerate entry of visual stimuli into awareness. *Journal of Vision*, *15*(8), 1–15. doi:10.1167/15.8.13

Ratcliff, R. (1985). Theoretical interpretations of the speed and accuracy of positive and negative responses. *Psychological Review*, *92*(2), 212–225. doi:10.1037//0033-295X.92.2.212

Ratcliff, R., & McKoon, G. (2008). The diffusion decision model: Theory and data for two-choice decision tasks. *Neural Computation*, *20*(4), 873–922. doi:10.1162/neco.2008.12-06-420

Ratcliff, R., & Rouder, J. N. (1998). Modeling response times for two-choice decisions. *Psychological Science*, *9*(5), 347–356. doi:10.1111/1467-9280.00067

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

Ridderinkhof, K. R., Van Den Wildenberg, W. P. M., Segalowitz, S. J., & Carter, C. S. (2004). Neurocognitive mechanisms of cognitive control: The role of prefrontal cortex in action selection, response inhibition, performance monitoring, and reward-based learning. *Brain and Cognition*, *56*(2), 129–140. doi:10.1016/j.bandc.2004.09.016

Rushworth, M. F. S., & Behrens, T. E. J. (2008). Choice, uncertainty and value in prefrontal and cingulate cortex. *Nature Neuroscience*, 11(4), 389–397. doi:10.1038/nn2066

Schouppe, N., Braem, S., De Houwer, J., Silvetti, M., Verguts, T., Ridderinkhof, K. R., & Notebaert, W. (2015). No pain, no gain: The affective valence of congruency conditions changes following a successful response. *Cognitive, Affective, & Behavioral Neuroscience*, *15*(1), 251–261. doi:10.3758/s13415-014-0318-3

Schultz, W., Dayan, P., & Montague, R. (1997). A neural substrate of prediction and reward. *Science (New York, N.Y.)*, *275*(5306), 1593–1599. doi:10.1126/science.275.5306.1593

Sergent, C., & Dehaene, S. (2004). Is consciousness a gradual phenomenon? Evidence for an all-or-none bifurcation during the attentional blink. *Psychological Science*, *15*(11), 720–728. doi:10.1111/j.0956-7976.2004.00748.x

Sevenster, D., Beckers, T., & Kindt, M. (2012). Retrieval per se is not sufficient to trigger reconsolidation of human fear memory. *Neurobiology of Learning and Memory*, *97*(3), 338–345. doi:10.1016/j.nlm.2012.01.009

Silvetti, M., Seurinck, R., & Verguts, T. (2011). Value and prediction error in medial frontal cortex: Integrating the single-unit and systems levels of analysis. *Frontiers in Human Neuroscience*, *5*(75). doi:10.3389/fnhum.2011.00075

Sklar, A. Y., Levy, N., Goldstein, A., Mandel, R., Maril, A., & Hassin, R. R. (2012). Reading and doing arithmetic nonconsciously. *Proceedings of the National Academy of Sciences*, *109*(48), 19614–19619.

Squire, L. R. (2004). Memory systems of the brain: A brief history and current perspective. *Neurobiology of Learning and Memory*, *82*(3), 171–177. doi:10.1016/j.nlm.2004.06.005

Stark, C. E. L., & Okado, Y. (2003). Making memories without trying: medial temporal lobe activity associated with incidental memory formation during recognition. *The Journal of Neuroscience : The Official Journal of the Society for Neuroscience*, *23*(17), 6748–53.

Summerfield, C., & de Lange, F. P. (2014). Expectation in perceptual decision making: Neural and computational mechanisms. *Nature Reviews Neuroscience*, *15*(11), 745–756. doi:10.1038/nrn3838

Summerfield, C., & Egner, T. (2009). Expectation (and attention) in visual cognition. *Trends in Cognitive Sciences*, *13*(9), 403–409. doi:10.1016/j.tics.2009.06.003

Sutton, R. S., & Barto, A. G. (1998). *Reinforcement Learning: An Introduction*. Cambridge, MA: MIT Press.

Sweeney-Reed, C. M., Zaehle, T., Voges, J., Schmitt, F. C., Buentjen, L., Kopitzki, K., … Rugg, M. D. (2016). Pre-stimulus thalamic theta power predicts human memory formation. *NeuroImage*. doi:10.1016/j.neuroimage.2016.05.042

Tsuchiya, N., & Koch, C. (2005). Continuous flash suppression reduces negative afterimages. *Nature Neuroscience*, *8*(8), 1096–1101. doi:10.1038/nn1500

Wiecki, T. V, Sofer, I., & Frank, M. J. (2013). HDDM: Hierarchical Bayesian estimation of the Drift-Diffusion Model in Python. *Frontiers in Neuroinformatics*, *7*, 1–10. doi:10.3389/fninf.2013.00014

Wilding, E., & Ranganath, C. (2011). Electrophysiological correlates of episodic memory processes. In S. J. Luck & E. S. Kappenman (Eds.), *The Oxford handbook of event-related potential components* (pp. 373–396). Oxford: Oxford University Press.

Xu, Y. (2009). Distinctive neural mechanisms supporting visual object individuation and identification. *Journal of Cognitive Neuroscience*, *21*(3), 511–518. doi:10.1162/jocn.2008.21024

Yordanova, J., Falkenstein, M., Hohnsbein, J., & Kolev, V. (2004). Parallel systems of error processing in the brain. *NeuroImage*, *22*(2), 590–602. doi:10.1016/j.neuroimage.2004.01.040

CHAPTER 6
NEDERLANDSTALIGE SAMENVATTING

In dit doctoraatsproefschrift beogen we dieper te onderzoeken hoe onze verwachtingen een invloed hebben op perceptie en leren. We vertrekken daarbij in **hoofdstuk 1** vanuit het *predictive coding* denkkader (Friston & Kiebel, 2009; Rao & Ballard, 1999; Summerfield et al., 2006). Dit theoretisch kader karakteriseert de hersenen als een voorspellend mechanisme dat constant voorspellingen maakt over de nabije toekomst. Het maakt daarvoor gebruik van een intern model van de wereld dat toelaat te anticiperen op binnenkomende informatie. Het enige wat achteraf nog verwerkt moet worden, is de niet-voorspelde input (i.e., de voorspellingsfouten). Op die manier stelt het voorspellend systeem de hersenen in staat informatie te verwerken op een energiezuinige wijze. In dat opzicht spelen voorspellingen en voorspellingsfouten een belangrijke rol bij het prioriteren van informatie binnen een breed scala van cognitieve processen.

Dit voorspellend mechanisme heeft twee centrale gevolgtrekkingen die we in dit proefschrift verder onderzoeken. Ten eerste spelen de voorspellingen die we maken op korte termijn een bepalende rol in hoe we de wereld om ons heen percipiëren. We onderzoeken daarbij hoe visuele verwerking beïnvloed wordt door deze voorspellingen en de resulterende voorspellingsfouten. Ten tweede fungeren voorspellingsfouten als een leersignaal zodat we ons intern model van de wereld kunnen optimaliseren op lange termijn. Vooral voorspellingsfouten rond beloningsfeedback (i.e., *reward prediction errors*, RPEs; Bush & Mosteller, 1951a, 1951b; Rescorla

& Wagner, 1972) spelen een cruciale rol tijdens leren. Hoewel de invloed van RPEs op leren reeds uitgebreid onderzocht werd binnen procedureel leren is er opmerkelijk weinig empirische evidentie voor een soortgelijke invloed op declaratief leren (Squire, 2004). Daarom focussen we in het tweede deel van dit proefschrift op het langetermijneffect van RPEs op declaratief leren.

Voor ons onderzoek kozen we bewust een systemische aanpak; vandaar de ondertitel van dit proefschrift, *a systems approach*. We vertrekken daarbij telkens van fundamenteel hersenonderzoek naar hoe neurale circuits informatie verwerken. Dit laat toe om de mechanismes waardoor voorspellingsfouten perceptie en leren beïnvloeden beter te begrijpen, met name door gebruik te maken van formele modellen. Deze formele modellen laten toe een onderscheid te maken tussen verschillende onderliggende processen. Dit kan door de verschillende processen te relateren aan unieke parameters die bovendien exact gekwantificeerd kunnen worden.

Hieronder geven we een korte toelichting van ons onderzoek naar het kortetermijneffect van voorspellingsfouten op perceptie en het langetermijneffect van voorspellingsfouten op leren. We gaan daarbij in op de belangrijkste resultaten en bespreken deze in het kader van de formele modellen gekoppeld aan elke onderzoeksvraag.

KORTETERMIJNEFFECT VAN VOORSPELLINGSFOUTEN OP PERCEPTIE

In eerst instantie richten we ons op perceptie en meer bepaald op de rol die voorspellingsfouten op korte termijn spelen binnen visuele

gewaarwording. De invloed van voorspellingen op visuele gewaarwording gaf reeds aanleiding tot een aanzienlijke onderzoekstraditie en wordt eveneens onderbouwd door een aantal theoretische kaders (Enns & Lleras, 2008; Hohwy, 2012; Rauss, Schwartz, & Pourtois, 2011; Seth, Suzuki, & Critchley, 2011; Summerfield et al., 2006). Zo zouden voorspellingen een belangrijke rol spelen bij het prioriteren van de input die we bewust waarnemen. Er is inderdaad heel wat evidentie die aangeeft dat voorspelde stimuli sneller en accurater worden waargenomen (Neely, 1991). Het specifieke mechanisme waardoor voorspellingen visuele gewaarwording beïnvloeden blijft echter betwist.

In **hoofdstuk 2** gingen we daarom na hoe de waarneming van een stimulus beïnvloedt wordt door voorafgaande informatie over de identiteit van die stimulus. Proefpersonen kregen de opdracht een stimulus te detecteren (i.e., de afbeelding van een huis of een gezicht) voorafgegaan door een icoon dat neutraal, congruent, of incongruent was met betrekking tot de identiteit van de afbeelding. De identiteit van de afbeelding was echter ongerelateerd aan het correcte antwoord, want proefpersonen werden enkel gevraagd de locatie van de stimulus te rapporteren (i.e., boven of onder het fixatiekruis) en niet de identiteit van de stimulus. Op die manier konden we het effect van voorspellingsfouten op visuele gewaarwording testen zonder een antwoord-bias te creëren. De resultaten van het experiment geven aan dat voorspellingsfouten met betrekking tot de identiteit van de stimulus (i.e., incongruente trials) leiden tot een tragere detectie van de stimuluslocatie. De voorspellingsfouten hadden echter geen invloed op de accuraatheid waarmee de detectietaak werd uitgevoerd.

Om de invloed van voorspellingsfouten op perceptie beter te vatten, maakten we gebruik van een diffusiemodel (i.e., het *drift diffusion model*, DDM; Bogacz, 2007; Gold & Shadlen, 2007; Ratcliff & McKoon, 2008; Ratcliff & Rouder, 1998; Ratcliff, 1985). Dit model beschrijft hoe tijdens een beslissingsproces (i.e., in dit geval een detectieproces) evidentie voor de verschillende antwoordalternatieven accumuleert door de tijd heen. Eens de geaccumuleerde evidentie een beslissingsdrempel bereikt (i.e., in dit geval een detectiedrempel), wordt het bijhorende antwoord gegeven. Voorspellingen kunnen op verschillende manieren het visuele detectieproces beïnvloeden. Enerzijds kunnen voorspellingen de accumulatiesnelheid van de (on)voorspelde informatie verhogen of verlagen. Anderzijds kan een voorspelling beïnvloeden hoeveel evidentie er nodig is voordat een beslissingsdrempel wordt bereikt.

De DDM-analyse wees uit dat voorspellingsfouten perceptie beïnvloeden via hun effect op de beslissingsdrempel; deze lag hoger op incongruente trials in vergelijking met congruente en neutrale trials. Met andere woorden, wanneer er een huis verschijnt terwijl je verwachtte een gezicht te zien, is er meer geaccumuleerde evidentie nodig vooraleer de afbeelding van het huis gelokaliseerd kan worden. De voorspellende informatie had echter geen effect op de snelheid waarmee evidentie in het detectieproces accumuleert.

Met dit experiment brachten we twee stromingen in de literatuur rond voorspellingsfouten en visuele gewaarwording voor het eerst samen. Voorheen werden diffusiemodellen reeds frequent ingezet in experimenten waarbij proefpersonen gevraagd werden de identiteit van een stimulus te bepalen (Domenech & Dreher, 2010; Dunovan, Tremel, & Wheeler, 2014;

Forstmann, Brown, Dutilh, Neumann, & Wagenmakers, 2010; Mulder, van Maanen, & Forstmann, 2014; Mulder, Wagenmakers, Ratcliff, Boekel, & Forstmann, 2012; Summerfield & de Lange, 2014). Dit *identificatieproces* is echter verschillende van het proces dat in het huidige experiment onderzocht werd. We gingen namelijk voor het eerste na hoe voorspellingen de visuele gewaarwording beïnvloeden in een detectietaak waarbij het lokaliseren van een stimulus centraal staat. Dit onderscheiden *individuatieproces* (i,e., het lokaliseren van een stimulus in de ruimte; Leslie, Xu, Tremoulet, & Scholl, 1998; Xu, 2009) staat centraal in onderzoek naar visueel bewustzijn (Baars, 1994; Overgaard & Sandberg, 2012; Sandberg, Timmermans, Overgaard, & Cleeremans, 2010; Sergent & Dehaene, 2004; Tononi & Koch, 2008). Met het huidige experiment bieden we bijgevolg een antwoord op de vraag hoe voorspellingen bepalen welke informatie prioritair het bewustzijn bereikt (Enns & Lleras, 2008; Hohwy, Roepstorff, & Friston, 2008; Kouider, de Gardelle, Sackur, & Dupoux, 2010). Het diffusiemodel biedt bovendien een dieper begrip van de processen die deze prioritering sturen en reikt de nodige handvaten aan om de neurale basis van dit proces verder te specificeren.

LANGETERMIJNEFFECT VAN VOORSPELLINGSFOUTEN OP LEREN

In tweede instantie onderzochten we hoe voorspellingsfouten op lange termijn een invloed hebben op leren. Een lange onderzoekstraditie wijst uit dat leren voornamelijk gedreven wordt door voorspellingsfouten met betrekking tot de mate waarin een ervaring of uitkomst als belonend wordt ervaren (i.e., een *reward prediction error*, RPE). Een ervaring kan meer belonend uitdraaien dan verwacht (positieve RPE) of minder belonend dan verwacht (negatieve RPE). Deze RPEs worden typisch gekwantificeerd als

het verschil tussen de verkregen beloning en de verwachte beloning. Uit hersenonderzoek weten we ook dat (positieve) RPEs een invloed hebben op leren dankzij hun stimulerend effect op de tijdelijke vrijzetting van dopamine die plaatsvindt in de substantia nigra en het ventraal tegmentaal gebied (Bayer & Glimcher, 2005; Schultz, Dayan, & Montague, 1997).

Hoewel het effect van positieve en negatieve RPEs uitgebreid onderzocht is in de literatuur rond procedureel leren (Squire, 2004), is er vooralsnog geen empirisch onderzoek naar de impact van RPEs op declaratief leren (bij mensen). Er is echter reeds heel wat evidentie uit onderzoek op dieren dat uitwijst dat dopamine een heel specifieke rol speelt bij consolidatieprocessen in de hippocampus. Zo stelt de neoHebbiaanse leertheorie van Lisman, Grace en Duzel (2011) dat een tijdelijke toename van dopaminevrijzetting bepaalt welke informatie door de hippocampus zal worden vastgelegd in het langetermijngeheugen en welke informatie zal worden vergeten. Gezien RPEs voor vergelijkbare toenames in dopaminevrijzetting zorgen, voorzagen we dat ook declaratief leren versterkt zou worden door RPEs.

Om dit te onderzoeken ontworpen we een leerparadigma waarin proefpersonen Nederlands-Swahili woordparen memoriseerden. Elk woordpaar ging tijdens het memoriseren gepaard met een positieve of negatieve RPE van gekende grootte. Deze RPEs werden gecreëerd door proefpersonen bij elk Nederlands woord eerst te laten gokken tussen een aantal Swahili opties. Door in dit keuzeproces per trial het aantal beschikbare opties te variëren (i.e., één, twee of vier opties) bepaalden we de verwachte kans op beloning. Daarnaast varieerden we ook of de proefpersonen positieve of negatieve beloningsfeedback kregen (i.e., een

geldelijke beloning of geen beloning). Het verschil tussen de verkregen en verwachte beloning biedt een exacte kwantificatie van de RPE die de proefpersonen ervoeren tijdens het memoriseren van het correcte Nederlands-Swahili woordpaar. Vervolgens testten we het geheugen voor de woordparen in een herkenningstest. De proefpersonen werden daarvoor opgedeeld in twee groepen waarbij de helft van de proefpersonen de herkenningstest meteen uitvoerde en de andere proefpersonen de volgende dag terugkwamen voor de herkenningstest. Volgens de neoHebbiaanse leertheorie zou het versterkend effect van dopaminevrijzetting op consolidatie immers toenemen in de tijd, en daardoor beter waarneembaar zijn op een uitgestelde test (e.g., na een dag of een week) dan op een onmiddellijke test (e.g., na twintig minuten).

In **hoofdstuk 3** rapporteren we de resultaten van onze eerste studie naar het effect van RPEs op declaratief leren. Deze wees uit dat woordparen vaker en met grotere zekerheid herkend worden naarmate er toenemend positieve RPEs optraden tijdens het leren. Gezien de geheugenperformantie toenam naarmate de RPEs toenemend positief werden en afnam naarmate de RPEs toenemend negatief werden, interpreteren we de RPEs hier als "beter-dan-verwacht" leersignalen (i.e., *signed reward prediction errors*, SRPEs). Bij procedureel leren daarentegen wordt leren frequent versterkt door zowel grote positieve als grote negatieve RPEs, gekend als "anders-dan-verwacht" leersignalen (i.e., *unsigned reward prediction errors*, URPEs). Zoals voorspeld door de neoHebbiaanse leertheorie was het effect van SRPEs op declaratief leren groter op de uitgestelde test dan op de onmiddellijke test. Bovendien gingen we na of de verbeterde herkenning ook verklaard kon worden door toegenomen verwerkingstijd (*time-on-task*) tijdens het

keuzeproces (i.e., kiezen tussen één, twee of vier opties). Dit bleek echter niet het geval te zijn.

Met deze studie leverden we de eerste empirische evidentie voor het effect van SRPEs op declaratief leren. De stabiliteit van het effect werd bovendien onderlijnd door een interne replicatie: SRPEs voorspelden de geheugenperformantie zowel bij de groep proefpersonen die de geheugentest onmiddellijk uitvoerde als bij de proefpersonen die de test pas een dag later uitvoerden. Het blijft echter de vraag of SRPEs een rechtstreekse invloed hebben op declaratief leren (i.e., door de invloed van dopaminevrijzetting op consolidatieprocessen in de hippocampus) of leren onrechtstreeks beïnvloeden, bijvoorbeeld door toegenomen aandacht na positieve beloningsfeedback (Chun, Turk-Browne, Tanaka, & Watanabe, 2007; Kruschke, 2001; Pearce & Hall, 1980; Wills, Lavric, Croft, & Hodgson, 2007).

Om een dieper inzicht te krijgen in de impact van SRPEs op het memoriseren van de woordparen voerden we een vervolgexperiment uit waarbij we tijdens het leren de hersenactiviteit registreerden met behulp van elektroencefalografie (EEG). Eerst en vooral fungeerde dit experiment als een dubbele replicatiestudie, met ook deze keer een aparte groep proefpersonen in de onmiddellijke en uitgestelde testconditie. Verder gingen we na of proefpersonen tijdens de beloningsfeedback effectief een "beter-dan-verwacht" leersignaal ervoeren. Om dit na te gaan pasten we een *time-frequency analyse* toe op de EEG-data en gingen we na of SRPEs een significante voorspeller waren voor de oscillatorische amplitude (i.e., *power*) in verscheidende frequentiebanden. Bovendien gingen we ook na of de power tijdens het leren een significante voorspeller was voor de

performantie op de herkenningstest. We maakten hierbij een belangrijk onderscheid tussen de power tijdens de beloningsfeedback en de power tijdens de daaropvolgende memorisatie van het correcte woordpaar. Op die manier konden we nagaan of SRPEs een rechtstreeks effect hebben op declaratief leren (i.e., powermodulaties tijdens de beloningsfeedback zijn voorspellend voor de geheugenperformantie) of een indirect effect hebben via toegenomen aandacht tijdens het daaropvolgend memorisatieproces (i.e., powermodulaties tijdens de memorisatie zijn voorspellend voor de geheugenperformantie).

In **hoofdstuk 4** rapporteren we de resultaten van deze EEG-vervolgstudie. Eerst en vooral boden de gedragsdata een volledige replicatie van de effecten die we in de eerste studie vonden: woordparen worden beter en met grotere zekerheid herkend naarmate een toenemend positieve RPE optrad tijdens het leren (i.e., "beter-dan-verwacht" leersignalen; SRPEs). Dit effect bleek opnieuw sterker te zijn in de uitgestelde test dan in de onmiddellijke test. De *time-frequency analyse* op de EEG data wees bovendien uit dat de proefpersonen effectief een SRPE ervoeren tijdens de beloningsfeedback aangezien de amplitude van de oscillaties in verschillende frequentiebanden significant voorspeld werd door de SRPEs. Dit bevestigt dat de kwantificatie van de SRPEs in onze studie overeenkwam met hoe proefpersonen de beloningsfeedback ervaren. Daarenboven kon de geheugenperformantie enkel significant voorspeld worden op basis van de power tijdens de beloningsfeedback en niet door de power tijdens de daaropvolgende memorisatiefase. Dit suggereert dat SRPEs een direct effect hebben op declaratief leren en biedt geen evidentie voor een indirect effect via toegenomen aandacht tijdens memorisatie.

Deze twee experimenten bieden de eerste empirische demonstratie van het effect van SRPEs op declaratief leren bij mensen. We toonden daarbij aan dat onze manipulatie van de SRPEs overeenkwam met de ervaring van de proefpersonen. Deze SRPEs bleken ook een zelfde neurale response uit te lokken (i.e., gelijkaardige powermodulaties) als typisch waargenomen tijdens RPEs in procedureel leren (Cavanagh, Cohen, & Allen, 2009; Gehring, Liu, Orr, & Carp, 2012; Luu, Tucker, & Makeig, 2004; Oliveira, McDonald, & Goodman, 2007; Yordanova, Falkenstein, Hohnsbein, & Kolev, 2004). Uit de resultaten bleek bovendien dat alternatieve processen zoals *time-on-task* tijdens het keuzeproces of toegenomen aandacht tijdens de memorisatiefase geen alternatieve verklaring kunnen bieden voor de toegenomen geheugenperformantie. Op deze manier biedt dit experiment een solide basis om verder te exploreren welke randvoorwaarden gelden bij het versterkend effect van SRPEs op declaratief leren. Hoe belangrijk is de temporele relatie of het causaal verband tussen de beloningsfeedback en het te leren materiaal? Wat is het precieze effect van negatieve RPEs? Gaat het versterkend effect van SRPEs op declaratief leren ook op voor andere bronnen van toegenomen dopaminevrijzetting zoals nieuwe informatie (*novelty*) of onzekerheid over het beloningsregime (*volatility*)? Ten slotte kan de voorgestelde neurale basis van het effect ook verder onderzocht worden door activatie in de hippocampale regio te registreren via intracraniale EEG of door de dopaminevrijzetting rechtstreeks te manipuleren in een farmacologische interventiestudie.

CONCLUSIE

Samengevat gingen we in het huidige proefschrift na hoe voorspellingsfouten een effect hebben op perceptie en leren. Door gebruik te maken van formele modellen waren we in staat specifiek na te gaan hoe voorspellingen bepalen welke informatie er prioritair verwerkt wordt en op welke manier deze prioriteit geïmplementeerd wordt. In de **algemene discussie** gaan we dieper in op de implicaties van onze resultaten voor de literatuur rond voorspellingsfouten, visuele gewaarwording en declaratief leren. Tot slot, evalueren we hoe onze systemische aanpak toekomstige mogelijkheden biedt om een gedetailleerd inzicht te krijgen in de prioriterende mechanismes waardoor voorspellingsfouten een impact hebben op perceptie en leren.

REFERENTIES

Baars, B. (1994). A global workspace theory of conscious experience. *Consciousness In Philosophy And Cognitive Neuroscience*, 149–171.

Bayer, H. M., & Glimcher, P. W. (2005). Midbrain dopamine neurons encode a quantitative reward prediction error signal. *Neuron*, *47*(1), 129–141. doi:10.1016/j.neuron.2005.05.020

Bogacz, R. (2007). Optimal decision-making theories: Linking neurobiology with behaviour. *Trends in Cognitive Sciences*, *11*(3), 118–125. doi:10.1016/j.tics.2006.12.006

Bush, R., & Mosteller, F. (1951a). A mathematical model for simple learning. *Psychological Review*, *58*(5), 313–123.

Bush, R., & Mosteller, F. (1951b). A model for stimulus generalization and discrimination. *Psychological Review*, *58*(6), 413–423.

Cavanagh, J. F., Cohen, M. X., & Allen, J. J. B. (2009). Prelude to and resolution of an error: EEG phase synchrony reveals cognitive control dynamics during action monitoring. *Journal of Neuroscience*, *29*(1), 98–105. doi:10.1523/JNEUROSCI.4137-08.2009

Chun, M. M., Turk-Browne, N. B., Tanaka, K., & Watanabe, T. (2007). Interactions between attention and memory. *Current Opinion in Neurobiology*, *17*, 177–184. doi:10.1016/j.conb.2007.03.005

Domenech, P., & Dreher, J.-C. (2010). Decision threshold modulation in the human brain. *The Journal of Neuroscience : The Official Journal of the Society for Neuroscience*, *30*(43), 14305–14317. doi:10.1523/JNEUROSCI.2371-10.2010

Dunovan, K., Tremel, J., & Wheeler, M. (2014). Prior probability and feature predictability interactively bias perceptual decisions. *Neuropsychologia*, *61*, 210–221. doi:10.1016/j.neuropsychologia.2014.06.024.Prior

Enns, J. T., & Lleras, A. (2008). What's next? New evidence for prediction in human vision. *Trends in Cognitive Sciences*, *12*(9), 327–333. doi:10.1016/j.tics.2008.06.001

Forstmann, B. U., Brown, S., Dutilh, G., Neumann, J., & Wagenmakers, E.-J. (2010). The neural substrate of prior information in perceptual decision making: A model-based analysis. *Frontiers in Human Neuroscience*, *4*, 40. doi:10.3389/fnhum.2010.00040

Friston, K., & Kiebel, S. (2009). Predictive coding under the free-energy principle. *Philosophical Transactions of the Royal Society of London. Series B, Biological Sciences*, *364*(1521), 1211–1221. doi:10.1098/rstb.2008.0300

Gehring, W., Liu, Y., Orr, J., & Carp, J. (2012). The error-related negativity (ERN/Ne). In S. J. Luck & E. S. Kappenman (Eds.), *The Oxford handbook of event-related potential components* (pp. 231–291). Oxford: Oxford University Press.

Gold, J. I., & Shadlen, M. N. (2007). The neural basis of decision making. *Annual Review of Neuroscience*, *30*, 535–574. doi:10.1146/annurev.neuro.29.051605.113038

Hohwy, J. (2012). Attention and conscious perception in the hypothesis testing brain. *Frontiers in Psychology*, *3*(96). doi:10.3389/fpsyg.2012.00096

Hohwy, J., Roepstorff, A., & Friston, K. (2008). Predictive coding explains binocular rivalry: An epistemological review. *Cognition*, *108*(3), 687–701. doi:10.1016/j.cognition.2008.05.010

Kouider, S., de Gardelle, V., Sackur, J., & Dupoux, E. (2010). How rich is consciousness? The partial awareness hypothesis. *Trends in Cognitive Sciences*, *14*(7), 301–307. doi:10.1016/j.tics.2010.04.006

Kruschke, J. K. (2001). Toward a unified model of attention in associative learning. *Journal of Mathematical Psychology*, *45*(6), 812–863. doi:10.1006/jmps.2000.1354

Leslie, A. M., Xu, F., Tremoulet, P. D., & Scholl, B. J. (1998). Indexing and the object concept: Developing "what" and "where" systems. *Trends in Cognitive Sciences*, *2*(1), 10–18. doi:10.1016/S1364-6613(97)01113-3

Lisman, J., Grace, A. A., & Duzel, E. (2011). A neoHebbian framework for episodic memory; role of dopamine-dependent late LTP. *Trends in Neurosciences*, *34*(10), 536–547. doi:10.1016/j.tins.2011.07.006

Luu, P., Tucker, D. M., & Makeig, S. (2004). Frontal midline theta and the error-related negativity: Neurophysiological mechanisms of action regulation. *Clinical Neurophysiology : Official Journal of the International Federation of Clinical Neurophysiology*, *115*(8), 1821–1835. doi:10.1016/j.clinph.2004.03.031

Mulder, M. J., van Maanen, L., & Forstmann, B. U. (2014). Perceptual decision neurosciences - a model-based review. *Neuroscience*, *277*, 872–884. doi:10.1016/j.neuroscience.2014.07.031

Mulder, M. J., Wagenmakers, E.-J., Ratcliff, R., Boekel, W., & Forstmann, B. U. (2012). Bias in the brain: A diffusion model analysis of prior probability and potential payoff. *The Journal of Neuroscience : The Official Journal of the Society for Neuroscience*, *32*(7), 2335–2343. doi:10.1523/JNEUROSCI.4156-11.2012

Neely, J. (1991). Semantic priming effects in visual word recognition: A selective review of current findings and theories. *Basic Processes in Reading: Visual Word Recognition*, *11*, 264–336.

Oliveira, F. T. P., McDonald, J. J., & Goodman, D. (2007). Performance monitoring in the anterior cingulate is not all error related: Expectancy deviation and the representation of action-outcome associations. *Journal of Cognitive Neuroscience*, *19*(12), 1994–2004. doi:10.1162/jocn.2007.19.12.1994

Overgaard, M., & Sandberg, K. (2012). Kinds of access: Different methods for report reveal different kinds of metacognitive access. *Philosophical Transactions of the Royal Society of London. Series B, Biological Sciences*, *367*(1594), 1287–1296. doi:10.1098/rstb.2011.0425

Pearce, J. M., & Hall, G. (1980). A model for Pavlovian learning: Variations in the effectiveness of conditioned but not of unconditioned stimuli. *Psychological Review*, *87*, 532–552. doi:10.1037//0033-295X.87.6.532

Rao, R. P., & Ballard, D. H. (1999). Predictive coding in the visual cortex: A functional interpretation of some extra-classical receptive-field effects. *Nature Neuroscience*, *2*(1), 79–87. doi:10.1038/4580

Ratcliff, R. (1985). Theoretical interpretations of the speed and accuracy of positive and negative responses. *Psychological Review*, *92*(2), 212–225. doi:10.1037//0033-295X.92.2.212

Ratcliff, R., & McKoon, G. (2008). The diffusion decision model: Theory and data for two-choice decision tasks. *Neural Computation*, *20*(4), 873–922. doi:10.1162/neco.2008.12-06-420

Ratcliff, R., & Rouder, J. N. (1998). Modeling response times for two-choice decisions. *Psychological Science*, *9*(5), 347–356. doi:10.1111/1467-9280.00067

Rauss, K., Schwartz, S., & Pourtois, G. (2011). Top-down effects on early visual processing in humans: A predictive coding framework. *Neuroscience & Biobehavioral Reviews*, *35*(5), 1237–1253. doi:10.1016/j.neubiorev.2010.12.011

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

Sandberg, K., Timmermans, B., Overgaard, M., & Cleeremans, A. (2010). Measuring consciousness: Is one measure better than the other? *Consciousness and Cognition*, *19*(4), 1069–1078. doi:10.1016/j.concog.2009.12.013

Schultz, W., Dayan, P., & Montague, R. (1997). A neural substrate of prediction and reward. *Science (New York, N.Y.)*, *275*(5306), 1593–1599. doi:10.1126/science.275.5306.1593

Sergent, C., & Dehaene, S. (2004). Is consciousness a gradual phenomenon? Evidence for an all-or-none bifurcation during the attentional blink. *Psychological Science*, *15*(11), 720–728. doi:10.1111/j.0956-7976.2004.00748.x

Seth, A. K., Suzuki, K., & Critchley, H. D. (2011). An interoceptive predictive coding model of conscious presence. *Frontiers in Psychology*, *2*, 395. doi:10.3389/fpsyg.2011.00395

Squire, L. R. (2004). Memory systems of the brain: A brief history and current perspective. *Neurobiology of Learning and Memory*, *82*(3), 171–177. doi:10.1016/j.nlm.2004.06.005

Summerfield, C., & de Lange, F. P. (2014). Expectation in perceptual decision making: Neural and computational mechanisms. *Nature Reviews Neuroscience*, *15*(11), 745–756. doi:10.1038/nrn3838

Summerfield, C., Egner, T., Greene, M., Koechlin, E., Mangels, J., & Hirsch, J. (2006). Predictive codes for forthcoming perception in the frontal cortex. *Science (New York, N.Y.)*, *314*(5803), 1311–1314. doi:10.1126/science.1132028

Tononi, G., & Koch, C. (2008). The neural correlates of consciousness: An update. *Annals of the New York Academy of Sciences*, *1124*, 239–261. doi:10.1196/annals.1440.004

Wills, A. J., Lavric, A., Croft, G. S., & Hodgson, T. L. (2007). Predictive learning, prediction errors, and attention: Evidence from event-related potentials and eye tracking. *Journal of Cognitive Neuroscience*, *19*(5), 843–854. doi:10.1162/jocn.2007.19.5.843

Xu, Y. (2009). Distinctive neural mechanisms supporting visual object individuation and identification. *Journal of Cognitive Neuroscience*, *21*(3), 511–518. doi:10.1162/jocn.2008.21024

Yordanova, J., Falkenstein, M., Hohnsbein, J., & Kolev, V. (2004). Parallel systems of error processing in the brain. *NeuroImage*, *22*(2), 590–602. doi:10.1016/j.neuroimage.2004.01.040

DATA STORAGE FACT SHEETS

In compliance with the UGent standard for research accountability, transparacy and reproducibility, the location of the datasets used in this dissertation are added below. For each of the empirical chapters (i.e., chapters 2 to 4) a separate Data Storage Fact Sheet is completed, detailing which data and analysis files are stored, where they are stored, who has access to the files and who can be contacted in order to request access to the files. In addition, the Data Storage Fact Sheets have been added to my public UGent Biblio account.

DATA STORAGE FACT SHEET FOR CHAPTER 2

Name/identifier study
Author: Esther De Loof
Date: 14-06-2016

1. Contact details

1a. Main researcher

- name: Esther De Loof
- address: Henri Dunantlaan 2 9000 Gent
- e-mail: estherdeloof@gmail.com or esther.deloof@ugent.be

1b. Responsible Staff Member (ZAP)

- name: Tom Verguts
- address: Henri Dunantlaan 2 9000 Gent
- e-mail: Tom.Verguts@ugent.be

If a response is not received when using the above contact details, please send an email to data.pp@ugent.be or contact Data Management, Faculty of Psychology and Educational Sciences, Henri Dunantlaan 2, 9000 Ghent, Belgium.

2. Information about the datasets to which this sheet applies

- Reference of the publication in which the datasets are reported:

  Chapter 2 of PhD dissertation (first empirical chapter): Predictive information speeds up visual awareness in an individuation task by modulating threshold setting, not processing efficiency

  De Loof, E., Van Opstal, F., & Verguts, T. (2016). Predictive information speeds up visual awareness in an individuation task by modulating threshold setting, not processing efficiency. *Vision Research, 121*, 104-112. http://dx.doi.org/10.1016/j.visres.2016.03.002

- Which datasets in that publication does this sheet apply to?:

  All data from the reported experiment (behavioral).

3. Information about the files that have been stored

3a. Raw data

- Have the raw data been stored by the main researcher? [x] YES / [ ] NO
  If NO, please justify:

- On which platform are the raw data stored?
  - [x]   researcher PC
  - [ ]   research group file server
  - [ ]   other (specify): ...

- Who has direct access to the raw data (i.e., without intervention of another person)?
  - [x]   main researcher
  - [ ]   responsible ZAP
  - [ ]   all members of the research group
  - [ ]   all members of UGent
  - [ ]   other (specify): ...

3b. Other files

- Which other files have been stored?
  - [ ]   file(s) describing the transition from raw data to reported results. Specify: ...
  - [ ]   file(s) containing processed data. Specify: ...
  - [x]   file(s) containing analyses. Specify: R scripts (containing all processing steps and their justification)
  - [x]   files(s) containing information about informed consent (printed informed consents stored in folder)
  - [ ]   a file specifying legal and ethical provisions
  - [ ]   file(s) that describe the content of the stored files and how this content should be interpreted. Specify: ...
  - [ ]   other files. Specify: ...

- On which platform are these other files stored?
  - [x]   individual PC
  - [ ]   research group file server
  - [ ]   other: ...

- Who has direct access to these other files (i.e., without intervention of another person)?
  - [x]   main researcher
  - [ ]   responsible ZAP
  - [ ]   all members of the research group
  - [ ]   all members of UGent
  - [ ]   other (specify): ...

4. **Reproduction**

- Have the results been reproduced independently?: [ ] YES / [x] NO

- If yes, by whom (add if multiple):
  - name:
  - address:
  - affiliation:
  - e-mail:

v0.2

DATA STORAGE FACT SHEET FOR CHAPTER 3

Name/identifier study
Author: Esther De Loof
Date: 14-06-2016

1. Contact details

1a. Main researcher

- name: Esther De Loof
- address: Henri Dunantlaan 2 9000 Gent
- e-mail: estherdeloof@gmail.com or esther.deloof@ugent.be

1b. Responsible Staff Member (ZAP)

- name: Tom Verguts
- address: Henri Dunantlaan 2 9000 Gent
- e-mail: Tom.Verguts@ugent.be

If a response is not received when using the above contact details, please send an email to data.pp@ugent.be or contact Data Management, Faculty of Psychology and Educational Sciences, Henri Dunantlaan 2, 9000 Ghent, Belgium.

2. Information about the datasets to which this sheet applies

- Reference of the publication in which the datasets are reported:

  Chapter 3 of PhD dissertation (second empirical chapter): Signed reward prediction errors drive declarative learning

- Which datasets in that publication does this sheet apply to?:

  All data from the reported experiment (behavioral).

3. <u>Information about the files that have been stored</u>

3a. <u>Raw data</u>

- Have the raw data been stored by the main researcher? [x] YES / [ ] NO
  If NO, please justify:

- On which platform are the raw data stored?
  - [x]   researcher PC
  - [ ]   research group file server
  - [ ]   other (specify): ...

- Who has direct access to the raw data (i.e., without intervention of another person)?
  - [x]   main researcher
  - [ ]   responsible ZAP
  - [ ]   all members of the research group
  - [ ]   all members of UGent
  - [ ]   other (specify): ...

3b. <u>Other files</u>

- Which other files have been stored?
  - [ ]   file(s) describing the transition from raw data to reported results. Specify: ...
  - [ ]   file(s) containing processed data. Specify: ...
  - [x]   file(s) containing analyses. Specify: R scripts (containing all processing steps and their justification)
  - [x]   files(s) containing information about informed consent (printed informed consents stored in folder)
  - [ ]   a file specifying legal and ethical provisions
  - [ ]   file(s) that describe the content of the stored files and how this content should be interpreted. Specify: ...
  - [ ]   other files. Specify: ...

- On which platform are these other files stored?
  - [x]   individual PC
  - [ ]   research group file server
  - [ ]   other: ...

- Who has direct access to these other files (i.e., without intervention of another person)?
  - [x] main researcher
  - [ ] responsible ZAP
  - [ ] all members of the research group
  - [ ] all members of UGent
  - [ ] other (specify): ...

4. **Reproduction**

- Have the results been reproduced independently?: [ ] YES / [x] NO

- If yes, by whom (add if multiple):
  - name:
  - address:
  - affiliation:
  - e-mail:

v0.2

DATA STORAGE FACT SHEET FOR CHAPTER 4

Name/identifier study
Author: Esther De Loof
Date: 14-06-2016

1. Contact details

1a. Main researcher

- name: Esther De Loof
- address: Henri Dunantlaan 2 9000 Gent
- e-mail: estherdeloof@gmail.com or esther.deloof@ugent.be

1b. Responsible Staff Member (ZAP)

- name: Tom Verguts
- address: Henri Dunantlaan 2 9000 Gent
- e-mail: Tom.Verguts@ugent.be

If a response is not received when using the above contact details, please send an email to data.pp@ugent.be or contact Data Management, Faculty of Psychology and Educational Sciences, Henri Dunantlaan 2, 9000 Ghent, Belgium.

2. Information about the datasets to which this sheet applies

- Reference of the publication in which the datasets are reported:

  Chapter 4 of PhD dissertation (third empirical chapter): The modulation of EEG oscillations by reward prediction errors drives declarative learning

- Which datasets in that publication does this sheet apply to?:

  All data from the reported experiment (behavioral + EEG).

3. <u>Information about the files that have been stored</u>

3a. <u>Raw data</u>

- Have the raw data been stored by the main researcher? [x] YES / [ ] NO
  If NO, please justify:

- On which platform are the raw data stored?
  - [x]   researcher PC
  - [x]   research group file server
  - [ ]   other (specify): ...

- Who has direct access to the raw data (i.e., without intervention of another person)?
  - [x]   main researcher
  - [x]   responsible ZAP
  - [ ]   all members of the research group
  - [ ]   all members of UGent
  - [ ]   other (specify): ...

3b. <u>Other files</u>

- Which other files have been stored?
  - [x]   file(s) describing the transition from raw data to reported results. Specify: Word document describing preprocessing steps for the EEG data
  - [x]   file(s) containing processed data. Specify: Folders with the substages of the preprocessing of the EEG data
  - [x]   file(s) containing analyses. Specify: R and MATLAB scripts (containing all processing steps and their justification)
  - [x]   files(s) containing information about informed consent (printed informed consents stored in folder)
  - [ ]   a file specifying legal and ethical provisions
  - [ ]   file(s) that describe the content of the stored files and how this content should be interpreted. Specify: ...
  - [ ]   other files. Specify: ...

- On which platform are these other files stored?
  - [x]   individual PC
  - [x]   research group file server
  - [ ]   other: ...

- Who has direct access to these other files (i.e., without intervention of another person)?
  - [x]   main researcher
  - [x]   responsible ZAP
  - [ ]   all members of the research group
  - [ ]   all members of UGent
  - [ ]   other (specify): ...

4. Reproduction

- Have the results been reproduced independently?: [ ] YES / [x] NO

- If yes, by whom (add if multiple):
  - name:
  - address:
  - affiliation:
  - e-mail:

v0.2