

Predicting beyond ourselves:

How expectations shape the perception of external
action-outcomes

Myrthel Dogge

ISBN: 978-94-6332-508-0

Lay-out: Ferdinand van Nispen, *my-thesis.nl*

Printed by: GVO drukkers & vormgevers, Ede, The Netherlands

Predicting beyond ourselves:

How expectations shape the perception of
external action-outcomes

Voorspellingen die verder reiken dan onszelf:
Hoe verwachting de perceptie van externe actie-uitkomsten beïnvloedt
(met een samenvatting in het Nederlands)

Proefschrift

ter verkrijging van de graad van doctor aan de Universiteit Utrecht
op gezag van de rector magnificus, prof.dr. H.R.B.M. Kummeling,
ingevolge het besluit van het college voor promoties in het openbaar te
verdedigen op maandag 3 juni 2019 des middags te 12.45 uur

door

Myrthel Dogge

geboren op woensdag 26 juli 1989
te Nieuwegein

Promotor: Prof. dr. H. Aarts

Copromotoren: Dr. R. Custers
Dr. D. Hofman

Dit proefschrift werd (mede) mogelijk gemaakt met financiële steun van de Nederlandse Organisatie voor Wetenschappelijk Onderzoek (NWO) [MaGW Onderzoekstalentbeurs: 406-14-039].

TABLE OF CONTENTS

Chapter 1	Introduction and overview	7
Chapter 2	Exploring the role of motor and non-motor predictive mechanisms in sensory attenuation: Perceptual and neurophysiological findings	21
Chapter 3	The influence of action-effect anticipation on bistable perception: Differences between onset rivalry and ambiguous motion	47
Chapter 4	Perception of action-outcomes is shaped by life-long and contextual expectations	73
Chapter 5	Cortical information flow during inferences of agency	95
Chapter 6	Moving forward: On the limits of motor-based forward models	123
Appendices		137
Summary in Dutch (Nederlandse samenvatting)		151
References		163
Acknowledgments (Dankwoord)		177
Biography		185

CHAPTER 1

Introduction and Overview

“The fundamental purpose of brains is to produce [the] future [...]. In order to cope, an organism must either armor itself (like a tree or a clam) and “hope for the best,” or else develop methods of getting out of harm’s way and into the better neighborhoods in its vicinity. If you follow this latter course, you are confronted with the primordial problem that every agent must continually solve: Now what do I do? In order to solve this problem, you need a nervous system, to control your activities in time and space. The juvenile sea squirt wanders through the sea searching for a suitable rock or hunk of coral to cling to and make its home for life. For this task, it has a rudimentary nervous system. When it finds its spot and takes root, it doesn’t need its brain anymore so it eats it!” – Daniel Dennett (1991)

Humans are nothing like sea squirts. Much unlike these peculiar creatures our lives are continuously colored by goal-directed actions through which we interact with our environment. Indeed, it has been argued that this faculty lies at the very essence of our humanness (Bandura, 2001). Nevertheless, Dennett’s illustration highlights a truth for human functioning that has been increasingly acknowledged over the past few decades: The ability to act has greatly affected, and continues to affect, human brain physiology and cognition (Herwig, 2015; Herwig, Beisert, & Prinz, 2013). One particularly intriguing example of such a side-effect of action is our conscious experience of this capacity. Specifically, whenever we act, we also feel and know that we are the cause of our movements and their resulting consequences. This pervasive phenomenon is known as our sense of agency, and is central to the current dissertation.

Although perhaps not immediately apparent, the sense of agency impacts many facets of our daily lives (for a review see: Moore, 2016). It is for instance tightly coupled to responsibility, and as such, in western society, seen as a prerequisite for punishment by law (Haggard, 2017; Haggard & Tsakiris, 2009). That is, individuals are only punished (or rewarded) for behavior that is performed volitionally and of which the consequences are understood. In addition, the sense of agency is highly relevant for the increasing number of our interactions that are mediated by technology. Through technological interfaces, the effects of our actions are no longer restricted to our immediate environment, nor require our physical presence (e.g., doctors can perform transatlantic surgeries (Marescaux et al., 2001) and we can “touch” loved ones that are miles away (e.g., Saadatian et al., 2014)). To ensure that these mediated actions feel as unmediated, and as closely related to the experience of controlling our own body as possible, the

mechanisms that give rise to agentic experiences need to be taken into account (Limerick, Coyle, & Moore, 2014).

The importance of the experience of agency also becomes apparent when considering the implications of its disturbance. Aberrant agentic processing is central to various psychiatric disorders. The by far most extensively studied example is schizophrenia, in which patients display grave impairments in their ability to recognize the source of their actions and resulting outcomes. These patients may both claim control over events they did not cause, as well as report a lack of control over events they did produce (Frith, Blakemore, & Wolpert, 2000a; Hur, Kwon, Lee, & Park, 2014). Likewise, depressed (Alloy & Abramson, 1979; Obhi, Swiderski, & Farquhar, 2013) and anxious patients (Gentsch, Endrass, & Kathmann, 2012; Oren, Friedmann, & Dar, 2016) tend to experience a loss of control over their behavior, resulting in feelings of helplessness and compulsive behaviors, respectively.

Considering these broad implications, understanding how the sense of agency arises is a key question, which I assess in my dissertation. Before turning to the specifics of the research I conducted, I will first provide a (by no means exhaustive) overview of the measures and explanatory models of agency experiences that have dominated the field.

RESEARCHING THE SENSE OF AGENCY

The sense of agency has been examined with a large variety of methods, which can be coarsely classified as focusing on one of two levels of agentic experience: *the judgement of agency* or *the feeling of agency* (Synofzik, Vosgerau, & Newen, 2008). The judgement of agency reflects the explicit attribution of actions and outcomes to ourselves or others. In daily life such explicit experiences mostly occur when there is uncertainty regarding the true cause of an event, such as when our actions do not lead to their expected or desired outcomes (e.g., when the room does not light up after turning on the light switch). Judgements of agency are generally assessed by means of explicit reports. Typically, participants are exposed to sensory action-outcomes of which the true causal origin is made ambiguous, either as result of spatiotemporal discrepancies between actions and effects (e.g., Sato & Yasuda, 2005), or due to the (implied) influence of another potential agent (e.g., Aarts, Custers, & Wegner, 2005; Renes, Haren, & Aarts, 2015). Participants are then asked to state to what extent they, as opposed to another source, caused the event to

occur. Hence, most (but not all) experiments on explicit agency judgements are inherently social.

On top of these high-level judgements of agency, another, rather different aspect of the sense of agency can be identified. In particular, during most of our actions we rarely stop in our tracks to think about whether or not we caused an event. Instead, the experience of agency has been described “phenomenally thin” (Haggard, 2005): we generally have an ongoing, “buzzing”, experience of causing the consequences of our actions – especially when we are alone and things unfold as intended. This more indirect feeling of agency is probed by implicit agency measures. Instead of directly asking participants about their experience, implicit measures assess the perceptual processing of action-outcomes, which are used as a proxy for experienced agency. Two perceptual phenomena that have been extensively used to study implicit agency experiences are intentional binding and sensory attenuation.

Intentional binding experiments have been inspired by Libet’s famous insight that conscious awareness of events surrounding voluntary action can be measured by means of a clock (Libet, Gleason, Wright, & Pearl, 1983). Borrowing from this idea, a typical intentional binding task requires participants to use the position of a rotating clock hand to judge the timing of a button press and the tone it produces (for a review see: Moore & Obhi, 2012). Participants tend to perceive their action as shifted towards the tone, and the tone as shifted towards the action, compared to conditions in which these events occur in isolation. That is, actions and effects are perceived as drawn together in time, which corresponds with the holistic experience we have of our actions and their consequences in daily life. Crucially, this binding effect is known to disappear in the absence of factual agency, such as when one’s actions are involuntary (Engbert, Wohlschläger, & Haggard, 2008; Haggard, Clark, & Kalogeras, 2002). This dependence on voluntary action is why the intentional binding effect was originally seen as a marker of agentic experience (a view that is not uncontested – see below).

In addition to the subjective timing, voluntary actions also affect the perceived intensity of action-outcomes. Specifically, self-produced effects tend to be suppressed compared to effects produced by an external source, both on a perceptual and a neurophysiological level. The most famous example of this sensory attenuation phenomenon is our inability to tickle ourselves (Blakemore, Wolpert, & Frith, 2000; Weiskrantz, 1971). Comparable attenuation effects have been observed in other sensory modalities. For instance, the neural activity that is evoked by the

sound of our own voice is suppressed compared to when passively listening to recordings of the same sound (Curio, Neuloh, Numminen, Jousmaki, & Hari, 2000; Heinks-Maldonado, Mathalon, Gray, & Ford, 2005). In addition, our visual sensitivity is decreased for self- versus externally produced stimuli (Cardoso-Leite, Mamassian, Schütz-Bosbach, & Waszak, 2010; Roussel, Hughes, & Waszak, 2013). As only self-produced effects are suppressed, these attenuation effects are thought to underlie our ability to attribute events to their proper cause (Frith, Blakemore, & Wolpert, 2000b).

Albeit elegant, implicit measures of agency have not escaped criticism. The main concern is to what extent these measures can really be seen as markers of the sense of agency. This concern is based on the fact that the evidence for a relationship between explicit agency reports on the one hand, and intentional binding and sensory attenuation on the other hand, has been mixed and inconclusive (Braun, Thorne, Hildebrandt, & Debener, 2014; Dewey & Knoblich, 2014; Ebert & Wegner, 2010; Imaizumi & Tanno, 2019; Pyasik, Burin, & Pia, 2018; Saito, Takahata, Murai, & Takahashi, 2015; Timm, Schönwiesner, Schröger, & SanMiguel, 2016). Importantly, however, the absence of such a direct relationship does not necessarily warrant the conclusion that implicit measures cannot tell us anything about agentic experiences. Instead, it might point to the possibility that the feeling of agency (arguably reflected by sensory attenuation and intentional binding) and the judgement of agency at least partially reflect distinct aspects of the sense of agency (Synofzik et al., 2008)¹.

Another potential reason for the absence of a clear relationship between implicit and explicit measures is that the methods generally used to assess implicit agency experiences might not capture the core aspects of our ongoing “buzzing” agentic experience. For instance, it is not directly clear how perceived attenuation of action-outcomes would lead to a positive experience of agency without the need of a rather counterintuitive additional interpretative step (for a critical discussion see: Christensen & Grünbaum, 2017; Haggard, 2017). In Chapter 3 of this dissertation I address this issue by introducing another, perhaps more intuitive, way to assess implicit agency processing.

1 This does not imply that the levels are completely independent of each other. Instead, it has been argued that the feeling of agency might serve as one of the cues, but not the only cue, affecting judgements of agency (and the other way around), which leaves room for occasional dissociations (Synofzik et al., 2008, 2013).

MECHANISMS UNDERLYING THE AGENTIC SENSE OF SELF

Much like the aforementioned dissidence on the best way to measure the sense of agency, there has been ample discussion on its underlying mechanisms. Two theories that have been particularly influential in the field are the comparator model and the theory of apparent mental causation. The comparator model was originally conceived in the context of sensorimotor control (Miall & Wolpert, 1996; Wolpert & Flanagan, 2001). According to this account, voluntary movements are accompanied by a copy of the motor command, which is used to generate predictions about the consequences of our movement. The predicted sensory input can in turn be compared to the desired outcomes to make rapid online adjustments to movements. In addition, by comparing predicted with actual action-consequences, input that is a mere by-product of movement can be suppressed. This suppression is thought to be beneficial as the effects that accompany our own movements (e.g., the sound of our voice or footsteps) are often far less informative than external input that could signal threat or reward (e.g., the sound of a child crying or a friend calling our name). Over the past few decades, this model has been increasingly extrapolated beyond simple sensorimotor control, and has been directly implicated in the experience of agency. Specifically, it has been proposed that a match between actual and predicted feedback underlies the conclusion that input must be self-caused, whereas a substantial discrepancy between what was predicted and what was observed will result in an external causal attribution (Feinberg, 1978; Frith et al., 2000a; Haggard & Chambon, 2012).

The view that the comparator model underlies agency processing is typically substantiated by the aforementioned perceptual differences between self-produced and externally-produced effects, both in terms of perceived intensity (i.e., sensory attenuation) and timing (i.e., intentional binding). That is, seeing that only outcomes that are self-produced can be predicted, only these outcomes (and not externally-produced outcomes) can be attenuated or temporally bound to actions. However, as has been explained in detail elsewhere (Hughes, Desantis, & Waszak, 2013a), the comparison between self-produced and externally-produced effects differs in more ways than merely the presence (or absence) of motor signals. For instance, when events are produced by our own actions we cannot only predict (with a certain probability) which sensory consequence will result from our action (motor identity prediction), but we can also predict at which point in time this event will occur (temporal prediction). As both intentional binding and attenuation are known to

be modulated by such alternative sources of information (e.g., Haggard et al., 2002; Lange, 2009; Ruess, Thomaschke, & Kiesel, 2017; Vroomen & Stekelenburg, 2010), the exact role of motor predictive signals still needs to be elucidated. This issue is addressed in Chapters 2 and 3 of this dissertation.

Another theoretical account, the theory of apparent mental causation, offers a rather different explanation of the origins of agentic experiences. In particular, this account proposes that the experience of agency results from the post-hoc comparison of the perceived event and previously activated thoughts or intentions (Wegner, 2002; Wegner & Wheatley, 1999). The extent to which agency is experienced depends on three key features: whether the thoughts about the action-outcome occurred before the action, whether the thoughts are consistent with the action (and the resulting outcome), and whether there are potential alternative causes of the event. Hence, in comparison to the comparator model, this account downplays the importance of sensorimotor signals, while higher-level, non-motor influences are presumed to carry a far greater weight. This idea has been supported by research showing that inferred agency over action-outcomes increased when the representations of these outcomes were pre-activated, either by asking participants to form the explicit goal to produce this outcome, or by means of priming (Aarts et al., 2005; Renes et al., 2015; Van der Weiden, Ruys, Aarts, & Ruys, 2013). In contrast to the comparator model, research on the mechanisms and neural substrates underlying these inferences of agency has been rather limited. In Chapter 5, we attempt to take a first step towards filling this gap.

On top of these two dominant models, cue-integration models have been proposed (Moore & Fletcher, 2012; Synofzik et al., 2008; Synofzik, Vosgerau, & Voss, 2013). As their name suggests, these accounts propose that the sense of agency results from multiple cues, including the motor-predictive signals that are central to the comparator model, as well as the cognitive, situational cues highlighted by the theory of apparent mental causation. The contribution of different cues has been proposed to rely on their relative availability and reliability. However, many questions about the details of their interaction remain.

CURRENT DISSERTATION

Taken together, the previous paragraphs illustrate the dependence of the sense of agency on both signals directly related to motor control, as well as on information

originating from non-motor sources. The present dissertation revolves around how these different sources of information jointly affect agentic processing. In Chapter 2, I first set out to establish whether action-outcome predictions based on motor and non-motor sources have comparable effects on experienced agency (as reflected by sensory attenuation). To this end, participants took part in a task in which they were exposed to associations between key presses (motor condition) and tones, or shape cues (non-motor condition) and tones. Subsequently it was examined whether tones that were compatible (versus incompatible) with learned associations were attenuated, both in terms of their perceived intensity, as well as in terms of evoked auditory potentials. Much to my surprise, however, no robust evidence was observed for the influence of motor- or non-motor predictions, neither on a perceptual nor on a neurophysiological level.

The findings from Chapter 2, and particularly the absence of a motor-prediction effect, were surprising given the strong weight that is generally given to the comparator account as an explanatory model for the sense of agency. In Chapter 3, I consider whether these null-findings might be related to shortcomings of the perceptual phenomena that are typically used to study motor-prediction effects. Specifically, the comparator account proposes that actions are accompanied by very specific predictions about the identity of upcoming action-effects. That is, when pressing a key on a piano, a pianist does not expect just any tone to occur, but instead anticipates hearing a tone with a specific pitch. Although both sensory attenuation and intentional binding reflect the processing of action-outcomes (i.e., their subjective timing and intensity), neither of these measures directly relates to “*what*” is predicted per se. In other words, stimulus identity is not central to these phenomena. Accordingly, these paradigms might not be optimally suitable to assess the processes that are specified in the comparator model.

In Chapter 3, I therefore utilize a phenomenon known as bistable perception, in which stimulus identity is of inherent importance. Bistable perception occurs when the brain is exposed to a stimulus that has multiple possible interpretations, and is characterized by the alternation of these interpretations over time (Blake & Logothetis, 2002; Leopold & Logothetis, 1999). One factor that is known to influence the interpretation of these stimuli is one’s prior expectations about what one will see (Scocchia, 2014). I designed a task in which bistable images were produced by actions, and in which only one of the potential interpretations of the image was compatible with predictions that were formed based on previously learned action-outcome associations. Across three experiments, mixed results were

obtained. While motor-predictions biased the perception of an ambiguous motion sphere², I observed no such effects when actions were followed by the presentation of rivalrous gratings³. One of the differences between these stimuli is that they arguably reflect different stages of visual competition. In particular, the conflict between rivalrous gratings might be resolved at a lower stage of visual processing, and therefore might be less susceptible to top-down modulation, than the conflict that is central to ambiguous motion. Although speculative, this interpretation points towards potential restrictions to motor-predictive influences on perception.

The findings discussed so far seem to indicate that motor-based signals at least affect perceptual processing of action-outcomes under some circumstances. In Chapter 4, I further zoom in on the characteristics of motor-predictions and examine to what extent they can be dynamically updated over time. (Motor)-predictions can develop across different timescales and contexts and are not always complementary to each other (Seriès & Seitz, 2013). For instance, when reverse parking a car with a trailer, one's overlearned expectation that the car will move in the same direction as the steering wheel, conflicts with the contextual expectation that the trailer will move in the opposite direction as one's action. In this chapter I examine how these long-term and short-term expectations jointly determine the perception of action-outcomes. Participants completed a task in which they initiated the motion of an ambiguous sphere with a rotary manual action. As a result of lifelong learning, participants tend to, by default, perceive the sphere as rotating in the same direction as their movement⁴. This bias was abolished after participants were shortly exposed to incompatible action-outcome associations (i.e., clockwise actions resulting in counterclockwise sphere motion). In contrast, exposure to compatible associations did not add anything to the already existing bias. These results suggest that both overlearned motor-predictions, as well as motor-predictions based on the current task-set, jointly affect perception of action-outcomes.

The results of Chapter 4 suggest that the context in which a task is performed can play an important role in how action-outcomes are perceived. The idea that

2 An ambiguous sphere is a type of bistable percept. In the experiment presented here, I used a sphere consisting of dots. Half of the dots moved in a leftward direction, and half of the dots moved in a rightward direction. This sphere can be perceived as rotating in a clockwise or counterclockwise direction depending on which dots (those moving in a rightward or a leftward direction) are perceived in the plane closest to the observer.

3 The presentation of rivalrous gratings entails the presentation of a different grating (often in terms of orientation) to each eye. Only one of these stimuli can dominate perception at the same time.

4 An example of a real-life action-outcome contingency that concurs with this bias is winding a watch, during which the clock hand can be seen to rotate in the same direction as the watch stem.

more general non-motor information, such as context, beliefs or intentions, are important for agency processing concurs with earlier research that I have conducted on explicit judgements of agency (Chapter 5). As reviewed above, the experience of agency is not only thought to result from the low-level motor processes that are described in the comparator model, but can also emerge via a different route. Specifically, it has been shown that participants can infer agency based on a correspondence between prior thoughts about outcomes and actual outcomes, even in the absence of actual action-outcome contingencies (Aarts et al., 2005). In contrast to motor-based experiences of agency, the neural processes underlying these inferences of agency are poorly understood. Moreover, most work examining the neural markers of agency has adopted a localization approach, while interactions between implicated brain areas are ignored (David, 2012). In Chapter 5, I address these issues by exploring EEG-based functional connectivity during agency inferences. Participants were exposed to action-outcomes that either matched or mismatched with pre-activated outcomes. Replicating previous work (Van der Weiden et al., 2013), I demonstrated that experienced agency was increased for action-outcomes that matched pre-activated information. In addition, evidence for (directed) connectivity between parietal and frontal areas was obtained during such match trials, which (roughly) concurs with connectivity patterns observed in previous work on motor-based agentic experiences (Nahab et al., 2011). Whilst preliminary, these results add to a more complete understanding of the neural processes underlying agency experiences, including those with a non-motor basis.

The work presented above demonstrates that evidence for the influence of motor-predictions on the perception of action-outcomes was difficult to obtain across a variety of measures. As a consequence, I have increasingly come to question the explanatory range of the comparator model over the past years. In Chapter 6 I reflect on this issue in more detail. Specifically, I evaluate to what extent the motor-predictions that are specified in the comparator model can be extrapolated to the processing of action-outcomes in the outside world - such as tones and visual stimuli that are generated by button presses. I put forward three lines of argumentation in support of the idea that the supposed role of motor signals in the prediction of such external action-effects is hitherto unjustified. First, although there is ample evidence for the role of prediction-based modulation of action-outcome perception, the presence of these prediction-effects in itself does not necessitate the involvement of motor-predictions. Second, prediction effects in humans are already observed after minimal learning of new action-outcome

associations, which does not match the slow formation of the neural pathways that are thought to underlie motor-prediction effects. Third, current empirical evidence (including the work presented in this dissertation) provides only weak evidence for a role of motor-prediction effects in the processing of external action-outcomes. Hence, whilst motor-predictions are likely to be implicated in the prediction of action-effects that are tightly coupled to preceding movements (such as during self-applied tickling), their influence seems questionable when predicting beyond ourselves.

READING GUIDE

The research outlined above is discussed in detail in the remaining chapters of this dissertation. All these chapters were written with the purpose of being published as an article in a scientific journal. As a consequence, some of the information in the chapters will overlap, and the order in which methods and results are discussed can vary depending on journal requirements. Note that all chapters can be read individually and in any desired order.

CHAPTER 2

Exploring the role of motor and non-motor predictive mechanisms in sensory attenuation: Perceptual and neurophysiological findings

This chapter is based on: Dogge, M., Hofman, D., Custers, R., & Aarts, H. (2019). Exploring the role of motor and non-motor predictive mechanisms in sensory attenuation: Perceptual and neurophysiological findings. *Neuropsychologia*, *124*, 216-225.
<https://doi.org/10.1016/j.neuropsychologia.2018.12.007>

ABSTRACT

Sounds that result from our own actions are perceptually and neurophysiologically attenuated compared to sounds with an external origin. This sensory attenuation phenomenon is commonly attributed to prediction processes implicated in motor control. However, accumulating evidence suggests that attenuation effects can also result from prediction processes beyond the motor domain. The aim of the present study was two-fold. First, we attempted to replicate the role of identity-specific motor predictions in sensory attenuation. Second, we set out to examine whether attenuation effects can be observed when tones cannot be predicted from preceding actions, but only from the non-motor cues accompanying them. Participants completed a two-alternative forced choice task on the loudness of tones whose pitch was congruent or incongruent with previously learned key-tone or cue-tone associations. No convincing evidence was observed for identity predictions on a perceptual level nor on a neurophysiological level. However, exploratory analyses revealed that attenuation was more pronounced for participants who first learned to rely on motor (instead of non-motor predictions). Together, these findings suggest that the role of motor identity predictions in sensory attenuation might have to be reconsidered.

INTRODUCTION

Self-produced sensations are perceived as less intense than externally produced sensations (Schafer & Marcus, 1973). This sensory attenuation effect is famously exemplified by our inability to tickle ourselves (Blakemore, Wolpert, & Frith, 2000; Weiskrantz, 1971) and is imperative for successful interaction with the environment. Reduced processing of action-effects not only frees up resources to deal with novel information, but is also thought to aid in the distinction of self and other produced effects (Frith et al., 2000b; Haggard & Tsakiris, 2009). Although sensory attenuation is commonly ascribed to predictive processes implicated in motor control (Blakemore et al., 2000; Blakemore, Frith, & Wolpert, 1999; Blakemore, Wolpert, & Frith, 1998; Frith et al., 2000b), accumulating evidence suggests that it can also be observed in the absence of any actions, when events are predictable from a different source (Hughes et al., 2013a; Schröger, Marzecová, & Sanmiguel, 2015). The present study examined how attenuation effects resulting from these more general (non-motor) predictive mechanisms compare to those resulting from action-based predictions.

While the differential processing of self-generated and externally generated effects is demonstrated across sensory modalities, the present study restricts itself to the auditory domain, which is most extensively studied (Hughes et al., 2013a). Sounds following one's own actions are systematically reduced in perceived loudness compared to sounds with an external origin (Sato, 2009; Weiss, Herwig, & Schütz-Bosbach, 2011a, 2011b; Weiss & Schütz-Bosbach, 2012). In addition, the N1 component of the auditory ERP, which is thought to reflect prediction error, has a smaller amplitude for self- versus externally induced sounds (Baess, Horváth, Jacobsen, & Schröger, 2011; Baess, Jacobsen, & Schröger, 2008; Schafer & Marcus, 1973; Timm, Sanmiguel, Saupe, & Schröger, 2013; Van Elk, Salomon, Kannape, & Blanke, 2014).

The aforementioned findings are generally alluded to as evidence for the role of motor prediction in sensory attenuation. Specifically, self-prepared movement is thought to be accompanied by a copy of the motor command (i.e., an efference copy) that can be used to predict action-effects and as such explain away, or attenuate, incoming sensory input. Seeing that only self-produced, but not externally-produced effects are accompanied by efference copies, only the former are attenuated (Frith et al., 2000b). Importantly, however, the nature of typically employed designs (i.e., comparing self-produced versus other produced tones)

obscures identification of the exact predictive mechanisms that underlie sensory attenuation. General interpretations of forward models presume that differences between the conditions result from *identity predictions*, that is, from predictions about the exact identity of an upcoming sound (Hughes et al., 2013a). However, other differences between the conditions, such as the predictability of the point in time in which the effect will occur (i.e., *temporal predictions*), might also account for the observations (see Hughes et al., 2013a for a review). Only a few studies have circumvented these issues and demonstrated the influence of isolated identity predictions by the direct comparison of self-produced tones that were either congruent or incongruent with predicted outcomes (Hughes, Desantis, & Waszak, 2013b; Kühn et al., 2011).

Sensory attenuation is not restricted to the motor domain and has also been observed as a function of non-motor predictions. For instance, N1 amplitudes are reduced for tones whose frequency can be predicted from a preceding pattern, compared to unpredicted (Lange, 2009) or mispredicted tones (Hsu, Le Bars, Hämäläinen, & Waszak, 2015). Similar attenuation effects were reported for tones whose identity could be inferred from simple contingencies, such as when the identity of the second tone of a pair is identical to the first tone (Hsu, Hämäläinen, & Waszak, 2014a). However, there are also some studies that failed to observe non-motor identity prediction effects. In one study, for example, the N1 amplitude did not differ between cued tones whose identity was fixed (i.e., predictable) versus randomly chosen (Hsu, Hämäläinen, & Waszak, 2013). Another study even observed an enhanced (instead of a reduced) N1 amplitude for attended predictable (vs. unpredictable) tones, whereas no prediction effect was observed when participants were not attending to the tones (Hsu, Hämäläinen, & Waszak, 2014b). Notwithstanding their diverging nature, the aforementioned findings suggest that sensory attenuation effects are broader than can be accounted for solely by forward models implicated in motor control.

While attenuation effects resulting from motor and non-motor predictions have been shown in isolation, only a few studies (across sensory modalities) have attempted to compare them in terms of quality and magnitude. Some of these studies have observed comparable attenuation effects as a result of both prediction sources (Desantis, Mamassian, Lisi, & Waszak, 2014), whereas other studies have failed to observe non-motor prediction effects (Bednark, Poonian, Palghat, McFadyen, & Cunnington, 2015; Cardoso-Leite, Mamassian, Schütz-Bosbach, &

Waszak, 2010; Richters & Eskew, 2009⁵). The findings of these studies are difficult to interpret as they employ different designs and varying indices of attenuation. For instance, motor and non-motor prediction effects have been compared both across samples (Cardoso-Leite et al., 2010; Richters & Eskew, 2009) and within the same sample (Bednark et al., 2015; Desantis, Mamassian, et al., 2014). Moreover, both neurophysiological (Bednark et al., 2015) and varying perceptual indices of attenuation have been employed, including measures of perceptual sensitivity (e.g., just noticeable difference, Desantis, Mamassian et al., 2014; and d' , Cardoso-Leite et al., 2010), as well as measures of perceptual intensity (e.g., point of subjective equality, Desantis, Mamassian et al., 2014). How these different indices of attenuation relate to each other is unclear, which complicates interpretations of the observed findings and the (dis)similarity between attenuation effects resulting from motor and non-motor prediction.

The aim of the present study was twofold. First, the current study served to determine the role of identity-specific motor predictions in attenuation. Second, we wished to examine whether non-motor predictions can take over and result in similar attenuation effects when motor predictions have no predictive value. To circumvent the aforementioned issues regarding design related differences across studies, we directly compared motor and non-motor prediction effects in the same sample of subjects, using frequently employed measures of neurophysiological attenuation (i.e., N1 amplitude) as well as perceptual intensity and perceptual sensitivity indices.

To examine the influence of predictive processes on attenuation, participants completed an adaptation of an existing auditory detection paradigm (Desantis, Mamassian, et al., 2014) consisting of a motor prediction and a non-motor prediction block. The order of these blocks was counterbalanced across participants. In the first phase of the task, participants were exposed to contingencies between freely chosen key-presses (motor prediction block) or geometrical stimuli accompanying these key presses (non-motor prediction block) and the pitch of a subsequent tone. In the subsequent test phase, we examined perceived loudness as a function of action-congruency (i.e., comparing tones that were congruent versus incongruent with previously learned relationships). Based on the studies outlined above, we at least expected lower perceived intensity, sensitivity and a decreased N1 amplitude for tones that were congruent (vs. incongruent) with previously

5 Bednark et al., 2015 did not observe any N1 attenuation related to identity-specific predictions, irrespective of prediction source.

learned action-effect associations. Whether or not a similar pattern would be observed for the non-motor prediction condition was more difficult to predict given the ambiguity of the existing literature regarding the effects of non-motor prediction cues on sensory attenuation.

METHODS

Participants

Twenty-four participants took part in the experiment ($M_{age} = 20.63$; $SD_{age} = 2.45$; 15 females; 19 right-handed)⁶. All participants had self-reported normal or corrected to normal vision and no hearing disabilities. In addition, none of the participants were smokers or recreational drug users and none reported current neurological conditions, mental illnesses or use of psychiatric medication. Participants were requested to refrain from caffeine consumption three hours prior to the experiment. All participants received written and oral information concerning the set-up of the experiment and signed an informed consent form. A monetary reimbursement was received in return for participation. The study received approval from the faculty's (Social and Behavioral Sciences) ethical board.

Procedure

Participants completed a modified version of an auditory detection paradigm as described in (Desantis, Mamassian, et al., 2014). The task consisted of a motor prediction block and a non-motor prediction block with a five-minute break in-between. The order of these blocks was counterbalanced between participants. Within each block participants completed ten acquisition phases (A), each consisting of 80 trials, and ten test phases (T), each consisting of 36 trials. These phases were presented in an interleaved (ATAT) order to reduce the likelihood of extinction effects. The acquisition phase served to learn associations between actions and tones (motor prediction block) or between visual cues and tones (non-motor prediction block). In the test phase, the effect of the learned associations on loudness perception was assessed. In order to diminish attention

6 Three participants were excluded prior to data analysis due to a technical error and one participant for not adhering to task instructions. Four new participants were recruited and assigned to the respective cells of the design. Demographics regard the final sample.

lapses, participants played Tetris (Pretris; Pfister, 2008⁷) for three minutes after completing half of the motor block and after half of the non-motor prediction block. Prior to the start of the motor and non-motor prediction block participants completed practice trials for both the acquisition phase (8 trials) and the test phase (4 trials). After the experiment, participants answered some general exit questions, including questions about handedness and demographic characteristics.

Task

Acquisition phase

Participants were instructed to produce a freely chosen right or left key press in response to a white fixation cross. Key presses were produced by pressing the left or right button on a Cedrus RB530 response pad (Cedrus Corporation, San Pedro, CA) with the corresponding index finger⁸. Participants were asked to aim for an equal response distribution. Every twenty trials feedback regarding the ratio of key presses was presented to assist participants in this attempt. In the motor prediction condition, each key press (motor cue) generated a tone after a 200-millisecond interval. For half of the participants, a left key press was associated with a low tone and a right key press was associated with a high tone. The opposite key-tone mapping applied to the other half of the participants. Participants were made explicitly aware of these associations prior to the start of each phase. In the non-motor prediction condition, key presses were immediately followed by either a white square or a white circle (non-motor cue) that was presented for 100 milliseconds. After a 100-millisecond interval a low or a high tone was presented. Importantly, the geometrical stimulus, and not the key press, predicted the tone pitch in this condition. There was no association between geometrical stimuli and key presses on a phase level. Similar to the motor prediction condition, cue-tone mappings were counterbalanced between participants. All tones were 100 millisecond lasting sine waves, including 10 millisecond onset and offset envelopes. The tones were presented binaurally at approximately 74 dB through foam in-earplugs (Earlink 3A Oty 50, Aearo Company Auditory Systems, Indianapolis, IN, USA). Two pairs comprising a low and a high tone frequency were counterbalanced between blocks

7 Traditional Tetris sounds were added to this version of Pretris.

8 One participant indicated to have used thumbs instead of index fingers to press the keys during both the acquisition and the test phase of the motor prediction block (first block for this participant). We decided against excluding this participant considering that the mapping of importance was hand-specific (e.g., left – low; right – high) and not necessarily finger-specific.

and participants to reduce the likelihood of spill-over of learning effects from the motor to the non-motor prediction block (or vice versa). Accordingly, for half of the participants a low tone of 750 Hz and a high tone of 900 Hz were presented in the motor prediction block and a low tone of 700 Hz and a high tone of 850 Hz were presented in the non-motor prediction block. The opposite ascription applied to the other half of the participants. All trials were separated by a 1000 millisecond inter-trial interval.

Each acquisition phase consisted of 80 trials (similar to Desantis, Mamassian, et al., 2014), including 20 percent of catch trials. The catch trials were identical to the main acquisition trials, with the exception that participants had to indicate the frequency of the presented tone (low or high) by pressing one of two foot pedals. In the motor prediction condition the labels of the foot pedals always matched learned associations. That is, if participants learned that a left key press was associated with a low tone, the label of the left pedal also corresponded to a low tone. In the non-motor prediction condition foot pedals labels were counterbalanced in a similar manner, such that for half of the participants the left pedal represented a low tone and the right pedal a high tone, whereas the opposite labeling applied to the other half of the participants.

Test phase

In the test phase participants were again instructed to generate freely chosen key presses as soon as a white fixation cross was presented. Similar to the acquisition phase the key press (motor) or the visual cue (non-motor) was followed by an approximately 74 dB tone. Importantly, however, the tones were now presented randomly such that the frequency of the tones was either congruent, or incongruent with learned action-effect or cue-effect associations. In addition, this (standard) tone was now followed by a second (sample) tone of equal frequency but varying loudness (~70-78 dB, with 1 dB intervals) after an interval of 1100 milliseconds. Participants completed a two-alternative forced choice task, in which they indicated whether the first or second tone was louder by using the foot pedals. The left foot pedal always indicated that the first tone was loudest, whereas the right foot pedal always indicated that the second tone was loudest. The ascription of frequency pairs to the motor and non-motor prediction block was identical to the acquisition phase. Figure 2.1 depicts the timeline of acquisition and test trials.

To ensure an approximately equal distribution of congruent and incongruent trials across congruency and sample tone levels, a list was pre-programmed

EEG recording

EEG was recorded with the BioSemi Active Two EEG system (BioSemi, Amsterdam) from 64 electrodes (sampling rate: 2048 Hz) that were positioned according to the international 10/20 system. An online Common Mode Sense-Driven Right Leg (CMS-DRL) was used as a reference. Electro-oculogram (EOG) was measured from electrodes placed on the suborbit and supraorbit of the right eye and on the outer canthi of both eyes.

Data pre-processing

Behavioral data: Test trials

Trials with erroneous responses (i.e., multiple key presses, multiple pedal presses, pressing a pedal when a key was supposed to be pressed or vice versa) were excluded from all analyses ($M = 4.81\%$, $SD = 2.95\%$). In addition, data inspection indicated that participants were occasionally very slow to respond to the fixation cross at the start of the trial, as well as to judge which of the two tones was louder. These delayed responses are problematic as, in the first case, participants might not have attended properly to the stimuli, whereas, in the latter case, information as to which tone was louder might no longer be accessible. For these reasons, we decided to reject trials when the onset time of key presses ($M = 1.85\%$, $SD = 0.51\%$) and/or pedal responses ($M = 2.00\%$, $SD = 0.54\%$) were more than 3 standard deviations above the mean of that participant (after first excluding trials with multiple responses). The mean number of trials per condition that was contained for the final analysis is presented in Table 2.1.

Table 2.1. Mean number of trials used for final analysis as a function of prediction type, congruency and sample tone magnitude

Prediction type	Congruency	Sample tone (dB)								
		70	71	72	73	74	75	76	77	78
Motor	Congruent	18.38 (1.66)	17.79 (1.22)	18.75 (1.78)	18.33 (1.83)	17.75 (1.70)	17.88 (1.26)	18.58 (1.41)	18.75 (1.42)	18.00 (1.89)
	Incongruent	17.58 (1.89)	18.42 (1.28)	18.33 (1.63)	17.83 (1.99)	18.13 (1.70)	18.21 (1.82)	18.50 (2.09)	18.92 (1.67)	18.38 (1.21)
Non-motor	Congruent	18.25 (1.26)	18.33 (1.76)	18.33 (1.40)	17.71 (1.90)	17.58 (1.38)	18.75 (1.48)	18.75 (1.54)	18.67 (1.69)	18.79 (1.25)
	Incongruent	18.92 (1.14)	18.38 (1.44)	18.08 (1.89)	17.88 (1.75)	17.83 (1.13)	18.25 (1.51)	18.58 (1.21)	18.96 (1.12)	18.92 (1.18)

Note. Numbers between parentheses represent standard deviations.

Behavioral data: Catch trials

Similar to the test trials, trials with multiple responses ($M = 4.39\%$, $SD = 2.60\%$) or extreme reaction times for hand ($M = 1.64\%$, $SD = 0.48\%$) and/or pedal responses ($M = 1.88\%$, $SD = 0.57\%$) were excluded prior to further analysis. The average percentage of the remaining trials was 92.37 % ($SD = 3.79\%$) for the motor prediction condition, and 92.32% ($SD = 3.10\%$) for the non-motor prediction condition.

EEG data

Offline, the data was downsampled to 256 Hz, bandpass filtered (1-25 Hz), re-referenced to the average reference and segmented into epochs from -200 milliseconds to 500 milliseconds relative to the onset of the first tone. Ocular artifacts were rejected using the EOG signal (Gratton, Coles, & Donchin, 1983). Trials with extreme and incorrect responses were excluded based on the criteria described in above. Remaining artifacts were excluded by automatically rejecting segments with signals exceeding ± 75 microvolts on channels of interest (see Data analysis section). Baseline correction was applied using an interval of 100 milliseconds prior to the first tone. An average percentage of approximately ninety percent of the trials was contained for final analyses in all the conditions (Motor prediction, congruent: $M = 89.55\%$, $SD = 4.61\%$; Motor prediction, incongruent: $M = 89.40\%$, $SD = 4.31\%$; Non-motor prediction, congruent: $M = 89.40\%$, $SD = 4.58\%$, Non-motor prediction, incongruent: $M = 90.25\%$, $SD = 3.61\%$)¹⁰.

Data analysis

Perceptual attenuation

The percentage of “second-tone-is-louder” responses was calculated separately for each combination of prediction type, congruency and magnitude of the second (sample) tone for each individual participant. These percentages were fitted with a psychometric function (cumulative Gaussian) to calculate indices of perceptual intensity (point of subjective equality) and perceptual sensitivity (just noticeable difference); see Figure 2.2. The point of subjective equality (PSE) represents the sample tone magnitude at which the sample tone is perceived as louder than the standard tone on fifty percent of the trials. Accordingly, a lower PSE value corresponds to

¹⁰ For one participant 21 trials were missing in the motor condition due to a technical malfunction. The reported percentages were calculated based on the remaining trials.

more attenuation of the standard tone (i.e., the first tone following the predictive cue). In addition to the PSE, previous studies on attenuation have often included the just noticeable difference (JND), which is half of the difference of the sample tone magnitude at which the sample tone is judged as louder than the standard tone on 75% of the trials and on 25% of the trials. This index is thought to represent perceptual sensitivity and reflects the variability of responses given by the participant.

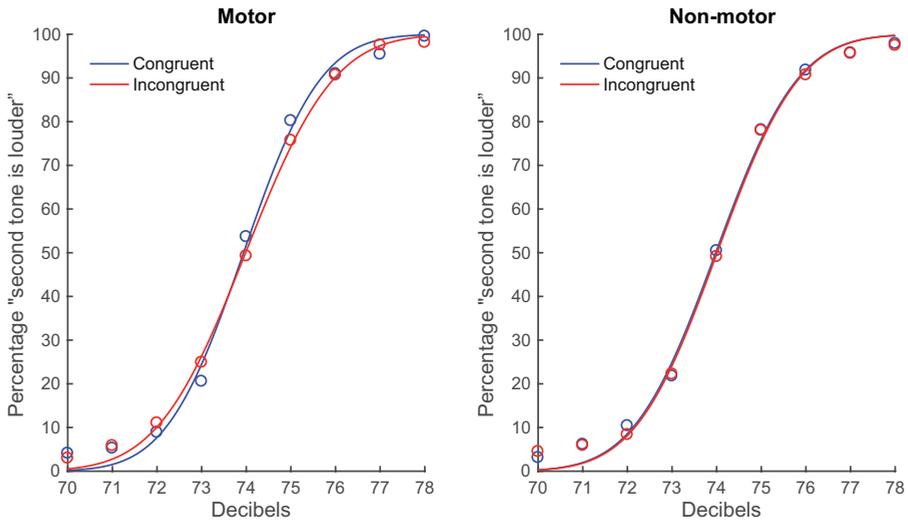


Figure 2.2. Average percentage of “second-tone-is-louder” responses for congruent and incongruent trials, as a function of sample tone magnitude and prediction type across all participants. The presented percentages were calculated excluding erroneous trials and outliers (see Data pre-processing section).

Neurophysiological attenuation

Given that the N1 is known to consist of several separate peaks, the ERP analysis focused on three peaks (N1a, N1b and N1c; Näätänen & Picton, 1987; Woods, 1995) that have previously examined in the context of motor prediction (c.f. Sanmiguel, Todd, & Schröger, 2013; Timm et al., 2013). The N1b peak maximizes over frontocentral electrodes, whereas the N1a and N1c peaks are maximal over bilateral mid-temporal electrodes (Woods, 1995). Given these differences in topography and latency, congruency and prediction type effects were assessed separately for each peak. Specifically, amplitudes were averaged across frontocentral electrodes (Cz, FCz and Fz) for the N1b peak, and across left (C5, FC5, FT7 and T7) and right (C6, FC6 FT8 and T8) mid-temporal electrodes for the N1a and N1c peaks. The time windows of interest were determined based on the observed grand

averages. The N1a and N1c peak were defined as the first (60 – 100 ms) and second (120 – 170 ms) negative peak on the temporal electrodes, respectively. The N1b was analyzed in a window stretching from 80 to 130 after tone onset. All peaks were quantified as the most negative amplitude for individual averages within the previously specified windows. Separate repeated measures ANOVA's were conducted for each component with prediction type (motor versus non-motor) and congruency (congruent versus incongruent), as independent variables.

RESULTS

Catch trial accuracy

Catch trial accuracy was high for all four tones: 700 Hz: $M_{acc} = .98$, $SD_{acc} = .04$; 750 Hz: $M_{acc} = .97$, $SD_{acc} = .03$; 850 Hz: $M_{acc} = .97$, $SD_{acc} = .06$; 900 Hz: $M_{acc} = .97$, $SD_{acc} = .03$. These results indicate that participants paid attention to the tones and were well able to identify them.

Perceptual attenuation

Point of subjective equality

To examine the hypothesized effects of prediction on perceptual intensity, PSE values were subjected to a repeated measures ANOVA with prediction type (motor versus non-motor) and congruency (congruent versus incongruent) as independent variables. This analyses yielded no main effects of prediction type, $F(1,23) = .13$, $p = .721$, $\eta_p^2 = .01$ and congruency, $F(1,23) = .16$, $p = .692$, $\eta_p^2 = .01$, nor an interaction between these two factors $F(1,23) = .03$, $p = .866$, $\eta_p^2 < .01$.

In order to examine to what extent these non-significant findings reflect evidence for the null hypothesis we calculated Bayes factors for the reported effects, using the R package Bain (Gu, Hooijink, Mulder, & Rosseel, 2018; <https://informative-hypotheses.sites.uu.nl/software/bain/>). This package differs from other, more generally known, software (e.g., JASP), in that it specifies the prior width based on a fraction of the data. In addition, instead of running omnibus tests, Bain allows one to specify the exact contrast of interest. Note that for the current study this results in separate evaluations of evidence for the main and interaction effects. The analyses for the main effects revealed that the observed data is about four times

as likely under the null hypothesis compared to the alternative hypothesis¹¹ for both prediction type ($BF_{01} = 4.59$) and congruency ($BF_{01} = 4.52$). In addition, the evidence for the absence of an interaction effect is almost five times as likely as the evidence for the presence of an interaction effect ($BF_{01} = 4.83$).

An additional, exploratory analysis was conducted to examine whether differences in preceding predictive context (e.g., the order in which the prediction type blocks were shown) might have affected the results. A three-way mixed ANOVA was executed, with order of the prediction type blocks (motor first versus non-motor first) as an additional between-subject variable. This analysis yielded a significant interaction between congruency and order, $F(1,22) = 5.66$, $p = .027$, $\eta_p^2 = .20$. As depicted in Figure 2.3, PSE values were descriptively lower for the congruent compared to the incongruent condition if participants started with the motor prediction block, $F(1,22) = 3.97$, $p = .059$, $\eta_p^2 = .15$, whereas an opposite, albeit weaker, pattern was observed for participants who started with the non-motor prediction block, $F(1,22) = 1.88$, $p = .184$, $\eta_p^2 = .08$ ¹². Notably, there was no three-way interaction between prediction type, order and congruency, $F(1,22) =$

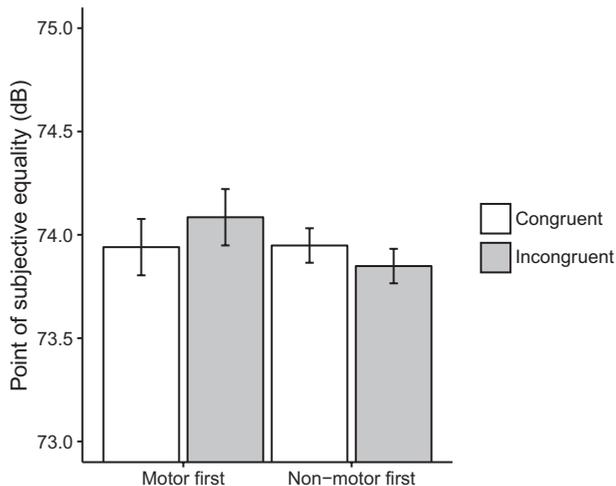


Figure 2.3. Point of subjective equality as a function of congruency and block order (collapsed over prediction type). Error bars reflect within-subject 95% confidence intervals calculated according to Morey's (2008) method.

- 11 For all the reported Bayesian analyses the null hypothesis refers to the absence of a main or interaction effect, whereas the alternative hypothesis refers to the presence of a main or interaction effect.
- 12 Considering the sensitivity of small samples for outliers, we further evaluated the simple main effects of congruency using separate Wilcoxon signed-ranked tests for each level of order. The same pattern of results was observed. That is, the effect of congruency was marginally significant for participants who started with the motor prediction block ($V = 20$, $p = .08$, one-tailed), whilst the effect for participants who started with the non-motor prediction block was not significant ($V = 58$, $p = .94$, one-tailed).

.31, $p = .583$, $\eta_p^2 = .01$, indicating that the observed order effect was similar in the motor and in the non-motor prediction condition.

Just noticeable difference

A separate 2 (prediction type: motor versus non-motor) \times 2 (congruency: congruent versus incongruent) repeated measures ANOVA was conducted on the JND values. This analysis yielded no main effects of prediction type, $F(1,23) = .20$, $p = .663$, $\eta_p^2 = .01$ and congruency, $F(1,23) = 2.76$, $p = .110$, $\eta_p^2 = .11$, nor an interaction between prediction type and congruency $F(1,23) = .30$, $p = .591$, $\eta_p^2 = .01$.

We further examined these non-significant effects by calculating Bayes factors. Bayesian analyses revealed that the observed data is about four times more likely in the absence of a main effect of prediction type, than in the presence of such an effect ($BF_{01} = 4.44$). The same is true for the interaction effect between prediction type and congruency ($BF_{01} = 4.22$). However, the evidence for an effect of congruency is inconclusive (i.e., neither evidence for the null or for the alternative hypothesis is obtained; $BF_{01} = 1.24$).

Contrary to the PSE analysis, the exploratory addition of order to the design did not yield an interaction between congruency and order, $F(1,22) < .01$, $p = .973$, $\eta_p^2 < .01$. The three-way interaction between prediction type, congruency and order, also did not reach significance, $F(1,22) = .03$, $p = .866$, $\eta_p^2 < .01$.

Neurophysiological attenuation

Grand average ERP's and voltage maps for the N1a, N1b and N1c components are shown separately for the motor- and non-motor prediction condition in Figure 2.4.

N1b

The analysis for the N1b time window revealed no significant main effect for congruency, $F(1,23) = 1.19$, $p = .287$, $\eta_p^2 = .05$, no significant main effect for prediction type, $F(1,23) = .29$, $p = .595$, $\eta_p^2 = .01$, and no significant interaction between congruency and prediction type, $F(1,23) = .02$, $p = .889$, $\eta_p^2 < .01$.

Similar to the analysis of the behavioral data, we conducted Bayesian analyses to evaluate to what extent the non-significant findings reflect evidence for the null hypothesis. The observed data is about four times as likely under the null hypothesis compared to the alternative hypothesis for both the main effect of prediction type, as well as for the interaction between prediction type and congruency (see Table 2.2). In addition, the data is almost three times as likely under the hypothesis that

there is no difference between congruent and incongruent conditions, than under the hypothesis that there is a difference between these conditions. Potential effects of predictive context were explored by adding order (motor-prediction condition

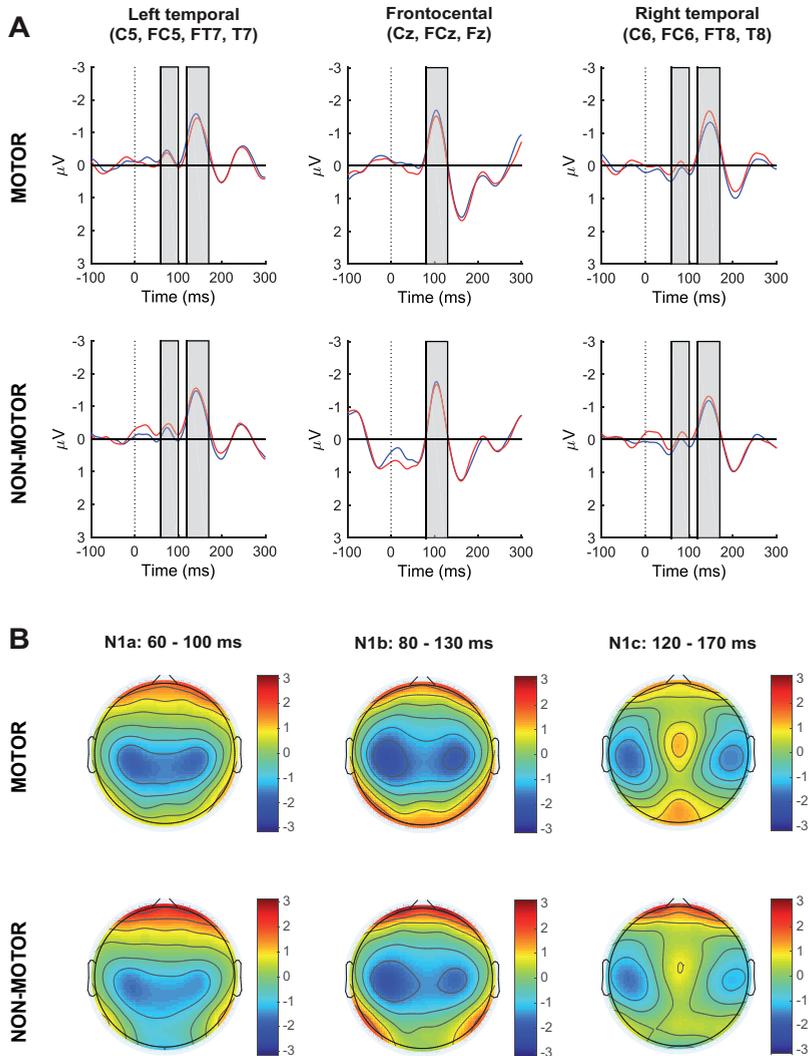


Figure 2.4. (A) Grand average ERPs over congruent (blue lines) and incongruent (red lines) trials in the motor- and the non-motor prediction condition. Grey areas represent time windows subjected to statistical analysis. (B) Topographic maps of the motor- and non-motor prediction conditions (collapsed over congruency).

Table 2.2. Evidence for the absence of effects (reflected by BF_{01}) for all N1 peaks

	N1b (80 – 130 ms)	N1a (60- 100 ms)	N1c (120-170 ms)
Prediction type	4.24	4.62	4.23
Congruency	2.71	1.29	2.36
Prediction type x Congruency	4.85	3.25	3.80

Table 2.3. Results of mixed-ANOVA's including order for all N1 peaks

	N1b (80 -130 ms)			N1a (60 – 100 ms)			N1c (120 – 170 ms)		
	<i>F</i>	<i>p</i>	η_p^2	<i>F</i>	<i>P</i>	η_p^2	<i>F</i>	<i>p</i>	η_p^2
Prediction type	0.29	.593	0.01	0.11	.740	<.01	0.29	.595	0.01
Congruency	1.28	.270	0.06	2.76	.111	0.11	1.50	.234	0.06
Prediction type x Congruency	0.02	.888	<0.01	0.95	.342	0.04	0.55	.468	0.02
Order	1.27	.271	0.06	1.81	.192	0.08	0.07	.802	<0.01
Order x Prediction type	1.36	.256	0.06	0.10	.750	0.01	0.64	.433	0.03
Order x Congruency	2.78	.109	0.11	1.83	.190	0.08	1.59	.220	0.07
Order x Congruency x Prediction type	1.21	.283	0.05	4.46	.046*	0.17	2.60	.121	0.11

first versus non-motor prediction condition first) to the design as a between subject factor. However, no significant interactions with this factor were observed (see Table 2.3).

N1a

Similar to the N1b time window, no significant main effect for prediction type, $F(1,23) = .12$, $p = .735$, $\eta_p^2 = .01$, or congruency, $F(1,23) = 2.66$, $p = .116$, $\eta_p^2 = .10$, nor a significant interaction between prediction type and congruency, $F(1,23) = .82$, $p = .374$, $\eta_p^2 = .03$, was observed for the N1a window. The main effect of laterality was significant, $F(1,23) = 4.46$, $p = .046$, $\eta_p^2 = .16$, showing a higher N1a amplitude for left electrodes ($M = -0.73$, $SE = 0.14$) compared to right electrodes ($M = -0.38$, $SE = 0.14$). Given that this main effect was not qualified by interactions with congruency, $F(1,23) = 3.05$, $p = .094$, $\eta_p^2 = .12$, prediction type, $F(1,23) = .26$, $p = .615$, $\eta_p^2 = .01$, or congruency and prediction type, $F(1,23) = .10$, $p = .751$, $\eta_p^2 < .01$, we decided to collapse the data across laterality levels prior to calculating Bayes factors (see Table 2.2) and exploring effects of order (see Table 2.3).

Bayesian analyses revealed that the data is about four times more likely under the hypothesis that there is no difference between the motor and the non-motor prediction condition, than that under the hypothesis there is a difference between these conditions. In contrast, the evidence for the main effect of congruency is

inconclusive. Finally, the data points towards the absence of an interaction effect (see Table 2.2).

The exploratory analyses including order yielded a significant three-way interaction between prediction type, congruency and order, $F(1,22) = 4.46$, $p = .046$, $\eta_p^2 = .17$. In order to further explore this interaction, we examined the simple interaction effect between prediction type and congruency at each level of order. As can be seen in Figure 2.5, the interaction between prediction type and congruency was more pronounced for participants who started the experiment with the motor prediction condition, $F(1,22) = 4.75$, $p = .040$, $\eta_p^2 = .18$, than for participants who started with the non-motor prediction condition, $F(1,22) = .65$, $p = .429$, $\eta_p^2 = .03$. For participants who started with the motor condition, an expected reduction in N1 amplitude for congruent versus incongruent trials was observed for the non-motor prediction condition, $F(1,22) = 7.42$, $p = .012$, $\eta_p^2 = .25$, but not for the motor prediction condition, $F(1,22) < .01$, $p = .965$, $\eta_p^2 < .01$ ¹³. Note that these results roughly mimic the observed pattern of the PSE values, in the sense that congruency effects were restricted to participants who started with the motor prediction condition. However, unlike perceptual attenuation effects, the order effects seem to be driven by the non-motor prediction condition on a neurophysiological level.

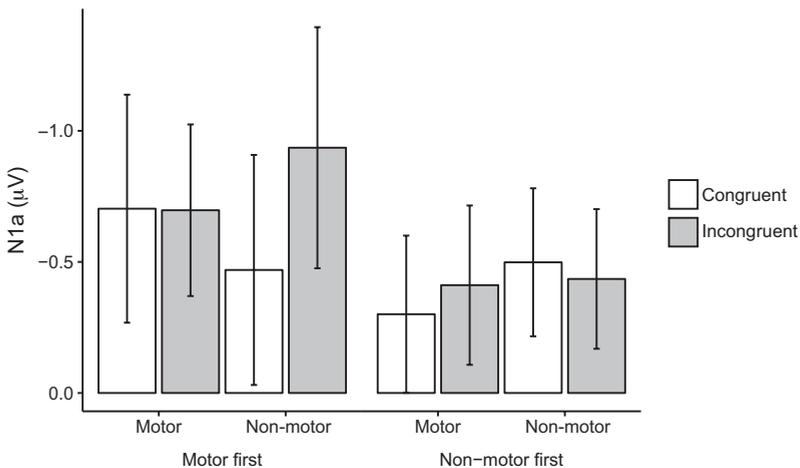


Figure 2.5. N1a amplitude as a function of prediction type, congruency and block order. Error bars reflect within-subject 95% confidence intervals calculated according to Morey's (2008) method.

13 Non-parametric Wilcoxon signed-ranked tests provided evidence for a similar pattern of second order simple main effects. Specifically, for participants who started with the motor condition, a significant effect of congruency was observed in the non-motor prediction condition ($V = 65$, $p = .02$, one-tailed), whereas there was no significant difference between congruent and incongruent trials in the motor prediction condition ($V = 40$, $p = .48$, one-tailed).

N1c

We observed no significant main effects for congruency, $F(1,23) = 1.46$, $p = .239$, $\eta_p^2 = .06$ and prediction type, $F(1,23) = .30$, $p = .592$, $\eta_p^2 = .01$, nor a significant interaction between congruency and prediction type, $F(1,23) = .51$, $p = .483$, $\eta_p^2 = .02$. In addition, we observed no main effect for laterality, $F(1,23) = .59$, $p = .449$, $\eta_p^2 = .03$, no interaction effect between prediction type and laterality, $F(1,23) = 1.43$, $p = .245$, $\eta_p^2 = .06$, and no interaction effect between congruency and laterality, $F(1,23) = 3.22$, $p = .086$, $\eta_p^2 = .12$. The interaction between prediction type, congruency and laterality also did not reach significance, $F(1,23) = 4.02$, $p = .057$, $\eta_p^2 = .15$. Given the absence of laterality effects, the data was collapsed prior to Bayesian analyses and exploratory analyses including order.

Bayesian analyses yielded substantial evidence for the absence of a main effect of prediction type and for an interaction effect between prediction type and congruency. Although the evidence for the absence of a main effect of congruency is not substantial, it is still stronger than evidence for the presence of such an effect. None of the interaction effects including order reach significance (see Table 2.3).

DISCUSSION

In an attempt to reconcile divergent findings in the literature regarding the similarity of attenuation effects resulting from motor and non-motor prediction, the present study directly compared these prediction effects in the same sample of subjects while employing frequently used indices of perceptual and neurophysiological attenuation. Overall, none of the indices of attenuation displayed the expected difference between tones that were prediction congruent (versus incongruent). Surprisingly, however, exploratory analyses revealed that the expected congruency effect could be observed on some of the measures (PSE and N1a component) for a subset of participants who started the experiment by learning action-effect (instead of cue-effect) relationships. Although these findings provide some support for the role of both motor- and non-motor identity prediction in attenuation, prediction-related effects were relatively weak and seemed to be conditional upon the person's learning history. Below we discuss our results in the context of the existing literature.

Perceptual attenuation

In contrast to our expectations, perceptual attenuation effects were only reflected in the point of subjective equality and were restricted to participants who completed the task in a specific block-order (i.e., starting with the motor prediction block). No statistical differences between prediction congruent and incongruent tones were observed for the just noticeable difference (JND). Importantly, however, Bayesian analyses revealed that the present data neither yields evidence for the alternative hypothesis, nor for the null hypothesis. According, the data do not allow for any further conclusions regarding the congruency effect on just noticeable difference values.

The finding that perceptual attenuation could only be observed for participants who completed the task in a specific block-order was not part of our initial expectations. Nevertheless, we wish to speculate about two subtle differences between the motor and non-motor prediction condition that might have contributed to this effect. Firstly, the mere presence of motor predictive cues was task relevant in the motor prediction block, as participants consistently had to choose which action to perform to produce a tone. In contrast, the non-motor predictive cues had no clear task relevance in the non-motor prediction block, where the identity of upcoming tones was determined by incidentally presented visual cues that occurred alongside non-predictive actions. As a consequence, attention towards actions and accompanying cues, as well as their relationship with subsequent effects, might have been reduced in non-motor prediction blocks compared to motor prediction blocks (see also: Hughes et al., 2013a). Since biases in information processing can be transferred from one task to another (Wylie & Allport, 2000), this difference in attentional tuning might not only have affected the acquisition of predictive relationships within each block, but might also have spilled over to the subsequent prediction block, explaining block order-dependent differences in attenuation.

Although speculative, this post-hoc explanation concurs with research on learned irrelevance, in which the associability of a cue has been shown to be impaired when it lacked predictive value in a preceding task (Kruschke & Blair, 2000; Le Pelley & McLaren, 2003). This phenomenon is commonly attributed to the idea that participants learn that a cue is irrelevant and consequently spend less attention to it, which in turn interferes with the formation of associations with new outcomes (Mackintosh, 1975). In a similar vein, it has been demonstrated that attention to predictive relationships in one task can affect learning of such

relationships in another task (Custers & Aarts, 2011). In particular, participants were more likely to learn unidirectional than bidirectional relationships when their attention was tuned to predictive (unidirectional) relationships in a prior unrelated task. Similarly, participants who started with the motor prediction block in the present study might have paid attention to predictive relationships, allowing them to pick up on the cue-tone associations in the non-motor prediction block. In contrast, participants who started with the non-motor prediction block presumably lacked this attentional tuning, which interfered with further predictive learning in the motor prediction block.

In addition to differences in attentional tuning, a second explanation for the observed block-order effect pertains to an asymmetry in the number of potential predictive sources that was included in the motor and non-motor prediction block. Whereas the motor prediction block only contained actions, the non-motor prediction block included both actions as well as visual cues. As a result of this asymmetry, the transition from the first to the second block differed in terms of prediction errors, depending on what was learned first. Specifically, for participants who started with the motor prediction block, the transition to the non-motor prediction block likely resulted in pitch misprediction, as actions that previously fully predicted the frequency of one specific (low or high) tone were now followed by both (low and high) tone frequencies. In contrast, when transitioning from a non-motor prediction block to a motor prediction block, previously predictive non-motor cues were absent, rendering prediction errors unlikely. Given the known role of prediction errors in new learning (e.g., Pearce & Hall, 1980), it is likely that participants who started with the motor prediction block were more attuned to seeking new predictors of the tones in the second block (in this case the non-motor cues) as opposed to participants who started with the non-motor prediction block. This active learning advantage might in turn have resulted in stronger predictions and attenuation effects for the former participants.

Altogether, the learning history of participants, and resulting effects on attention allocation, might thus be moderating conditions for the contribution of identity prediction to attenuation. The potential importance of attention to predictive relationships is supported by recent research suggesting that attenuation results from a shift of attention to surprising (incongruent) events, instead of cancellation of expected congruent events (Yon & Press, 2017). Future work incorporating indices of learning and attention would be valuable to further

examine to what extent these factors might explain the mixed results regarding prediction effects in the literature.

Notably, the observed order effect was present in both the motor and the non-motor prediction condition. In other words, no difference between action-based and cue-based expectations was observed. This concurs with previous research that has observed comparable effects for both prediction sources (Desantis, Mamassian, et al., 2014), and suggests that attenuation effects might reflect a more general predictive mechanism rather than self-specific motor predictions. In line with this idea, a recent study demonstrated that the common neurophysiological attenuation of self- versus externally produced tones disappeared when the onset of the tone was made predictable by a preceding visual countdown (Kaiser & Schütz-Bosbach, 2018, although see: Weiss & Schütz-Bosbach, 2012). Combined with the current findings, these observations call into question the frequently proposed contribution of sensory attenuation to the sense of agency in general and self-other distinction in particular (e.g., Haggard & Tsakiris, 2009).

Neurophysiological attenuation

The perceptual results were roughly mimicked on a neurophysiological level. Firstly, no general statistically significant congruency effects were observed on any of the three N1 components. However, although the Bayesian analyses generally indicated that the data provided evidence for the null hypothesis, this evidence was not substantial for the N1b and N1c peak, and inconclusive for the N1a peak. These findings therefore have to be interpreted with caution. Secondly, a similar order effect as observed for the PSE was observed on the N1a peak. That is, participants who started with the motor prediction block showed more attenuation for congruent versus incongruent tones than participants who started with the non-motor prediction block. Unlike perceptual attenuation, however, this effect was restricted to cue-based predictions.

The presence of the order effect on a neurophysiological level supports the previously proposed importance of learning history for sensory attenuation. The attenuation of the N1a peak itself is however hard to interpret as it does not coincide with previous work that has examined subcomponents of the N1 wave (Sanmiguel et al., 2013; Timm et al., 2013). In contrast to the present work, these previous studies only observed attenuation of the N1b and N1c peaks. Considering the exploratory nature of the current analyses, the above findings are therefore difficult to interpret.

The fact that we did not observe a generally reduced N1 amplitude for prediction-congruent versus prediction-incongruent tones was particularly surprising in case of the motor prediction condition. At first glance, this finding might seem to clash with previous demonstrations of neurophysiological attenuation following self- versus other-produced movements (Baess et al., 2011; Martikainen, 2005; Schafer & Marcus, 1973; Timm et al., 2013; Van Elk et al., 2014). It is important to note, however, that these studies examined the general influence of operational actions (i.e., actions that produce sounds) on sensory attenuation, whereas the present study scrutinized attenuation effects resulting from action-based predictions on the precise identity of a sound. The few existing studies that have hitherto examined such motor identity prediction effects have yielded mixed evidence (e.g., Baess et al., 2008; Bednark et al., 2015; Hughes et al., 2013b; Kühn et al., 2011)

The absence of clear identity-prediction effects in the current study also coincides with observations on other implicit measures of action-outcome perception, such as intentional binding. Similar to sensory attenuation, intentional binding (the temporal attraction of between self-produced actions and outcomes; Haggard, Clark, & Kalogeras, 2002) is generally attributed to motor predictive mechanisms (c.f., Moore & Haggard, 2008). However, studies directly comparing effects that are congruent or incongruent with action-based predictions have not found any evidence for this notion (Bednark et al., 2015; Desantis et al., 2012; Haering & Kiesel, 2014). This suggests that observed differences in perceptual processing of self-produced and externally produced effects are unlikely to be driven by specific motor-predictions as specified by forward models, but rather result from the more general heightened predictability of self-produced effects (Hughes et al., 2013a; Kaiser & Schütz-Bosbach, 2018). Taken together then, our results add to the existing evidence that both motor prediction and non-motor prediction effects on neurophysiological attenuation are not unequivocal and are possibly sensitive to variations in predictability manipulations across studies (for a similar argument, see: Bednark et al., 2015)

CONCLUSIONS

Sensory attenuation is a fascinating phenomenon that is generally thought to play a pivotal role in our ability to distinguish the events that we cause ourselves from those that are caused by external sources. This proposition builds on the assumption that attenuation is particularly driven by specific *action-dependent* predictions about upcoming sensory input. The present study set out to examine this assumption by directly comparing attenuation of tones whose identity could be predicted from motor or non-motor sources. Overall only weak evidence for identity-prediction effects was observed. In fact, the only (small) attenuation effects that we observed were conditional upon participant's learning history within the task (i.e., effects were only present for participants who started with the motor prediction condition). At most, the current data therefore suggest that motor-prediction processes might facilitate causal learning. Importantly however, the underlying process of attenuation does not appear to be action-dependent. Specifically, the small observed perceptual attenuation effects were comparable for action-based and cue-based predictions, suggesting that attenuation is more likely to result from a more generic predictive mechanism.

CHAPTER 3

The influence of action-effect anticipation on bistable perception: Differences between onset rivalry and ambiguous motion

This chapter is based on: Dogge, M., Gayet, S., Custers, R., & Aarts, H. (2018). The influence of action-effect anticipation on bistable perception: Differences between onset rivalry and ambiguous motion. *Neuroscience of Consciousness*, 2018, 1, niy004. <https://doi.org/10.1093/nc/niy004>

ABSTRACT

Perception is strongly shaped by the actions we perform. According to the theory of event coding, and forward models of motor control, goal-directed action preparation activates representations of desired effects. These expectations about the precise stimulus identity of one's action-outcomes (i.e., identity predictions) are thought to selectively influence perceptual processing of action-contingent effects. However, the existing evidence for such identity-prediction effects is scarce and mixed. Here, we developed a new paradigm to capture such effects, and examined whether action-outcome predictions can bias the perception of binocular onset rivalry (Study 3.1a and 3.1b) and bistable motion (Study 3.2). Participants performed learning tasks in which they were exposed to action-outcome associations. On test trials, actions were followed by bistable stimuli that could be perceived as being either congruent or incongruent with the aforementioned associations (i.e., rivalrous oriented gratings in Study 3.1a and 1b and spheres with ambiguous rotation directions in Study 3.2). Across three studies, we show that, whilst exposure to action-effect associations can bias the apparent motion direction of ambiguous spheres, it fails to influence perceptual selection of grating orientations in binocular onset rivalry. This pattern of results extends previous work on ambiguous motion by demonstrating that action-induced modulations do not generalize to all types of bistable percepts.

INTRODUCTION

Humans are active agents that can intentionally bring about desired effects in the world. This capacity inherently relies on predictions we make about the consequences of our actions. Indeed, it has been proposed that we can only behave in a goal-direct manner because we represent our actions in terms of their specific sensory effects (Aarts & Dijksterhuis, 2000; Hommel, Müsseler, Aschersleben, & Prinz, 2001), and, accordingly, automatically pre-activate the *identity* of these effects whenever we prepare to act (Miall & Wolpert, 1996). These specific expectations about the identity of one's action-outcomes (i.e., identity predictions) may help us to select the appropriate actions to reach our goals (Elsner & Hommel, 2001; Prinz, 1997), but are also thought to color the subjective perception of our action outcomes (Waszak, Cardoso-Leite, & Hughes, 2012). That is, if specific action-outcomes are pre-activated, actual, prediction-consistent effects, should be selectively affected (e.g., prediction-consistent, but not inconsistent, effects should have facilitated access to conscious awareness). However, as of yet, the existing evidence for such identity-prediction effects is scarce and mixed. In the present studies, we tap into this issue and examine whether anticipating a specific action-outcome can affect the very content of what is consciously perceived.

The view that predicted action-outcomes elicit pre-activation of sensory representations is typically substantiated by two perceptual phenomena: Compared to sensory outcomes with an external origin (e.g., tones, haptic feedback), self-produced sensory input is perceived as occurring earlier (intentional binding; Haggard, Clark, & Kalogeras, 2002) and as suppressed in perceptual intensity and neurophysiological activity (sensory attenuation; Schafer & Marcus, 1973). For instance, neural responses during self-produced speech are smaller compared to when listening to a recording of yourself speaking (Curio et al., 2000).

Importantly, however, these observations typically involve the comparison between self-produced and externally produced outcomes, which differ in more ways than the mere availability of action-outcome predictions. For instance, when preparing to speak, you do not only have information about what is going to be said, but also about when you will start speaking. This type of temporal control and temporal predictability is not present when listening to a tape or to someone else speaking. As such, typical intentional binding and sensory attenuation effects cannot be unequivocally attributed to anticipations about the *content* (or *identity*) of action-effects (see Hughes, Desantis, & Waszak, 2013a for a review). A few

studies have attempted to circumvent the aforementioned issue by manipulating the prediction congruency of self-produced outcomes (e.g., by comparing the perceived intensity of tones that were congruent or incongruent with previously learned action-tone associations); hence isolating identity predictions. The results of these studies were, however, mixed and inconclusive. Specifically, some studies have reported increased sensory attenuation for prediction congruent versus incongruent effects (Cardoso-Leite et al., 2010; Desantis, Mamassian, et al., 2014; Hughes et al., 2013b), while other studies did not find (clear) evidence for such prediction effects (Bednark et al., 2015; Dogge, Hofman, Custers, & Aarts, 2019). In contrast, intentional binding seems unaffected by identity predictions (Bednark et al., 2015; Desantis et al., 2012; Haering & Kiesel, 2014). Taken together, direct evidence for the influence of action-based identity predictions on the perception of action-outcomes is scarce.

Arguably, a more potent approach for capturing the hypothesized influence of identity-predictions on conscious perception, is to capitalize on the phenomenon of perceptual bistability. Bistable images are single stimuli that yield two conflicting perceptual interpretations (identities), one of which can match the expected perceptual outcome (Blake & Logothetis, 2002; Leopold & Logothetis, 1999). Considering that the perceptual conflict in such paradigms can be a conflict of stimulus identity, bistable phenomena are likely to be more susceptible to modulation by action-based predictions compared to intentional binding and sensory attenuation. A handful of studies has provided results compatible with this idea (Beets et al., 2010; Di Pace & Saracini, 2014; Maruya, Yang, & Blake, 2007; Mitsumatsu, 2009; Wallis & Backus, 2016; Wohlschläger, 2000). In these studies, voluntary movement of participants biased the content of bistable perception (e.g., ambiguous motion was perceived in the same direction as the action). Although these results are promising, and generally explained in terms of the pre-activation of action-effects, their reliance on these mechanisms is not unequivocal. In the aforementioned studies participants viewed the bistable stimuli while moving at the same time. This is problematic considering that the execution of an action is not only accompanied by a prediction of upcoming effects, but also by tactile and proprioceptive feedback. Both these types of information are known to bias predominant percepts in favor of the congruent visual interpretation (Lunghi & Alais, 2013; Lunghi, Binda, & Morrone, 2010; Salomon, Lim, Herbelin, Hesselmann, & Blanke, 2013). Accordingly, the aforementioned studies cannot clearly distinguish the effects of outcome-anticipation from other action-related

sources of information (see Wohlschläger, 2000, Experiment 4 for an important exception). The role of identity prediction processes in conscious perception in general - and in bistable perception specifically - thus remains to be elucidated.

Across three experiments, we investigated whether action-outcome anticipation can bias the content of conscious perception. We did so, by utilizing different forms of bistable perception. Study 3.1a and 1b examined the influence of outcome prediction on binocular rivalry; a particular type of bistable perception in which each eye is presented with a conflicting image (e.g., house versus face, or vertical versus horizontal gratings), causing conscious perception to alternate between the two eyes' images in periods of exclusive dominance (Blake, 2001). The percept that initially becomes dominant during exposure to binocular rivalry stimuli (i.e., onset rivalry; Carter & Cavanagh, 2007; Stanley, Forte, Cavanagh, & Carter, 2011) is of particular interest, as action-outcome predictions already emerge as early as the late stages of motor preparation (Desantis, Roussel, & Waszak, 2014). A focus on onset rivalry thus provides a direct test of the proposed mechanisms underlying identity-prediction effects: If action-preparation results in early pre-activation of specific action-outcome representations, this specific outcome-identity should reach the threshold of conscious awareness faster (Waszak et al., 2012), and win the battle for initial dominance.

In Study 3.2, we built upon the few previously mentioned bistable perception studies involving actions (Beets et al., 2010; Di Pace et al., 2014; Maruya et al., 2007; Mitsumatsu, 2009; Wallis & Backus, 2016; Wohlschläger, 2000), and employed bistable motion. Relative to stationary ambiguous stimuli, perception of motion direction is known to be more susceptible to factors like voluntary control (Hugrass & Crewther, 2012). In addition, bistable motion can be modulated by identity predictions, such as expectations based on learned cue-movement associations (Haijiang, Saunders, Stone, & Backus, 2006; Schmack, Weilhhammer, Heinzle, Stephan, & Sterzer, 2016), and hence seems suitable to pick up action prediction effects.

To address our research question, we devised a novel learning task consisting of two types of trials. On induction trials, participants were exposed to associations between a manual rotary movement and a subsequent unambiguous percept (i.e., one of two grating orientations in Study 3.1a and 1b, and one of two motion directions in Study 3.2). On test trials, this same rotary movement was followed by ambiguous percepts (i.e., different gratings in each eye in Study 3.1a and 1b, and

an ambiguously rotating sphere in Study 3.2). Importantly, the rotary movements were temporally separated from the bistable percepts allowing us to examine the role of outcome-anticipation while minimizing influences of other action-related sources of information. We hypothesized that participants would be more likely to report percepts that were consistent with action-outcome anticipations, both when presented with rivalrous stimuli and ambiguous motion.

STUDY 3.1A

Methods and materials

The current line of studies was conducted in line with the recommendations of the ethics board of Social and Behavioral Sciences at Utrecht University and all participants provided informed consent prior to participating in the study. Participants received money and/or course credit in exchange for their participation.

In all studies, Bayesian analyses were used to allow for both the assessment of the evidence in favor of the alternative hypothesis that action predictions affect conscious perception, as well as for the null hypothesis that action predictions do not affect conscious perception (Wagenmakers et al., 2018). Bayesian analyses were performed in JASP version 0.8.2.0, using the default Cauchy prior width of 0.707. Sample size was determined by a pre-specified optional Bayesian stopping rule: Data collection was terminated as soon as a Bayes factor of 6, favoring either the null or the alternative hypothesis, was obtained (i.e., reflecting substantial evidence; Jeffreys, 1961; Lee & Wagenmakers, 2013). However, we decided to collect a minimum of 20 participants (after exclusions based on strong color- or eye-dominance, or failure to adhere to task instructions - see data exclusion). Note that, unlike Frequentist test statistics, Bayesian posterior probabilities and Bayes factors can be recomputed and re-evaluated at any point during the data collection process without error inflation (see Rouder, 2014, for a more detailed explanation). Frequentist statistics are also reported to promote comparison with other studies, but should be treated with caution due the usage of an optional Bayesian stopping rule.

Participants

Thirty-two participants (28 females, 1 left-handed, $M_{age} = 21.84$, $SD_{age} = 2.53$) took part in the study.

Stimuli and apparatus

Participants viewed the stimuli through a mirror stereoscope consisting of two mirrors at a 45 degree angle, each reflecting one of two linearized 23-inch LCD monitors (Dell UZ2315H; resolution: 1920 x 1080; refresh rate: 60 Hz) that were facing each other (see Brascamp & Naber, 2017 for a more detailed description). A chin rest and forehead rest were used to ensure a stabilized head position at a viewing distance of approximately 82 cm.

The rivalrous stimuli comprised of circular square-wave gratings (diameter of 1.48 degrees of visual angle (dva)) that were surrounded by a gray annulus (2.5 Cd/m², 1.67 dva¹⁴; thickness: 0.11 dva) with four diagonal rectangular junctures (see Figure 3.1). In addition, a vertical rectangle, which had to be aligned with two of these junctures (see procedure), was superimposed on the grating. Both the rectangle as well as the grating were colored gray on induction trials and red-gray or blue-gray on rivalrous test trials. To promote stereoscopic fusion of the complementary images presented to each eye, the stimuli were presented within identical circular frames of 2.59 dva in diameter (1.1 Cd/m²), which were in turn encapsulated by a Brownian noise rectangle (2.96 dva by 3.33 dva) on a black background (< 1 Cd/m²). All gratings had a spatial frequency of 6.75 cycles per degree, and a Michelson contrast of 0.2 between minimum (1.6 Cd/m²) and maximum (2.5 Cd/m²) luminance values.

Procedure and design

The experimental task was framed as a “bridge building game” in which participants turned vertical gratings into oblique (tilted) “bridges” by rotating a custom-made cylindrical knob with their dominant hand. The required rotation direction was indicated by a visual cue presented to the left or right side of the grating. Specifically, participants were instructed to rotate the vertical grating in such a way that an imaginary ball would be able to roll towards the cue following the line-orientation of the grating (i.e., a cue on the right side required a leftward rotation); see Figure 3.1.

The task featured two types of trials: induction trials and test trials. In induction trials, rotary movements were immediately followed by presentation of the corresponding (unambiguous) grating to both eyes. These trials served as

14 Despite matching the calibration of both displays with the Spyder3Elite (Datacolor, Lawrenceville, NJ), there were slight luminance discrepancies between the screens (mean discrepancy was approximately 9%). All reported luminance values are as measured on the left screen.

exposure to action-effect associations between rotary movements and oblique grating orientations. In test trials the perceptual outcome following the rotary movement was made ambiguous by presenting each eye with a different -45 or $+45$ degree grating orientation; one of which was consistent with the preceding movement. Accordingly, test trials allowed us to examine whether the action-effect associations that were learned in induction trials biased perceptual selection between rivalrous orthogonal gratings. Note that no changes in grating orientation occurred while the rotary movement was being executed. Once the movement was completed the vertical gratings immediately changed to the oblique outcome gratings.

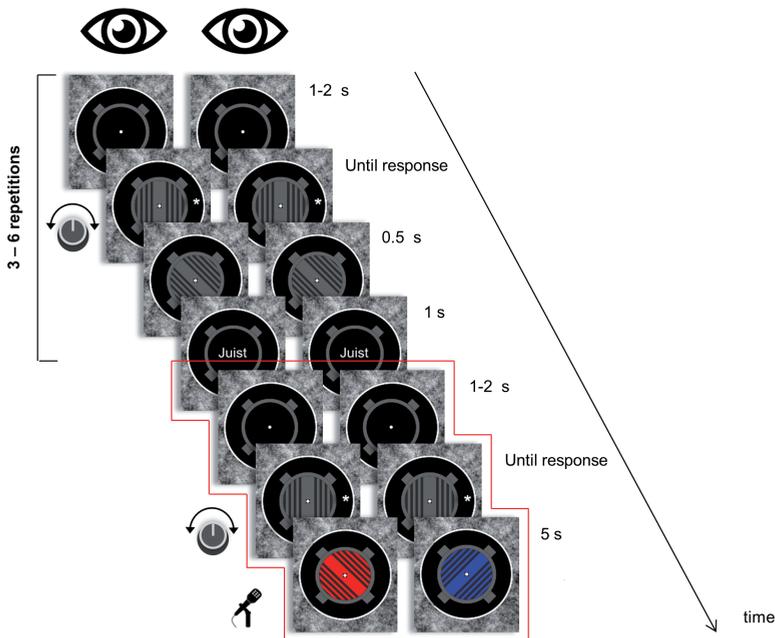


Figure 3.1. Schematic representation of trial events in induction and rivalry trials of Study 3.1a. Contrasts and colors deviate from those used in the experiment for illustrative purposes. Participants were instructed to rotate a vertical grating in such a way that the lines of the grating were pointing towards the side of a cue (*). In a sequence of three to six induction trials, rotary movements were followed by the presentation of the same (action-coherent) oblique grating in both eyes, and by feedback on the accuracy of the action (i.e., the Dutch word for correct (“juist”) or incorrect (“onjuist”)). In the rivalry trials (outlined in red), each eye was presented with a different grating whose orientation was either consistent or inconsistent with the preceding movement. Participants were asked to verbally report the initially dominant percept and all subsequent switches. Cue position could vary within a sub-block.

Each trial started with the presentation of a fixation point for a randomly chosen duration of 1000, 1250, 1500, 1750, or 2000 milliseconds. Subsequently, participants viewed a vertical grating with a movement cue on the left or right side, presented at an eccentricity of 1.05 degrees of visual angle. In response to this cue, participants were required to execute a clockwise or counterclockwise manual rotation in order to rotate the grating by 45 degrees towards the required ending position. In order to help participants to execute the proper movement (i.e., rotate the grating to the proper ending position), four diagonal junctures were shown on the annulus surrounding the grating. The vertical superimposed rectangle presented in the center of the grating had to be aligned with two of these diagonal junctures. In addition, the rotary button produced a soft but audible click when reaching a rotation angle of 45 degrees. In induction trials, this movement resulted in the presentation of an action-coherent outcome (500 milliseconds), as well as textual feedback on the accuracy of the movement (i.e., the Dutch words for “correct” or “incorrect”; 1000 milliseconds). To motivate participants to pay attention to the cue position, participants could earn money with correct rotations and lose money with incorrect rotations on induction trials. In the test trials a different grating was presented to each eye, at corresponding retinal locations and participants were instructed to verbally report the color of the first dominant percept and of all subsequent percepts during five seconds. The verbal responses were recorded by a microphone. Note that by asking participants to report a feature of the grating that was orthogonal to the feature of interest (i.e., color instead of orientation), we reduced the likelihood that prediction consistent responses are caused by demand characteristics; see Attarha and Moore (2015) for a similar approach.

Before the experiment participants completed pre-screening tests to check for color vision (Ishihara, 1917), visual acuity (customized, digital Landolt-C test), and stereoscopic vision. Participants did not take part in the experimental phase if they displayed abnormal color vision, failed to distinguish a gap of 0.07 degrees of visual angle in the Landolt-C task¹⁵ or experienced diplopia. In addition, participants were screened for eye dominance (i.e., a bias for perceiving percepts in a particular eye), and color dominance (i.e., a bias for seeing one of the two grating colors) as both of these biases reduce the amount of variance that can be explained by prediction consistency effects. As part of this screening participants

15 This width corresponds to one line (i.e., half a sinus cycle) in the gratings, and therefore reflects the minimally required visual acuity for distinguishing between the two grating orientations.

completed 16 trials in which they passively viewed a vertical grey grating that changed into a colored rivalrous grating (i.e., the same gratings that were used in the actual experiment). This rivalrous grating was presented for five seconds while participants indicated their dominant percept by means of the “q” and “p” key on the keyboard (representing red and blue respectively). Orientation placement (left-eye -45 and right-eye $+45$ vs. left-eye -45 and right-eye $+45$), and Color placement (left-eye red and right-eye blue vs. left-eye blue and right-eye red) was randomized over trials¹⁶. Participants who showed either eye or color onset-dominance on more than 85% of the trials did not participate in the study (the first four trials were not included in this calculation). After successfully completing the pre-screening participants completed several practice rounds to become familiar with the rotary button and task-setup.

Induction and test trials were presented in sub-blocks consisting of a sequence of three to six induction trials and one final test trial. Length of sub-blocks was varied to prevent anticipation of test trials. The experiment was divided into four blocks that each contained sixteen sub-blocks. The four possible sub-block lengths were each presented four times in a randomized order within each block. This resulted in a total of 288 induction trials and 64 rivalrous test trials per participant. Cue position (left vs. right), orientation placement (left-eye -45 and right-eye $+45$ vs. left-eye -45 and right-eye $+45$), and color placement (left-eye red and right-eye blue vs. left-eye blue and right-eye red) were manipulated within participants. Each combination was shown two times in each block in a randomized order.

Results

Data exclusion

Despite successfully completing the aforementioned pre-screening tests, data sets from ten participants had to be excluded because eye- or color dominance exceeded 85% during the task¹⁷. The remaining sample consisted of twenty-two participants (18 females, 1 left-handed, $M_{age} = 21.68$, $SD_{age} = 2.08$). In addition, trials were excluded from analysis when the button was incorrectly handled, such

16 For six of the inclusions, the prevalence of different factor levels was not perfectly equated due to a programming error. As this pre-screening procedure was only used to assess whether participants were allowed to participate in the experiment in the first place (with the aim to minimize post-experiment exclusion), and color and eye dominance were also assessed during the main experiment, this is unlikely to have affected the reported results.

17 For one participant the task was terminated after two experimental blocks in which only the red percept reached initial dominance.

as when it was rotated too fast (before the presentation of the cue), too slow (not within the response limit of 2 seconds), not far enough, back and forth or more than once; $M = 0.92\%$, $SD = 1.33\%$). In addition, trials in which the button was rotated in the wrong direction (i.e., not in the direction that was indicated by the movement cue; $M = 0.14\%$, $SD = 0.46\%$) and trials in which the first spoken response after presentation of the rivalry stimulus was something other than “red” or “blue” were also excluded from analysis ($M = 0.79\%$, $SD = 1.06\%$).

Perceptual selection

Only the first verbal responses to the rivalry stimuli were considered in the analyses. The proportion of prediction-consistent percepts of all participants (see Figure 3.2) were submitted to a directional one-sample t-test against chance (i.e., a proportion of 0.5 prediction-consistent responses). This analysis revealed that perceptual selection of prediction-consistent percepts ($M = 0.51$, $SD = .07$) was not likely to be greater than what is expected by chance ($BF_{+0} = 0.59$; $t(21) = 1.02$, $p = .16$, $d = .22$). Instead the data were 1.7 times more likely to reflect a null effect (no influence of prediction) than to reflect greater prevalence of prediction-consistent responses.

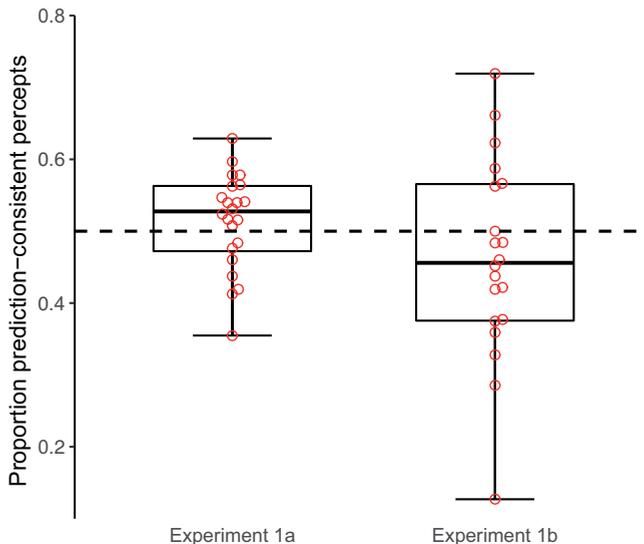


Figure 3.2. Proportion of prediction-consistent percepts for all participants in Study 3.1a and 1b. Lower and upper box limits represent the 25th and 75th percentiles respectively. Whiskers extend to the most extreme values that fall within 1.5 times the interquartile distance from the hinges of the box. Central lines and red data points represent the median and individual participants, respectively. The dotted line represents chance.

Discussion

Study 3.1a did not show convincing evidence for the influence of action-based predictions on perceptual selection during onset rivalry. These results have to be interpreted with caution, considering that the outcome of the Bayesian analysis was inconclusive and therefore also failed to provide support for the null hypothesis. We opted to end the experiment preliminary (i.e., prior to reaching all of the stopping rule requirements) as sequential analysis of the Bayes factor failed to consistently diverge to either of the pre-specified, stopping-rule thresholds. Participant debriefings suggested that a potential reason for the absence of the prediction effects might be related to suboptimal task design. Specifically, participants indicated that they used the position of the cue to determine the correct rotation direction, rather than the required position of the grating. In other words, rather than envisioning which oblique orientation would allow the ball to roll towards the cue, they used an easier stimulus-response strategy (e.g., “if the cue is presented on the left, I have to rotate to the right”). This strategy clearly reduces the task relevance of the grating orientation, which might have prevented participants from picking up on the associations between these stimulus features and preceding actions. Indeed, it has been shown that the acquisition (and/or retrieval) of action-effect associations is hampered for non-salient, task-irrelevant outcomes (Dutzi & Hommel, 2009). As expectations about upcoming action-outcomes cannot be formed in the absence of these associations, the reduced relevance of grating orientations might explain the absence of the prediction effects in the present study.

In Study 3.1b we tackled this issue by discouraging participants from using aforementioned stimulus-response strategies. Instead of always starting in a vertical position, the grating could now also be presented in a horizontal position. Accordingly, the position of the cue itself was no longer sufficient to determine the required rotation direction (i.e., a cue presented on the left required a rotation to the right for a vertical grating and a rotation to the left for a horizontal grating). In comparison to Study 3.1a, this approach promotes visualization of the action outcome. In addition, we would like to point out that, in contrast to Study 3.1a, each oblique grating could now result from both a clockwise and a counterclockwise rotation, which precludes the possibility that potential prediction effects are caused by mere action execution rather than anticipated action-outcomes.

Another design element that might have accounted for the absence of the predicted effect is the difference between gratings presented in the induction trials and test trials. While participants are presented with chromatic gratings in the test trials, they learn to produce (and presumably predict) gray oblique gratings in the induction trials. These predicted perceptual outcomes in grayscale might not (fully) transfer to chromatic gratings. There is some evidence that is compatible with this notion from sensory attenuation research. Specifically, neurophysiological suppression has shown to be largest for the action-outcomes that most closely match predicted outcome features (Baess et al., 2008; Heinks-Maldonado et al., 2005). The partial mismatch in the predicted and actual features of the gratings in the present study might therefore have reduced modulation of onset rivalry by predictions of action outcome. In Study 3.1b we reduced this mismatch by utilizing identical grayscale gratings in the induction and test trials.

On top of the aforementioned adjustments, several other changes were implemented to further increase the likelihood of observing prediction effects. Firstly, participants used keys instead of verbal reporting in Study 3.1b. Verbal reports are significantly slower, which means that participants sometimes found themselves in the process of verbalizing a percept, while a switch to another percept already occurred. This possibly noise-inducing factor is less likely when using key responses. Secondly, participants reported significant fatigue and loss of concentration while doing the task. We therefore shortened the length of the task by showing the rivalrous stimuli for one second instead of five seconds (see Carter and Cavanagh (2007) for a similar approach to measure onset rivalry).

STUDY 3.1B

Methods and materials

Participants

Thirty participants (23 females, 3 left-handed, $M_{age} = 23.17$, $SD_{age} = 2.88$) completed the study. None of the participants participated in Study 3.1a.

Stimuli and apparatus

Stimuli were the same as in Study 3.1a, with a few exceptions. Only gray gratings were used. Moreover, feedback text was replaced by a red or green

fixation point (i.e., incorrect versus correct responses) to facilitate a focus on the fixation point throughout the trial. We also sought to reduce eye dominance effects by adjusting the contrast of the grating presented to the non-dominant eye for each participant, so as to equate perceptual dominance of the two eyes. To do so, two adaptive, accelerated stochastic approximation staircases (Kesten, 1958) of 40 trials were used in which the contrast of the dominant eye was held constant (Michelson contrast of 0.2 between background (9.5 Cd/m^2) and minimum/maximum luminance), while the contrast of the non-dominant eye (determined by means of a hole-in-the-card task) was adjusted based on the initially dominant percept (starting at a Michelson contrast of 0.2 or 0.8). The mean contrast value of the last eleven trials of both staircases was the contrast used for the non-dominant eye in the main experiment ($M = 0.27$, $SD = 0.14$).

Procedure and design

The procedure was the same as Study 3.1a with a few exceptions. In order to emphasize the association between the rotation of the knob and the rotation of the gratings, the instructions were slightly altered. Participants were told the gratings were visual representations of the knob that had to be set in the correct position. Specifically, the gratings had to be rotated in such a way that all the lines would point towards a cue (*) presented slightly above the left or right bottom juncture on the annulus surrounding the grating (see Figure 3.3). In addition, the starting position of the grating could now be either vertical or horizontal, the duration of rivalry stimuli was shortened to one second, and participants reported their dominant percept by means of a button press. Specifically, participants used the upward or downward arrow key to select one of two vertically stacked images depicting either a -45 degree or a $+45$ -degree grating. Vertical rather than horizontal response options were chosen such as to minimize response bias elicited by the preceding leftward or rightward rotary action. In addition, these responses were executed with the non-dominant hand, whereas the dominant hand was used for rotating the knob.

The aforementioned changes resulted in a design comprising starting position (vertical vs. horizontal), cue position (left vs. right) and orientation placement (left-eye -45 and right-eye $+45$ vs. left-eye -45 and right-eye $+45$) as within subject factors, and response mapping (top -45 and bottom $+45$ versus top $+45$

and bottom -45) as a between subject factor¹⁸. Each combination of within subject factors was shown two times in each block in a randomized order.

Before the experiment participants completed pre-screening tests to check for stereoscopic vision and eye dominance. Since a staircase was used to diminish eye dominance effects, participants were only excluded if they displayed diplopia.

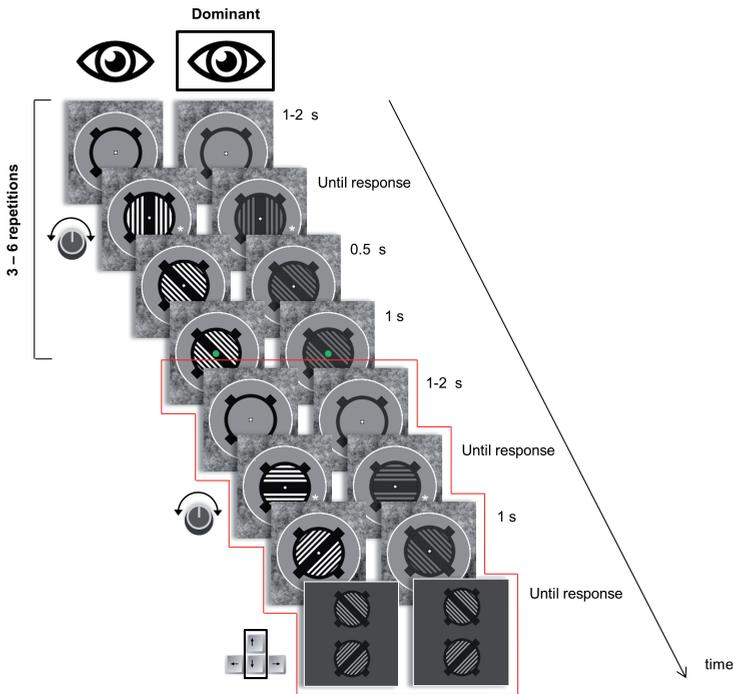


Figure 3.3. Schematic representation of trial events in induction and rivalry trials of Study 3.1b. Contrasts and colors deviate from those used in the experiment for illustrative purposes. Participants completed sub-blocks consisting of three to six induction trials and one rivalry trial (outlined in red). Cue position could vary within a sub-block.

18 Due to participant exclusion, response mapping was not equally divided between groups (13 participants received the first response mapping and 7 participants received the second response mapping). A Bayesian ANOVA with mapping as a between-subject factor did not provide conclusive evidence for (nor against) a difference in prediction consistency between these two groups ($BF_{10} = 1.16$). However, we would like to stress that there is no reason to assume that one of the response options would lead to more prediction consistent responses.

Results

Data exclusion

For eight participants eye dominance exceeded 85% during the task. The datasets of these participants were excluded from further analyses, as motivated in Study 3.1a. The fact that we had to exclude participants for excessive eye dominance, despite the use of the staircases, concurs with previous work demonstrating that equating stimulus strength of rivalry stimuli does not necessarily prevent dominance biases in onset rivalry (Stanley, Carter, & Forte, 2011). In addition, the datasets of two participants were excluded for failing to keep the head rested in the chinrest and for consistently executing multiple instead of one knob rotation, respectively. The remaining sample consisted of twenty participants (15 females, 1 left-handed, $M_{age} = 23.65$, $SD_{age} = 3.18$). Trials in which the button was incorrectly handled (i.e., rotated too fast, too slow, not far enough, back and forth or more than once; $M = 1.88\%$, $SD = 2.61\%$) or rotated in the wrong direction (i.e., not in the direction that was indicated by the movement cue; $M = 0.39\%$, $SD = 0.69\%$) were also excluded from analysis.

Perceptual selection

The proportion of initially dominant prediction-consistent percepts was submitted to a directional one-sample t-test against chance (i.e., a proportion of .50). In contrast to our hypothesis, this analysis provided substantial evidence *against* the hypothesis that participants are more likely than chance to perceive prediction-consistent percepts ($M = 0.46$, $SD = 0.14$; $BF_{+0} = 0.12$; $t(19) = -1.25$, $p = .89$, $d = -0.28$). That is, the observed results were roughly eight times as likely to be observed under the null hypothesis compared to the alternative hypothesis (see Figure 3.2).

Discussion

Onset rivalry was not influenced by exposure to action-effect associations. Unlike the previous study, the Bayesian analysis now yielded substantial evidence in favor of the null hypothesis. Even a complete match between gratings on induction and rivalry trials, as well as more relevance of the grating orientation was not sufficient to bias perceptual selection. Together, the combined findings of Study 3.1a and 1b suggest that, at least under the current experimental conditions, onset rivalry is insensitive to action-based predictions.

STUDY 3.2

One possible reason for the absence of prediction effects in Study 3.1 might be that onset rivalry involves competition that is not, or only weakly, influenced by top-down factors. Indeed, onset rivalry is very sensitive to bottom-up stimulus features (e.g., contrast, luminance), and early trait-like influences such as eye dominance, while the role of top-down processes has not yet been extensively examined (Stanley, Forte, et al., 2011).

In Study 3.2 we attended to this issue by examining whether action-outcome anticipation can influence perception of a structure from motion sphere; a bistable percept comprised of moving dots that can be perceived as a globe rotating in a leftward or a rightward direction (Andersen & Bradley, 1998). Perception of ambiguous spheres has been shown to be sensitive for modulation from a variety of top-down factors including beliefs (Schmack et al., 2013; Sterzer, Frith, & Petrovic, 2008), working memory (Scocchia, Valsecchi, Gegenfurtner, & Triesch, 2013), associative learning (Schmack et al., 2016), and voluntary control (Brouwer & Van Ee, 2006). By keeping all other aspects of the task as similar as possible to the previous experiments, Study 3.2 therefore allows us to examine whether the influence of action-based predictions can be observed when using a type of bistable perception that is known to be susceptible to modulation by top-down factors.

Methods and materials

Participants

Twenty participants (15 females, 2 left-handed, $M_{age} = 23$, $SD_{age} = 2.99$) participated in the study.

Stimuli and apparatus

Visual stimuli were structure-from-motion spheres (diameter = 2.22 dva) consisting of 240 white squared “dots” (0.06 dva; 113 Cd/m²), which each had a fifty percent chance of moving in a rightward or a leftward direction. In the center, the dots moved with a speed of 45 degrees/s, which decreased to a speed of zero near the edges of a circular aperture following a sinusoidal profile, thereby eliciting the percept of a sphere. Dot lifetime was one second (to avoid that all dots would be replaced simultaneously, the initial dot ‘age’ was randomly chosen between zero and one seconds). The rotation direction of the sphere was either unambiguous (induction

trials) or ambiguous (test trials). Ambiguous spheres consisted of two identical 2-D sphere projections that were presented to both eyes. Due to the sinusoidal speed profile, observers tend to see the dots moving in one direction as being in a different depth plane than the dots moving in the other direction, causing the percept of a leftward or rightward rotating globe (depending on whether the leftward or rightward dots are perceived in the front plane respectively). Unambiguous spheres were created by adding binocular disparity to the dots (i.e., adding a horizontal offset between the two eyes; maximum offset from fixation = 0.04 dva), inducing a stable depth percept (e.g., by presenting the leftward moving dots on the near depth plane, and the rightward moving dots on the far depth plane). Importantly, the monocular stimulation induced by presentation of the unambiguous and ambiguous spheres were virtually identical, while the percept elicited by the unambiguous spheres could still be reliably manipulated experimentally. All stimuli were presented on a black background ($< 1 \text{ Cd/m}^2$) and viewed through the same stereoscope as described for the first two studies.

Procedure and design

The experiment featured two types of trials: 1) unambiguous induction trials in which only the prediction-consistent rotation was presented and 2) ambiguous test trials in which perception could be of either the prediction-consistent or prediction-inconsistent rotation direction. Each trial started with a gray fixation point (56.2 Cd/m^2 ; 0.13 dva) that was presented for a randomly chosen duration of 1000,1250,1500,1750 or 2000 milliseconds. Subsequently, participants observed a stationary 2D sphere with a fixation point in the center and a red movement cue (23.3 Cd/m^2 ; 0.13 dva) presented left or right of the sphere at an eccentricity of 1.48 degrees of visual angle. Participants were instructed to keep their gaze fixated on the fixation point and to use a rotary button to rotate the front of the sphere towards the cue (i.e., a cue on the right side required a leftward rotation). To motivate participants to pay attention to the cue position, they were told that correct rotations would be rewarded with a monetary compensation. Similar to the previous onset rivalry experiments, participants could only earn money on the induction trials. This rule was not explicitly mentioned in the present study to not draw attention to differences between induction and test trials. After correctly executing the rotary movement within response time limit (two seconds), a rotating ambiguous or unambiguous sphere was shown for one second. Participants reported the perceived rotation direction by selecting one of two vertically stacked images depicting either

a leftward or rightward rotation. Images were selected by means of the upward or downward arrow key (rather than the leftward or rightward keys) to minimize bias by the preceding leftward or rightward rotary action (see Figure 3.4).

Induction and test trials were presented in the same sub-blocks as described in Study 3.1. Cue position (left vs. right) was manipulated within participants. On half of the induction trials (36 trials) and half of the test trials (8 trials) in each block a target was presented on the left side of the sphere and on the other half of the trials a target was presented on the right side of the sphere. The order of the targets was randomized within each block. Response mapping (upward arrow = leftward rotation and downward arrow = rightward rotation vs. upward arrow =

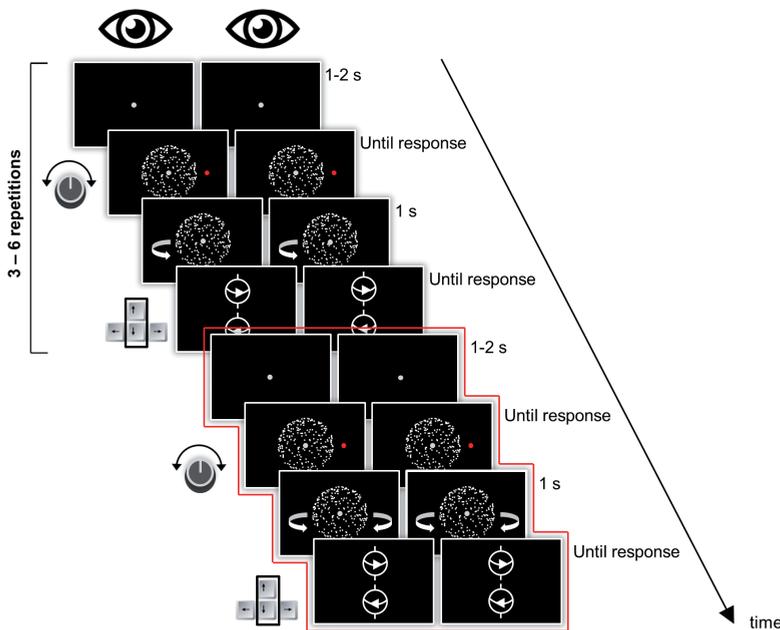


Figure 3.4. Schematic representation of trial events in induction and test trials of Study 3.2. Participants were instructed to rotate the front of a sphere towards a red dot presented on its left or right side. In induction trials, movements resulted in an unambiguous, action-coherent, rotation. In test trials, participants were presented with ambiguous bistable spheres that could be perceived as rotating in either a leftward or rightward direction. Participants indicated the perceived rotation direction by selecting the corresponding image with the arrow keys. Cue position could vary within a sub-block.

rightward rotation and downward arrow = leftward rotation) was counterbalanced between participants.

Before the experimental phase participants were pre-screened for stereoscopic

vision by inducing varying retinal disparity. Participants who experienced diplopia did not take part in the study.

Results

Data exclusion

One participant showed a remarkably low accuracy level in the unambiguous induction trials (accuracy = 0.62; z -score = -3.68). Considering that this resulted in significantly reduced exposure to correct action-effect associations, this participant was excluded from further analyses. The remaining sample consisted of nineteen participants (14 females, 2 left-handed, $M_{age} = 22.68$, $SD_{age} = 2.71$). Trials in which participants incorrectly handled the button (i.e., rotated the button too fast, too slow, not far enough, back and forth or performed more than one rotation; $M = 0.58\%$, $SD = 1.07\%$), as well as trials in which participants rotated the button in the wrong direction (i.e., not in the direction that was required by the movement cue; $M = 0.82\%$, $SD = 1.59\%$), were excluded from further analyses.

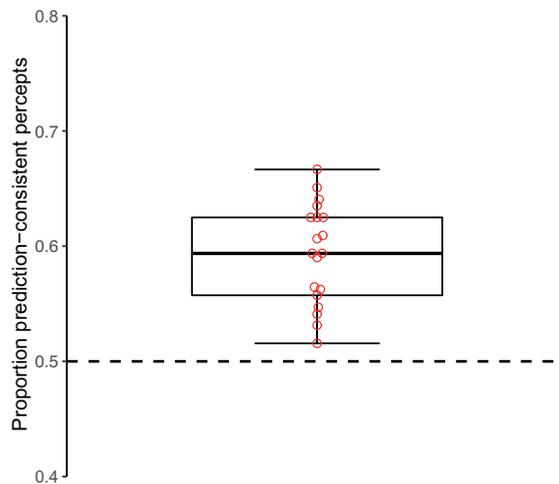


Figure 3.5. Proportion of prediction-consistent percepts for all participants in Study 3.2. The lower and upper box limit represent the 25th and 75th percentile respectively. Whiskers extend to the most extreme values that fall within 1.5 times the interquartile distance from the hinges of the box. The central line and red data points represent the median and individual participants, respectively. The dotted line represents chance.

Perceived motion

The proportion of test trials in which the perceived rotation corresponded with learned action-effect associations was calculated for all participants for both induction and rivalrous trials. Mean accuracy on induction trials

was high ($M = 96.50\%$, $SD = 4.54\%$) indicating that the rotation directions of unambiguous spheres were clear and easily distinguishable. Mean proportion of prediction-consistent percepts on rivalrous trials were submitted to a directional one-sample t-test. This test indicated that the observed data was more than a million times more likely under the alternative hypothesis (prediction-consistency > chance), compared to the null hypothesis (prediction-consistency = chance; $M = 0.59$, $SD = 0.04$; $BF_{+0} = 1.16 \times 10^6$; $t(18) = 9.41$, $p < .001$, $d = 2.16$)¹⁹. Numerically, every single participant had a higher proportion of prediction-consistent than prediction-inconsistent responses (see Figure 3.5).

GENERAL DISCUSSION

Across three studies we demonstrate that while action-outcome anticipation does not affect perceptual selection during onset rivalry, it does modulate the perceived rotation direction of ambiguous spheres. These latter results are in line with previous demonstrations of identity prediction effects in sensory attenuation (Cardoso-Leite et al., 2010; Desantis, Mamassian, et al., 2014; Hughes et al., 2013b). Here we move beyond this work by showing that action-outcome anticipation cannot only modulate perceptual processing (i.e., the perceived *intensity* of stimuli), but can also affect perceptual content (i.e., the perceived *identity* of stimuli). Moreover, the observed prediction effect on bistable motion corroborates previous work showing that actions tend to bias perception of ambiguous motion in the direction of movement (Mitsumatsu, 2009; Wallis & Backus, 2016; Wohlschläger, 2000). However, in contrast to these previous studies, we separated expectations about action-outcomes from other action-related information sources (e.g., proprioception), and, as such, provide more direct evidence for the proposed role of outcome-anticipation in perceptual processing (Waszak et al., 2012).

Interestingly, action-identity prediction effects do not appear to generalize to all types of bistable perception, as no effects were observed for onset rivalry. These diverging results raise the question of whether predictive processes might have been weaker or absent in Study 3.1a and 1b compared to Study 3.2. We deem this unlikely considering that the exposure to action-effect associations, in terms of number of trials, was identical across studies. There are, however, several other

¹⁹ Including the dataset of the excluded participant in the analysis (see “Data exclusion”) did not change this pattern of results: $BF_{+0} = 1.54 \times 10^5$.

important differences between the studies that might explain the divergent effects; the most apparent being the variation in the type of bistability that was employed. One possible, albeit speculative, explanation of the results is that these different levels of bistable perception might reflect different levels of visual processing, with varying susceptibility for top down modulation. Specifically, we propose that (onset) rivalry might be resolved at a lower stage of visual processing than structure-from-motion, which might be less permeable to top-down influences (see Meng & Tong, 2004 for a similar argument). There are several arguments that can be put forward in support of this line of reasoning.

With respect to permeability by top-down influences, binocular rivalry has been shown to be less susceptible to top-down modulation compared to other forms of bistable competition or perceptual suppression. For instance, Meng and Tong (2004) showed a reduced ability of participants to selectively switch between rivalrous stimuli as opposed to the two possible orientations of a Necker cube. In addition, perception of ambiguous structure-from-motion spheres is biased towards the concurrent content of visual working memory (Scocchia et al., 2013), whereas perception during binocular onset rivalry is not (Gayet, Brascamp, Van der Stigchel, & Paffen, 2015). Also, only basic visual processing can be applied to an image that is binocularly suppressed (for reviews, see Gayet, Van der Stigchel, & Paffen, 2014; Lin & He, 2009), compared to other forms of perceptual suppression that allow for more elaborate processing of the suppressed image (Breitmeyer, 2014).

In addition, there are at least two ways that rivalry seems different from structure-from-motion spheres in terms of the level of visual processing. First, the “perceptual manifestation” of binocular rivalry differs from that of other forms of bistable perception, such as structure-from-motion (Brascamp, Sterzer, Blake, & Knapen, 2018). Specifically, in an ambiguous structure-from-motion sphere, one of the two percepts is effectively perceptually suppressed (e.g., the leftwards moving percept), but the visual elements that constitute this percept are nonetheless consciously accessible (i.e., the dots that make up the sphere). In binocular rivalry suppression, however, the observer not only fails to see the suppressed percept of interest (i.e., the grating orientation), but also lacks conscious access to image-parts that constitute the suppressed percept, suggesting more complete suppression. This “deeper” suppression might render rivalry stimuli less accessible to top-down influences, potentially explaining the lack of observed prediction effects in Study 3.1a and 1b. In contrast, structure-from-motion spheres are more similar to other

forms of bistable perception, like Necker cubes, whose visual elements remain accessible and, are therefore modifiable by factors such as voluntary attention (e.g., voluntarily attending to certain ribs of the cube can be used to induce switches; Kawabata, 1986).

Second, there are some indications that rivalry can be resolved at very early stages of visual processing, such as V1 (Lee & Blake, 2002) and even at the level of the lateral geniculate nucleus (Haynes, Deichmann, & Rees, 2005). Whilst rivalry has also been associated with activity of higher fronto-parietal areas (Sterzer et al., 2009), the extent to which these areas play a role in modulating interocular competition, or are involved in cognitive processes that arise after the interocular competition is resolved (e.g., task demands, or perceptual reports), remains a topic of vehement debate (Brascamp et al., 2018; Giles, Lau, & Odegaard, 2016; Zaretskaya & Narinyan, 2014). In addition, there is some recent evidence suggesting that, although frontal-parietal involvement might sometimes be observed during rivalry, it is not necessary for perceptual reversals to occur (Brascamp, Blake, & Knäpen, 2015; Zou, He, & Zhang, 2016). These latter findings are in line with our proposal that rivalry reflects a relatively low stage of visual competition (although note studies showing modulations of perceptual switching as a function of parietal-TMS; Carmel, Walsh, Lavie, & Rees, 2010; Zaretskaya, Thielscher, Logothetis, & Bartels, 2010).

At first glance, the idea that binocular rivalry occurs too early in the visual processing hierarchy for top-down modulation by action-outcome anticipation might seem to contradict recent findings showing that onset rivalry can be biased by motion predictions (Attarha & Moore, 2015; Denison, Piazza, & Silver, 2011). In these studies, participants were more likely to select the rivalrous grating with an orientation that logically followed the preceding motion sequence than a grating with a motion-inconsistent orientation. Although these results might be interpreted as evidence for the influence of top-down predictions on (early) visual competition, an alternative explanation has been offered. Specifically, Attarha and Moore (2015) posit that instead of reflecting high-level predictive influences, these results might also have resulted from more local, low-level, priming mechanisms, in which motion-consistent percepts are primed by synchronized activity between subcortical areas and the visual cortex. This interpretation is in line with the observation that both predictions based on semantic context and more complex pattern motion failed to bias onset rivalry (Attarha & Moore, 2015), as well as with our failure to observe an influence of outcome-predictions on perceptual selection.

Taken together, there is some, although not unequivocal, evidence to support the idea that the difference in findings between the two present studies might be the result of differences in the stages of visual competition that were involved. We like to stress that this notion is speculative, and point out that the current results and explanation might be restricted to the initial dominance percept of rivalry. Indeed, there is evidence to suggest that the underlying mechanisms of onset rivalry and ongoing rivalry might be at least partially independent (e.g., Stanley, Forte, et al., 2011). Additional research comparing various forms of bistable perception (and binocular rivalry), with varying depths of suppression, could provide valuable insights into these issues.

There are several alternative explanations, relating to subtle procedural differences between the studies, that need to be considered. One of these procedural differences pertains to the extent of spatial correspondence between hand movements and grating rotations. In the rivalry studies, the rotation plane of the knob was perpendicular to the movement of the grating on the screen. In contrast, in Study 3.2, the sphere rotated on the same plane as the rotary button. Although previous work has suggested that spatial correspondence between action and motion displays facilitates action-based modulation of bistable perception (Wohlschläger, 2000; Experiment 1), these effects have also been observed in the absence of such correspondence. Specifically, when action and motion display share a cognitive dimension, a congruent spatial relation has shown not to be a prerequisite for the effects (Wohlschläger, 2000; Experiment 3). Considering that such cognitive coupling between hand movements and bistable stimuli was equally induced across studies, we are confident that action-identity prediction occurred in all studies.

Another procedural factor that could be argued to account for the observed differences between Study 3.1 and 2 is demand characteristics. That is, the results of Study 3.2 might not be a reflection of outcome prediction effects, but instead reflect the propensity of participants to report their explicit expectations concerning what they should perceive instead of their actual percepts. We assume that such demand characteristics are especially probable in the absence of a clear sensory signal. In contrast, if participants are very certain about what they see, they will tend to rely on (and report) their percepts. If demand characteristics are fully responsible for the effect we observed in Study 3.2, the question remains why they would not have been present in Studies 1a and 1b. We have reasons to believe that such demand characteristics were, in fact, more likely to play a part in Studies

1a and 1b (in which no effect of action-prediction was observed). One indicator of perceptual uncertainty in bistable perception paradigms is the number of reported mixed percepts, referring to situations in which participants do not exclusively perceive one of the competing percepts, but rather a mixture of the two (i.e., parts of each eye's stimulus during binocular rivalry (Blake, 2001), and two convex or concave half spheres during ambiguous motion (Hol, Koene, & Van Ee, 2003). While participants in our lab regularly report the former, the perception of mixed sphere percepts is rare. Considering the absence of prediction effects in the first two studies, in which perceptual uncertainty was higher, it is unlikely that demand characteristics can exclusively account for the observed effects in Study 3.2.

CONCLUSION

We extend previous work by showing that action-outcome anticipation cannot merely modulate perceptual processing, but can also bias the actual content of our conscious perception. Importantly, prediction effects were only observed when actions were followed by bistable motion, but not by rivalrous gratings. The present results coincide with previous work suggesting that binocular rivalry in general, and onset rivalry in particular, might be less susceptible to top-down influences. This discrepancy might point towards differences in underlying mechanisms. Further examination of the commonalities and discrepancies between distinct forms of bistable perception, as well as the possible restrictions to the influence of top-down biases, will provide important insights in the role of predictive processing on awareness.

CHAPTER 4

Perception of action-outcomes is shaped by life-long and contextual expectations

This chapter is based on: Dogge, M. Custers, R., Gayet, S., Hoijsink, H., and Aarts, H. (2019). Perception of action-outcomes is shaped by life-long and contextual expectations. *Scientific Reports*, 9, 5225. <https://doi.org/10.1038/s41598-019-41090-8>

ABSTRACT

The way humans perceive the outcomes of their actions is strongly colored by their expectations. These expectations can develop over different timescales and are not always complementary. The present work examines how long-term (structural) expectations – developed over a lifetime - and short-term (contextual) expectations jointly affect perception. In two studies, including a pre-registered replication, participants initiated the movement of an ambiguously rotating sphere by operating a rotary switch. In the absence of any learning, participants predominantly perceived the sphere to rotate in the same direction as their rotary action. This bias toward structural expectations was abolished (but not reversed) when participants were exposed to incompatible action-effect contingencies (e.g., clockwise actions causing counterclockwise percepts) during a preceding learning phase. Exposure to compatible action-effect contingencies, however, did not add to the existing structural bias. Together, these findings reveal that perception of action-outcomes results from the combined influence of both long-term and immediate expectations.

INTRODUCTION

Perception is a crucial faculty for moving organisms. In order to support fast and efficient interaction with the environment, perceptual processing cannot merely rely on bottom-up sensory input, but instead requires top-down knowledge to shape and understand incoming signals (Hommel et al., 2001). As a consequence, perception is strongly affected by what one expects to perceive (Clark, 2013; De Lange, Heilbron, & Kok, 2018). Perceptual expectations can be derived from a myriad of sources (De Lange et al., 2018). This includes not only the spatial and temporal regularities in the sensory signal itself, but also the actions of the observer (e.g., when turning a steering wheel, drivers tend to expect their car to turn in the same direction). In line with this idea, it has been proposed that action preparation is accompanied by the prediction of the sensory consequences of movement (Miall & Wolpert, 1996; Wolpert & Flanagan, 2001). The neural computations involved in programming and comparing these predicted action-outcomes with actual input are thought to be an important signal for how humans experiences themselves as causal agents (Frith et al., 2000b; Moore, 2016). Hence, elucidating how motor expectations are formed and updated when interacting with a dynamic physical, and social, environment is crucial for understanding how the agentic sense of self emerges.

Expectations about upcoming events can develop across different timescales, and have been broadly divided into two categories (Seriès & Seitz, 2013). Firstly, humans form sensory expectations by exposure to relatively stable statistical regularities in the environment across their lifetime. Such long-term (structural) expectations can for instance be observed in expert piano players, who are more likely to perceive an ambiguous tone pair as going up or down in pitch after pressing keys on a keyboard in a left-right order or right-left order, respectively (Repp & Knoblich, 2007). Along similar lines, tones are generally judged as lasting longer when preceded by a movement of a longer duration (Yon, Edey, Ivry, & Press, 2017), the perceived number of visual events is biased by the number of preceding key presses (Kunde & Kiesel, 2006), and ambiguous motion percepts are perceived as moving in the same direction as accompanying bodily actions (a phenomenon known as action capture; Wohlschläger, 2000). Secondly, perception is also affected by more temporary action-outcome expectations, that are more easily updated through contextual demands (Seriès & Seitz, 2013). For instance, the perceived motion direction of ambiguous action-outcomes can be biased

by acquiring new, arbitrary, associations between actions and following effects (Kawabe, 2013; Wallis & Backus, 2016).

In everyday life humans are regularly exposed to circumstances in which contextual expectations conflict with overlearned (structural) expectations. For instance, when reversing a car with a trailer into a parking space, turning the wheel has the opposite result to regular driving conditions (i.e., turning the wheel in a rightward direction will now result in a leftward turn). To date, little is understood about how fast expectations are updated in these kinds of circumstances, and more specifically, how long-term and short-term expectations jointly determine the predicted perceptual outcome of an action. Although a few studies have provided evidence for the notion that structural expectations can be modulated by contextual expectations (Adams, Graf, & Ernst, 2004; Sotiropoulos, Seitz, & Seriès, 2011), such updating effects are not always observed (Gekas, Seitz, & Seriès, 2015). In addition, action-based expectations have hitherto largely been ignored. One important exception is a recent study on the amenability of structural motion biases (i.e., the tendency to see ambiguous motion in line with one's own movement; Wallis & Backus, 2016). However, whilst this study reports clear contextual learning effects, no consistent structural expectations were observed. This leaves open the question whether similar contextual updating would be observed in the face of more robust long-term expectations.

Another key issue that remains unclear from existing work is how expectations developed over different timescales persist over time. Specifically, existing studies do not address the question whether and how flexibly, structural expectations overtake contextual expectations after the latter expectations are no longer valid or reinforced by evidence.

Here, we present two studies in which we assessed the joint influence of structural and contextual expectations on the perception of action-outcomes, using ambiguous motion displays. In Study 4.1 participants manually operated a rotary switch, which initiated the presentation of a bistable rotating sphere that could be perceived as rotating in a clockwise or counterclockwise direction. In order to assess the interaction between structural and contextual expectations, a preceding learning phase was completed in which participants were exposed to action-motion contingencies that were either compatible (i.e., a clockwise movement resulting in a clockwise rotation of the sphere) or incompatible with structural expectations (i.e., a clockwise movement resulting in a counterclockwise rotation of the sphere). Finally, we included a baseline condition, in which test blocks were

not preceded by a learning phase, thus allowing us to assess pre-existing (i.e., structural) action-effect anticipations.

To ensure that participants' reported rotation directions genuinely reflected their perceived rotation direction (rather than a percept-unrelated response tendency), we devised a response-task that was orthogonal to our measure-of-interest: The sphere motion was divided into two successive epochs, and participants were asked to report whether or not a reversal of rotation direction had occurred. Because the rotation direction of the sphere in the second epoch was always unambiguous, we could infer the perceived rotation direction of the sphere in the first epoch (i.e., the initial percept) from the presence or absence of a reversal report, while at the same time alleviating a potential response bias.

Based on the research reviewed above, we firstly hypothesized to observe a general bias to perceive the initial rotation direction of the sphere in the same direction as one's movement (i.e., following structural expectations). On top of that we expected a tendency to perceive the sphere to rotate in the most recently learned direction (i.e., contextual expectations). Finally, we expected the strength of this contextual bias to weaken as a function of time due to extinction of the learned contingencies. In Study 4.2 we conducted a pre-registered replication of Study 4.1.

RESULTS

Study 4.1

Perceptual selection

The mean proportion of percepts that matched structural expectations in test trials (hereafter: action-consistent percepts) was submitted to a directional Bayesian one-sample t-test against chance (i.e., .50) in all three conditions (JASP version 0.9.1, default Cauchy prior width = 0.707). The proportion of action-consistent percepts exceeded chance level in the baseline ($M = 0.54$, $SD = 0.05$, $BF_{+0} = 87.94$, Cohen's $d = 0.92$), and compatible conditions ($M = 0.53$, $SD = 0.05$, $BF_{+0} = 24.70$, Cohen's $d = 0.70$), but did not deviate from chance in the incompatible condition ($M = 0.50$, $SD = 0.06$, $BF_{-0} = 0.16$, Cohen's $d = 0.08$). Subsequent two-sided pairwise comparisons with the baseline condition showed that the proportion of action-consistent percepts was smaller in the incompatible condition than in the

baseline condition ($BF_{10} = 3.00$, Cohen's $d = 0.74$). In contrast, the evidence for the *absence* of a difference between the baseline and compatible condition was about three times as large as the evidence for a difference in the proportion of action-consistent percepts ($BF_{10} = 0.36$, Cohen's $d = 0.19$). However, note that the evidence for both of these between-group differences is relatively weak. No adjustments for multiple testing were applied, as type I and type II errors do not apply to Bayesian tests (Berger, 2003; Hoijtink, Mulder, Van Lissa, & Gu, 2018). The proportion of action-consistent percepts for all conditions is depicted in Figure 4.1.

We refrained from statistically testing assumptions of normality and homogeneity of variance as these tests would have little power as a result of the relatively small sample sizes. Importantly, however, visual inspection of the data (see Figure 4.1) suggests the absence of any gross assumption violations.

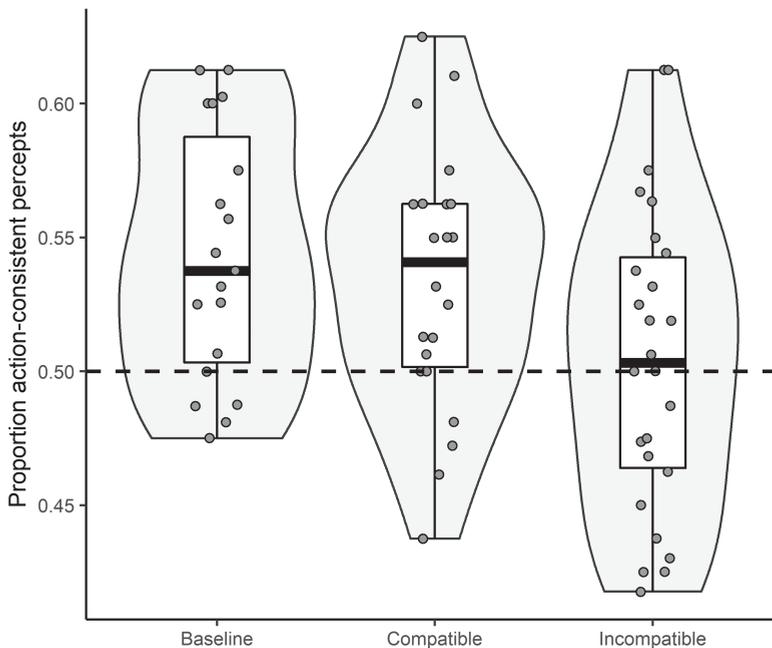


Figure 4.1. Proportion of action-consistent percepts for all three conditions in Study 4.1. Lower and upper box limit represent the 25th and 75th percentile respectively. Whiskers extend to the most extreme values that fall within 1.5 times the interquartile distance from the hinges of the box. Central lines represent the median. Data points represent individual participants.

Effect over time

In order to examine whether the observed biases would change over time, the proportion of action-consistent percepts was compared between the first and the second half of the test phase in all conditions (i.e., compatible learning, incompatible learning, and baseline). Directed paired sampled t -tests were used for the compatible and incompatible conditions as we expected potential extinction effects of the learned contingencies. Specifically, we expected the proportion of action-consistent percepts to reduce in the compatible condition and to increase in the incompatible condition (i.e., moving towards the baseline). In the baseline condition, a non-directional paired sampled Bayesian t -test was used. Unexpectedly, time effects were observed in none of the conditions (Baseline: $BF_{10} = 0.38$; Compatible: $BF_{+0} = 0.20$; Incompatible: $BF_{-0} = 0.11$, see Table 4.1).

Table 4.1. Descriptives of action-consistent percepts over time in Study 4.1

	Time 1 (first half)		Time 2 (second half)		Cohen's d
	Mean	SD	Mean	SD	
Baseline	0.53	0.06	0.56	0.08	-0.24
Compatible	0.53	0.08	0.54	0.08	-0.04
Incompatible	0.52	0.07	0.50	0.08	0.22

Note. Cohen's d represents the effect size for time 1 \neq time 2 for the baseline condition, time1 > time 2 for the compatible condition and time 1 < time 2 for the incompatible condition.

Induction trials

The propensity of participants to perceive action-outcomes in line with their expectations was also evaluated in the induction trials. In contrast to test trials, induction trials contained objective switches in the rotation direction of the sphere, and as such the perceived switches on these trials could be objectively (in)correct. Performance (in terms of the proportion of accurate trials) was near perfect for both learning conditions, and did not depend on whether the initial sphere motion was compatible (compatible learning condition: $M_{acc} = 0.98$, $SD_{acc} = 0.05$, incompatible learning condition: $M_{acc} = 0.97$, $SD_{acc} = 0.07$) or incompatible (compatible learning condition: $MD_{acc} = 0.98$, $SD_{acc} = 0.06$, incompatible learning condition: $M_{acc} = 0.97$, $SD_{acc} = 0.03$) with the preceding rotary action. Note that these accuracy levels are calculated after excluding error trials (see “Data exclusion” section) and are based on a different number of trials for the compatible and incompatible condition (see “Procedure and design section”). The near-ceiling performance demonstrates that participants understood the instructions and paid attention to the task.

Study 4.2

Perceptual selection

Our expectations for Study 4.2 followed from the observed findings of Study 4.1. We expected participants to be more likely than chance to perceive the bistable stimulus in line with structural (i.e., action-consistent) expectations in the test trials of the baseline and compatible condition. In addition, we expected the proportion of action-consistent percepts in both these conditions to be larger than in the incompatible condition. To test whether we could replicate this specific pattern of results, we formulated the following informative hypothesis (H-inf): $\mu_{\text{baseline}} > 0.5$, $\mu_{\text{compatible}} > 0.5$, $\mu_{\text{baseline}} > \mu_{\text{incompatible}}$, $\mu_{\text{compatible}} > \mu_{\text{incompatible}}$. The evidence for this hypothesis was compared against its complement H-c (i.e., not H-inf). This analysis yielded strong evidence for the expected pattern ($\text{BF}_{\text{inf,c}} = 208.62$), suggesting the overall observed pattern of Study 4.1 was replicated ($M_{\text{base}} = 0.52$, $SD_{\text{base}} = 0.04$; $M_{\text{comp}} = 0.53$, $SD_{\text{comp}} = 0.05$; $M_{\text{incomp}} = 0.49$, $SD_{\text{incomp}} = 0.04$; see Figure 4.2).

To further examine this pattern of results, several pre-registered follow-up analyses were executed. Firstly, to further evaluate the specific evidence for the presence of a structural expectation, we compared the informative hypothesis (H-inf) $\mu_{\text{baseline}} > 0.5$, against the competing hypothesis (H-c) $\mu_{\text{baseline}} = 0.5$. This analysis yielded evidence for a proportion of action-consistent percepts that exceeds chance level ($\text{BF}_{\text{inf,c}} = 4.97$, Cohen's $d = 0.69$). In addition, we observed evidence for the *absence* of a difference between the baseline and compatible condition (H-inf: $\mu_{\text{baseline}} = \mu_{\text{compatible}}$, tested against H-c: $\mu_{\text{baseline}} < \mu_{\text{compatible}}$; $\text{BF}_{\text{inf,c}} = 5.35$, Cohen's $d = 0.13$ (reflecting the effect size for $\mu_{\text{baseline}} < \mu_{\text{compatible}}$)). The evidence for a difference between the baseline and the incompatible condition went in the same direction as in Study 4.1, but was not conclusive (H-inf: $\mu_{\text{baseline}} > \mu_{\text{incompatible}}$, tested against H-c: $\mu_{\text{baseline}} = \mu_{\text{incompatible}}$; $\text{BF}_{\text{inf,c}} = 2.49$, Cohen's $d = 0.75$). Finally, similar to Study 4.1, the incompatible condition was not different from chance (H-inf: $\mu_{\text{incompatible}} = 0.5$, tested against H-c: $\mu_{\text{incompatible}} \neq 0.5$; $\text{BF}_{\text{inf,c}} = 6.61$, Cohen's $d = 0.13$ (reflecting the effect size for $\mu_{\text{incompatible}} \neq 0.5$)). No explicit assumption tests nor adjustments for multiple comparisons were performed.

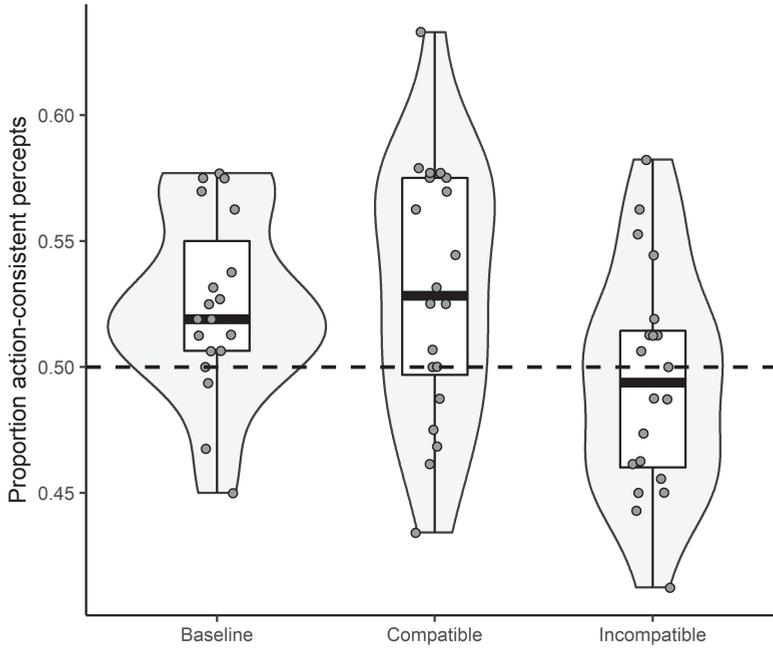


Figure 4.2. Proportion of action-consistent percepts as a function of experimental condition in Study 4.2. Lower and upper box limits represent the 25th and 75th percentile respectively. Whiskers extend to the most extreme values that fall within 1.5 times the interquartile distance from the hinges of the box. Central lines represent the median. Data points represent individual participants.

Effect over time

No evidence for a difference in the proportion of action-consistent percepts as a function of time was observed for any of the conditions in Study 4.1. To test this specific pattern of results we compared the evidence for the informative hypothesis H-inf: $\mu_{\text{baseline1}} = \mu_{\text{baseline2}}, \mu_{\text{compatible1}} = \mu_{\text{compatible2}}, \mu_{\text{incompatible1}} = \mu_{\text{incompatible2}}$, with the competing hypothesis H-c: $\mu_{\text{baseline1}} \neq \mu_{\text{baseline2}}, \mu_{\text{compatible1}} \neq \mu_{\text{compatible2}}, \mu_{\text{incompatible1}} \neq \mu_{\text{incompatible2}}$. Coinciding with the results of Study 4.1, we found evidence for the *absence* of time effect ($BF_{\text{inf,c}} = 4.26$; see also Table 4.2).

Table 4.2. Descriptives of action-consistent effects over time in Study 4.2

	Time 1 (first half)		Time 2 (second half)		Cohen's d (time1 ≠ time2)
	Mean	SD	Mean	SD	
Baseline	0.52	0.04	0.53	0.05	0.02
Compatible	0.53	0.08	0.53	0.09	0.05
Incompatible	0.48	0.08	0.51	0.05	0.25

Induction trials

Similar to Study 4.1, accuracy levels were near ceiling in both the compatible and incompatible induction blocks, and did not depend on whether the initial sphere motion was compatible (compatible learning condition: $M_{acc} = 0.98$, $SD_{acc} = 0.04$, incompatible learning condition: $M_{acc} = 0.96$, $SD_{acc} = 0.09$) or incompatible with the preceding action (compatible learning condition: $M_{acc} = 0.97$, $SD_{acc} = 0.09$, incompatible learning condition: $M_{acc} = 0.97$, $SD_{acc} = 0.05$).

DISCUSSION

Across two studies we observed that both structural and contextual expectations affect the perception of action-outcomes. Interestingly, these effects did not seem to operate in a simple additive manner. That is, the bias induced by already existing expectations was not further enhanced by compatible contextual expectations. In addition, incompatible learning abolished, but did not reverse, existing structural biases. All the observed effects were stable over time. The implications of these findings are discussed below.

The modulation of perception by action-based predictions has been examined extensively and is generally attributed to internal forward models implicated in motor control (Schütz-Bosbach & Prinz, 2007; Waszak et al., 2012). Specifically, when one prepares an action, forward models are thought to use a copy of the motor command to predict the sensory consequences of the action, which in turn influences perceptual processing (Miall & Wolpert, 1996; Wolpert & Flanagan, 2001). These forward models are generally assumed to be highly adaptable in order to allow for efficient behavior in a constantly changing environment. Indeed, research on motor adaptation suggests that people can concurrently acquire multiple internal models, and flexibly switch between these models based on task demands (e.g., Imamizu et al., 2007). To the best of our knowledge, however, it thus far remains unknown how action-based expectations that developed across different timescales are combined.

In the present studies we show that pre-existing, overlearned expectations, are quickly abolished by short exposure to disconfirming evidence. Similar updating effects have been observed in research examining the opposite side of the action-perception link. Specifically, the generally automatic (structural) tendency to imitate an observed action (Rizzolatti, Fogassi, & Gallese, 2001), can be abolished

through learning (Heyes, Bird, Johnson, & Haggard, 2005), or even be replaced by a tendency to perform complementary actions as a result of instructed task rules (Newman-Norlund, Van Schie, Van Zuijlen, & Bekkering, 2007; Van Schie, Van Waterschoot, & Bekkering, 2008). Together these findings point towards a strong top-down influence of contextual information on how the motor and perceptual system interact. However, the mechanism that underlies such contextual influences is still a matter of debate. According to the MOSAIC model, people can switch between multiple, motor-specific, internal models that exist alongside each other (Wolpert & Kawato, 1998). In contrast, more recent accounts suggest that there is only one, context-sensitive, predictive model that is not specific to motor predictions, but rather involves predictions about sensory events in general (Adams, Shipp, & Friston, 2013; Pickering & Clark, 2014). The observed flexibility in our current studies can be explained by both models, and future research is necessary to elucidate the exact processes via which these updating effects come about.

The present results are also in accordance with previous work on the updating of structural, action-based, predictions (Wallis & Backus, 2016). Similar to the present results, this study demonstrated that the pre-existing bias to perceive ambiguous motion in the same direction as one's action can be altered by learning. Unlike the present findings, however, the structural bias was enhanced by compatible learning and reversed by incompatible learning. One potential explanation for these diverging results, is that no consistent effect of structural expectations was observed in the previous study. In the absence of such a (strong) structural bias, perception might be more amendable by recently acquired contextual expectations. In contrast, the clear structural bias observed in the present study could be at the limits of the malleability of perception, restricting further enhancement by compatible learning. Likewise, this strong structural bias might have been relatively difficult to break down by incompatible learning, explaining why incompatible contextual learning was sufficient for abolishing but not for reversing the extant structural bias.

This line of reasoning is corroborated by another recent study examining the joint influence of long-term and short-term expectations on action-outcome monitoring (Wirth, Steinhauser, Janczyk, Steinhauser, & Kunde, 2018). In this study participants performed a simple ideomotor task in which one of two keys could be used to add a puzzle piece to the top or bottom of another piece presented in the center of a computer screen. Actions could either result in puzzle pieces appearing at the intended (spatially compatible) location or at an incompatible

location. Incompatible outcomes were shown to reduce response speed in a subsequently performed, unrelated task. Importantly, this interference effect was modulated by the type of expectations that participants could rely on. Specifically, in some blocks, outcomes were produced at random, which meant that monitoring could only be affected by the pre-existing (long-term) expectation that actions generally lead to compatible outcomes. In contrast, in other blocks participants were exposed to probabilistic action-outcome contingencies that were either compatible (high-compatible blocks) or incompatible (high-incompatible blocks) with this long-term bias. Whilst a strong interference effect was observed in both random and high-compatible blocks, this effect disappeared in blocks in which the majority of trials was incompatible. In other words, much like the present study, incompatible short-term expectations led to the disappearance, but not a reversal, of a pre-existing (structural) processing bias.

Contrary to our expectations, the influence of contextual learning on the perception of action-outcomes did not diminish over time. The absence of such an extinction effect might suggest that the observed influences of learning did not emerge as a result of repeated exposure to contingencies. Instead these effects might have resulted from an all-or-none adherence to the instructed task-set (i.e., the sphere will rotate in the same or a different direction as your action) that was expected to remain stable over the course of the experiment. The absence of a time effect has to be interpreted with some caution, though, as the observed evidence in Study 4.2 was not robust to variations in prior width (see Supplementary Results section in Appendix A). Nevertheless, the findings are in line with research showing temporal stability of general instruction effects on bistable perception. For instance, participants who were made to believe that a specific set of glasses would change the apparent rotation direction of a sphere, were biased to perceive a subsequent bistable sphere in the instructed direction - and continued to do so in the absence of any further evidence as long as the glasses were worn (Sterzer et al., 2008). Likewise, research on action-effect learning has shown that instructed action-outcome beliefs are remarkably resistant against disconfirming evidence after one's expectations have been consolidated by initial evidence (Pilditch & Custers, 2017). In the present study the transition of the induction to the test phase was very subtle. Similar to the aforementioned work it is therefore likely that participants held on to the beliefs that were formed based on the combination of structural expectations and initial contextual evidence (i.e., my action either does or does not cause the sphere to move in the same direction as my movement) throughout the experiment.

In conclusion, the present studies demonstrate that the perception of action-outcomes is jointly affected by long-term and short-term expectations. Structural motor expectations that are likely to be formed over the duration of one's lifetime were abolished after short exposure to incompatible action-effect contingencies. Although this malleability of perception appears to be restricted by the strength of the existing bias (i.e., no reversal effect was observed) these findings suggest a remarkable adaptability of even strong and solidified motor expectations. Considering that humans live in a constantly changing world in which perceptual expectations regularly need to be updated (e.g., when reverse parking a car with a trailer, or when required to drive on the opposite side of the road in a foreign country), this flexibility of our predictive system represents a crucial function for efficient behavior.

METHODS

Study 4.1

Sample size determination

Sample size was determined by means of a pre-specified stopping rule. This rule prescribed termination of data collection as soon as substantial evidence (i.e., a Bayes factor of 6; Jeffreys, 1961) was reached for the presence or absence of action-consistent effects, after inclusion of a minimum of twenty participants per condition (Rouder, 2014; Schönbrodt & Wagenmakers, 2018). The stopping rule for the compatible condition slightly deviated from the other conditions (see “Sample size determination” section in Appendix A). We defined action-consistent percepts as percepts that are in line with structural expectations (i.e., a clockwise sphere motion following a clockwise rotary action).

Participants

A total of seventy-one participants took part in one of three experimental conditions, which were run as sequential experiments. Twenty participants were assigned to the baseline condition ($M_{age} = 21.20$, $SD_{age} = 1.67$, three left handed, sixteen females), twenty-five participants were assigned to the compatible condition ($M_{age} = 22.52$, $SD_{age} = 2.42$, two left handed, nineteen females), and twenty-six were assigned to the incompatible condition ($M_{age} = 21.50$, $SD_{age} = 1.86$,

four left handed, twenty-four females). None of the participants took part in more than one condition. Prior to the experiment, a pre-screening for stereoscopic vision was completed (see “Pre-screening Procedure” section in Appendix A). The studies were conducted in accordance with the Declaration of Helsinki and the used paradigm is approved by the ethics board of the Faculty of Social Sciences at Utrecht University. Participants provided informed consent prior to the start of all studies and received money and/or course credit in exchange for participation.

Apparatus

Participants were seated in front of a mirror stereoscope consisting of two mirrors at a 45-degree angle, which each reflected one of two opposite linearized 23-inch LCD monitors (Dell UZ2315H; resolution: 1920 x 1080; refresh rate: 60 Hz). This set-up has been described in detail elsewhere (Brascamp & Naber, 2017). A stabilized head position and constant viewing distance of approximately 82 centimeters was ensured by a chin and forehead rest.

Stimuli

Participants were exposed to disambiguated and ambiguous visual structure-from-motion spheres. The spheres were adapted from those used in previous research (Dogge, Gayet, Custers, & Aarts, 2018), and consisted of 240 white, square “dots”, half of which moved in a leftward direction, and half of which moved in a rightward direction. The speed of the dots was 45 degrees of visual angle (dva) per second in the center and decreased to zero toward the edges of a circular aperture following a sinusoidal, thus eliciting the percept of a rotating sphere. Each dot had a lifetime of one second after which it was replaced by a new, randomly positioned dot. Initial dots were assigned a random ‘age’ between zero and one seconds to prevent simultaneous replacement of all dots. The rotation direction of the resulting structure-from-motion spheres was intrinsically ambiguous. Dots moving in the same direction are typically interpreted as belonging to a similar depth plane, and depending on whether the leftward dots or the rightward dots appear closer to the observer, the sphere is perceived as rotating leftwards or rightwards, respectively. Stimuli were presented using the Psychophysics Toolbox 3 (Brainard, 1997; Pelli, 1997) in Matlab R2016b (The Mathworks, Natick, MA).

In this study, both spheres with an ambiguous rotation direction (in test trials), and spheres with a disambiguated rotation direction (in induction trials) were used. Similar spheres were separately presented to the two eyes of the participants,

and different horizontal offsets between the eyes (i.e., disparities) were added to the leftward and rightward moving dots. This caused dots moving in one direction to appear closer to the observer than dots moving in the other direction, hence disambiguating the rotation direction of the spheres. Disparities ranged between 0 and 0.04 dva relative to fixation. To create spheres with an ambiguous rotation direction the disparity was set to 0 (i.e., the two eyes' images were identical). Intermittent presentation of ambiguous spheres is known to induce perceptual stabilization effects. Specifically, the perceptual interpretation of intermittently presented bistable stimuli is strongly biased towards the previous percept (Pearson & Brascamp, 2008), potentially causing the initial percept to persist across successive trials for more than ten minutes (Leopold & Logothetis, 1999). In order to reduce these influences, the location of both ambiguous and unambiguous spheres varied randomly between one of four quadrants around the fixation point (Chen & He, 2004).

Procedure and design

Throughout the experiment, participants initiated the rotation of a sphere by rotating a custom-made rotary switch. This switch was designed to induce a strong mapping between action and motion as the rotary movements of participants were parallel to the sphere's rotation plane.

Participants took part in one of three conditions: the baseline (i.e., no-learning) condition, the compatible learning condition or the incompatible learning condition. The baseline condition was designed to assess the influence of participants' structural action-outcome expectations on perception (i.e., the tendency to see a clockwise rotating sphere after rotating the switch in a clockwise direction). Accordingly, this condition only comprised of a test phase, in which the rotation direction of the sphere was ambiguous. Additionally, the compatible and incompatible learning conditions were designed to manipulate participants' contextual action-outcome expectations. In these two conditions, the test phase was preceded by a learning phase (40 trials), in which participants used the rotary switch to set the sphere into motion with an unambiguous rotation direction. Rotary actions were followed by sphere motion directions that were either 80% compatible or 80% incompatible with existing structural expectations. Participants were explicitly informed on the direction of these contingencies. In the subsequent test phase (80 trials), all rotary actions were followed by a sphere with an ambiguous rotation direction.

The timeline of both induction and test trials is depicted in Figure 4.3. Each trial started with the presentation of a fixation point in the center of the screen. After a random duration of 1000 to 2000 milliseconds (with intervals of 250 milliseconds) a stationary sphere appeared. As soon as the sphere appeared on the screen, the fixation point was replaced by an action-cue (a clockwise or counterclockwise arrow), which indicated the direction in which the rotary switch had to be rotated. Participants handled the switch with their dominant hand. If participants rotated the switch within the time limit of three seconds, the stationary sphere started to move. Incorrect rotations were followed by an error message. In order to motivate participants to perform the correct actions, they could earn money for correct rotations and lose money for incorrect rotations. Unbeknownst to the participants, the earned amount was rounded up and all participants were paid the same amount.

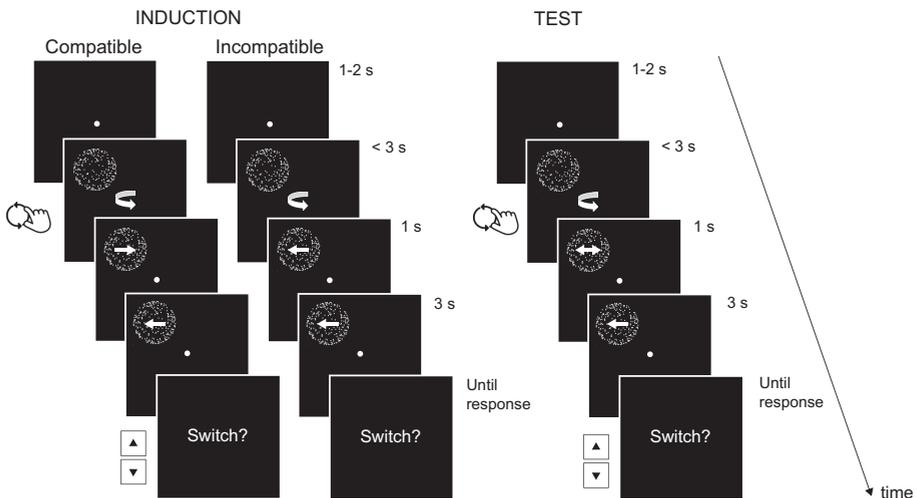


Figure 4.3. Schematic representation of trial events in the induction phase and the test phase.

In induction trials, the initial rotation direction of the sphere was either compatible or incompatible with the direction in which the rotary switch was moved by the participants. After one second, the sphere either continued to rotate in the same direction for three seconds, or switched and rotated in the opposite direction for three seconds (crucially, only the initial rotation direction was conditional upon the participant's action). Participants were asked to indicate whether they perceived a switch by pressing the "up" arrow key for "yes" and the "down" arrow

key for “no”. In test trials, the rotation direction of the sphere was ambiguous in the first second, and was unambiguous in the last three seconds (50% clockwise, 50% counterclockwise). Registering switch reports (rather than percept reports) allowed us to infer the perceived initial rotation direction, while minimizing demand characteristics (i.e., participant’s tendency to report what they think they should have seen, instead of what they actually perceived).

The task consisted of five blocks of eight induction trials (in the compatible and incompatible learning conditions) and ten blocks of eight test trials (in all three conditions). Each possible combination of the required rotation of the switch (clockwise versus counterclockwise) and the second unambiguous rotation period of the sphere (clockwise versus counterclockwise) was shown twice in each induction and test block in a randomized order. Action-outcome contingency was kept at 80/20 for the compatible condition and at 20/80 for the incompatible condition, across all induction blocks. All four sphere locations were each shown ten times in the induction phase and twenty times in the test phase. The order of the location was randomized, with the restriction that no location could be shown twice in a row. Prior to the experiment participants completed two practice rounds to get familiarized with the switch (8 trials) and to practice with the task (20 trials in the learning conditions and 8 trials in the baseline condition).

Data exclusion

Induction trials

Trials were excluded from analysis when the switch was rotated in the wrong direction (i.e., not in the direction that was indicated by the action-cue; Compatible: $M = 0.20\%$, $SD = 0.69\%$; Incompatible: $M = 0.48\%$, $SD = 2.00\%$). On top of that, we also excluded trials in which the switch was incorrectly handled, such as when it was rotated too early (before the presentation of the cue), too late (not within the response time limit), not far enough, too far, back and forth, or more than once; Compatible: $M = 0.90\%$, $SD = 1.22\%$; Incompatible: $M = 1.73\%$, $SD = 2.09\%$).

Accuracy levels for the perceived switches in motion direction were calculated over the remaining induction trials. Two participants were excluded from further data analysis for performing at (or below) chance level in the compatible condition. In addition, another participant in this same condition was excluded from analysis for scoring significantly lower than the overall group (z -score = -3.81 , calculated after excluding the aforementioned participants). The low accuracy score of these

participants suggests that they were not able to distinguish between the different disambiguated rotation directions. One additional participant was excluded in the baseline condition due to insufficient stereoscopic vision (see “Pre-screening procedure” section in Appendix A).

Test trials

Similar to the induction trials, test trials were excluded when the switch was rotated in the opposite direction as the action-cue (Baseline: $M = 0.13\%$, $SD = 0.39\%$; Compatible: $M = 0.06\%$, $SD = 0.27\%$; Incompatible: $M = 0.72\%$, $SD = 1.84\%$), and when the switch was incorrectly handled (Baseline: $M = 1.18\%$, $SD = 1.79\%$; Compatible: $M = 0.97\%$, $SD = 2.25\%$; Incompatible: $M = 1.01\%$, $SD = 2.21\%$). The reported percentages were calculated after participant exclusion.

Study 4.2

To establish the robustness of the observed results, we replicated Study 4.1 in a pre-registered study, of which the pre-registration form can be accessed at osf.io/bmxwv. The pre-registration comprises an exhaustive description of all planned data pre-processing and analyses, so that no parameters could be adjusted after data collection was completed (Wagenmakers, Wetzels, Borsboom, Van der Maas, & Kievit, 2012). Note that this replication attempt also allowed us to confirm whether the observed pattern of results in Study 4.1 was indeed caused by the experimental manipulations, rather than by pre-existing, between-group differences. Study 4.2 was identical to the previous experiment with a few exceptions. Firstly, the contingency between the participants' action and the ensuing sphere rotation was now kept at 80% *within* each individual block of the induction phase (rather than across blocks, as in Study 4.1), resulting in a fully balanced design. Secondly, evidence for the expected pattern of results was tested by means of Bayesian model selection (Hojtink, 2012), using the R package *Bain* (Gu, Mulder, & Hoijtink, 2018). This approach allowed us to compare a model that reflects our predictions for all conditions simultaneously, against a model in which our specific set of predictions are not met (see Results section for details). *Bain's* Markov chain Monte Carlo settings are chosen such that the program itself determines whether or not the sampling procedures have converged, and accurate estimates of complexity, fit and Bayes factors have been obtained (for details see: Gu, Hoijtink, Mulder, & Rosseel, 2018; Gu, Mulder, et al., 2018; Hoijtink, Gu, & Mulder, 2018).

Participants

Sixty-three participants took part in the experiment ($M_{age} = 22.27$, $SD_{age} = 2.20$, nine left handed, forty-eight females).

Data exclusion

Induction trials

Similar to Study 4.1, trials were excluded when the switch was rotated in the wrong direction relative to the action-cue (Compatible: $M = 0.57\%$, $SD = 2.67\%$; Incompatible: No errors). On top of that, trials in which the rotary switch was incorrectly handled were also excluded (Compatible: $M = 1.59\%$, $SD = 1.97\%$; Incompatible: $M = 1.67\%$, $SD = 2.54\%$).

In addition, in line with the pre-registered exclusion criteria, three participants (two in the compatible condition and one in the incompatible condition) were replaced before data analysis for showing suboptimal accuracy in perceiving switches in motion direction ($< 80\%$) in the induction phase. For practical reasons these participants were not replaced immediately, as described in the pre-registration. Instead, the participants were replaced in a random order after running twenty participants per condition. This resulted in a final sample of twenty participants in each condition. After completing data collection, one additional participant was excluded from the baseline condition for being an outlier according to pre-registered criteria (i.e., the mean proportion of action-consistent percepts exceeded 1.5 times the interquartile range from the 1st quartile for the baseline condition).

Test trials

Trials were excluded when participants rotated the switch in the opposite direction of the action-cue (Baseline: $M = 0.59\%$, $SD = 1.21\%$; Compatible: $M = 0.38\%$, $SD = 1.68\%$; Incompatible: $M = 0.38\%$, $SD = 1.00\%$). In addition, trials in which the rotary switch was incorrectly handled were also excluded (Baseline: $M = 0.72\%$, $SD = 1.05\%$; Compatible: $M = 1.38\%$, $SD = 1.57\%$; Incompatible: $M = 1.00\%$, $SD = 0.96\%$). The reported percentages were calculated after participant exclusion.

Code availability

Code for data pre-processing and analysis can be accessed at <https://osf.io/9xq6v/>. Matlab code for running the experiment is available upon reasonable request from M.D.

Data availability

Data to reproduce the reported results can be accessed at <https://osf.io/9xq6v/>. This includes data to reproduce Figure 4.1 and 4.2.

CHAPTER 5

Cortical information flow during inferences of agency

This chapter is based on: Dogge, M., Hofman, D., Boersma, M., Dijkerman, H. C., & Aarts, H. (2014). Cortical information flow during inferences of agency. *Frontiers in Human Neuroscience*, 8, 609. <https://doi.org/10.3389/fnhum.2014.00609>

ABSTRACT

Building on the recent finding that agency experiences do not merely rely on sensorimotor information but also on cognitive cues, this exploratory study uses electroencephalographic recordings to examine functional connectivity during agency inference processing in a setting where action and outcome are independent. Participants completed a computerized task in which they pressed a button followed by one of two color words (red or blue) and rated their experienced agency over producing the color. Before executing the action, a matching or mismatching color word was pre-activated by explicitly instructing participants to produce the color (goal condition) or by briefly presenting the color word (prime condition). In both conditions, experienced agency was higher in matching vs. mismatching trials. Furthermore, increased electroencephalography (EEG)-based connectivity strength was observed between parietal and frontal nodes and within the (pre)frontal cortex when color-outcomes matched with goals and participants reported high agency. This pattern of increased connectivity was not identified in trials where outcomes were pre-activated through primes. These results suggest that different connections are involved in the experience and in the loss of agency, as well as in inferences of agency resulting from different types of pre-activation. Moreover, the findings provide novel support for the involvement of a fronto-parietal network in agency inferences.

INTRODUCTION

Humans generally feel in control of their actions and the events that follow from them. This sense of agency plays a key role in self-awareness as well as social interaction (Haggard & Tsakiris, 2009; Ruys & Aarts, 2012). Although experiences of self-agency arise naturally in most individuals, abnormalities in agency processing, such as feeling in control over externally generated outcomes, or, oppositely, experiencing a loss of control over outcomes that one did produce, have been observed in a variety of psychiatric and neurological disorders (Blakemore, Wolpert, & Frith, 2002). Examining the neural substrates underlying self-agency in the healthy brain thus is an important step to comprehend the origin of disturbed agency experiences, and eventually uncover possible ways to alleviate them.

The experience of agency has primarily been studied from the perspective of comparator models that are part of the motor control system (Frith et al., 2000a). These models often rely on paradigms in which visual, tactile, or auditory feedback of the participants' action is manipulated (e.g., David, 2012; Sperduti, Delaveau, Fossati, & Nadel, 2011). According to the comparator model, the execution of an action is accompanied by the prediction of sensory action-outcomes based on internal copies of movement-predicting signals (i.e., efference copies) generated by the motor system. Because internal motor predictions are generally fast and reliable, sensory outcomes are readily perceived as self-produced when these predictions correspond with the actual outcome (Frith et al., 2000a). This motor prediction process of agency has been found to be associated with brain activity in various areas, including the superior temporal gyrus, the inferior parietal lobe, as well as motor regions such as the pre-supplementary motor area and the cerebellum (for an overview, see: Sperduti et al., 2011).

According to the comparator model, experiences of agency are less likely to occur when the motor system cannot produce an efference copy (i.e., when acts are not self-generated) or when these signals are weak or noisy, such as when there is no clear causal relationship between an action and an effect. However, recent research has demonstrated that people can feel in control over externally generated events (Wegner, Sparrow, & Winerman, 2004) and in the absence of high action-effect contingency (Moore, Lagnado, Deal, & Haggard, 2009; Van der Weiden, Aarts, & Ruys, 2011). These findings strongly suggest that agency experiences can also emerge via a different route. This alternative route, specified by the inference model, involves cognitive inferences of the correspondence

between action outcomes and prior activation of information about the outcome (Wegner, 2002). Despite the role of these inference processes in the emergence of agency (Moore et al., 2009; Sato, 2009; Synofzik et al., 2008, 2013), their neural basis has hitherto received relatively little empirical attention. Another issue that remains unclear from prior work is how brain regions associated with agency interact and exert influence over each other. The present study builds on recent advancement in the quantification of neural communication to examine the interactions between cortical regions during inferences of agency.

Agency inferences

The inference model proposes that upon observing an event, people determine whether or not it has resulted from their actions by comparing the outcome with prior activated information or thoughts about action-effects. If there is a match, they ascribe the action-outcome to themselves, whereas, in case of a mismatch, the effect is ascribed to an external cause. Although this account involves a predictive element regarding action-outcomes similar to the comparator model, the prior expectations specified in this model only minimally depend on motor signals. Instead, these expectations pertain to cognitive priors such as intentions and beliefs (Synofzik et al., 2013). Moreover, even though predictive elements are involved in inference processes, the critical information is provided by the action outcome (Synofzik et al., 2013).

It is important to note that inferences of agency are normally thought to result from intentions. That is, if an intention to produce a certain outcome matches the actual sensory consequences following one's action, people tend to experience causal responsibility for these consequences, whereas if the intention mismatches with the observed outcome, a reduced sense of agency is experienced (Wegner, 2002). Intriguingly, however, recent research suggests that prior knowledge regarding action outcomes does not necessarily need to be explicitly activated for agency inferences to occur, but can also consist of outcome primes as a source of agency (Aarts et al., 2005; Belayachi & Van der Linden, 2010; Dannenberg, Förster, & Jostmann, 2012; Jones, de-Wit, Fernyhough, & Meins, 2008; Linser & Goschke, 2007; Ruys & Aarts, 2012). This evidence possibly accounts for the emergence of experienced agency in everyday situations where people do not produce action-outcomes themselves or lack awareness of the actual causes of their behavior.

Although goals and primes give rise to similar inferences of agency, there is some preliminary evidence to suggest that the two sources produce qualitatively different effects (Van der Weiden et al., 2013). Specifically, pursuing a goal instigates a control process that causes people to focus on the specific outcome one wants to reach and, at the same time, to inhibit all other possible outcomes (Aarts, 2012; Fisbach & Ferguson, 2007; Förster, Liberman, & Friedman, 2007). Consequently, inference processes based on goals are very specific and reliable in the detection of deviations from intended outcomes. These goal-directed control processes are less likely to occur in case of outcome priming, because outcome priming is assumed to merely enhance the accessibility of the outcome representations and other information associated with it (Van der Weiden et al., 2013). This implies that agency inferences based on priming are less sensitive to deviations and hence have a noisier processing mechanism than goal-directed processes (Van der Weiden et al., 2013). Based on these qualitative differences between inferences resulting from goals and primes, we not only examined the neural communication between cortical regions underlying goal-based inferences but, for exploratory purposes, also investigated these neural processes during prime-based inferences.

Neural communication and agency inferences

Cognitive functioning, including inferences of agency, is dependent on the integration of information within and between functionally specialized brain sites (Stam & Van Straaten, 2012a; Varela, Lachaux, Rodriguez, & Martinerie, 2001). There is increasing agreement that this integration, or more precisely, the communication between neurons, arises from synchronization of neural activity (Buzsaki & Draguhn, 2004; Fries, 2005; Salinas & Sejnowski, 2001; Sauseng & Klimesch, 2008; Schnitzler & Gross, 2005; Varela et al., 2001). Specifically, neurons' responsiveness has the property to oscillate, referring to fluctuations in excitability of their membrane potential (Buzsaki & Draguhn, 2004). These fluctuations create time windows in which a neuron is most responsive to signals by other neurons (Buzsaki & Draguhn, 2004). Hence, for two neurons to successfully exchange information, their excitability period needs to be aligned, which happens whenever they oscillate in phase (Buzsaki & Draguhn, 2004; Fries, 2005). In contrast, when phase synchronization between the oscillations of two neurons is absent, their communication is inhibited (Fries, 2005). Accordingly, the neural networks underlying agency processing can be studied by examining the

synchronization of neural activity, and thus the exchange of information between local and distant groups of neurons (Varela et al., 2001).

Two recent studies using functional magnetic resonance imaging (fMRI) have provided some insights into the neural networks underlying agency processing. In one study, participants were asked to indicate perceived control over actions based on congruent or incongruent movement feedback (David et al., 2007). Increased connectivity was observed between the pre-motor cortex, cerebellum, and posterior parietal cortex (PPC) when movements were correctly identified as externally generated, and between the insula and somatosensory cortex when movements were correctly classified as self-generated (David, 2012; David et al., 2007). In another study, leading and lagging networks were identified during experiences of loss of control in response to incongruent visual feedback (Nahab et al., 2011). The leading network consisted, among others, of the inferior parietal lobe and the insula and was shown to send information to a lagging network consisting of several areas in the posterior parietal and prefrontal lobe. The authors interpreted the leading network as being involved in the comparison of motor predictions with actual effects, whereas the lagging network (in particular the pre-frontal lobe) was thought to be responsible for the translation of the outcome of this comparison into higher order processing of agency, such as the conscious awareness of this experience.

Although the aforementioned connectivity studies provide a first glimpse into neural networks underlying experiences of agency and to the direction of information flow between them, they deal with agency processes informed by motor predictive signals, and not by cognitive inferences processes per se. A recent fMRI study addressed this notion by examining the neural substrates of goal-based agency inferences (Renes, Vermeulen, Kahn, Aarts, & Van Haren, 2013). During the ascription of outcomes to oneself, activation was observed in the inferior parietal lobe, the superior frontal cortex and the medial prefrontal cortex, implying that the lagging network identified by Nahab and colleagues (2011) might indeed be involved in the previously mentioned higher order agency processing.

The present study

In the present study we further examine and extend these findings by analyzing the pattern of information flow during inferences of agency using measures of (directed) phase synchronization. By doing so, we not only build on recent calls

for a shift from localization to network perspectives on agency processing (David, 2012), but also expand prior work on the connectivity underlying the sense of agency by employing a more direct measure of neural communication.

To explore the cortical interactions underlying inferences of agency we used an action-outcome task in which participants perform an action (pressing a key) that is followed by a sensory effect (the color word red or blue presented on the computer screen) that either matches or mismatches with pre-activated knowledge of this outcome. After observing the outcome, self-agency over producing the outcome is reported. Importantly, participants learn that the outcome they observe is not always caused by their actions but can be determined by the computer as well. As a consequence, sensorimotor predictive processes are unreliable in this task, allowing us to pinpoint agency experiences that are informed by inferences. Furthermore, pre-activation of knowledge about outcomes is manipulated by explicitly instructed goals to produce the outcome or by briefly presented primes of the outcome, thus allowing us to study goal-based and prime-based agency inferences.

To examine the neural communication of agency inferences we used the electroencephalogram (EEG), which has a temporal resolution that is sufficient to non-invasively examine phase synchronization (Sauseng & Klimesch, 2008; Stam & Van Straaten, 2012a). Based on prior work we are particularly interested in coupling strength between parietal and frontal regions and the direction of information flow between them.

METHODS

Participants

Thirty right-handed participants ($M_{\text{age}} = 21.03$, $SD_{\text{age}} = 3.20$; 22 females) who indicated no current neurological condition, mental illness or use of psychiatric medication took part in the experiment. Participants were asked to refrain from the consumption of caffeine 3 hours prior to the experiment. All participants received course credit or a monetary reward in exchange for their participation. The study received approval from our internal faculty board (Social and Behavioral Sciences) at Utrecht University. Furthermore, written informed consent of each participant was obtained.

Agency inference task

The agency inference task was adapted from Renes, Van Haren and Aarts (2015). Similar to playing a slot machine, this task required participants to stop a sequence of rapidly presented information to produce a particular outcome (i.e., the color word red or blue) on the computer screen. Specifically, participants pressed a key in response to a cue while viewing alternating letter strings. Upon pressing this key, the stream of letter strings stopped and the color word “red” or “blue” was presented. This outcome could either match or mismatch with prior knowledge regarding the action-effect (i.e., goals or outcome primes; see below). In addition, participants learned that the computer could have caused the presented outcome as well. In other words, the cause of the observed effect was ambiguous (Aarts et al., 2005; Sato, 2009). After viewing the sensory effect following their key press, participants reported experienced agency over causing the perceived effect.

Each trial consisted of five different phases: an exposure phase, a filler interval, an action phase, an outcome phase and a rating phase (see Figure 5.1). The last four phases were identical for all trials. During the filler interval, participants attended to rapidly alternating letter strings. This interval served as a delay between exposure to pre-activated information and the action that was also present in previous work on agency inferences (Van der Weiden et al., 2013). In the action phase, participants responded to a circle (the letter “o” presented in Arial 24 pt. at an approximate visual angle of 2.10°) that appeared above or below the letter strings, by pressing the corresponding upper or lower key on a response box with their right index finger. This action cue was included to ensure that participants paid attention to the outcome prime or goal presented amidst of the letter strings. The interval in which a response could be given lasted 800 milliseconds. If participants pressed the key within this interval, the strings continued to alternate until the end of a 960-millisecond lasting interval, whereas if they pressed too late, an error message occurred and the trial was processed as missing.

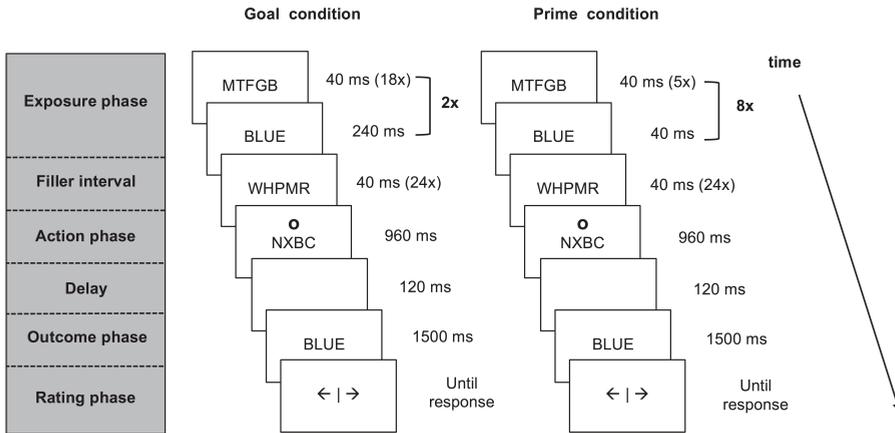


Figure 5.1. Schematic presentation of a match trial in the agency inference task for the goal condition and the prime condition. Both goal and prime trials start with the pre-activation of a color word that is presented within a stream of letter strings. In the goal trials participants are instructed to produce this outcome. In the prime trials participants are merely exposed to the prime words. After a short interval participants press a key in response to an action cue appearing above or below the letter strings. Upon this key press the stream of information stops and a color word matching or mismatching the pre-activated word is presented. Participants are asked to report experienced self-agency over this outcome.

Following the action phase, the color word “red” or “blue” (counterbalanced between trials) was shown for 1500 milliseconds, after a short delay of 120 milliseconds. To ensure that participants would maintain looking at the letter strings, participants were told that pressing the key during the presentation of a string containing the letter R (e.g., MWRT) would cause the word “red” to appear, whereas a key press during the presentation of a string containing the letter “B” (e.g., BTSZW) was followed by the word “blue.” In reality, the computer determined the presentation of color words.

After each trial, experienced agency was assessed during a rating phase by asking participants to what extent they felt their key press caused the presented color word to occur. They could respond by moving a square on a 9-point analog scale ranging from the Dutch word “niet” (in this context roughly corresponding to: “Not at all” to “wel” (“Very much”). The square had to be moved at least one position to the left or the right of the scale, starting in the center (i.e., answer “5”). This caused the data to consist of split responses (i.e., data ranging from 1 to 4 and 6 to 9). In order to form a continuous scale ranging from 1 to 8, the agency ratings were recoded (i.e., 9 = 8, 8 = 7, 7 = 6, and 6 = 5).

Pre-activated knowledge about outcomes

As mentioned earlier, the exposure phase was not identical for all trials. Specifically, in this phase knowledge regarding the outcome was activated by either goals or by primes.

In goal trials, participants were exposed to a series of letter strings followed by a color word that was clearly presented on the screen for 240 milliseconds. This sequence was repeated twice (see Figure 5.1), using the same color word. Participants were instructed to form the goal to produce the color word that appeared within the series of letter strings.

In outcome prime trials, participants viewed five random letter strings followed by a briefly presented color word (40 ms). This sequence of events was repeated eight times, resulting in a total of eight identical primes during a 1920 millisecond period (see Figure 5.1). Importantly, participants were not instructed to formulate a goal in the prime trials.

Note that, in contrast to prior studies (Van der Weiden et al., 2013), the duration and moment of the exposure phase was identical for both types of pre-activated outcome information. Accordingly, differences between prime and goal-based inferences could be examined in a more controlled manner.

The goal trials and outcome prime trials were presented in two separated blocks which each consisted of 64 randomly presented trials. All participants started with the prime condition to prevent transference of instructions from the goal condition to the prime condition (i.e., to prevent participants from using the primed information to form a goal). In half of the trials, pre-activated color words corresponded with the actual outcome, whereas in the other half of the trials they did not correspond with this outcome. Participants practiced for both blocks before the onset of the experiment (eight trials per practice block). After completing these practice trials participants completed the outcome-priming block, followed by a practice block for the goal condition (four trials) and the actual goal block. In between the two blocks participants were allowed to have a break. In addition, participants paused for 30 s after completing the first half (i.e., 32 trials) of each block.

EEG recording and pre-processing

EEG was recorded at a sampling rate of 2048 Hz during the entire agency inference task from 32 electrodes positioned according to the international 10/20 system using the BioSemi Active Two EEG system (BioSemi). The Electro-

oculogram (EOG) was measured from electrodes placed on the suborbit and supraorbit of the right eye and on the outer canthi of both eyes. Raw EEG data was band pass filtered offline (0.5–50 Hz) with a roll-off of 48 dB/oct and a 50 Hz Notch filter. Time series were re-referenced against an average reference. In order to correct for eye movements, Gratton and Cole's method (Gratton et al., 1983) was used. A semi-automated artifact correction tool (Brain Vision Analyzer software package; Version 2.0), allowing a maximum difference of 50 $\mu\text{V}/\text{ms}$, was employed to detect further artifacts. The corrected data was chunked down to 128 trial-specific segments that started at the onset of the outcome presentation and ended after 1000 milliseconds. This time window corresponds to the interval of interest used in prior work on the neural basis of agency inferences (Renes et al., 2013).

Functional connectivity

EEG was employed to assess both bidirectional and directional neural communication during agency inferences; quantified by the phase lag index (PLI; Stam, Nolte, & Daffertshofer, 2007) and directed phase lag index (dPLI; Stam & Van Straaten, 2012b) respectively.

Phase lag index

The PLI identifies statistical interdependency of two time series based on the level of asymmetry of the distribution of their phase differences (for mathematical details see: Stam et al., 2007). Since the PLI only reflects correlations between signals of which the phase difference deviates from zero, it is less affected by common source problems and amplitude changes than other connectivity measures (however, see: Muthukumaraswamy & Singh, 2011). The PLI ranges from 0 to 1, with a score of zero indicating no coupling or coupling that might result from common source problems, and a score of 1 indicating perfect coupling.

BRAINWAVE software (version 9.75) was used to compute the instantaneous phase (using a Hilbert transformation) and PLI between all pairs of electrodes for each trial in the broadband (2–50 Hz), delta band (2–4 Hz), theta band (4–8 Hz), alpha band (8–12 Hz), beta band (13–30 Hz), and gamma band (30–40 Hz). By doing so, trial specific 32x32 connectivity matrices were created. Given that the present study aims to examine functional connectivity associated with agency experiences emerging from inferences, rather than connectivity as a function of

task conditions, we decided to examine the low vs. high agency contrast within each task condition (i.e., as a function of matching and type of pre-activation). In line with prior research (Renes et al., 2013), the aforementioned 32 x 32 trial matrices were sorted into two groups based on agency ratings (Low agency: rating 4, High agency: rating 5). The frequency distributions of agency ratings for matching and pre-activation cells are presented in Supplementary Figure S5.1 (Appendix B). To allow for group comparison, average matrices were created for each possible combination of type of pre-activation, matching and level of agency. This resulted in eight average 3D group matrices comprising PLI values for all possible pairs of electrodes per participant.

Nonparametric permutation tests adapted from (Boersma et al., 2012) were used to test for differences in PLI for all possible electrode pairs between low agency and high agency for match and mismatching conditions (i.e., low agency vs. high agency for goals matching the outcome, low agency vs. high agency for goals mismatching the outcome, low agency vs. high agency for primes matching the outcome and low agency vs. high agency for primes mismatching the outcome). These tests involved a resampling method with replacement, which was used to generate ten thousand random pairs of groups from the two originally specified observations (i.e., low and high agency), across participants²⁰. By comparing the mean PLI values for all electrode pairs between these random groups, a distribution of differences for all pairwise connections was created. The position of the original difference value in this distribution was used to determine p-values for each contrast. Significant differences (alpha 0.05) were visualized using a modified version of the topoplot function in the EEGLab toolbox (Delorme & Makeig, 2004). Specifically, networks were plotted in which each node is represented by an EEG electrode and the links between the nodes correspond to a significant difference in connectivity between low and high agency.

Directed phase lag index

dPLIs were calculated to examine the direction of information flow of pairwise connections. Similar to the PLI the dPLI is a measure of the asymmetry of the

20 The authors are aware that the within-subject nature of the data violates the exchangeability assumption of the permutation tests and thus increases the likelihood of false positives. However, given that the design of the study leads to varying numbers of trials in the low and high agency condition (precluding within-subject permutation), as well as our aspiration to visualize networks, this analytical procedure was deemed most appropriate. Nevertheless, caution is advised when interpreting the results.

distribution of phase differences of two signals (Stam & Van Straaten, 2012b). However, dPLI also assesses the direction of the asymmetry (i.e., the probability that the phase of the signal measured at electrode X is smaller than the phase of the signal measured at electrode Y), whereas PLI merely determines the presence of absolute asymmetry. The direction of the asymmetry allows one to infer whether a signal recorded from a node is phase leading (i.e., sending information) or phase lagging (i.e., receiving information) compared to the signal recorded from all the other nodes (Stam & Van Straaten, 2012b). Specifically, time series measured from a node with a dPLI score larger than 0.5 are thought to be leading in phase, whereas a dPLI score smaller than 0.5 indicates the opposite pattern (Stam & Van Straaten, 2012b). In the present study a modified version of the BRAINWAVE software (version 9.70) was used to assess directional connectivity between pairs of nodes.

For each trial, dPLI matrices for all electrode pairs and average group matrices corresponding to each possible combination of matching, type of pre-activation and agency were constructed. These average matrices were used to obtain dPLI values for all participants for each connection that significantly differed in PLI between groups. Exploratory one-sample t-tests were used to examine whether dPLI values of the connections significantly differed from 0.5. Corrections for multiple comparisons were made by means of Benjamini and Hochberg's (1995) false discovery rate procedure. These analyses were performed using SPSS (version 20).

RESULTS

Behavioral data

Agency ratings

One hundred and eleven trials (2.89% of the total amount) were excluded from the analyses due to the absence of a key press within the interval of the action phase. Mean agency ratings were calculated for matches and mismatches in the goal trials and in the prime trials. Visual inspection of the data as well as normality tests indicated non-normality of the data. However, considering the robustness of ANOVA for these departures from normality, we refrained from the use of non-parametric alternatives. The mean ratings were submitted to a 2 (type of pre-activation: goal vs. prime) \times 2 (matching: mismatch vs. match) repeated

measures ANOVA. This analysis yielded a main effect of matching, $F(1, 29) = 13.06, p = .001, \eta_p^2 = .31$, indicating higher agency experiences when pre-activated outcome information was consistent as opposed to inconsistent with the actual outcome. Moreover, an interaction between type of pre-activation and matching was observed, $F(1, 29) = 5.39, p = .03, \eta_p^2 = .16$. The main effect for type of pre-activation was not significant, $F(1, 29) = 1.25, p = .27, \eta_p^2 = .04$.

To gain further insight into the interaction, simple main effects using Bonferroni correction (corrected alpha 0.0125) were calculated. These analyses yielded higher agency ratings for matching vs. mismatching in both the goal, $F(1, 29) = 11.36, p = .002, \eta_p^2 = 0.28$, and outcome priming condition, $F(1, 29) = 10.74, p = .003, \eta_p^2 = 0.27$. A marginally significant simple main effect of type of pre-activation was observed within match trials, $F(1,29) = 6.66, p = .02, \eta_p^2 = 0.19$, but not in mismatch trials $F(1,29) = 2.22, p = .15, \eta_p^2 = 0.07$. The means of the cells are depicted in Figure 5.2.

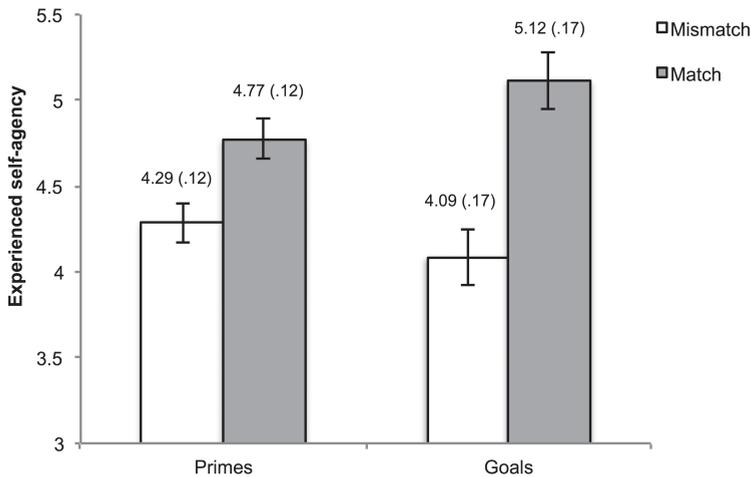


Figure 5.2. Reported experience of agency as a function of type of pre-activation and matching. The numbers above the bars represent the mean and standard error (also reflected by the error bars) of the corresponding condition

Key-press reaction times

To check whether participants responded differently to the action cue by pressing the key as a function of the type of trial, mean reaction times were submitted to a 2 (type of pre-activation: goal vs. prime) x 2 (matching: mismatch vs. match) repeated measures ANOVA. This analysis yielded a non-significant main effect of type of pre-activation, $F(1, 29) = 2.68, p = .11, \eta_p^2 = .09$, indicating no difference

in reaction time between goal trials ($M = 436.25$, $SE = 11.85$) and prime trials ($M = 449.28$, $SE = 9.75$). In addition, no difference in reaction time was observed between match ($M = 445.43$, $SE = 9.77$) and mismatch trials ($M = 440.10$, $SE = 10.69$), as evidenced by a non-significant main effect of matching, $F(1, 29) = 2.38$, $p = .13$, $\eta_p^2 = .08$. Finally, the interaction effect between type of pre-activation and matching was not significant, $F(1, 29) = 0.05$, $p = .83$, $\eta_p^2 = .002$.

Agency rating times

The time participants took to report experienced agency was also assessed by submitting mean rating times (in milliseconds) to a 2 (type of pre-activation: goal vs. primes) \times 2 (matching: mismatch vs. match) repeated measures ANOVA. Although the data was non-normally distributed, we refrained from using non-parametric alternatives for previously mentioned reasons. Participants reported experienced agency faster in goal trials ($M = 1470.69$, $SE = 129.57$) than in prime trials ($M = 1669.35$, $SE = 131.58$), $F(1, 29) = 6.51$, $p = .02$, $\eta_p^2 = .18$. The differences in reaction time between mismatch trials ($M = 1554.70$, $SE = 122.13$) and match trials ($M = 1585.33$, $SE = 130.79$), $F(1, 29) = .49$, $p = .49$, $\eta_p^2 = .02$, as well as the interaction effect between type of pre-activation and matching, $F(1, 29) = 1.95$, $p = .17$, $\eta_p^2 = .06$, were non-significant.

In short, the behavioral data shows two notable findings. First, participants report higher agency experiences when the observed effect matches vs. mismatches with pre-activated outcome information. This effect tends to be more pronounced in case of goal-based agency inferences than in case of prime-based agency inferences. Moreover, participants provided faster ratings concerning their feeling of agency in goal trials as opposed to prime trials.

EEG data

Data exclusion

Visual data inspection led to the detection of noisy data on one or more channels for five participants. These participants were excluded from further EEG analyses to retain the option of analyzing all 32×32 channel pairs. In addition, four participants had no trials left in one or more cells that were created by splitting the data in low and high agency ratings; these participants were also excluded. Hence, the total sample for EEG analysis consisted of 21 participants ($M_{\text{age}} = 21.43$, $SD_{\text{age}} = 3.37$; 17 females). In addition, 1.67% of the trials were

excluded based on semi-automated visual artifact rejection. Finally, trials that were characterized as missing in the agency inference task (i.e., trials in which the key was not pressed within the action interval) were omitted from analyses (2.75%)²¹.

Connectivity

Figure 5.3 provides an overview of connectivity for the contrast between low and high agency as a function of matching and type of pre-activation. Re-running the permutation analyses can result in marginal variation in the null distribution of mean differences. As a result, inclusion of connections with PLI differences near the significance threshold (alpha 0.05; two-tailed) is subject to similar variation. Dashed lines ($.02 \leq p \leq .03$) are used to discriminate these connections from those that are more robust over different runs (i.e., solid lines; $p < .02$).

Goal trials

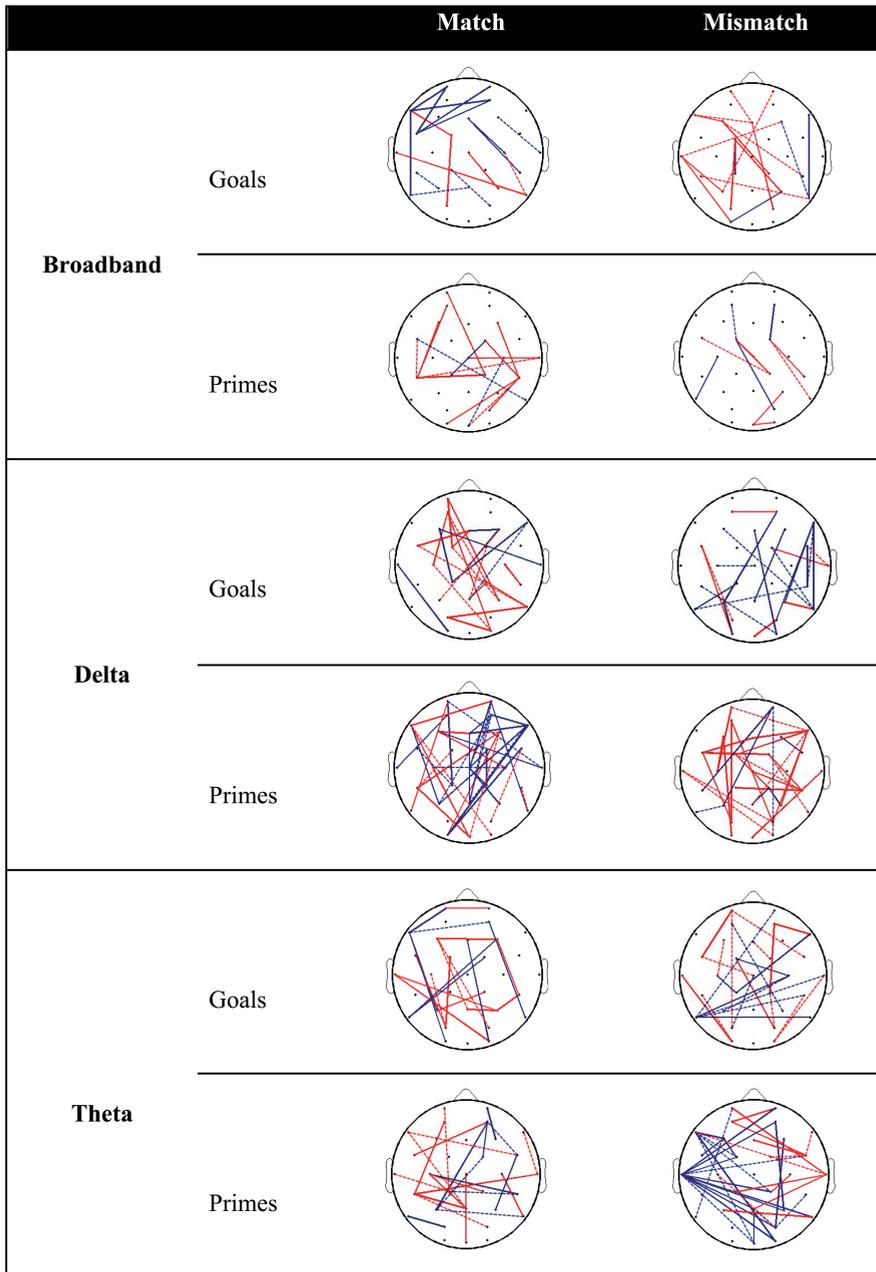
The behavioral data suggests that people experience more self-agency when a goal matches the observed outcome vs. when it does not. In other words, matches are more likely to be associated with high agency, whereas mismatches are associated with low agency²². When examining connectivity associated with high agency (vs. low agency) experiences in trials in which goals match with the outcome (Figure 5.3), increased connectivity is observed between parietal and frontal regions as well as within the frontal cortices in the broadband. With regard to specific frequency bands, high agency experiences during match trials seem particularly governed by increased connectivity in the beta band.

Different connections emerge during low-agency experiences in trials in which goals mismatch with the outcome. Specifically, in the broadband frequency increased connectivity for low agency experiences (compared to high agency experiences) is observed within and between parietal and frontal areas. In

21 Although the analyses of behavioral data and EEG data differ (in the sense that different contrasts are assessed), we checked whether the reported findings for behavioral data change when excluding the artifact trials and participants that are excluded from the EEG analysis. The repeated measures ANOVA yielded a main effect of matching, $F(1, 20) = 8.36, p = .009, \eta_p^2 = .29$, showing that matching outcomes corresponded with higher agency experiences than mismatching outcomes. The main effect of type of pre-activation, $F(1, 20) = .02, p = .89, \eta_p^2 = .001$ and the interaction between type of pre-activation and matching, $F(1, 20) = 1.27, p = .27, \eta_p^2 = .06$ were not significant.

22 Note that the behavioral data indicates that participants sometimes report low agency in a match trial and high agency in a mismatch trial. These experiences, and corresponding connectivity, are not likely to reflect agency processes emerging from prime-based or goal-based inferences—and are thus not explicitly discussed in the result section.

addition, this increased connectivity can particularly be observed in alpha and gamma bands.



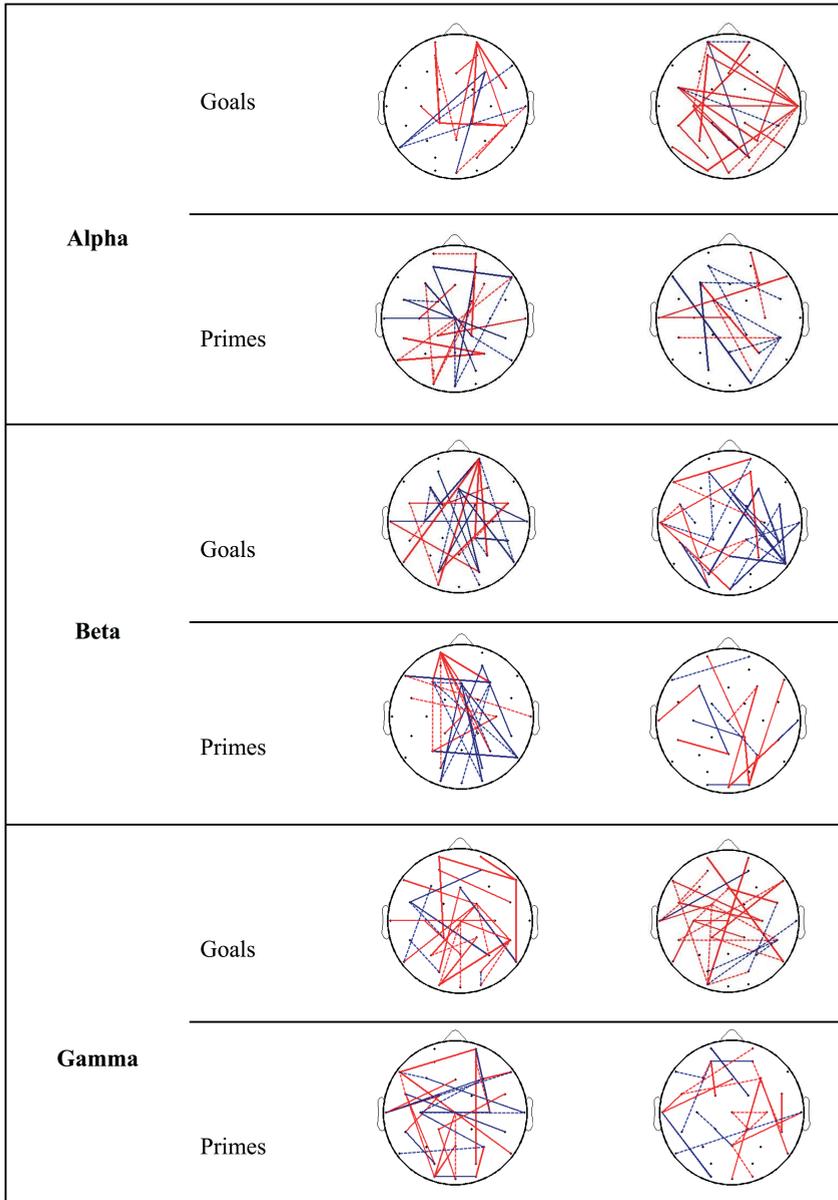


Figure 5.3. Significant differences in PLI between low and high agency as a function of matching and type of pre-activation across frequency bands. Red lines indicate that $PLI_{low\ agency} > PLI_{high\ agency}$, whereas blue lines represent the opposite pattern. Dashed lines represent PLI differences near the significance threshold ($0.02 \leq p \leq 0.03$) whereas solid lines represent connections that are more robust across re-runs of the permutation analysis ($p < 0.02$).

Prime trials

During experiences of high agency (vs. low agency) in trials in which primes match with the observed effect, increased connectivity can be observed between parietal and frontal regions in the broadband. Increased fronto-parietal connectivity is also present to a larger extent in the delta, theta, alpha, and beta band.

Reports of low agency (as opposed to high agency) during primed mismatch trials are associated with enhanced coupling between parietal and frontal areas. With regard to specific frequency bands, increased connectivity between parietal, and frontal regions during experiences of low agency (vs. high agency) is especially apparent in the delta band.

Direction of information flow

The connectivity pattern that was observed in the broadband during high agency experiences in trials in which goals matched action-effects, is in line with previous findings on the neural basis of agency (e.g., Nahab et al., 2011). To explore whether the information flow between the identified nodes is also consistent with prior work (i.e., directed from parietal to frontal lobes), directed phase lag indices were calculated for all conditions in this frequency range (see Table 5.1 for mean PLI values). As can be seen in Table 5.2, the signal measured from the left parietal electrode is leading in phase compared to the signal at the left frontal electrode in trials in which goals match the outcome, suggesting that there is a trend of anteriorly directed information flow in these trials. In the other conditions no clear direction of information flow could be observed. It should be noted that the reported effects are not corrected for multiple comparisons. After implementing this correction, none of the dPLI values were different from 0.5 at the conventional significance level of $p < 0.05$.

Table 5.1. Mean PLI during high (HA) and low experiences of agency (LA) in trials in which (A) goals matched the outcomes, (B) goals mismatched the outcomes, (C) primes matched the outcomes, and (D) primes mismatched the outcomes in broadband frequency.

	Connection	PLI (LA)		PLI (HA)	
		<i>M</i>	<i>SD</i>	<i>M</i>	<i>SD</i>
(A)					
	PO3_FC1	0.18	0.04	0.15	0.03
	FC1_F7	0.16	0.04	0.14	0.03
	P7_F7	0.12	0.03	0.15	0.03
	F7_AF4	0.11	0.04	0.14	0.04
	FC5_FP1	0.12	0.03	0.14	0.03
	FC5_FP2	0.11	0.03	0.14	0.03
	FC5_AF4	0.11	0.04	0.14	0.03
	F7_FP1	0.11	0.03	0.14	0.04
	CP6_Fz	0.12	0.03	0.15	0.02
(B)					
	PO4_Fz	0.16	0.04	0.13	0.03
	PO3_FC1	0.15	0.04	0.12	0.04
	P4_FC1	0.16	0.03	0.13	0.04
	CP2_F3	0.16	0.03	0.13	0.04
	F3_F7	0.15	0.04	0.12	0.04
	P8_F8	0.13	0.03	0.16	0.05
	CP1_FC1	0.16	0.03	0.18	0.03
(C)					
	PO4_CP6	0.16	0.03	0.13	0.03
	CP6_FC2	0.17	0.04	0.14	0.03
	CP6_F4	0.16	0.05	0.13	0.03
	CP2_FP1	0.14	0.04	0.11	0.03
	CP2_CP5	0.16	0.03	0.14	0.02
	CP5_F3	0.15	0.03	0.13	0.04
	CP5_AF3	0.15	0.04	0.12	0.03
	CP1_FC2	0.14	0.02	0.16	0.03
(D)					
	CP2_FC1	0.16	0.03	0.14	0.03
	CP6_FC2	0.16	0.04	0.13	0.03
	PO4_FC1	0.13	0.03	0.16	0.04
	FC2_AF4	0.13	0.03	0.15	0.04

Note. Only connections between parietal and frontal electrodes (based on positions in the international 10/20 system) with robust significant PLI group differences are shown (i.e., solid lines in Figure 5.3).

Table 5.2. Results of one-sample t-tests for dPLI values during high (HA) and low experiences of agency (LA) in trials in which (A) goals matched the outcomes, (B) goals mismatched the outcomes, (C) primes matched the outcomes, and (D) primes mismatched the outcomes in broadband frequency.

	Connection	<i>T</i>	<i>p</i>	Direction effect
(A)				
	LA_PO3_FC1	-.21	.83	-
	LA_FC1_F7	-.65	.53	-
	HA_P7_F7	2.46	.02	↑
	HA_F7_AF4	2.11	.05	↑
	HA_FC5_FP1	3.28	.004	↑
	HA_FC5_FP2	3.23	.004	↑
	HA_FC5_AF4	1.40	.18	-
	HA_F7_FP1	2.74	.01	↑
	HA_CP6_Fz	-1.09	.29	-
(B)				
	LA_PO4_Fz	-.04	.97	-
	LA_PO3_FC1	-.59	.56	-
	LA_P4_FC1	.30	.77	-
	LA_CP2_F3	-.13	.90	-
	LA_F3_F7	-.81	.43	-
	HA_P8_F8	1.73	.099	↑
	HA_CP1_FC1	.37	.72	-
(C)				
	LA_PO4_CP6	-.53	.60	-
	LA_CP6_FC2	-2.89	.009	↓
	LA_CP6_F4	-2.25	.04	↓
	LA_CP2_FP1	-1.91	.07	↓
	LA_CP2_CP5	2.02	.06	↑
	LA_CP5_F3	-.89	.38	-
	LA_CP5_AF3	.15	.88	-
	HA_CP1_FC2	.62	.54	-
(D)				
	LA_CP2_FC1	-.35	.73	-
	LA_CP6_FC2	-1.40	.18	-
	HA_PO4_FC1	.09	.93	-
	HA_FC2_AF4	-1.16	.26	-

Note. Only connections between parietal and frontal electrodes (based on positions in the international 10/20 system) with robust significant PLI group differences are shown (i.e., solid lines in Figure 5.3). The arrows represent the direction of the information flow between the two nodes (i.e., electrodes) specified in the contrast. Upward arrows indicate that the first node is sending information to the second node, whereas downward arrows indicate that the first node is receiving information of the second node. Arrows are presented for *p*-values < 0.10.

DISCUSSION

Building on recent interest in neural networks underlying agency processing (David, 2012), the present study examined cortical information flow during inferences of agency. Whereas some insights into the networks underlying agency processing have been provided by previous studies employing fMRI (David et al., 2007; Nahab et al., 2011), here we offered a first attempt to investigate this connectivity by tapping into the mechanism that is proposed to underlie neural communication (i.e., phase synchronization).

The role of inference processes in self-agency experiences is supported by the current behavioral data. In line with the inference model (Wegner, 2002) and previous work (Aarts et al., 2005; Van der Weiden et al., 2011, 2013; Wegner & Wheatley, 1999), participants reported higher agency experiences when pre-activated knowledge was congruent vs. incongruent with actual outcomes. Importantly, these results cannot be easily accounted for by the comparator account, as predictive motor processes were unreliable (or even absent) due to the experimental set-up. Specifically, there was no causal relation between the key press of participants and the presentation of the outcome, which restricts the motor system in its prediction of sensory action consequences (Sato, 2009). Accordingly, the reported experiences of agency are likely to be informed by cognitive inferences formed upon the occurrence of the outcome.

The results of the EEG data provide insight into neural connectivity underlying agency inferences during matches and mismatches. Increased coupling between parietal and frontal cortices, as well as within frontal areas, was identified in the broadband during high agency experiences in trials in which outcomes matched prior goals. These regions have been associated with agency processing in general (David, Newen, & Vogeley, 2008; Sperduti et al., 2011) and agency inferences in particular (Renes et al., 2013). The PPC has been implicated in the detection of congruence between motor predictions and sensory action consequences, and has mainly been activated during mismatches (David, 2010). Nevertheless, Renes and colleagues (2013) have also identified activity in this region during matches, suggesting that it might be involved in more general comparative processes between outcome expectations and action-effects. Activity in prefrontal areas has been linked to a conscious monitoring function (i.e., the conscious experience of having caused an outcome or not; David, 2010). Although the observed fronto-parietal connectivity concurs with this prior research, it is important to note that

observed connectivity during agency inferences was not restricted to these areas, as can be seen in Figure 5.3.

Connectivity between parietal and frontal areas in the broadband was also observed during low agency experiences in trials in which goals mismatched with the outcome. Notably, however, the coupling within frontal areas that was observed during high agency in match trials was not detected during low agency experiences in mismatch trials. A possible explanation for this finding is that this frontal network is especially involved in the ascription of outcomes to oneself as opposed to external sources. Some indirect support for this idea comes from research on self-referential processing showing increased activity of the medial prefrontal cortex when participants judged personality traits as self-descriptive vs. not self-relevant (Moran, Macrae, Heatherton, Wyland, & Kelley, 2006; Rameson, Satpute, & Lieberman, 2010). This fits with our observation that the frontal network was not involved in case of mismatching outcomes that were not ascribed to oneself (i.e., that were deemed to be non-relevant).

Beyond the mere presence of increased coupling, a trend of directionality pointing toward information flow from parietal to frontal cortices was observed in the broadband during high agency experiences following from outcomes matching goals. This finding is in line with results by Nahab and colleagues (2011) who speculated that the PPC serves as a low-level congruence detection network that transmits mismatch information to prefrontal cortices in order to give rise to higher order agency processing (i.e., a conscious experience of agency). Although this observation is exciting, it is important to note that the observed directionality in the current study was relatively weak (in terms of statistical significance) and absent in trials in which goals mismatched the actual outcome. That is, whereas increasing coupling between parietal and frontal regions was observed during low agency experiences in these trials, parietal nodes were not leading in phase compared to frontal nodes. Therefore, interpretations with regard to direction of information flow should be made with caution. More generally, it is important to note that there is no unique relationship between the time series recorded by EEG and their underlying source, allowing only crude interpretations concerning underlying brain regions. Importantly, however, the main interest of the present study was to elucidate connectivity between frontal and more posterior parts of the brain, rather than to relate specific brain areas to agency inferences.

In addition to neural communication between cortical regions in the broadband frequency, interactions in specific frequency bands were assessed.

Intriguingly, fronto-parietal connections were present across frequency bands, while none of the bands seemed particularly involved in agency inferences as a whole. These observations might be attributable to the complex nature of agency processing, in the sense that it encompasses functions that have been related to specific bands, such as keeping outcome representations active in working memory (associated with theta band oscillations; Klimesch, Schack, & Sauseng, 2005) and, in the case of goal-based inferences, the prioritizing of top-down influence (i.e., goals) over novel events (associated with beta band oscillations; Engel & Fries, 2010). Accordingly, the observed connectivity in the variety of bands might be a reflection of the different dimensions of the integration process involved in agency inferences (Varela et al., 2001).

Recent findings suggest that agency experiences can result from goal-based inferences as well as from prime-based inferences (Van der Weiden et al., 2013). Based on these findings, we examined the neural communication involved in both type of agency inferences. When comparing connectivity patterns between goals and primes in the broadband, frontal connections were observed during high agency experiences in trials in which goals matched the outcome that were absent in trials in which primes matched the outcome. Similarly, more fronto-parietal coupling was observed in goal trials than in prime trials during low agency experiences in mismatch trials. This general decrease in connectivity associated with primes (vs. goals) might be explained by differences in the process underlying the two types of pre-activation (Van der Weiden et al., 2013). In contrast to goals, mere priming of outcome information is not assumed to install an attentional control process that maintains the specific outcome active in mind, while inhibiting other irrelevant (but associated) items at hand. Therefore, the activation of the outcome representation by priming (compared to goals) might be more transient and less stable. The behavioral data provides evidence for this notion. First, the difference in agency experiences resulting from matches and mismatches tends to be more predominantly expressed in goal trials than in prime trials. In addition, participants were significantly faster to report experienced self-agency in the former (vs. the latter) trials. These findings are in line with the notion that agency inferences occurring via priming processes are less stable and noisier than goal-based inferences, which may account for the reduced connectivity associated with the former processes.

This line of reasoning might shed light onto the recent observation that patients suffering from schizophrenia show specific disturbances in prime-based

inferences processes whereas their goal-based inferences seem intact (Renes et al., 2013). Schizophrenia has been related to reduced structural connectivity between various brain regions, including reduced integrity of white matter tracts connecting parietal and frontal nodes (Ellison-Wright & Bullmore, 2009; Voineskos et al., 2010; Whitford, Kubicki, & Shenton, 2011). Given that anatomical connections restrict the functional networks that can be formed (Fries, 2005), agency inferences that rest on fronto-parietal functional connectivity are likely to be disturbed as well. The present study suggests that functional connectivity related to prime-based inferences is weaker compared to goal-based inferences. When taking into account that only the prime-driven processes are disturbed in schizophrenia patients, it can be speculated that the relatively strong functional connectivity pattern underlying inferences based on goals, might allow schizophrenic patients to experience agency despite decreased anatomical fronto-parietal connectivity. In contrast, inferences based on primes are already associated with weaker functional connectivity and accordingly might not be able to overcome these structural abnormalities. However, the notion that primed-based agency inferences are reduced in schizophrenic patients as a result of the quality of fronto-parietal anatomical connectivity awaits further testing.

There are several methodological limitations that warrant consideration when interpreting the present results. By examining connectivity on the scalp we cannot exclude the possibility that observed differences between conditions have been affected by spontaneous and systematic changes of distant sources. That is, due to the discontinuity PLI, noise induced by these sources can shift phase leads to phase lags, which in turn might give rise to spurious differences or, oppositely, mask real differences in connectivity (Vinck, Oostenveld, Van Wingerden, Battaglia, & Pennartz, 2011). Future studies incorporating source-localization procedures might provide additional insight into the influence of distant sources. Another factor that might affect PLI measurements is the number of trials used to estimate this index. When this number is small, as in the current study, PLI values tend to be overestimated, especially in case of small PLI values (Vinck et al., 2011). Note, however, that this overestimation of PLI would be expected in both low and high agency conditions. As such, the connectivity difference of interest is relatively unaffected by this issue. A final confounding factor in the present study is the multiple comparisons problem. Statistical analysis of EEG data inherently copes with testing of condition effects at a large number of pairs, across multiple frequency bands. Although there are methods to correct for multiple

testing, these methods are either overly conservative when a large number of tests is conducted, or focused on networks rather than individual connections. Given the exploratory aim of the present research, an uncorrected comprehensive overview of connectivity is provided. Accordingly, observed connectivity has to be interpreted with some caution.

CONCLUSION

To conclude, we have demonstrated the potential of recent methodological advances in the quantification of brain dynamics to elucidate the neural basis underlying inferences of agency. In particular, we were able to extend prior research that has mainly focused on localized activation and provide preliminary support for the existence of fronto-parietal interactions involved in sending information from parietal to frontal areas to arrive at the conscious experience of agency. By doing so, we hope that the present results will encourage future research to move beyond mere snapshots of the brain and to further explore the neural networks underlying agentic self-awareness.

CHAPTER 6

Moving forward: On the limits of motor-based forward models

This chapter is based on an invited submission to Trends in Cognitive Sciences: Dogge, M., Custers, R., Aarts, H. (under review).
Moving forward: On the limits of motor-based forward models.

ABSTRACT

The human ability to anticipate the consequences that result from action is an essential building block for cognitive, emotional and social functioning. A dominant view is that this faculty is based on predictive motor-signaling, in which a forward model uses a copy of the motor command to predict imminent sensory action-consequences. Although this model was originally conceived to explain the processing of action-outcomes that are tightly coupled to bodily movements, it has been increasingly extrapolated to effects beyond the body. Here we critically evaluate this generalization and argue that, although there is ample evidence for the role of predictions in the processing of external action-effects, there is hitherto little reason to assume that these predictions result from motor-based forward models.

HUMANS AS ANTICIPATORY AGENTS

Actions are an essential part of our everyday life. They allow us to break free from mere enslavement to our surroundings, and to interact with our environment in a goal-directed manner. Although this ability provides us with countless opportunities, it is also a source of unique challenges. Unlike stationary beings, we are not only exposed to input from the environment, but also to the sensory consequences that directly result from our own bodily-movements (e.g., the sound of our own footsteps). The resulting overflow of information could greatly interfere with our ability to track progress towards goal-attainment, as well as with quick online corrections of behavior. In order to deal with these issues, we, and a plethora of other acting animals, have evolved the powerful ability to anticipate the consequences of movement before they actually occur (Crapse & Sommer, 2008a; Straka, Simmers, & Chagnaud, 2018; Waszak et al., 2012). However, how exactly do our brains foresee these consequences?

A dominant and widely accepted view is that predictions about action-outcomes are generated by forward models (Miall & Wolpert, 1996; Wolpert & Flanagan, 2001). These forward models can be seen as simulations of the motor system that use a copy of the motor command, known as an **efferece copy**, to predict the sensory consequences that are expected to result from the action in question (known as corollary discharge). Such predictions have been associated with various adaptive functions, including more efficient perceptual processing of self-produced input (Box 1). The forward model account has gained extensive support, and has been implicated in the prediction of a wide variety of action-outcomes, ranging from effects close to the effector (such as **proprioceptive** changes; see Glossary) to effects in the outside world (such as the visual changes that result from turning on a light switch). Here, we critically examine this explanatory range of forward models. We propose that while the prediction of body-related effects may rely on these models, predicting effects beyond the body often requires more general predictive mechanisms, such as posited by, for instance, **predictive coding** accounts (Friston, 2010).

PREDICTIONS ABOUT THE BODY AND BEYOND

The idea that perception is influenced by motor-based predictions has a long history. Although generally attributed to Helmholtz (Helmholtz, 1962), the development of this notion can be traced back to ancient times – including observations made by the likes of Aristotle (Grüsser, 1995). Nevertheless, the forward model account did not gain traction until 1950, when two independent research groups coined the terms “efference copy” and “corollary discharge” that are still central to the account today (Sperry, 1950; Von Holst & Mittelstaedt, 1971). Both groups validated the influence of motor-predictions by means of illustrative effects in animals, such as the behavior that is observed after surgical eye alterations. Specifically, after rotating the eyes (or the entire head) by 180 degrees, fish and flies show a curious response: they start to move in indefinite circles. This **optokinetic response** was attributed to the idea that the animals predict the retinal motion that results from movement. Under normal circumstances, such predictions can be used to cancel out the influence of action on perception. However, after surgery, the retinal motion will be opposite to what is expected, resulting in a prediction error. As animals pursue perceptual stability, they will start to turn in the same direction as the retinal motion, which only magnifies the error and causes them to get caught in a self-reinforcing behavioral loop.

As is evident from this example, early work on motor-predictive processes focused on action-outcomes that are narrowly tied to the preceding action. Similar body-related prediction effects have been observed in work on sensory attenuation in humans (Box 2). A well-known example is the inability of humans to tickle themselves, which is thought to result from the prediction and cancellation of self-produced **somatosensory input** (Blakemore et al., 2000). Aside from such body-related action-outcomes, motor-based forward models have also been increasingly associated with the processing of external effects (Box 2 and 3). Such effects include input that is produced by button presses, ranging from abstract stimuli such as tones (e.g., Baess, Widmann, Roye, Schröger, & Jacobsen, 2009; Baess et al., 2008; Haggard et al., 2002; Hughes et al., 2013b; Jones, Hughes, & Waszak, 2013; Sato, 2008; Timm, Sanmiguel, Keil, Schröger, & Schönwiesner, 2014; Weiss & Schütz-Bosbach, 2012) and **Gabor patches** (Cardoso-Leite et al., 2010) to more complex visual stimuli like faces and houses (Hughes & Waszak, 2014). Unlike body-related effects, which have a tight and consistent coupling with the preceding actions, the relationship between these action-outcomes and the preceding action

is relatively arbitrary and needs to be learned. In this paper we will refer to these latter outcomes as environment-related (a term adopted from Pfister, 2019). Given the differences between outcomes that affect the body, versus those that affect the environment, it is not self-evident that the predictions of these outcomes rely on the same mechanism. Indeed, the involvement of motor-based forward models in the predictions of (especially environment-related) action-outcomes is disputable. Below we will discuss three lines of argumentation in support of this view.

CONVERGENT EVOLUTION OF MOTOR-PREDICTIVE SIGNALING

Evidence for motor-based forward models has been observed across the animal kingdom, ranging from invertebrates to primates (Crapse & Sommer, 2008a, 2008b; Straka et al., 2018). Despite being ubiquitous, these predictive processes can express themselves in remarkably different ways. One well-studied example is the male cricket, who sings to attract and entice future mates (Poulet & Hedwig, 2006). The cricket's song, which is generated by rubbing the forewings together, is very loud (> 100 dB) and would result in overstimulation of the auditory system. In order to circumvent this issue, crickets use information from motor commands to inhibit the central auditory pathway in phase with wing movements and resulting chirps. A rather different and more advanced way to take advantage of motor-predictions has been observed in electric fish, such as mormyrids (Bell, 1981). Electric fish generate electric fields and assess perturbations caused by objects and other animals to navigate their environment. In order to distinguish between stimulation of the electroreceptors that is caused by themselves and by other fish, mormyrids use motor commands that are generated by the electric organ to predict the precise input (both in terms of timing and content) that will be sensed by their electroreceptors. As the predicted input is opposite in sign, but equal in duration to self-produced input, only electric stimulation that is self-caused will be cancelled out. Hence, although both crickets and mormyrids clearly use motor-based forward models, the specific expression of the predictive mechanism depends on the unique challenges the species faces (Crapse & Sommer, 2008a; Straka et al., 2018).

The evidence for motor-based models in animals, as well as the fact that similar modulations of perceptual processing are observed in humans, is often taken as evidence for the idea that these effects result from similar mechanisms.

However, humans have unique cognitive capacities that allow them to anticipate the future in ways that transcend the predictive mechanisms in many other animals (Gilbert & Wilson, 2007). It is therefore conceivable that they have evolved alternative routes to predict action-outcomes. Indeed, there is ample evidence for such non-motor predictions, as we will discuss in more detail later. Accordingly, the mere observation of prediction effects on the processing of action-outcomes does not necessitate the involvement of motor-based models, as predictions could potentially result from various distinctive mechanisms. Although this reasoning holds for both body-related and environment-related outcomes, the involvement of motor-predictions seems particularly unlikely for the latter.

FORMATION AND NEURAL ENCODING OF MOTOR-BASED FORWARD MODELS

The fact that evidence for motor-predictive signaling has been observed across the animal kingdom is particularly valuable as animal work, and its amenability for the use of relatively intrusive techniques (such as **intracellular recordings**), has resulted in detailed mappings of the neural circuits that underlie motor-to-perception interactions. One of the most precisely described circuits is the neural pathway underlying auditory suppression in mice (Reznik & Mukamel, 2019; Schneider & Mooney, 2018). Like humans, mice suppress auditory consequences of bodily-movement, such as vocalizing and locomotion. This attenuation has at least partially been attributed to inhibitory **interneurons** in the auditory cortex, which in turn suppress auditory **pyramidal cells** (Schneider, Nelson, & Mooney, 2014). Crucially, the interneurons receive input from neurons in the secondary motor cortex (M2). **Optogenetic stimulation** of M2 axon terminals in the auditory cortex leads to comparable suppression effects as actual movement – hence providing evidence for direct motor modulation.

The motor-auditory circuit in mice is of particular interest to the present discussion, as recent research has examined how suppression via this pathway can be adaptively updated by experience to new sounds for which no previously overlearned relationship exists (Rummell, Klee, & Sigurdsson, 2016; Schneider, Sundararajan, & Mooney, 2018). In these studies, mice learned to associate environment-related actions-outcomes (i.e., tone bursts) with either lever presses (Rummell et al., 2016) or treadmill running (Schneider et al., 2018).

Movement-related suppression was observed that was specific for the learned **reafferent** frequency. Importantly, the process was shown to be very gradual. While one hour of learning was not enough to produce specific suppression of the reafferent frequency, this effect became increasingly more apparent after extensive practice over the course of several days (Schneider et al., 2018).

The idea that suppression effects might depend on the strength of learned associations concurs with recent work in humans showing that larger attenuation effects have been observed for more intuitive action-outcome links. Specifically, sounds that result from button presses were more attenuated than sounds resulting from **saccades** and less attenuated compared to sounds that resulted from mouth-movements (i.e., blowing into a microphone) (Mifsud, Beesley, Watson, & Whitford, 2016; Mifsud & Whitford, 2017). A similar pattern was observed for flashes, such that saccades lead to more attenuation than button presses (Mifsud et al., 2018). Hence, more action-based attenuation was observed for action-outcome links that are more frequently encountered during one's lifetime.

Together these studies suggest that the formation of the neural pathways underlying motor-related modulation takes practice. In contrast, most motor-prediction effects in humans appear after short acquisition periods, even for environment-related effects for which no tight action-outcome association exists. Indeed, the acquisition phases of most experiments often do not surpass the one hour of learning (Horváth, 2015) that was shown to be insufficient for mice. Considering that the timescale on which formation of neural connections takes place is likely to be comparable for mice and humans, the fast modulations in our species suggests the reliance on predictions generated via an alternative route.

EVALUATING EVIDENCE FOR MOTOR-PREDICTION EFFECTS BEYOND THE BODY

The idea that the extrapolation of motor-based forward models to outcomes beyond the body is not straightforward is further evidenced by evaluating existing support for such effects. In particular, there are two issues that can be identified, relating to the experimental designs employed, as well as the strength of the empirical evidence.

With regard to the first issue, most empirical studies use a design that cannot exclude the contribution of alternative (non-motor) factors. In particular, the

influence of motor-predictive processes is generally examined by comparing a condition in which participants execute actions with a condition in which participants merely observe effects. This comparison builds on the idea that forward models can only generate action-outcome predictions for voluntary actions that are accompanied by an efference copy, whereas no such predictions are present in the absence of action. However, as has been reviewed in detail elsewhere (Horváth, 2015; Hughes et al., 2013a), these conditions differ in more ways than simply the presence of motor-based predictions about upcoming effects. For instance, self-produced and externally-produced effects differ in temporal predictability (i.e., the predictability of the moment in time at which effects occur), as well as in the attention that is allocated to them.

One way to control for these factors is by directly comparing how predicted versus unpredicted or mispredicted effects of actions are processed. As these conditions only differ in the match between predicted and actual action-effects, predictive action-mechanism are isolated from the aforementioned confounding influences. In addition, this comparison allows one to assess the extent to which the prediction effects are specific to the precise identity of the upcoming action-outcome. Such specificity is necessary for effects to be attributed to motor-based forward models, as precise predictions about upcoming effects are required for their involvement in motor-control. For body-related action-outcomes, such prediction-specificity has indeed been observed. For instance, in the auditory domain, the neurophysiological attenuation of one's own voice greatly diminishes when its playback has an altered pitch (Behroozmand, Liu, & Larson, 2011; Heinks-Maldonado et al., 2005; Heinks-Maldonado, Nagarajan, & Houde, 2006), as well as when the sound one utters differs from one's general median production (i.e., the predicted utterance) of that sound (Niziolek, Nagarajan, & Houde, 2013). Similarly, when participants tickle themselves by means of a robotic arm, reported ticklishness increases when spatial perturbations are introduced (Blakemore et al., 1999).

However, the strength of evidence for specific motor-predictions in environment-related effects is not as convincing. On the one hand, some studies have observed evidence in line with motor-based forward models. For instance, reduced neurophysiological attenuation was observed for tones whose identity could be predicted compared to tones whose identity was unpredictable (Baess et al., 2008) or did not match with action-outcome expectations (Hughes et al., 2013b). Likewise, reduced contrast sensitivity was observed for visual stimuli, such as letters (Roussel et al., 2013) and Gabor patches (Cardoso-Leite et al.,

2010), when the identity of these stimuli was congruent (versus incongruent) with learned action-effect associations. On the other hand, studies that have isolated motor-based predictions have failed to observe such effects. In work on sensory attenuation, several recent studies failed to obtain (clear) evidence for perceptual and neurophysiological attenuation for predicted compared unpredicted or mispredicted tones (Bednark et al., 2015; Dogge, Hofman, et al., 2019). In addition, studies examining intentional binding (Box 3) have failed to show an influence of specific motor-predictions. To be more precise, similar binding was perceived for all outcomes following actions, regardless of whether their specific identity is predictable, or in line with previously learned action-outcome associations (Bednark et al., 2015; Desantis et al., 2012; Haering & Kiesel, 2014). These studies therefore cast doubt on the influence of motor-based identity predictions on perception.

In recent work we further assessed this issue by capitalizing on **bistable perception** (Dogge et al., 2018). This perceptual phenomenon occurs when the brain is exposed to a single stimulus that can yield two conflicting perceptual interpretations (Blake & Logothetis, 2002; Leopold & Logothetis, 1999; Sterzer et al., 2009). Considering that conflict of different stimulus identities is at the heart of bistable perception, we hypothesized that it could be more susceptible to modulation by motor-based forward models than sensory attenuation or intentional binding. Participants executed an action that initiated a bistable stimulus, of which one possible interpretation matched previously learned action-outcome expectations. The observed results were mixed and depended on the type of bistable stimulus that was used. In particular, while motor-predictions biased the perception of an **ambiguous motion sphere**, no such effect was observed for **rivalrous gratings**. One important difference between these two stimuli is that they arguably reflect different stages of visual competition. Although this view can be contested (see Dogge et al., 2018 for a discussion), it can be argued that the conflict that is central to binocular rivalry is resolved at a lower stage of visual processing, and is therefore less sensitive to top-down modulation, than the conflict related to ambiguous motion. If correct, this might mean that the influence of motor-predictions is restricted to certain levels of perceptual processing. Whether this is indeed the case remains to be investigated.

Taken together, the support for the extension of forward models to effects beyond the body is hitherto not convincing. This is evidenced by both the confounds in typically employed experimental designs, as well as the mixed evidence resulting from studies that try to control for these confounds.

CONCLUDING REMARKS AND MOVING FORWARD

We have reviewed the role of motor-based forward models in the processing of action-outcomes. While evidence for these models is ubiquitous across species, there are three reasons why the extrapolation of these models to effects beyond the body hitherto seems unjustified. First, the observation of prediction effects does not in itself necessitate the involvement of motor predictions. Second, the relatively slow formation of action-outcome connections that has been observed in animal research does not match with the rapid changes in prediction effects that have been observed in humans – suggesting that a different mechanism might underlie the latter effects. Third, the evidence for the involvement of motor-predictive signals in the processing of environment-related effects is relatively weak both due to methodological confounds as well as the presence of both positive and negative results.

The idea that motor-based forward models are unlikely to be involved in the prediction of environment-related action-outcomes, raises the question which other mechanisms might underlie predictions of these effects. Importantly, the anticipation of sensory input is known to be strongly utilized outside of the motor system, and to result from a variety of sources (Bubic, von Cramon, & Schubotz, 2010; De Lange et al., 2018). An example of such a non-motor prediction effect is the neurophysiological attenuation of tones whose pitch can be predicted from a preceding tone-sequence, as opposed to tones that are unpredictable from the preceding tone-set (Lange, 2009). How these non-motor predictions compare to motor-predictions is unclear. Whilst a few recent studies have directly compared the magnitude of perceptual modulation resulting from motor and non-motor prediction sources (Bednark et al., 2015; Cardoso-Leite et al., 2010; Dogge, Hofman et al., 2019; Stekelenburg & Vroomen, 2015), the results from these studies were mixed, and are hard to reconcile due to varying manipulations of predictability, as well as various measures of perceptual processing (for a discussion see: Dogge, Hofman et al., 2019). In addition, non-motor expectations are generally assessed in the absence of any actions, which means their role in the processing of action-outcomes remains to be elucidated.

In the past few years there has been an increasing interest for more direct non-motor influences on motor-perception interactions. For example, research on instruction effects has shown that instructions about stimulus-response or response-stimulus mappings can become automatized without any practice (Liefoghe, Wenke, & De Houwer, 2012; Meiran, Liefoghe, & De Houwer, 2017). In particular, responses that are incongruent with instructions lead to impaired performance

during tasks in which the instructions are irrelevant. Although these studies do not directly examine expectation effects, their findings are highly relevant for the present discussion as they suggest that action-outcome relations can be activated without the presence of a motor-act, and hence without the involvement of a motor-based forward model. The fact that these instruction effects are observed without any practice is particularly interesting in light of the generally short duration of studies on environment-related action-outcomes. That is, processes similar to the ones that underlie instruction effects might explain how expectations about action-outcomes can be formed in the absence of extensive learning.

In our own recent work (Dogge, Custers, Gayet, Hoijsink, & Aarts, 2019), we have also examined the influence of instructed task sets on the perceptual processing of action-outcomes. In this experiment, participants initiated the movement of a bistable rotating sphere. In the absence of any learning, participants tend to perceive the sphere as rotating in the same direction as their action. This motion bias was abolished after participants learned (both through instructions and short exposure) that sphere motion would be opposite to the preceding action. In contrast, compatible learning did not add anything to the already existing bias. In other words, both the instructed task set, as well as the existing bias, affected the perceived percepts. It remains to be seen whether instructions in the absence of any practice would lead to similar expectation effects (see Outstanding Questions). While indirect, these joint findings provide evidence for the idea that more general prediction mechanisms could be involved in expectations of environment-related action-effects (see also: Kaiser & Schütz-Bosbach, 2018).

The idea that the perceptual processing of action-effects might result from more general prediction mechanisms coincides with the previously mentioned predictive coding account. According to recent extensions of this account there is no separate mechanism that we use to predict the consequences of our actions (Friston, 2010; Pickering & Clark, 2014). Instead, expectations about action-outcomes are based on a general predictive mechanism that holds for all sensory events (Friston, 2010; Pickering & Clark, 2014). More specifically, it is proposed that expectations about sensory action-outcomes result in prediction errors that can be diminished by executing the movement that will bring about the expected outcome (known as **active inference**). As such, there is no need for efference copies. Note that this is true for both body-related and environment-related action-effects. Here we argue for a less extreme, hybrid model (see also: Pickering & Clark, 2014). That is, instead of discarding motor-based forward models altogether, we

suggest that the contribution of these models is likely to vary as a function of the type of action-outcome. Whilst motor-based models are likely to be involved in predictions of body-related and overlearned outcomes, these models might contribute less to predictions about external events²³. Importantly, we do not imply a strict dichotomy with regards to the type of model that is involved in processing different type of action-effects. In some circumstances, the models might interact, such as when the timing of the upcoming action-outcome is important and cognitive models could benefit from using motor-based predictions as input.

In conclusion, although motor-based forward models are of undoubted value for the perceptual processing of action-outcomes, we argue that their explanatory range has been overestimated, while the influence of other, cognitive, sources has been mostly overlooked. The notion that we can use cognitive processes to anticipate the consequences of our actions is hardly controversial or novel. Indeed, this presumption is central to classic models of goal-directed behavior (Aarts & Custers, 2010; Carver & Scheier, 2000; Miller, Galanter, & Pribram, 1960). In these models actions are driven by an attempt to reduce the discrepancy between current and desired (or predicted) states, which resonates with recent predictive coding accounts. Here we have attempted to both reignite as well as expand these ideas, by emphasizing the importance of cognitive mental representations, not only for directing action, but also for the sensory processing of action-effects themselves.

23 As of yet, it cannot be ruled out that even predictions of body-related effects could result from more general predictive mechanisms. Whether the literature on motor-based forward models can be reinterpreted in terms of predictive coding accounts is an empirical question that can only be answered by empirical work (see also: Pickering & Clark, 2014).

APPENDICES

APPENDIX A: SUPPLEMENTARY SECTION OF CHAPTER 4

Supplementary methods

Sample size determination

Study 4.1

Sample size was based on a pre-specified stopping rule (see Methods section main document). For the compatible condition (which was run first) this rule was slightly different than for the other two conditions. Specifically, data collection was planned to be terminated as soon as substantial evidence was found for (or against) a learning effect in the *first half* of the test trials. This rule was based on the expected extinction of contextual learning over time. Data collection was stopped despite not meeting this pre-specified criterion for several reasons. First, after twenty-five participants the Bayes factor still failed to exceed one of the evidence thresholds (i.e., a BF of 6 or 1/6). Second, the proportion of action-consistent percepts was constant over time, and substantial evidence for a prediction effect was observed across all trials. Based on these observations, the stopping rule was changed for subsequent experiments. Note that this posteriori choice is unlikely to have affected the presented results, which is further evidenced by the replication in Study 4.2.

Study 4.2

The stopping rule for Study 4.2 has been described in detail in the pre-registration document (osf.io/bmxwv).

Pre-screening procedure

Study 4.1

In order to check whether participants had normal stereoscopic vision, they took part in a simple disparity task. In this task, participants were exposed to a 3D black annulus in the center of the screen. Stereoscopic depth information was provided by adding a horizontal offset to the annulus that was presented to each eye, relative to fixation (i.e., binocular disparity). The amount of horizontal offset was varied by means of two adaptive, stochastic approximation staircases (Kesten, 1958) that converged to an accuracy level of 80 percent after 30 trials.

On each trial, participants were asked to report the depth plane in which the annulus was perceived. Correct and incorrect responses resulted in a decrease and an increase of the offset on the next trial, respectively. The minimum offset that resulted in a stable depth percept was determined by averaging the last eleven offsets for both staircases. Only participants for whom this minimum offset was equal to or smaller than 0.04 dva were invited to participate in the actual experiment. We initially used a cut-off of 0.07 dva for inclusion. However, during data collection in a parallel project it became clear that a stricter threshold of 0.04 dva is preferable to ensure that participants reliably observe unambiguous sphere directions. All but one of the already included participants met this criterion. This participant was excluded from analyses as described in the result section of Study 4.1.

Study 4.2

In Study 4.2 an adjusted screening procedure was applied. The pre-screening consisted of two short computer tasks checking for visual acuity and stereoscopic vision. The visual acuity task was a computerized variant of the Landolt-C task, in which the letter “C” was presented in the center of the screen at an angle of 0, 90, 180 or 270 degrees. Participants were tasked to indicate the position of the gap by means of the arrow keys, which caused the letter to change in size on the subsequent trial. The size of the letter was determined via an adaptive staircase procedure that converged to an accuracy level of 80 percent after 40 trials. The minimum gap that was perceived correctly on at least 80% of the trials was calculated by averaging the last ten gap widths. Participants who were not able to distinguish a gap equal to or smaller than 0.04 dva were not invited to participate in the main experiment.

To test for stereoscopic vision, participants were exposed to the actual stimuli used in the task (instead of to an annulus as was the case for Study 4.1). Accordingly, they viewed twenty spheres with a disambiguated motion direction. Participants were asked to report the rotation direction of the sphere by means of the arrow keys (e.g., pressing the left key when the front of the sphere was moving in a leftward direction). Participants who were not able to correctly indicate the correct rotation direction on at least 80% of the trials, were not invited to participate in the main experiment.

Supplementary results

Sensitivity analyses

Study 4.1

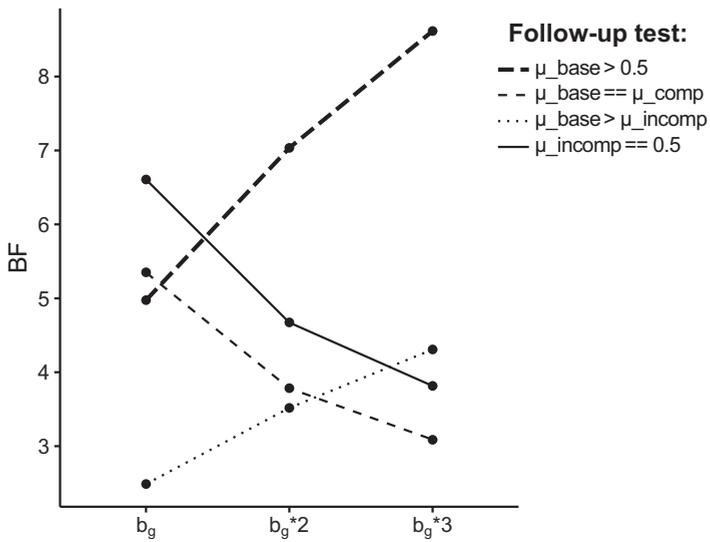
In order to check the robustness of the drawn conclusions under changing prior models, robustness regions (RR) were calculated. These regions reflect the range of Cauchy prior widths under which the same scientific inferences can be drawn (we used $BF \geq 3$ in favor of the reported conclusion). The regions for all tests are depicted in Supplementary Table 4.1. Note that JASP does not allow the entry of a Cauchy prior width greater than 2. Accordingly, prior widths exceeding this threshold are not evaluated.

Supplementary Table 4.1. Robustness ranges for Study 4.1

	Main effect	Time effect
Baseline condition	RR[0.01, 2]	RR[0.85, 2]
Compatible condition	RR[0.02, 2]	RR[0.37, 2]
Incompatible condition	RR[0.29, 2]	RR[0.18, 2]
Baseline vs. Compatible	RR[0.78, 2]	NA
Baseline vs. Incompatible	RR[0.39,0.70]	NA

Study 4.2

In contrast to JASP, the variance of the prior distribution that is used in Bain is estimated based on a fraction of the data. Accordingly, a different method of determining the robustness of drawn inferences is appropriate. Specifically, a sensitivity analysis can be executed in which the fraction that is used to estimate the prior is varied (for more information visit: <https://informative-hypotheses.sites.uu.nl/software/bain/>; Hoijtink, Mulder et al., 2018). Considering that the variance of the prior only affects hypotheses that include equality constraints (Gu, Mulder, et al., 2018), sensitivity analyses will not be reported for tests restricted to inequality constraints. The sensitivity analyses were not pre-registered. Supplementary Figure 4.1 depicts the sensitivity analyses for the follow-up tests that were described in detail in the main document. In addition, we also conducted a sensitivity analysis for the time effect across the three conditions. The evidence for the absence of a time effect becomes inconclusive as the prior width decreases ($BF_{\text{inf},c} = 1.51$ and $BF_{\text{inf},c} = 0.82$ for a b_g of two and three respectively).



Supplementary Figure 4.1. Sensitivity analyses for the follow-up tests of Study 4.2. The Bayes factor is plotted as a function of the fraction of the data that was used to estimate the prior variance (b_g). This variance decreases as the fraction increases.

APPENDIX B: SUPPLEMENTARY SECTION OF CHAPTER 5

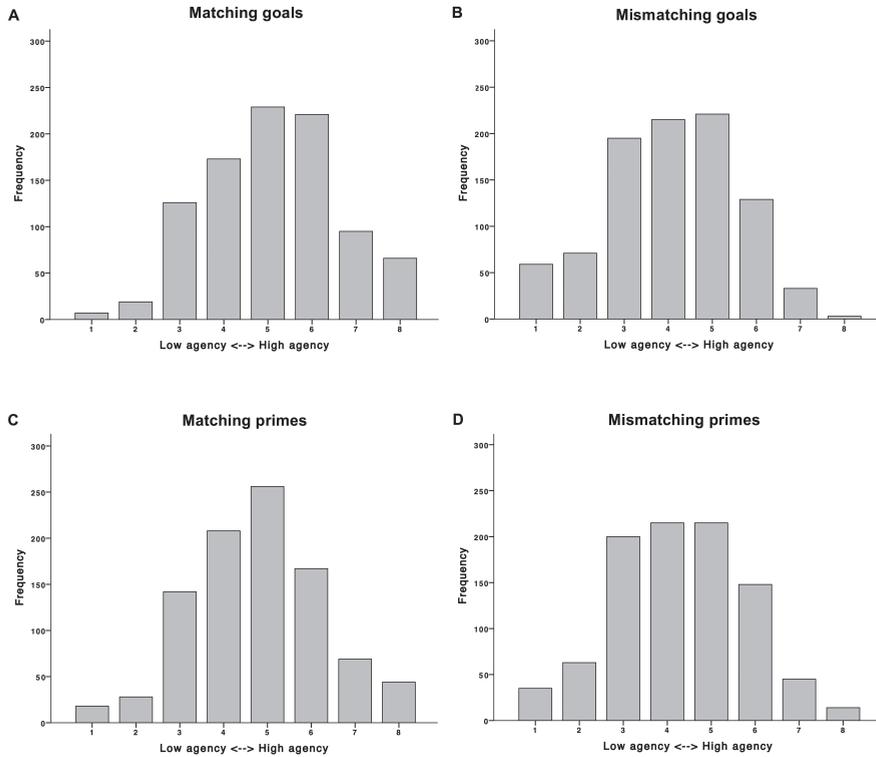


Figure S5.1. Frequency distribution of self-agency ratings in trials in which (A) goals matched the outcome, (B) goals mismatched the outcome, (C) primes matched the outcome and (D) primes mismatched the outcome. In the goal condition, high agency ratings were reported on 65.3% of the match trials and on 41.7% of the mismatch trials, whereas in the prime condition high agency was reported on 57.5% of the match trials and on 45.1% of the mismatch trials.

APPENDIX C: TEXT BOXES AND GLOSSARY OF CHAPTER 6

Box 1: Forward models and their functional role

According to the forward model account, the brain uses forward models to simulate the motor system in response to motor commands. In doing so it can predict both the upcoming state of the body, as well as accompanying sensory consequences. Such forward models have been associated with various adaptive functions (for a review see: Miall & Wolpert, 1996).

One essential use of forward models is control during calibrated or fast movement (Davidson & Wolpert, 2005; Desmurget & Grafton, 2000; Kawato, 1999; Wolpert & Ghahramani, 2000). In order to execute coordinated actions (especially actions involving multiple limbs, such as clapping), the brain needs to know the position and velocity of the effectors involved. This information cannot be obtained from our senses due to the relatively long processing time of sensory feedback. These same delays also prevent sensory feedback from assisting in the execution of fast movements, such as saccades. Forward models solve these issues by providing predictions that can be used before feedback is received.

Forward models also play an important role in skill acquisition. When learning a new skill, such as playing the piano, there will be discrepancies between predicted and actual actions and action-outcomes. In order to learn from these prediction errors, the observed sensory discrepancies need to be translated into motor errors. Forward models are thought to be essential in this translation as they capture the relationship between actions and following sensory effects (a phenomenon known as distal supervised learning) (Jordan, 2003). In a similar way, forward models underlie our ability to mentally rehearse actions without actually performing them (Miall & Wolpert, 1996).

In addition to the proposed use in motor processing, forward models have also been implicated in perception, which is the central topic of this paper. In particular, by predicting the sensory consequences of action, actual self-produced effects can be attenuated or cancelled out (e.g., Blakemore et al., 2000; Wolpert & Flanagan, 2001; see also Box 2). This is adaptive as it allows for more resources to be spend on the processing of input generated by external sources, which generally hold more informative value. In addition, predictions of forward models are thought to underlie the ability to attribute events to their proper cause. That is, while one is likely to be the cause of an event when there

is a match between predicted and actual input, considerable prediction errors generally justify external attributions (Feinberg, 1978; Frith et al., 2000a).

Box 2: Sensory attenuation of body-related and environment-related action-outcomes

Sensory attenuation refers to the perceptual and neurophysiological attenuation of sensory input that is produced by actions as opposed to external sources. Perceptual attenuation is generally measured by means of explicit intensity ratings (such as the experienced ticklishness of touch; Blakemore et al., 2000; Claxton, 1975; Weiskrantz, 1971), or by more implicit psychophysical measures, such as the point of subjective equality (e.g., Dogge, Hofman et al., 2019; Sato, 2008; Weiss et al., 2011b; Weiss & Schütz-Bosbach, 2012). On a neurophysiological level, sensory attenuation is typically measured by means of event-related potentials, such as the auditory N1 - the amplitude of which is diminished for self-produced (versus externally-produced) effects (see: Horváth, 2015; Schröger et al., 2015 for recent reviews).

Attenuation effects have been observed both for body-related and environment-related action-outcomes. Examples of the former are self-applied tickles (Claxton, 1975; Weiskrantz, 1971) and constant force (Bays, Wolpert, & Flanagan, 2005; Shergill, Bays, Frith, & Wolpert, 2003), which are perceived as less intense and evoke less activity in the somatosensory cortex (as measured by fMRI; Blakemore et al., 2000; Blakemore et al., 1998), compared to tactile input produced by another person or by a robotic device. Likewise, listening to self-produced vocalizations during speech evokes smaller auditory evoked responses compared to passively listening to the same sounds (Curio et al., 2000; Heinks-Maldonado et al., 2005; Wang et al., 2014). In addition to these body-related outcomes, attenuation effects have also been observed for environment-related action-outcomes for which the connection with preceding actions is far more arbitrary, such as for visual stimuli (Cardoso-Leite et al., 2010; Hughes & Waszak, 2014; Roussel, Hughes, & Waszak, 2014; Roussel et al., 2013) or tones resulting from button presses (e.g., Baess et al., 2009; Baess et al., 2008; Hughes et al., 2013b; Jones et al., 2013; Sato, 2008; Timm et al., 2014; Weiss & Schütz-Bosbach, 2012).

Although action-related attenuation effects are generally ascribed to motor-based forward models, the precise underlying mechanism remains a topic of discussion. The most common explanation is the cancellation account

that is based on original descriptions by von Holst and Mittelstaedt (Von Holst & Mittelstaedt, 1971). This account proposes that motor predictions can be seen as “photo negatives” that cancel out self-produced input. Alternatively, the pre-activation account (Roussel et al., 2013; Waszak et al., 2012) proposes that motor-predictions result in heightened baseline activity in the sensory areas that represent expected action-outcomes. As this heightened baseline level is harder to distinguish from activity that is subsequently induced by actual input, self-produced effects are experienced as relatively attenuated. Finally, sharpening accounts suggest that the observed suppression of self-produced effects actually reflects a sharpening response (Yon, Gilbert, De Lange, & Press, 2018). According to this account, activity is only increased for the minority of neurons that are tuned to the expected (self-produced) effects, whilst signals inconsistent with the predictions are suppressed. This will result in a decrease in the net population response and hence in an overall attenuation effect. This account might explain why motor-predictions sometimes result in enhanced rather than suppressed perceptual processing of expected effects (for further discussion see: De Lange et al., 2018; Reznik & Mukamel, 2019).

Box 3: Intentional binding of environment-related action-outcomes

Similar to sensory attenuation (see Box 2), intentional binding is a perceptual phenomenon that has frequently been used to assess the influence of motor-based predictions on perceptual processing (for a review see: Moore & Obhi, 2012). Intentional binding refers to the perceived temporal attraction of actions and subsequent effects, compared to when these events occur in isolation. This effect disappears when action-outcomes are externally triggered either through TMS or through passive movement (Engbert et al., 2008; Engbert, Wohlschläger, Thomas, & Haggard, 2007; Haggard & Clark, 2003; Haggard et al., 2002), demonstrating its reliance on intentionality. The most commonly used method to assess binding borrows from the work of Libet (Libet et al., 1983) who realized that conscious awareness of actions could be captured by means of a clock. In the standard binding paradigm participants attend to a rotating clock hand and indicate its position at the time of action-execution or the sensory action-outcome (Haggard et al., 2002). While the temporal awareness of the action is generally shifted towards the outcome, the awareness of the outcome is generally shifted towards the preceding action. Another popular method to assess intentional binding is to directly

ask participants to estimate the interval between their actions and subsequent outcomes (e.g., Caspar, Christensen, Cleeremans, & Haggard, 2016; Engbert et al., 2008, 2007), with reduced intervals being observed for self-produced outcomes.

In contrast to sensory attenuation, research on intentional binding has been restricted to environment-related effects. This is a consequence of the temporal separation of actions and outcomes that is inherent to the aforementioned measures. A few apparent exceptions are studies that have examined binding of somatic action-outcomes, such as the passive displacement of a finger (Engbert et al., 2008; Tsakiris & Haggard, 2003). Crucially, however, these somatic effects were produced in the same artificial and arbitrary way as environment-related action-effects (i.e., by means of button presses). As such, the overlearned, one-on-one relationship that is present for what we deem to be *true* bodily-related effects is absent in these studies.

Intentional binding is generally attributed to similar motor-based forward models as sensory attenuation (but see: Moore & Fletcher, 2012). However, how exactly motor-predictive mechanisms would lead to the aforementioned perceptual attraction often remains unspecified. An important exception is the pre-activation account (Waszak et al., 2012), which offers a specific explanation for the shift in the perceived onset time of the action-outcome. According to this account, motor-predictions lead to the pre-activation of an expected sensory template. As a result of this elevated baseline, the relative amount of time needed to reach the threshold for conscious awareness is reduced for self-produced effects, as opposed to externally-produced effects for which no such pre-activation can occur.

Outstanding questions box

- What is the time course of the formation of motor-based forward models, and the underlying neural motor-perceptual circuit, in humans?
- Can instructions about action-outcome associations in the absence of learning lead to similar perceptual modulation (e.g., sensory attenuation) as learned associations?
- Can instructions equally affect the formation of predictions regarding body-related and environment-related action-outcomes?
- What are the different characteristics of body-related and environment-related predictions, in terms of for instance, their resistance to evidence-based updating (i.e., flexibility) or the relative importance or weight of temporal predictability (i.e. “when” an effect might occur might be less important for

environment-related effects that are not as closely linked to the preceding action).

- Are predictions about body-related and environment-related action-outcomes equally disrupted in disorders that are commonly associated with aberrant predictive mechanisms (such as schizophrenia)?

GLOSSARY

Active inference: an extension of predictive coding models that includes motor control. In particular, active inference presumes that prediction-errors regarding expected sensory input cannot only be resolved by updating one's model of the world, but also by actively bringing about the expected sensory input through one's actions.

Ambiguous motion sphere: a rotating sphere with an ambiguous motion direction. The sphere consists of dots, half of which move in a leftward direction, and half of which move in a rightward direction. The sphere can be perceived as rotating in a clockwise or a counterclockwise direction depending on whether the leftward or the rightward moving dots are perceived to be located in the front plane (i.e., closer to the observer).

Bistable perception: a single stimulus that can result in two perceptual interpretations that tend to alternate over time.

Gabor patch: a visual stimulus to which the primary visual cortex is particularly responsive, consisting of a sinusoidal grating with a circular Gaussian amplitude envelope.

Intracellular recordings: a technique that allows one to measure electrical currents on the level of a single neuron by sticking a microelectrode inside the cell.

Interneurons: a type of nerve cell with a predominantly inhibitory function

Optogenetic stimulation: a relatively non-invasive technique that combines genetic manipulation and optics to control the activity of neurons. In particular,

genetic modification can be used to express light-sensitive proteins in cells. These proteins will in turn cause an influx of positively-charged or negatively-charged ions in response to light and as such activate or deactivate the neuron, respectively.

Optokinetic response: reflexive movement of the eyes (as well as the head and body in some animals) in response to the full motion of one's visual field. The rotation is in the same direction as the movement to minimize retinal motion and maintain perceptual stability.

Predictive coding: theoretical framework on perception in which the brain forms internal models about sensory input, which it continuously updates by minimizing prediction errors that result from the comparison between predicted and actual input.

Proprioception: the sense of one's body in space (e.g., in terms of position and velocity).

Pyramidal cells: a common type of excitatory cell that can be found in the cortex of mammals.

Reafferent: resulting from one's own movement.

Rivalrous gratings: the presentation of unique gratings (often in terms of their orientation) to each eye. Instead of perceiving a fused image, perception will alternate between the two possible interpretations over time (a phenomenon known as binocular rivalry).

Saccades: fast eye movements that shift gaze from one fixation point to another.

Somatosensory input: sensory input resulting from skin, muscles and joints (e.g., touch, temperature etc.).

SUMMARY IN DUTCH

(Nederlandse samenvatting)

ACHTERGROND

Acties zijn een essentieel onderdeel van ons dagelijks leven. Door te handelen zijn we niet enkel passief onderworpen aan de input van de wereld om ons heen, maar kunnen we onze omgeving op een doelgerichte manier beïnvloeden. Dit vermogen heeft een enorme impact op hoe we de wereld waarnemen en ervaren. Een intrigerend voorbeeld is het gevoel van controle dat we hebben tijdens het uitvoeren van intentionele bewegingen. Dat wil zeggen, wanneer we een actie uitvoeren, weten en voelen we doorgaans dat wij de oorzaak van deze actie en de opvolgende consequenties zijn. Dit doordringende gevoel van zelf-causatie, ook wel bekend als *sense of agency* staat centraal in het huidige proefschrift.

De ervaring van zelf-causatie beïnvloedt veel facetten van ons dagelijks leven. Zo is deze ervaring bijvoorbeeld nauw verbonden met strafrechtelijke verantwoordelijkheid: Individuen worden enkel gestraft voor gedrag dat vrijwillig is uitgevoerd en waarvan de gevolgen worden begrepen (Haggard, 2017; Haggard & Tsakiris, 2009). In het verlengde hiervan is de ervaring van zelf-causatie ook van belang voor het grote (en groeiende) aantal van onze acties dat gemedieerd wordt door technologische interfaces. Denk hierbij bijvoorbeeld aan de toenemende automatisering van auto's en vliegtuigen, waarin onze objectieve controle steeds meer uit handen wordt gegeven. Een belangrijke kwestie is hoe ervaren controle en daarmee verbonden verantwoordelijkheid desondanks gewaarborgd kunnen blijven.

Het belang van het ervaren van zelf-causatie wordt ook duidelijk wanneer men kijkt naar mentale stoornissen waarin deze ervaring is verstoord. Zo zijn er patiënten met schizofrenie die er van overtuigd zijn dat hun acties en gedachten door externe bronnen worden bestuurd (Frith et al., 2000a). Daarnaast leiden ervaringen van verminderde controle in depressieve (Alloy & Abramson, 1979; Obhi, Swiderski, & Farquhar, 2013) en angstige (Gentsch, Endrass, & Kathmann, 2012; Oren, Friedmann, & Dar, 2016) patiënten tot gevoelens van hulpeloosheid of compulsief gedrag. Deze brede implicaties illustreren hoe belangrijk het is om te begrijpen hoe ervaringen van zelf-causatie tot stand komen.

MECHANISMEN DIE TEN GRONDSLAG LIGGEN AAN ZELF-CAUSATIE ERVARINGEN

Er bestaat veel discussie over de onderliggende mechanismen van de ervaring van zelf-causatie. Eén van de meest dominante modellen in het veld veronderstelt dat deze ervaring gebaseerd is op voorspellingen vanuit ons motorische systeem (Miall & Wolpert, 1996; Wolpert & Flanagan, 2001). Volgens dit model zorgt de continue en levenslange blootstelling aan acties en opvolgende consequenties ervoor dat we leren om deze gebeurtenissen met elkaar te associëren. Door middel van deze associaties kan ons brein de zintuiglijke gevolgen van een actie (zoals het geluid van een deurbel) vervolgens al voorspellen op basis van de motorische commando's die voorafgaand aan een actie worden gegenereerd.

Onderzoek naar motorische voorspellingen was oorspronkelijk met name gericht op ons vermogen om vrijwillige acties te controleren. Om vloeiend en snel te kunnen bewegen, moeten we onze acties namelijk gedurende de beweging kunnen aansturen. Hiervoor kunnen we niet wachten op de relatief langzame terugkoppeling van onze zintuigen. In plaats daarvan voorspelt ons motorische systeem de consequenties van onze acties. Door deze voorspellingen te vergelijken met verlangde uitkomsten kunnen acties tijdens de uitvoering worden gecorrigeerd. In de afgelopen decennia is er voorgesteld dat dezelfde voorspellingen betrokken zijn in onze subjectieve ervaring van zelf-causatie (Feinberg, 1978; Frith et al., 2000a; Haggard & Chambon, 2012). Het idee is hierbij dat deze ervaring voortkomt uit een vergelijking tussen verwachte zintuiglijke input en daadwerkelijk ondervonden zintuiglijke input. Wanneer verwachtingen en ervaringen overeenstemmen zijn we geneigd om controle te ervaren over onze acties en daaropvolgende consequenties. Substantiële discrepanties leiden daarentegen doorgaans tot externe attributies van causaliteit.

Het idee dat motorische voorspellingen ten grondslag liggen aan de ervaring van zelf-causatie wordt vaak onderbouwd door middel van onderzoek dat gebruik maakt van indirecte maten van deze gevoelens. In plaats van participanten direct te vragen of ze controle ervaren, richten deze maten zich op de perceptuele verwerking van actie-uitkomsten. De perceptie van deze uitkomsten wordt vervolgens als een indicator voor meer expliciete zelf-causatie gevoelens gebruikt. Twee aspecten van perceptuele verwerking die veelvuldig zijn toegepast om ervaringen van zelf-causatie te onderzoeken zijn intensiteit en timing. Gebeurtenissen die het gevolg zijn van onze eigen acties worden doorgaans als minder intens (Schafer & Marcus,

1973) en als eerder in tijd (Haggard et al., 2002) waargenomen dan gebeurtenissen die veroorzaakt zijn door een externe bron. Dit is in lijn met het idee dat motorische voorspellingen de perceptie van actie-uitkomsten beïnvloeden. Immers kunnen er alleen voorspellingen worden gemaakt op basis van motorische signalen wanneer gebeurtenissen door acties worden veroorzaakt. Het is echter belangrijk om op te merken dat zelf-veroorzaakte en extern veroorzaakte effecten op meer punten verschillen dan enkel de aanwezigheid van motorische signalen (Hughes et al., 2013a). Wanneer we een effect veroorzaken door middel van een actie kunnen we niet alleen voorspellen wat de uitkomst van deze actie gaat zijn (in termen van identiteit), maar bijvoorbeeld ook wanneer deze gebeurtenis plaats zal vinden (temporele voorspelbaarheid). Aangezien eerder onderzoek heeft aangetoond dat de perceptie van intensiteit en timing beïnvloed worden door dit soort alternatieve informatie (e.g., Haggard et al., 2002; Lange, 2009; Ruess et al., 2017; Vroomen & Stekelenburg, 2010), is de unieke bijdrage van motorische signalen aan de ervaring van zelf-causatie nog onduidelijk.

Naast de invloed van motorische signalen zijn er ook andere verklarende mechanismen voor zelf-causatie ervaringen voorgesteld. Een dominant alternatief is het idee dat de ervaring van zelf-causatie resulteert uit een post-hoc vergelijking van een waargenomen gebeurtenis en eerder geactiveerde gedachten of intenties (Wegner, 2002; Wegner & Wheatley, 1999). Motorische voorspellingen spelen maar een minimale rol in dit model. In plaats daarvan is de mate waarin we acties en opvolgende consequenties aan onszelf toekennen afhankelijk van drie factoren: de mate waarin gedachten over de actie voorafgaan aan de actie, de mate waarin gedachten consistent zijn met de actie en actie-uitkomsten, en de aanwezigheid van potentiële alternatieve oorzaken van zintuigelijke input. Bewijs voor dit model komt uit onderzoek waarin gedachten over actie-uitkomsten vooraf werden geactiveerd. Dit werd bijvoorbeeld gedaan door participanten actief het doel te laten vormen om een bepaalde uitkomst te veroorzaken. Participanten gaven aan meer controle te ervaren over de gevolgen van hun acties wanneer deze gevolgen overeenkwamen met dit doel (Aarts et al., 2005). In vergelijking met onderzoek naar motorische voorspellingen is het onderzoek naar de mechanismen en neurale substraten die ten grondslag liggen aan dit model echter beperkt.

HUIDIG PROEFSCHRIFT

De voorgaande paragrafen illustreren dat de ervaring van zelf-causatie een product is van zowel voorspellingen op basis van motorische signalen als van informatie die voortkomt uit meer hogere orde, cognitieve bronnen. In het huidige proefschrift wordt er onderzocht hoe deze verschillende informatiebronnen ervaren zelf-causatie beïnvloeden.

In Hoofdstuk 2 wordt er bekeken of voorspellingen die voortkomen uit verschillende bronnen een vergelijkbare impact hebben op hoe intens actie-uitkomsten worden waargenomen (een impliciete maat van zelf-causatie). Specifiek wordt er in dit hoofdstuk onderzocht of de intensiteit waarmee participanten de effecten van hun acties waarnemen afneemt wanneer ze de identiteit van deze effecten vooraf kunnen voorspellen vanuit (A) motorische signalen of (B) vanuit andere actie-ongelateerde bronnen. Hiertoe namen participanten deel aan een taak waarin ze door middel van simpele drukknoppen tonen met variërende toonhoogtes konden produceren. Er werd op verschillende manieren een verwachting aangeleerd over de opkomende toonhoogte. In één conditie leerden participanten om tonen te voorspellen aan de hand van hun acties (bijvoorbeeld: de linker toets werd gevolgd door een lage toon en de rechtertoets door een hoge toon). In een andere conditie werden de acties van participanten vergezeld door de presentatie van abstracte visuele symbolen op het computerscherm. In deze conditie werd de toonhoogte van de toon bepaald door deze symbolen in plaats van door de acties (bijvoorbeeld: een vierkant werd gevolgd door een lage toon en cirkel door een hoge toon). Na deze leerfasen werden participanten blootgesteld aan tonen die consistent of inconsistent waren met de aangeleerde actie-toon en symbool-toon relaties. Op basis van de veronderstelde rol van motorische voorspellingen werd er in ieder geval verwacht dat consistente (voorspelbare) tonen zowel op perceptueel en neurofysiologisch niveau onderdrukt zouden worden ten opzichte van inconsistente tonen in de actie conditie. Er werd echter in geen van de condities robuuste evidentie voor dit soort suppressie effecten gevonden.

Deze bevindingen zijn opvallend gezien de dominante rol die doorgaans wordt toegekend aan motorische voorspellingen als verklarend mechanisme voor ervaringen van zelf-causatie. In Hoofdstuk 3 worden de nul bevindingen uit Hoofdstuk 2 nader onderzocht. Specifiek overweeg ik hier of de afwezigheid van de verwachte effecten een gevolg zou kunnen zijn van tekortkomingen van

veelgebruikte paradigma's. De voorspellingen vanuit het motorische systeem worden namelijk verondersteld heel specifiek te zijn. Wanneer een pianist bijvoorbeeld een toets indrukt op een piano dan verwacht hij niet zomaar elke toon, maar een toon met een specifieke toonhoogte. De twee populaire indirecte maten van zelf-causatie reflecteren echter enkel de verwerking van actie-uitkomsten (in termen van timing en intensiteit), maar zijn niet direct gerelateerd aan "wat" er wordt voorspeld. Met andere woorden, de stimulusidentiteit staat niet centraal binnen deze perceptuele fenomenen, waardoor ze mogelijk niet optimaal geschikt zijn om invloeden van motorische voorspellingen op te pikken.

In Hoofdstuk 3 maak ik daarom gebruik van een fenomeen dat bekend staat als bistabiele perceptie, waarvoor stimulus identiteit van inherent belang is. Bistabiele perceptie komt voor wanneer het brein wordt blootgesteld aan een stimulus die op meerdere manieren geïnterpreteerd kan worden, waardoor de waarneming over tijd alterneert tussen deze verschillende interpretaties (Blake & Logothetis, 2002; Leopold & Logothetis, 1999). Een bekend voorbeeld hiervan is een ronddraaiende ballerina, die afwisselend als naar links of naar rechtsdraaiend kan worden waargenomen. Eerder onderzoek heeft aangetoond dat de waargenomen stimulus identiteit van dit soort stimuli gevoelig is voor verwachtingen over wat men gaat zien (Scocchia, 2014). In dit hoofdstuk bouw ik voort op dit onderzoek en introduceer ik een taak waarin simpele manuele acties van participanten werden gevolgd door bistabiele afbeeldingen. Slechts één van de interpretaties van deze afbeeldingen kwam overeen met de voorspellingen die gevormd konden worden op basis van eerder aangeleerde associaties tussen acties en daaropvolgende effecten. Als op actie gebaseerde voorspellingen inderdaad invloed hebben op perceptuele verwerking van actie-uitkomsten, dan zouden participanten meer geneigd moeten zijn om de voorspelde stimulus identiteit waar te nemen. Bijvoorbeeld: als een participant heeft geleerd dat het roteren van een draaiknop naar links gevolgd wordt door een naar links roterende bol, dan zou deze participant vervolgens ook meer geneigd moeten zijn om een bol met een ambigue draairichting naar links te zien draaien na dezelfde actie. Dit idee werd onderzocht voor verschillende typen bistabiele stimuli. De verwachte effecten werden niet gevonden voor bistabiele stimuli waarvan wordt gedacht dat het conflict tussen de verschillende stimulus interpretaties op een relatief laag niveau van visuele verwerking wordt opgelost. Participanten waren echter wél meer geneigd om de verwachte effecten waar te nemen wanneer acties gevolgd werden door bistabiele stimuli waarvan het conflict waarschijnlijk op een hoger niveau van visuele verwerking plaatsvindt. Hoewel

speculatief, suggereren deze resultaten daarmee dat de invloed van motorische voorspellingen mogelijk enkel invloed hebben op een relatief hoog stadium van visuele verwerking.

De resultaten in de voorafgaande hoofdstukken laten zien dat voorspellingen over actie-uitkomsten in ieder geval onder sommige omstandigheden de perceptuele verwerking van deze uitkomsten kan beïnvloeden. In Hoofdstuk 4, zoom ik verder in op de eigenschappen van motorische voorspellingen en onderzoek ik de mate waarin deze voorspellingen over de tijd kunnen veranderen. (Actie-)gerelateerde voorspellingen kunnen zich ontwikkelen over verschillende tijdschalen en verschillende voorspellingen zijn niet altijd compatibel met elkaar (Seriès & Seitz, 2013). Dit wordt bijvoorbeeld duidelijk wanneer je probeert om een auto met een aanhangwagen achteruit in te parkeren. De over lange tijd opgebouwde verwachting dat de auto doorgaans dezelfde kant op draait als het stuur, conflicteert in dit geval namelijk met de tijdelijke verwachting dat de aanhangwagen in deze specifieke context precies in de tegenovergestelde richting zal bewegen. In dit hoofdstuk bekijk ik hoe dit soort korte- en lange termijn verwachtingen gezamenlijk de perceptie van actie-uitkomsten beïnvloeden. Participanten namen deel aan een taak waarin ze door middel van een draaiknop een bistabiele bol in beweging brachten die zowel als rechtsom draaiend en als linksom draaiend kon worden waargenomen. In het dagelijks leven draaien objecten meestal dezelfde richting op als voorafgaande acties (bijvoorbeeld wanneer we een horloge opwinden). Door de constante blootstelling aan dit soort associaties zijn participanten zonder training geneigd om de bol in dezelfde richting als hun actie te zien roteren. In de studies die worden gepresenteerd in dit hoofdstuk laat ik zien dat deze bestaande bias teniet wordt gedaan wanneer participanten kort bloot worden gesteld aan contextuele actie-effect associaties die incompatibel zijn met deze bias (d.w.z. wanneer participanten een bol te zien kregen die in 80% van de gevallen de andere kant opdraaide dan de richting waarin de participanten de draaiknop hadden gedraaid). De blootstelling aan compatibele actie-effect associaties droeg echter niets bij aan de al bestaande bias. Deze resultaten suggereren dat de perceptie van actie-uitkomsten zowel beïnvloed wordt door verwachtingen die voortkomen uit langdurige blootstelling aan actie-effect relaties, als door verwachtingen die gebaseerd zijn op de huidige context of taak.

Naast voorspellingen op basis van motorische signalen, is er ook evidentie voor het belang van voorspellingen die gebaseerd zijn op non-motorische bronnen

(zoals overtuigingen, doelen en intenties). De rol van dit soort cognitieve cues is met name onderzocht in studies waarin participanten expliciet naar ervaren zelf-causatie worden gevraagd. Zo is er bijvoorbeeld aangetoond dat participanten ook controle over zintuigelijke input rapporteren wanneer er in werkelijkheid geen relatie is tussen acties en opvolgende effecten - zolang de effecten maar overeenkomen met intenties of gedachten die vooraf gingen aan de actie (Aarts et al., 2005). Deze bevindingen worden vaak verklaard aan de hand van de eerder beschreven post-hoc inferentie processen. De neurale processen die ten grondslag liggen aan deze inferenties zijn maar beperkt bestudeerd. Daarnaast richt het meeste bestaande breinonderzoek zich op het lokaliseren van specifieke gebieden die actief zijn tijdens zelf-causatie ervaringen, terwijl de communicatie tussen gebieden wordt genegeerd (David, 2012). In Hoofdstuk 5 speel ik in op deze tekortkomingen en meet ik de informatiestroom tussen hersengebieden tijdens ervaringen van zelf-causatie die voortkomen uit non-motorische voorspellingen. Specifiek namen participanten deel aan een taak waarin gedachten over de gevolgen van acties (in dit geval de presentatie van het woord “rood” of “blauw” op het computerscherm) vooraf werden geactiveerd. Dit werd gedaan door participanten vooraf het doel te laten vormen om één van deze woorden te veroorzaken of door participanten voor een hele korte periode één van de woorden te laten zien. In overeenstemming met eerder werk rapporteerden participanten een sterker gevoel van controle over de daadwerkelijke gevolgen van hun acties wanneer deze consistent (versus inconsistent) waren met eerder geactiveerde gedachten. Daarnaast werd er evidentie gevonden voor communicatie tussen pariëtale en frontale hersengebieden tijdens ervaringen van zelf-causatie. In Hoofdstuk 5 beschouw ik deze resultaten in de context van eerder werk naar de neurale substraten van ervaren zelf-causatie.

De bovengenoemde studies laten veelal kleine of weinig robuuste effecten van motorische voorspellingen op de perceptie van actie-uitkomsten zien. Als gevolg hiervan kan men zich afvragen in hoeverre voorspellingen vanuit het motorische systeem daadwerkelijk een rol spelen in de (impliciete) ervaring van zelf-causatie. Meer specifiek kunnen er vraagtekens gezet worden bij de mate waarin motorische voorspellingen een rol spelen in de perceptie van *alle* gevolgen van onze acties. Er kan een onderscheid gemaakt worden tussen (A) lichaam-gerelateerde actie-uitkomsten en (B) externe actie-uitkomsten. Lichaam-gerelateerde actie-uitkomsten hebben een duidelijke, aangeboren of ingeprente link met voorafgaande acties. Voorbeelden hiervan zijn de tactiele sensatie die je ervaart wanneer je jouw

eigen huid aanraakt en het geluid van jouw stem wanneer je spreekt. Voor externe actie-uitkomsten is de link tussen acties en opvolgende zintuiglijke gevolgen daarentegen arbitrair, zoals voor de tonen en bistabiele stimuli die gebruikt zijn in de studies voor dit proefschrift. Initieel onderzoek naar motorische voorspellingen heeft zich met name gericht op de lichaam-gerelateerde actie-uitkomsten die belangrijk zijn voor actie-controle (bijv. de huidige positie van een arm). In veel recent onderzoek wordt er echter gesuggereerd dat het motorische systeem ook voorspellingen kan genereren over externe actie-uitkomsten. In Hoofdstuk 6 evalueer ik deze generalisatie en concludeer ik, mede op basis van de resultaten van de voorgaande hoofdstukken, dat er op dit moment niet voldoende evidentie is om dit te rechtvaardigen. Ik stel voor dat het motorische systeem waarschijnlijk met name een rol speelt bij de perceptie van actie-uitkomsten die een duidelijke één-op-één relatie hebben met de voorafgaande actie. Wanneer het echter gaat om voorspellingen die verder reiken dan onszelf is het waarschijnlijk dat cognitieve (non-motorische) bronnen een grotere rol spelen dan op dit moment wordt gewaardeerd.

REFERENCES

- Aarts, H. (2012). Goals, motivated social cognition, and behavior. In S. T. Fiske & C. N. Macrae (Eds.), *The SAGE handbook of social cognition* (pp. 75–95). London: Sage Publications Ltd.
- Aarts, H., & Custers, R. (2010). The unconscious will : How the pursuit of goals operates outside of conscious awareness. *Science*, 329, 47–51.
- Aarts, H., Custers, R., & Wegner, D. M. (2005). On the inference of personal authorship: Enhancing experienced agency by priming effect information. *Consciousness and Cognition*, 14, 439–458.
- Aarts, H., & Dijksterhuis, A. (2000). Habits as knowledge structures: Automaticity in goal-directed behavior. *Journal of Personality and Social Psychology*, 78(1), 53–63.
- Adams, R. A., Shipp, S., & Friston, K. J. (2013). Predictions not commands: Active inference in the motor system. *Brain Structure and Function*, 218(3), 611–643.
- Adams, W. J., Graf, E. W., & Ernst, M. O. (2004). Experience can change the “light-from-above” prior. *Nature Neuroscience*, 7(10), 1057–1058.
- Alloy, L. B., & Abramson, L. Y. (1979). Judgment of contingency in depressed and nondepressed students: Sadder but wiser ? *Journal of Experimental Psychology: General*, 108(4), 441–485.
- Andersen, R. A., & Bradley, D. C. (1998). Perception of three-dimensional structure from motion. *Trends in Cognitive Sciences*, 2(6), 222–228.
- Attarha, M., & Moore, C. M. (2015). Onset rivalry: Factors that succeed and fail to bias selection. *Attention, Perception, & Psychophysics*, 77(2), 520–535.
- Baess, P., Horváth, J., Jacobsen, T., & Schröger, E. (2011). Selective suppression of self-initiated sounds in an auditory stream: An ERP study. *Psychophysiology*, 48(9), 1276–1283.
- Baess, P., Jacobsen, T., & Schröger, E. (2008). Suppression of the auditory N1 event-related potential component with unpredictable self-initiated tones: Evidence for internal forward models with dynamic stimulation. *International Journal of Psychophysiology*, 70(2), 137–143.
- Baess, P., Widmann, A., Roye, A., Schröger, E., & Jacobsen, T. (2009). Attenuated human auditory middle latency response and evoked 40-Hz response to self-initiated sounds. *European Journal of Neuroscience*, 29(7), 1514–1521.
- Bandura, A. (2001). Social cognitive theory: An agentic perspective. *Annual Review of Psychology*, 52, 1–26.
- Bays, P. M., Wolpert, D. M., & Flanagan, J. R. (2005). Perception of the consequences of self-action is temporally tuned and event driven. *Current Biology*, 15(12), 1125–1128.
- Bednark, J. G., Poonian, S. K., Palghat, K., McFadyen, J., & Cunnington, R. (2015). Identity-specific predictions and implicit measures of agency. *Psychology of Consciousness: Theory, Research, and Practice*, 2(3), 253–268.
- Beets, I. A. M., Hart, B. M. t, Rösler, F., Henriques, D. Y. P., Einhäuser, W., & Fiehler, K. (2010). Online action-to-perception transfer: Only percept-dependent action affects perception. *Vision Research*, 50(24), 2633–2641.
- Behroozmand, R., Liu, H., & Larson, C. R. (2011). Time-dependent neural processing of auditory feedback during voice pitch error detection. *Journal of Cognitive Neuroscience*, 23(5), 1205–1217.
- Belayachi, S., & Van der Linden, M. (2010). Feeling of doing in obsessive-compulsive checking. *Consciousness and Cognition*, 19(2), 534–546.
- Bell, C. C. (1981). An efference copy which is modified by reafferent input. *Science*, 214(4519), 450–453.
- Benajmini, Y., Hochberg, Y., Benajmini, Y., & Hochberg, Y. (1995). Controlling the false discovery rate: A practical and powerful approach to multiple testing. *Journal of the Royal Statistical Society: Series B*, 57(1), 289–300.
- Berger, J. O. (2003). Could Fisher, Jeffreys and Neyman have agreed on testing? *Statistical Science*, 18(1), 1–32.
- Blake, R. (2001). A primer on binocular rivalry, including current controversies. *Brain and Mind*, 2(1), 5–38.
- Blake, R., & Logothetis, N. K. (2002). Visual competition. *Nature Reviews Neuroscience*, 3(1), 13–21.
- Blakemore, S. J., Wolpert, D. M., & Frith, C. D. (2000). Why can't you tickle yourself? *NeuroReport*, 11(11), R11–R16.
- Blakemore, S. J., Frith, C. D., & Wolpert, D. M. (1999). Spatio-temporal prediction modulates the perception of self-produced stimuli. *Journal of Cognitive Neuroscience*, 11(5), 551–559.
- Blakemore, S. J., Wolpert, D. M., & Frith, C. D. (1998). Central cancellation of self-produced tickle sensation. *Nature Neuroscience*, 1(7), 635–640.

- Blakemore, S. J., Wolpert, D. M., & Frith, C. D. (2002). Abnormalities in the awareness of action. *Trends in Cognitive Sciences*, 6(6), 237–242.
- Boersma, M., Kemner, C., De Reus, M. A., Collin, G., Snijders, T. M., Hofman, D., ... Van den Heuvel, M. P. (2012). Disrupted functional brain networks in autistic toddlers. *Brain Connectivity*, 3(1), 41–49.
- Brainard, D. H. (1997). The Psychophysics Toolbox. *Spatial Vision*, 10(4), 433–436.
- Brascamp, J., Blake, R., & Knapen, T. (2015). Negligible fronto-parietal BOLD activity accompanying unreportable switches in bistable perception. *Nature Neuroscience*, 18(11), 1672–1678.
- Brascamp, J., Sterzer, P., Blake, R., & Knapen, T. (2018). Multistable perception and the role of frontoparietal cortex in perceptual inference. *Annual Review of Psychology*, 69, 77–103.
- Brascamp, J. W., & Naber, M. (2017). Eye tracking under dichoptic viewing conditions: A practical solution. *Behavior Research Methods*, 49(4), 1303–1309.
- Braun, N., Thorne, J. D., Hildebrandt, H., & Debener, S. (2014). Interplay of agency and ownership: The intentional binding and rubber hand illusion paradigm combined. *PLoS ONE*, 9(11), e111967.
- Breitmeyer, B. G. (2014). *The visual (un)conscious and its (dis)contents: A microtemporal approach*. Oxford: Oxford University Press.
- Brouwer, G. J., & Ee, R. Van. (2006). Endogenous influences on perceptual bistability depend on exogenous stimulus characteristics. *Vision Research*, 46, 3393–3402.
- Bubic, A., von Cramon, D. Y., & Schubotz, R. I. (2010). Prediction, cognition and the brain. *Frontiers in Human Neuroscience*, 4, 25.
- Buzsaki, G., & Draguhn, A. (2004). Neuronal oscillations in cortical networks. *Science*, 304, 1926–1929.
- Cardoso-Leite, P., Mamassian, P., Schütz-Bosbach, S., & Waszak, F. (2010). A new look at sensory attenuation: Action-effect anticipation affects sensitivity, not response bias. *Psychological Science*, 21(12), 1740–1745.
- Carmel, D., Walsh, V., Lavie, N., & Rees, G. (2010). Right parietal TMS shortens dominance durations in binocular rivalry. *Current Biology*, 20(18), R799–R800.
- Carter, O., & Cavanagh, P. (2007). Onset rivalry: Brief presentation isolates an early independent phase of perceptual competition. *PLoS ONE*, 2(4), e343.
- Carver, C. S., & Scheier, M. F. (2000). On the structure of behavioral self-regulation. In M. Boekaerts, P. R. Pintrich, & M. Zeidner (Eds.), *Handbook of self-regulation* (pp. 41–84). San Diego, CA: Academic Press.
- Caspar, E. A., Christensen, J. F., Cleeremans, A., & Haggard, P. (2016). Coercion changes the sense of agency in the human brain. *Current Biology*, 26(5), 585–592.
- Chen, X., & He, S. (2004). Local factors determine the stabilization of monocular ambiguous and binocular rivalry stimuli. *Current Biology*, 14, 1013–1017.
- Christensen, M. S., & Grünbaum, T. (2017). Sense of moving : Moving closer to the movement. In M. S. Christensen & T. Grünbaum (Eds.), *Sensation of Movement* (pp. 64–84). Abingdon: Routledge.
- Clark, A. (2013). Whatever next? Predictive brains, situated agents, and the future of cognitive science. *Behavioral and Brain Sciences*, 36(3), 181–204.
- Claxton, G. U. Y. (1975). Why can't we tickle ourselves? *Perceptual and Motor Skills*, 41, 335–338.
- Crapse, T. B., & Sommer, M. A. (2008a). Corollary discharge across the animal kingdom. *Nature Reviews Neuroscience*, 9(8), 587–600.
- Crapse, T. B., & Sommer, M. A. (2008b). Corollary discharge circuits in the primate brain. *Current Opinion in Neurobiology*, 18(6), 552–557.
- Curio, G., Neuloh, G., Numminen, J., Jousmaki, V., & Hari, R. (2000). Speaking modifies voice-evoked activity in the human auditory cortex. *Human Brain Mapping*, 9(4), 183–191.
- Custers, R., & Aarts, H. (2011). Learning of predictive relations between events depends on attention, not on awareness. *Consciousness and Cognition*, 20(2), 368–378.
- Dannenberg, L., Förster, J., & Jostmann, N. B. (2012). “If only...”: When counterfactual thoughts can reduce illusions of personal authorship. *Consciousness and Cognition*, 21, 456–463.
- David, N. (2010). Functional anatomy of the sense of agency: Past evidence and future directions. In M. Balconi (Ed.), *Neuropsychology of the Sense of Agency* (pp. 69–80). Springer Milan.
- David, N. (2012). New frontiers in the neuroscience of the sense of agency. *Frontiers in Human Neuroscience*, 6, 161.

- David, N., Cohen, M. X., Newen, A., Bewernick, B. H., Shah, N. J., Fink, G. R., & Vogeley, K. (2007). The extrastriate cortex distinguishes between the consequences of one's own and others' behavior. *NeuroImage*, 36(3), 1004–1014.
- David, N., Newen, A., & Vogeley, K. (2008). The “sense of agency” and its underlying cognitive and neural mechanisms. *Consciousness and Cognition*, 17(2), 523–534.
- Davidson, P. R., & Wolpert, D. M. (2005). Widespread access to predictive models in the motor system: A short review. *Journal of Neural Engineering*, 2(3), S313–S319.
- De Lange, F. P., Heilbron, M., & Kok, P. (2018). How do expectations shape perception? *Trends in Cognitive Sciences*, 22(9), 764–779.
- 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.
- Denison, R. N., Piazza, E. A., & Silver, M. A. (2011). Predictive context influences perceptual selection during binocular rivalry. *Frontiers in Human Neuroscience*, 5, 166.
- Dennett, D. C. (1991). *Consciousness explained*. Boston, MA: Little, Brown and Company.
- 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.
- Desantis, A., Mamassian, P., Lisi, M., & Waszak, F. (2014). The prediction of visual stimuli influences auditory loudness discrimination. *Experimental Brain Research*, 232(10), 3317–3324.
- Desantis, A., Roussel, C., & Waszak, F. (2014). The temporal dynamics of the perceptual consequences of action-effect prediction. *Cognition*, 132(3), 243–250.
- Desmurget, M., & Grafton, S. (2000). Forward modeling allows feedback control for fast reaching movements. *Trends in Cognitive Sciences*, 4(11), 1–9.
- Dewey, J. A., & Knoblich, G. (2014). Do implicit and explicit measures of the sense of agency measure the same thing? *Plos One*, 9(10), e110118.
- Di Pace, E., & Saracini, C. (2014). Action imitation changes perceptual alternations in binocular rivalry. *PLoS ONE*, 9(5), e98305.
- Dogge, M., Custers, R., Gayet, S., Hoijtink, H., & Aarts, H. (2019). Perception of action-outcomes is shaped by life-long and contextual expectations. *Scientific Reports*, 9, 5225.
- Dogge, M., Gayet, S., Custers, R., & Aarts, H. (2018). The influence of action-effect anticipation on bistable perception: Differences between onset rivalry and ambiguous motion. *Neuroscience of Consciousness*, 4, niy004.
- Dogge, M., Hofman, D., Custers, R., & Aarts, H. (2019). Exploring the role of motor and non-motor predictive mechanisms in sensory attenuation: Perceptual and neurophysiological findings. *Neuropsychologia*, 124, 216–225.
- Dutzi, I. B., & Hommel, B. (2009). The microgenesis of action-effect binding. *Psychological Research*, 73(3), 425–435.
- Ebert, J. P., & Wegner, D. M. (2010). Time warp : Authorship shapes the perceived timing of actions and events. *Consciousness and Cognition*, 19(1), 481–489.
- Ellison-Wright, I., & Bullmore, E. (2009). Meta-analysis of diffusion tensor imaging studies in schizophrenia. *Schizophrenia Research*, 108(1–3), 3–10.
- Elsner, B., & Hommel, B. (2001). Effect anticipation and action control. *Journal of Experimental Psychology: Human Perception and Performance*, 27(1), 229–240.
- Engbert, K., Wohlschläger, A., & Haggard, P. (2008). Who is causing what? The sense of agency is relational and efferent-triggered. *Cognition*, 107(2), 693–704.
- Engbert, K., Wohlschläger, A., Thomas, R., & Haggard, P. (2007). Agency, subjective time, and other minds. *Journal of Experimental Psychology: Human Perception and Performance*, 33(6), 1261–1268.
- Engel, A. K., & Fries, P. (2010). Beta-band oscillations-signalling the status quo? *Current Opinion in Neurobiology*, 20(2), 156–165.
- Feinberg, I. (1978). Efference copy and corollary discharge: Implications for thinking and its disorders. *Schizophrenia Bulletin*, 4(4), 636–640.
- Fisbach, A., & Ferguson, M. J. (2007). The goal construct in social psychology. In A. W. Kruglanski and E. T. Higgins (Eds.), *Social psychology: Handbook of basic principles* (pp. 490–515). New York, NY: Guilford Press.

- Förster, J., Liberman, N., & Friedman, R. S. (2007). Seven principles of goal activation: A systematic approach to distinguishing goal priming from priming of non-goal constructs. *Personality and Social Psychology Review, 11*(3), 211–233.
- Fries, P. (2005). A mechanism for cognitive dynamics: Neuronal communication through neuronal coherence. *Trends in Cognitive Sciences, 9*(10), 474–480.
- Friston, K. (2010). The free-energy principle: A unified brain theory? *Nature Reviews Neuroscience, 11*(2), 127–138.
- Frith, C. D., Blakemore, S. J., & Wolpert, D. M. (2000a). Explaining the symptoms of schizophrenia: Abnormalities in the awareness of action. *Brain Research Reviews, 31*, 357–363.
- Frith, C. D., Blakemore, S. J., & Wolpert, D. M. (2000b). Abnormalities in the awareness and control of action. *Philosophical Transactions of the Royal Society B: Biological Sciences, 355*(1404), 1771–1788.
- Gayet, S., Brascamp, J. W., Van der Stigchel, S., & Paffen, C. L. E. (2015). Cogito ergo video: Task-relevant information is involuntarily boosted into awareness. *Journal of Vision, 15*(5), 3.
- 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*, 460.
- Gekas, N., Seitz, A. R., & Seriès, P. (2015). Expectations developed over multiple timescales facilitate visual search performance. *Journal of Vision, 15*(9), 10.
- Gentsch, A., Endrass, T., & Kathmann, N. (2012). Dysfunctional forward model mechanisms and aberrant sense of agency in obsessive-compulsive disorder. *Biological Psychiatry, 71*(7), 652–659.
- Gilbert, D. T., & Wilson, T. D. (2007). Propection: Experiencing the future. *Science, 317*(5843), 1351–1355.
- Giles, N., Lau, H., & Odegaard, B. (2016). What type of awareness does binocular rivalry assess? *Trends in Cognitive Sciences, 20*(10), 719–720.
- Gratton, G., Coles, M. G. H., & Donchin, E. (1983). A new method for off-line removal of ocular artifact. *Electroencephalography and Clinical Neurophysiology, 55*(4), 468–484.
- Grüsser, O. J. (1995). On the history of the ideas of efference copy and reafference. In C. Debru (Ed.), *Essays in the history of the physiological sciences, Clío Medica 33* (pp. 35–55).
- Gu, X., Hoijtink, H., Mulder, J., & Rosseel, Y. (2018). Bain: A program for Bayesian testing of order constrained hypotheses in structural equation models. *Unpublished*.
- Gu, X., Mulder, J., & Hoijtink, H. (2018). Approximated adjusted fractional Bayes factors: A general method for testing informative hypotheses. *British Journal of Mathematical and Statistical Psychology, 71*(2), 229–261.
- Haering, C., & Kiesel, A. (2014). Intentional binding is independent of the validity of the action effect's identity. *Acta Psychologica, 152*, 109–119.
- Haggard, P. (2005). Conscious intention and motor cognition. *Trends in Cognitive Sciences, 9*(6), 290–295.
- Haggard, P. (2017). Sense of agency in the human brain. *Nature Reviews Neuroscience*.
- Haggard, P., & Chambon, V. (2012). Sense of agency. *Current Biology, 22*(10), R390–R392.
- Haggard, P., & Clark, S. (2003). Intentional action: Conscious experience and neural prediction. *Consciousness and Cognition, 12*(4), 695–707.
- Haggard, P., Clark, S., & Kalogeras, J. (2002). Voluntary action and conscious awareness. *Nature Neuroscience, 5*(4), 382–385.
- Haggard, P., & Tsakiris, M. (2009). The experience of agency: Feelings, judgments, and responsibility. *Current Directions in Psychological Science, 18*(4), 242–246.
- Haijiang, Q., Saunders, J. A., Stone, R. W., & Backus, B. T. (2006). Demonstration of cue recruitment: Change in visual appearance by means of Pavlovian conditioning. *Proceedings of the National Academy of Sciences, 103*(2), 483–488.
- Haynes, J. D., Deichmann, R., & Rees, G. (2005). Eye-specific effects of binocular rivalry in the human lateral geniculate nucleus. *Nature, 438*(7067), 496–499.
- Heinks-Maldonado, T. H., Mathalon, D. H., Gray, M., & Ford, J. M. (2005). Fine-tuning of auditory cortex during speech production. *Psychophysiology, 42*(2), 180–190.
- Heinks-Maldonado, T. H., Nagarajan, S. S., & Houde, J. F. (2006). Magnetoencephalographic evidence for a precise forward model in speech production. *NeuroReport, 17*(13), 1375–1379.

- Helmholtz, H. (1962). *Treatise on physiological optics*. (J. P. C. Southall, Ed.) (Vol. 3). New York: Dover. (Original work published in 1867).
- Herwig, A. (2015). Linking perception and action by structure or process? Toward an integrative perspective. *Neuroscience and Biobehavioral Reviews*, 52, 105–116.
- Herwig, A., Beisert, M., & Prinz, W. (2013). Action science emerging: Introduction and leitmotifs. In W. Prinz, M. Beisert, & A. Herwig (Eds.), *Action science: Foundations of an emerging discipline* (pp. 1–33). Cambridge, MA: MIT Press.
- Heyes, C., Bird, G., Johnson, H., & Haggard, P. (2005). Experience modulates automatic imitation. *Cognitive Brain Research*, 22(2), 233–240.
- Hojitink, H. (2012). *Informative hypotheses: Theory and practice for behavioral and social scientists*. Boca Raton, FL: Chapman and Hall/CRC.
- Hojitink, H., Gu, X., & Mulder, J. (2018). Bayesian evaluation of informative hypotheses for multiple populations. *British Journal of Mathematical and Statistical Psychology*, 3, 14.
- Hojitink, H., Mulder, J., Van Lissa, C., & Gu, X. (2018). Tutorial: Testing hypotheses using the Bayes factor. *Psychological Methods*.
- Hol, K., Koene, A., & Van Ee, R. (2003). Attention-biased multi-stable surface perception in three-dimensional structure-from-motion. *Journal of Vision*, 3(7), 486–498.
- Hommel, B., Müsseler, J., Aschersleben, G., & Prinz, W. (2001). The Theory of Event Coding (TEC): A framework for perception and action planning. *Behavioral and Brain Sciences*, 24(5), 849–878.
- Horváth, J. (2015). Action-related auditory ERP attenuation: Paradigms and hypotheses. *Brain Research*, 1626, 54–65.
- Hsu, Y. F., Hämäläinen, J. A., & Waszak, F. (2013). Temporal expectation and spectral expectation operate in distinct fashion on neuronal populations. *Neuropsychologia*, 51(13), 2548–2555.
- Hsu, Y. F., Hämäläinen, J. A., & Waszak, F. (2014a). Repetition suppression comprises both attention-independent and attention-dependent processes. *NeuroImage*, 98, 168–175.
- Hsu, Y. F., Hämäläinen, J. A., & Waszak, F. (2014b). Both attention and prediction are necessary for adaptive neuronal tuning in sensory processing. *Frontiers in Human Neuroscience*, 8, 152.
- Hsu, Y., Le Bars, S., Hämäläinen, J. A., & Waszak, F. (2015). Distinctive representation of mispredicted and unpredicted prediction errors in human electroencephalography. *The Journal of Neuroscience*, 35(43), 14653–14660.
- Hughes, G., Desantis, A., & Waszak, F. (2013a). Mechanisms of intentional binding and sensory attenuation: The role of temporal prediction, temporal control, identity prediction, and motor prediction. *Psychological Bulletin*, 139(1), 133–151.
- Hughes, G., Desantis, A., & Waszak, F. (2013b). Attenuation of auditory N1 results from identity-specific action-effect prediction. *European Journal of Neuroscience*, 37(7), 1152–1158.
- Hughes, G., & Waszak, F. (2014). Predicting faces and houses: Category-specific visual action-effect prediction modulates late stages of sensory processing. *Neuropsychologia*, 61, 11–18.
- Hugrass, L., & Crewther, D. (2012). Willpower and conscious percept: Volitional switching in binocular rivalry. *PLoS ONE*, 7(4), e35963.
- Hur, J. W., Kwon, J. S., Lee, T. Y., & Park, S. (2014). The crisis of minimal self-awareness in schizophrenia: A meta-analytic review. *Schizophrenia Research*, 152(1), 58–64.
- Imazumi, S., & Tanno, Y. (2019). Intentional binding coincides with explicit sense of agency. *Consciousness and Cognition*, 67, 1–15.
- Imamizu, H., Sugimoto, N., Osu, R., Tsutsui, K., Sugiyama, K., Wada, Y., & Kawato, M. (2007). Explicit contextual information selectively contributes to predictive switching of internal models. *Experimental Brain Research*, 181(3), 395–408.
- Ishihara S. (1917). Test for Colour-Blindness. Tokyo: Hongo Harukicho.
- Jeffreys, H. (1961). *Theory of probability* (3rd ed.). Oxford, UK: Oxford University Press.
- Jones, A., Hughes, G., & Waszak, F. (2013). The interaction between attention and motor prediction. An ERP study. *NeuroImage*, 83, 533–541.
- Jones, S. R., de-Wit, L., Fernyhough, C., & Meins, E. (2008). A new spin on the Wheel of Fortune: Priming of action-authorship judgements and relation to psychosis-like experiences. *Consciousness and Cognition*, 17(3), 576–586.

- Jordan, M. (2003). Forward models: Supervised learning with a distal teacher. *Cognitive Science*, 16(3), 307–354.
- Kaiser, J., & Schütz-Bosbach, S. (2018). Sensory attenuation of self-produced signals does not rely on self-specific motor predictions. *European Journal of Neuroscience*, 47(11), 1303–1310.
- Kawabata, N. (1986). Attention and depth perception. *Perception*, 15(5), 563–572.
- Kawabe, T. (2013). Side effect of acting on the world: Acquisition of action-outcome statistic relation alters visual interpretation of action outcome. *Frontiers in Human Neuroscience*, 7, 610.
- Kawato, M. (1999). Internal models for motor control and trajectory planning. *Current Opinion in Neurobiology*, 9(6), 718–727.
- Kesten, H. (1958). Accelerated stochastic approximation. *Annals of Mathematical Statistics*, 29(1), 41–59.
- Klimesch, W., Schack, B., & Sauseng, P. (2005). The functional significance of theta and upper alpha oscillations. *Experimental Psychology*, 52(2), 99–108.
- Kruschke, J. K., & Blair, N. J. (2000). Blocking and backward blocking involve learned inattention. *Psychological Bulletin & Review*, 7(4), 636–645.
- Kühn, S., Nenchev, I., Haggard, P., Brass, M., Gallinat, J., & Voss, M. (2011). Whodunnit? Electrophysiological correlates of agency judgements. *PLoS ONE*, 6(12), e28657.
- Kunde, W., & Kiesel, A. (2006). See what you've done! Active touch affects the number of perceived visual objects. *Psychonomic Bulletin and Review*, 13(2), 304–309.
- Lange, K. (2009). Brain correlates of early auditory processing are attenuated by expectations for time and pitch. *Brain and Cognition*, 69(1), 127–137.
- Le Pelley, M. E., & McLaren, I. P. L. (2003). Learned associability and associative change in human causal learning. *The Quarterly Journal of Experimental Psychology*, 56(1), 67–79.
- Lee, M. D., & Wagenmakers, E. J. (2013). *Bayesian data analysis for cognitive science: A practical course*. New York: Cambridge University Press.
- Lee, S., & Blake, R. (2002). V1 activity is reduced during binocular rivalry. *Journal of Vision*, 2(9), 618–626.
- Leopold, D. A., & Logothetis, N. K. (1999). Multistable phenomena: Changing views in perception. *Trends in Cognitive Sciences*, 3(7), 254–264.
- Libet, B., Gleason, C. A., Wright, E. W., & Pearl, D. K. (1983). Time of conscious intention to act in relation to onset of cerebral activity (readiness-potential). *Brain*, 106, 623–642.
- Liefvooghe, B., Wenke, D., & De Houwer, J. (2012). Instruction-based task-rule congruency effects. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 38(5), 1325–1335.
- Limerick, H., Coyle, D., & Moore, J. W. (2014). The experience of agency in human-computer interactions: A review. *Frontiers in Human Neuroscience*, 8, 643.
- Lin, Z., & He, S. (2009). Seeing the invisible: The scope and limits of unconscious processing in binocular rivalry. *Progress in Neurobiology*, 87(4), 195–211.
- Linser, K., & Goschke, T. (2007). Unconscious modulation of the conscious experience of voluntary control. *Cognition*, 104(3), 459–475.
- Lunghi, C., & Alais, D. (2013). Touch interacts with vision during binocular rivalry with a tight orientation tuning. *PLoS ONE*, 8(3), e58754.
- Lunghi, C., Binda, P., & Morrone, M. C. (2010). Touch disambiguates rivalrous perception at early stages of visual analysis. *Current Biology*, 20(4), R143–R144.
- Mackintosh, N. J. (1975). A theory of attention: Variations in the associability of stimuli with reinforcement. *Psychological Review*, 82(4), 276–298.
- Marescaux, J., Leroy, J., Gagner, M., Rubino, F., Mutter, D., Vix, M., ... Smith, M. K. (2001). Transatlantic robot-assisted telesurgery. *Nature*, 413(6854), 379–380.
- Martikainen, M. H. (2005). Suppressed responses to self-triggered sounds in the human auditory cortex. *Cerebral Cortex*, 15(3), 299–302.
- Maruya, K., Yang, E., & Blake, R. (2007). Voluntary action influences visual competition. *Psychological Science*, 18(12), 1090–1098.
- Meiran, N., Liefvooghe, B., & De Houwer, J. (2017). Powerful instructions: Automaticity without practice. *Current Directions in Psychological Science*, 26(6), 509–514.

- Meng, M., & Tong, F. (2004). Can attention selectively bias bistable perception? Differences between binocular rivalry and ambiguous figures. *Journal of Vision*, 4(7), 539–551.
- Miall, R. C. ., & Wolpert, D. M. (1996). Forward models for physiological motor control. *Neural Networks*, 9(8), 1265–1279.
- Mifsud, N. G., Beesley, T., Watson, T. L., Elijah, R. B., Sharp, T. S., & Whitford, T. J. (2018). Attenuation of visual evoked responses to hand and saccade-initiated flashes. *Cognition*, 179, 14–22.
- Mifsud, N. G., Beesley, T., Watson, T. L., & Whitford, T. J. (2016). Attenuation of auditory evoked potentials for hand and eye-initiated sounds. *Biological Psychology*, 120, 61–68.
- Mifsud, N. G., & Whitford, T. J. (2017). Sensory attenuation of self-initiated sounds maps onto habitual associations between motor action and sound. *Neuropsychologia*, 103, 38–43.
- Miller, G. A., Galanter, E. H., & Pribram, K. H. (1960). *Plans and the structure of behavior*. New York, NY: Holt, Rinehart, & Winston, Inc.
- Mitsumatsu, H. (2009). Voluntary action affects perception of bistable motion display. *Perception*, 38(10), 1522–1535.
- Moore, J., & Haggard, P. (2008). Awareness of action: Inference and prediction. *Consciousness and Cognition*, 17(1), 136–144.
- Moore, J. W. (2016). What is the sense of agency and why does it matter? *Frontiers in Psychology*, 7, 1272.
- Moore, J. W., & Fletcher, P. C. (2012). Sense of agency in health and disease: A review of cue integration approaches. *Consciousness and Cognition*, 21(1), 59–68.
- Moore, J. W., Lagnado, D., Deal, D. C., & Haggard, P. (2009). Feelings of control: Contingency determines experience of action. *Cognition*, 110(2), 279–283.
- Moore, J. W., & Obhi, S. S. (2012). Intentional binding and the sense of agency: A review. *Consciousness and Cognition*, 21(1), 546–561.
- Moran, J. M., Macrae, C. N., Heatherton, T. F., Wyland, C. L., & Kelley, W. M. (2006). Neuroanatomical evidence for distinct cognitive and affective components of self. *Journal of Cognitive Neuroscience*, 18(9), 1586–1594.
- Morey, R. D. (2008). Confidence intervals from normalized data: A correction to Cousineau (2005). *Tutorials in Quantitative Methods for Psychology*, 4(2), 61–64.
- Muthukumaraswamy, S. D., & Singh, K. D. (2011). A cautionary note on the interpretation of phase-locking estimates with concurrent changes in power. *Clinical Neurophysiology*, 122(11), 2324–2325.
- Näätänen, R., & Picton, T. (1987). The N1 wave of the human electric and magnetic response to sound: A review and an analysis of the component structure. *Psychophysiology*, 24(4), 375–425.
- Nahab, F. B., Kundu, P., Gallea, C., Kakareka, J., Pursley, R., Pohida, T., ... Hallett, M. (2011). The neural processes underlying self-agency. *Cerebral Cortex*, 21(1), 48–55.
- Newman-Norlund, R. D., Van Schie, H. T., Van Zuijlen, A. M. J., & Bekkering, H. (2007). The mirror neuron system is more active during complementary compared with imitative action. *Nature Neuroscience*, 10(7), 817–818.
- Niziolek, C. A., Nagarajan, S. S., & Houde, J. F. (2013). What does motor efference copy represent? Evidence from speech production. *Journal of Neuroscience*, 33(41), 16110–16116.
- Obhi, S. S., Swiderski, K. M., & Farquhar, R. (2013). Activating memories of depression alters the experience of voluntary action. *Experimental Brain Research*, 229(3), 497–506.
- Oren, E., Friedmann, N., & Dar, R. (2016). Things happen: Individuals with high obsessive-compulsive tendencies omit agency in their spoken language. *Consciousness and Cognition*, 42, 125–134.
- 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(6), 532–552.
- Pearson, J., & Brascamp, J. (2008). Sensory memory for ambiguous vision. *Trends in Cognitive Sciences*, 12(9), 334–341.
- Pelli, D. G. (1997). The VideoToolbox software for visual psychophysics: Transforming numbers into movies. *Spatial Vision*, 10(4), 437–442.
- Pfister, R. (2008). Pretris – Tetris for presentation. *Arch. Neurobehav. Exp. Stimuli* 198.
- Pfister, R. (2019). Effect-based action control with body-related effects: Implications for empirical approaches to ideomotor action control. *Psychological Review*, 126(1), 153–161.

- Pickering, M. J., & Clark, A. (2014). Getting ahead: Forward models and their place in cognitive architecture. *Trends in Cognitive Sciences*, 18(9), 451–456.
- Pilditch, T. D., & Custers, R. (2017). Communicated beliefs about action-outcomes: The role of initial confirmation in the adoption and maintenance of unsupported beliefs. *Acta Psychologica*, 184, 46–63.
- Poulet, J. F. A., & Hedwig, B. (2006). The cellular basis of a corollary discharge. *Science*, 311, 518–522.
- Prinz, W. (1997). Perception and action planning. *European Journal of Cognitive Psychology*, 9(2), 129–154.
- Pyasik, M., Burin, D., & Pia, L. (2018). On the relation between body ownership and sense of agency: A link at the level of sensory-related signals. *Acta Psychologica*, 185, 219–228.
- Rameson, L. T., Satpute, A. B., & Lieberman, M. D. (2010). The neural correlates of implicit and explicit self-relevant processing. *NeuroImage*, 50(2), 701–708.
- Renes, R. A., Van Haren, N. E. M., & Aarts, H. (2015). Attentional control and inferences of agency: Working memory load differentially modulates goal-based and prime-based agency experiences. *Consciousness and Cognition*, 38, 38–49.
- Renes, R. A., Vermeulen, L., Kahn, R. S., Aarts, H., & Van Haren, N. E. M. (2013). Abnormalities in the establishment of feeling of self-agency in schizophrenia. *Schizophrenia Research*, 143(1), 50–54.
- Repp, B. H., & Knoblich, G. (2007). Action can affect auditory perception. *Psychological Science*, 18(1), 6–7.
- Reznik, D., & Mukamel, R. (2019). Motor output, neural states and auditory perception. *Neuroscience & Biobehavioral Reviews*, 96, 116–126.
- Richters, D. P., & Eskew, R. T. (2009). Quantifying the effect of natural and arbitrary sensorimotor contingencies on chromatic judgments. *Journal of Vision*, 9(4), 27.
- Rizzolatti, G., Fogassi, L., & Gallese, V. (2001). Neurophysiological mechanisms underlying the understanding and imitation of action. *Nature Reviews Neuroscience*, 2(9), 661–670.
- Rouder, J. N. (2014). Optional stopping: no problem for Bayesians. *Psychonomic Bulletin & Review*, 21(2), 301–308.
- Roussel, C., Hughes, G., & Waszak, F. (2013). A preactivation account of sensory attenuation. *Neuropsychologia*, 51(5), 922–929.
- Roussel, C., Hughes, G., & Waszak, F. (2014). Action prediction modulates both neurophysiological and psychophysical indices of sensory attenuation. *Frontiers in Human Neuroscience*, 8, 115.
- Ruess, M., Thomaschke, R., & Kiesel, A. (2017). The time course of intentional binding. *Attention, Perception, and Psychophysics*, 79(4), 1123–1131.
- Rummell, B. P., Klee, J. L., & Sigurdsson, X. T. (2016). Attenuation of responses to self-generated sounds in auditory cortical neurons. *Journal of Neuroscience*, 36(47), 12010–12026.
- Ruys, K. I., & Aarts, H. (2012). I didn't mean to hurt you! Unconscious origins of experienced self-agency over others' emotions. *Emotion*, 12(1), 132–141.
- Saadatian, E., Samani, H., Parsani, R., Pandey, A. V., Li, J., Tejada, L., ... Nakatsu, R. (2014). Mediating intimacy in long-distance relationships using kiss messaging. *International Journal of Human-Computer Studies*, 72(10–11), 736–746.
- Saito, N., Takahata, K., Murai, T., & Takahashi, H. (2015). Discrepancy between explicit judgement of agency and implicit feeling of agency: Implications for sense of agency and its disorders. *Consciousness and Cognition*, 37, 1–7.
- Salinas, E., & Sejnowski, T. J. (2001). Correlated neuronal activity and the flow of neural information. *Nature Reviews Neuroscience*, 2(8), 539–550.
- 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.
- Sanmiguel, I., Todd, J., & Schröger, E. (2013). Sensory suppression effects to self-initiated sounds reflect the attenuation of the unspecific N1 component of the auditory ERP. *Psychophysiology*, 50(4), 334–343.
- Sato, A. (2008). Action observation modulates auditory perception of the consequence of others' actions. *Consciousness and Cognition*, 17(4), 1219–1227.
- Sato, A. (2009). Both motor prediction and conceptual congruency between preview and action-effect contribute to explicit judgment of agency. *Cognition*, 110(1), 74–83.

- Sato, A., & Yasuda, A. (2005). Illusion of sense of self-agency: Discrepancy between the predicted and actual sensory consequences of actions modulates the sense of self-agency, but not the sense of self-ownership. *Cognition*, *94*, 241–255.
- Saupe, K., Widmann, A., Trujillo-Barreto, N. J., & Schröger, E. (2013). Sensorial suppression of self-generated sounds and its dependence on attention. *International Journal of Psychophysiology*, *90*(3), 300–310.
- Sauseng, P., & Klimesch, W. (2008). What does phase information of oscillatory brain activity tell us about cognitive processes? *Neuroscience and Biobehavioral Reviews*, *32*(5), 1001–1013.
- Schafer, E., & Marcus, M. (1973). Self-stimulation alters human sensory brain responses. *Science*, *181*(4095), 175–177.
- Schmack, K., Gómez-Carrillo de Castro, A., Rothkirch, M., Sekutowicz, M., Rössler, H., Haynes, J.-D., ... Sterzer, P. (2013). Delusions and the role of beliefs in perceptual inference. *The Journal of Neuroscience*, *33*(34), 13701–13712.
- Schmack, K., Weilhhammer, V., Heinzle, J., Stephan, K. E., & Sterzer, P. (2016). Learning what to see in a changing world. *Frontiers in Human Neuroscience*, *10*, 263.
- Schneider, D. M., & Mooney, R. (2018). How movement modulates hearing. *Annual Review of Neuroscience*, *41*, 553–572.
- Schneider, D. M., Nelson, A., & Mooney, R. (2014). A synaptic and circuit basis for corollary discharge in the auditory cortex. *Nature*, *513*(7517), 189–194.
- Schneider, D. M., Sundararajan, J., & Mooney, R. (2018). A cortical filter that learns to suppress the acoustic consequences of movement. *Nature*, *561*(7723), 391–395.
- Schnitzler, A., & Gross, J. (2005). Normal and pathological oscillatory communication in the brain. *Nature Reviews Neuroscience*, *6*(4), 285–296.
- Schönbrodt, F. D., & Wagenmakers, E. J. (2018). Bayes factor design analysis: Planning for compelling evidence. *Psychonomic Bulletin & Review*, *25*(1), 128–142.
- Schröger, E., Marzecová, A., & Sanmiguel, I. (2015). Attention and prediction in human audition: A lesson from cognitive psychophysiology. *European Journal of Neuroscience*, *41*(5), 641–664.
- Schütz-Bosbach, S., & Prinz, W. (2007). Perceptual resonance: Action-induced modulation of perception. *Trends in Cognitive Sciences*, *11*(8), 349–355.
- Scocchia, L. (2014). Top-down influences on ambiguous perception: The role of stable and transient states of the observer, 1–18.
- Scocchia, L., Valsecchi, M., Gegenfurtner, K. R., & Triesch, J. (2013). Visual working memory contents bias ambiguous structure from motion perception. *PLoS ONE*, *8*(3), e59217.
- Seriès, P., & Seitz, A. R. (2013). Learning what to expect (in visual perception). *Frontiers in Human Neuroscience*, *7*, 668.
- Shergill, S. S., Bays, P. H., Frith, C. D., & Wolpert, D. M. (2003). Two eyes for an eye: The neuroscience of force escalation. *Science*, *301*(5630), 187.
- Sotiropoulos, G., Seitz, A. R., & Seriès, P. (2011). Changing expectations about speed alters perceived motion direction. *Current Biology*, *21*(21), R883–R884.
- Sperduti, M., Delaveau, P., Fossati, P., & Nadel, J. (2011). Different brain structures related to self- and external-agency attribution: A brief review and meta-analysis. *Brain Structure and Function*, *216*(2), 151–157.
- Sperry, R. (1950). Neural basis of the spontaneous optokinetic response produced by visual inversion. *Journal of Comparative and Physiological Psychology*, *43*(6), 482–489.
- Stam, C. J., Nolte, G., & Daffertshofer, A. (2007). Phase lag index: Assessment of functional connectivity from multi channel EEG and MEG with diminished bias from common sources. *Human Brain Mapping*, *28*(11), 1178–1193.
- Stam, C. J., & Van Straaten, E. C. W. (2012a). The organization of physiological brain networks. *Clinical Neurophysiology*, *123*(6), 1067–1087.
- Stam, C. J., & Van Straaten, E. C. W. (2012b). Go with the flow: Use of a directed phase lag index (dPLI) to characterize patterns of phase relations in a large-scale model of brain dynamics. *NeuroImage*, *62*(3), 1415–1428.
- Stanley, J., Carter, O., & Forte, J. (2011). Color and luminance influence, but can not explain, binocular rivalry onset bias. *PLoS ONE*, *6*(5), e18978.

- Stanley, J., Forte, J., Cavanagh, P., & Carter, O. (2011). Onset rivalry: The initial dominance phase is independent of ongoing perceptual alternations. *Frontiers in Human Neuroscience*, 5, 140.
- Stekelenburg, J. J., & Vroomen, J. (2015). Predictive coding of visual–auditory and motor–auditory events: An electrophysiological study. *Brain Research*, 1626, 88–96.
- Sterzer, P., Frith, C., & Petrovic, P. (2008). Believing is seeing: Expectations alter visual awareness. *Current Biology*, 18(16), R697–R698.
- Sterzer, P., Kleinschmidt, A., & Rees, G. (2009). The neural bases of multistable perception. *Trends in Cognitive Sciences*, 13(7), 310–318.
- Straka, H., Simmers, J., & Chagnaud, B. P. (2018). A new perspective on predictive motor signaling. *Current Biology*, 28(5), R193–R194.
- Synofzik, M., Vosgerau, G., & Newen, A. (2008). Beyond the comparator model: A multifactorial two-step account of agency. *Consciousness and Cognition*, 17(1), 219–239.
- Synofzik, M., Vosgerau, G., & Voss, M. (2013). The experience of agency: An interplay between prediction and postdiction. *Frontiers in Psychology*, 4, 127.
- Timm, J., Sanmiguel, I., Keil, J., Schröger, E., & Schönwiesner, M. (2014). Motor intention determines sensory attenuation of brain responses to self-initiated sounds. *Journal of Cognitive Neuroscience*, 26(7), 1481–1489.
- Timm, J., Sanmiguel, I., Saupé, K., & Schröger, E. (2013). The N1-suppression effect for self-initiated sounds is independent of attention. *BMC Neuroscience*, 14, 2.
- Timm, J., Schönwiesner, M., Schröger, E., & SanMiguel, I. (2016). Sensory suppression of brain responses to self-generated sounds is observed with and without the perception of agency. *Cortex*, 80, 5–20.
- Tsakiris, M., & Haggard, P. (2003). Awareness of somatic events associated with a voluntary action. *Experimental Brain Research*, 149(4), 439–446.
- Van der Weiden, A., Aarts, H., & Ruys, K. I. (2011). Prime and probability: Causal knowledge affects inferential and predictive effects on self-agency experiences. *Consciousness and Cognition*, 20(4), 1865–1871.
- Van der Weiden, A., Ruys, K. I., Aarts, H., & Ruys, K. I. (2013). A matter of matching: How goals and primes affect self-agency experiences. *Journal of Experimental Psychology: General*, 142(3), 954–966.
- Van Elk, M., Salomon, R., Kannape, O., & Blanke, O. (2014). Suppression of the N1 auditory evoked potential for sounds generated by the upper and lower limbs. *Biological Psychology*, 102, 108–117.
- Van Schie, H. T., Van Waterschoot, B. M., & Bekkering, H. (2008). Understanding action beyond imitation: Reversed compatibility effects of action observation in imitation and joint action. *Journal of Experimental Psychology: Human Perception and Performance*, 34(6), 1493–1500.
- Varela, F., Lachaux, J. P., Rodriguez, E., & Martinerie, J. (2001). The brainweb: Phase synchronization and large-scale integration. *Nature Reviews Neuroscience*, 2(4), 229–239.
- Vinck, M., Oostenveld, R., Van Wingerden, M., Battaglia, F., & Pennartz, C. M. A. (2011). An improved index of phase-synchronization for electrophysiological data in the presence of volume-conduction, noise and sample-size bias. *NeuroImage*, 55(4), 1548–1565.
- Voineskos, A. N., Lobaugh, N. J., Bouix, S., Rajji, T. K., Miranda, D., Kennedy, J. L., ... Shenton, M. E. (2010). Diffusion tensor tractography findings in schizophrenia across the adult lifespan. *Brain*, 133(5), 1494–1504.
- Von Holst, H., & Mittelstaedt, H. (1971). The principle of reafference: Interactions between the central nervous system and the peripheral organs. In P. C. Dodwell (Ed.), *Perceptual processing: Stimulus equivalence and pattern recognition* (pp. 41–71). New York, NY: Appleton-Century-Crofts. (Original work published in 1950).
- Vroomen, J., & Stekelenburg, J. (2010). Visual anticipatory information modulates multisensory interactions of artificial audiovisual stimuli. *Journal of Cognitive Neuroscience*, 22(7), 1583–1596.
- Wagenmakers, E.-J., Marsman, M., Jamil, T., Ly, A., Verhagen, J., Love, J., ... Morey, R. D. (2018). Bayesian inference for psychology. Part I: Theoretical advantages and practical ramifications. *Psychonomic Bulletin & Review*, 25(1), 35–57.
- Wagenmakers, E. J., Wetzels, R., Borsboom, D., Van der Maas, H. L. J., & Kievit, R. A. (2012). An agenda for purely confirmatory research. *Perspectives on Psychological Science*, 7(6), 632–638.

- Wallis, G. M., & Backus, B. T. (2016). When action conditions perception: Evidence of cross-modal cue recruitment. *Journal of Vision, 16*(14), 6.
- Wang, J., Mathalon, D. H., Roach, B. J., Reilly, J., Keedy, S. K., Sweeney, J. A., & Ford, J. M. (2014). Action planning and predictive coding when speaking. *NeuroImage, 91*, 91–98.
- Waszak, F., Cardoso-Leite, P., & Hughes, G. (2012). Action effect anticipation: Neurophysiological basis and functional consequences. *Neuroscience & Biobehavioral Reviews, 36*(2), 943–959.
- Wegner, D. M. (2002). *The illusion of conscious will*. Cambridge, MA: MIT Press.
- Wegner, D. M., Sparrow, B., & Winerman, L. (2004). Vicarious agency: Experiencing control over the movements of others. *Journal of Personality and Social Psychology, 86*(6), 838–848.
- Wegner, D. M., & Wheatley, T. (1999). Apparent mental causation: Sources of the experience of will. *American Psychologist, 54*(7), 480–492.
- Weiskrantz, L. (1971). Preliminary observations on tickling oneself. *Nature, 230*, 598–599.
- Weiss, C., Herwig, A., & Schütz-Bosbach, S. (2011a). The self in action effects: Selective attenuation of self-generated sounds. *Cognition, 121*(2), 207–218.
- Weiss, C., Herwig, A., & Schütz-Bosbach, S. (2011b). The self in social interactions: Sensory attenuation of auditory action effects is stronger in interactions with others. *PLoS ONE, 6*(7), e22723.
- Weiss, C., & Schütz-Bosbach, S. (2012). Vicarious action preparation does not result in sensory attenuation of auditory action effects. *Consciousness and Cognition, 21*(4), 1654–1661.
- Whitford, T. J., Kubicki, M., & Shenton, M. E. (2011). Diffusion tensor imaging, structural connectivity, and schizophrenia. *Schizophrenia Research and Treatment, 709523*.
- Wirth, R., Steinhauser, R., Janczyk, M., Steinhauser, M., & Kunde, W. (2018). Long-term and short-term action-effect links and their impact on effect monitoring. *Journal of Experimental Psychology: Human Perception and Performance, 44*(8), 1186–1198.
- Wohlschläger, A. (2000). Visual motion priming by invisible actions. *Vision Research, 40*(8), 925–930.
- Wolpert, D. M., & Flanagan, J. R. (2001). Motor prediction. *Current Biology, 11*(18), R729–R732.
- Wolpert, D. M., & Ghahramani, Z. (2000). Computational principles of movement neuroscience. *Nature Neuroscience, 3* (Suppl.), 1212–1217.
- Wolpert, D. M., & Kawato, M. (1998). Multiple paired forward and inverse models for motor control. *Neural Networks, 11*, 1317–1329.
- Woods, D. L. (1995). The component structure of the N1 wave of the human auditory evoked potential. *Electroencephalography and Clinical Neurophysiology. Supplement, 44*, 102–109.
- Wylie, G., & Allport, A. (2000). Task switching and the measurement of “switch costs.” *Psychological Research, 63*(3–4), 212–233.
- Yon, D., Edey, R., Ivry, R. B., & Press, C. (2017). Time on your hands: Perceived duration of sensory events is biased towards concurrent actions. *Journal of Experimental Psychology, 146*(2), 182–193.
- Yon, D., Gilbert, S. J., De Lange, F. P., & Press, C. (2018). Action sharpens sensory representations of expected outcomes. *Nature Communications, 9*(1), 4288.
- Yon, D., & Press, C. (2017). Predicted action consequences are perceptually facilitated before cancellation. *Journal of Experimental Psychology: Human Perception and Performance, 43*(6), 1073–1083.
- Zaretskaya, N., & Narinyan, M. (2014). Introspection, attention or awareness? The role of the frontal lobe in binocular rivalry. *Frontiers in Human Neuroscience, 8*, 527.
- Zaretskaya, N., Thielscher, A., Logothetis, N. K., & Bartels, A. (2010). Disrupting parietal function prolongs dominance durations in binocular rivalry. *Current Biology, 20*(23), 2106–2111.
- Zou, J., He, S., & Zhang, P. (2016). Binocular rivalry from invisible patterns. *Proceedings of the National Academy of Sciences, 113*(30), 8408–8413.

ACKNOWLEDGEMENTS

(Dankwoord)

“Expectations were like fine pottery. The harder you held them, the more likely they were to crack” – Brandon Sanderson

Hoewel de bovenstaande uitspraak niet helemaal strookt met de huidige wetenschappelijke consensus over verwachtingen in het brein, weet Sanderson mijn ervaring van de afgelopen jaren wél goed samen te vatten: Mijn PhD project verliep, om het zacht uit te drukken, niet helemaal zoals geanticipeerd. Aan de start van het project had ik grootse plannen en twijfelde ik er geen moment over dat alles zich precies zou ontfouwen zoals we zo mooi in onze beursaanvraag hadden opgeschreven. Niets was echter minder waar.

De realisatie hiervan was niet altijd makkelijk en ik heb me regelmatig afgevraagd of er überhaupt ooit een proefschrift zou komen. Het feit dat er nu toch een boekje ligt is het gevolg van de onvoorwaardelijke steun van een bijzondere groep mensen.

Allereerst mijn promotor en wetenschapspapa: **Henk**. Tien jaar geleden heb je mij onder je vleugels genomen tijdens het honorstraject van mijn bachelor. In de hierop volgende jaren heb je mij met aanstekelijke passie alle kneepjes van het vak geleerd en ook in mij de drang om onderzoek te gaan doen ontvlamd. Ondanks dat de laatste jaren niet helemaal liepen zoals we het vooraf voor ogen hadden, ben ik je ontzettend dankbaar voor het grenzeloze vertrouwen, de deur die altijd openstond en alle mogelijkheden die je voor me hebt gecreëerd. Dat ik “het nest” moet gaan verlaten vind ik nog steeds een beetje onwerkelijk, maar ik put er vertrouwen uit dat je me alle handvatten hebt gegeven om nu op eigen benen te gaan staan. Dankjewel voor alles.

Daarnaast wil ik natuurlijk ook mijn copromotoren bedanken. **Ruud**, door de vele uren die je in mijn project hebt gestopt vergeet ik soms dat je niet vanaf het begin betrokken bent geweest. Ik denk met een glimlach terug aan onze gedeelde liefde voor priegelwerk en oneindige discussies over de meest minuscule stimulus details. In mijn ervaring zijn er maar weinig mensen die dat soort discussies met zo'n oprecht enthousiasme kunnen voeren. Ik heb veel geleerd van jouw scherpe inzichten en ben je dankbaar voor jouw relativerende humor in tijden van tegenslag. **Dennis**, nu ik zelf mijn PhD aan het afronden ben, realiseer ik me pas hoe bijzonder de begeleiding was die ik al van jou ontving toen je zelf nog aan het promoveren was. Je kwam op mij altijd over alsof je al jaren een dokterstitel op zak had. Daarnaast denk ik niet dat

ik hetzelfde geduld had kunnen opbrengen als er de afgelopen maanden constant een onbegripvolle student met 101 vragen in mijn kamer had gestaan. Zonder jouw onuitputbare kennis en hulp bij het bouwen van ambitieuze opstellingen waren geen van mijn projecten mogelijk geweest.

Dit proefschrift zou er niet liggen zonder mijn uitzonderlijke paranimf duo. Waar menig promovendus pas in de laatste stadia van de proefschrift afrondingen wat taken aan de paranimfen toebedeeld, zijn mijn paranimfen funest geweest voor het tot stand komen van mijn proefschrift. **Robert**, het aantal keren dat je me in de afgelopen jaren letterlijk en figuurlijk van de grond van onze huiskamer hebt moeten vegen kan ik al niet eens meer tellen. Door jouw nuchtere wereldblik en (de)motiverende woorden (“alles komt goed”) heb ik de eindstreep dan eindelijk gehaald. Dankjewel voor al je geduld tijdens mijn eindeloze tirades over onderzoek perikelen en tijdens mijn herhaalde pogingen om in een “*semi-sentient*” staat ons appartement af te branden. **Surya**, mensen die jou niet kennen geloven me nooit als ik je omschrijf en mensen die je wel kennen zijn het er allemaal over eens dat je eigenlijk ook niet echt menselijk bent. Niet alleen ben je een van de meest begaafde mensen die ik ken, maar daarnaast ben je ook nog eens gezegend met een bewonderingswaardige portie vriendelijkheid en engelengeduld waar je U tegen zegt. De 352 keer dat je het concept dispariteit opnieuw aan me uit moest leggen waren altijd met dezelfde rust en positiviteit – ongeacht het tijdstip. Ook voor alle “visuele perceptie voor dummies” lessen kan ik je niet genoeg bedanken. Al mag ik het van jou nooit zeggen: Zonder jou was dit allemaal *echt* niet gelukt. Ik kijk er naar uit om in het komende jaar met je samen te werken.

Naast mijn officiële begeleiding heb ik ook veel additionele hulp mogen ontvangen van een team van inspirerende mensen. **Roy**, jij hebt zo goed als al mijn opstellingen gebouwd en ondersteund. Van fotodiodes tot spraakherkenningsprogramma's, alles wat ik kon verzinnen kon jij bouwen. Dit deed je vaak nog dezelfde dag (of nacht) dat ik weer met een nieuw voorstel aan kwam zetten. Je was verder ook nooit te beroerd om al mijn naïeve programmeervragen te beantwoorden (ook als ik weer eens “alles kapot had gemaakt”). Ik ben je ontzettend dankbaar en weet zeker dat je in Amsterdam voor veel mensen net zo onmisbaar zal zijn. **Herbert**, ik wens vaak dat ik je eerder tijdens mijn project had mogen ontmoeten. Ondanks dat ik op aansporen van Surya huiverig begon aan mijn eerste Bayesiaanse analyses in JASP, ben ik inmiddels helemaal overtuigd van de waarde van Bayes factoren. Jouw positieve instelling en wel-

willendheid om mijn eindeloze vragen te beantwoorden hebben daar een grote rol in gespeeld. Het voelt stiekem een beetje alsof ik een extra begeleider heb gehad. Ik vind het ontzettend leuk dat ik de ontwikkeling van BAIN van dichtbij mee heb kunnen maken en ben er trots op dat het een onderdeel is van ons gezamenlijk gepubliceerde artikel. **Ron**, jij blijft voor mij een grote bron van inspiratie. Ik bewonder de precisie en passie waarmee je werkt en ook zeker de moed om een andere weg in te durven slaan. Jouw constructieve kritische blik heeft me vaak geholpen, zowel op academisch vlak als daarbuiten. Ik hoop dat we ook de komende jaren contact kunnen houden. **Chris P.**, dankjewel voor alle hulp in de prille stadia van mijn allereerste binoculaire rivaliteit studies. Ik moet er nog vaak om lachen hoe naïef ik was in mijn inschatting dat ik dat wel even snel onder de knie kon krijgen. Zonder jouw toolbox was ik waarschijnlijk nergens gekomen. **Maarten**, hetzelfde geldt voor jouw onmisbare mini-lezingen over psychofysica en curvefitting. Ik vind het nog steeds heel bijzonder dat je altijd bereid was om mijn vragen te beantwoorden alsof ik “één van jullie” was. **Jim**, ik sta regelmatig versteld van jouw eindeloze positiviteit en behulpzaamheid. Je was altijd bereid om jouw expert-ogen even uit te lenen voor een pilot of om me van goed advies te voorzien. Jouw high-five streak heeft de laatste paar weken van mijn PhD een stuk dragelijker gemaakt. Daar kan geen kinderchocolade tegenop.

The **goallab (/SCAN/SPAM)** group of course also needs a central place in this list of acknowledgments. **Hans**, I still aspire to adopt a fraction of your calm and level-headed attitude that so often helped me to put things into proper perspective. **Loek**, I was thrilled to find a “predictive coding” buddy when you joined the lab. I still have many happy memories of our trip to Edinburgh and cherished our bi-weekly statistics and programming meetings. **Chris, H.**, our pseudoproject was one of the most fun projects I have done, and I wished there would have been time to do more. I’m still truly impressed by the orderly manner in which you work and am thankful for all the tips you have given me along the way. It goes without saying that I count on you to make Jeffreys proud when I’m gone. **Samy**, I always thought we would end up collaborating on something, but it somehow never happened. Nevertheless, I greatly value all of our academic (and non-academic) discussions, conference trips (even the ones on which not all of my organs were functioning) and plant adopting expeditions. **Manuel**, I really enjoyed our statistics discussions and short Gutenberg breaks. I hope we can finally sit down for a relaxed cup of tea to celebrate that we both made it and are (sort of) alive in the near future. **Tom**, it is a shame that the agency reading club is no longer in existence, I have happy

memories of our mutual passionate rants about the newest controversial papers. **Yuru**, I still remember our very first trip to the EEG course in Amsterdam. I think we have both grown a lot since then and I have enjoyed learning about the sense of agency together. **Dan** and **Kevin**, thank you for the interesting lab discussions and for introducing me to the amazing Chinese cuisine (it appears that our Dutch version doesn't even come close to doing it any justice).

Ook de SCIPers kunnen in dit rijtje natuurlijk niet overgeslagen worden. **Neeltje**, ik vind het nog steeds jammer dat onze plannen voor studies bij schizofreniepatiënten nooit werkelijkheid zijn geworden. Zonder jouw warme steun en de vrijheid die je me hebt gegeven in mijn jaar bij het UMCU hadden we nooit de financiering voor dit project ontvangen. Ik ben blij dat je nu als commissielid toch ook bij de afronding betrokken kunt zijn. **Anouk**, komende uit hetzelfde nest heb je me over de jaren heen bij een verscheidenheid aan koffiezetautomaten van veel gewaardeerde wijze woorden voorzien. **Merel** en **Martijn**, ik heb veel warme herinneringen aan onze SCIP-diners en hoop dat we de traditie nog een lange tijd voort kunnen blijven zetten.

In al mijn jaren op de SHOP-afdeling zijn er natuurlijk nog veel meer mensen die me van steun hebben voorzien. **Reine** en **Madelijn**, zoals ik jullie laatst al vertelde zijn jullie mijn grote voorbeeld van mensen die de sleutel tot de ultieme work-life balance hebben ontdekt. Dankjulliewel voor alle zorgzame woorden en inspirerende rust. **Theo**, de afgelopen tijd heb je me veel vaker langs je kantoor zien rennen dan dat er tijd was voor een praatje. Ik hoop dat we de gezellige diners of interessante culturele avonden in de toekomst weer op kunnen pakken. **Tina**, ik moet soms nog steeds grinniken als ik terugdenk aan onze verwoede poging om het proefpersoonuren systeem te veranderen. Ik had de oprechte en (achteraf) naïeve overtuiging dat we dat wel even zouden regelen. Dat het uiteindelijk niet mocht baten heeft in ieder geval niet aan jouw aanstekelijke strijd lust gelegen. **Iлона**, je hebt me als kamergenoot in het laatste half jaar vooral in halfbewuste staat meegemaakt, maar ik heb blije herinneringen aan ons gedeelde ongeloof over de afwezigheid van gepaste gang-etiquette. Verder wil ik alle andere PhD-lotgenoten, zoals **Melissa**, **Lianne**, **Laurens** en **Onur** bedanken voor de fijne en relativerende gesprekken tijdens de lunchpauzes en bij de koffiezetautomaten. Het heeft meer voor me betekent dan ik heb kunnen laten blijken.

Tenslotte wil ik alle mensen bedanken die me buiten de werkvloer hebben opgevangen. **Niek**, ik prijs mezelf nog regelmatig gelukkig om hoezeer wij twee handen op één buik zijn. Het is enorm fijn als iemand direct snapt waar emoties of gedachten vandaan komen zonder dat dit tot in den treure uitgelegd moet worden, onafhankelijk van afstand of tijd. Jouw rustige en logische advies tijdens onze telefoongesprekken maakt altijd een groot verschil. **Pascal**, mijn dag is altijd beter als ik jouw pretoogjes door mijn glazen kantoordeur zie gluren. Dankjewel dat je altijd langs blijft komen voor kopjes koffie en de nieuwste updates over wat voor geniaals je nu weer hebt gebouwd. **Mike**, ik heb ontzettend veel respect voor jouw veerkrachtigheid en voor hoe je in het leven staat. Onze korte wandelingen naar de Gutenberg waren altijd precies lang genoeg om me weer nieuwe moed te geven om door te gaan. Ik hoop dat we elkaar de komende periode meer kunnen gaan zien. **Mettie, Daniëlle, Diane, Kim, en Ninja**, dankjewel voor alle lieve woorden en kneedkamerknuffels. **Peter**, ik zou willen dat je kon zien dat het uiteindelijk allemaal toch gelukt is. Ik mis jouw passie, ongepaste memes en goede raad, maar ben dankbaar dat we de bijzondere tijd die we hadden mochten delen.

Daarnaast wil ik ook mijn familie bedanken. **Mama**, dankjewel voor de rust waarmee je mijn frustraties de afgelopen jaren hebt aangehoord. Ook de schoonmaaktornado's die je regelmatig door mijn appartementen hebt laten woeden werden zeer gewaardeerd. **Papa**, onze zondagmiddaglunches zijn inmiddels een traditie. Ik koester de momenten waarop je jouw nieuwste levenslessen of top 100 met mij deelt. **Albert**, zelfs als we elkaar wat minder zien, weet ik dat ik altijd op je kan rekenen. Dankjewel voor je onvoorwaardelijke steun en vertrouwen. **Mathijs**, van klein aapje tot volwassen broer. Ik heb je met bewondering gadeslagen en ben trots op hoe jij je staande houdt tussen de niet altijd even makkelijke groep jongeren waarmee je werkt. Jouw enthousiasme is aanstekelijk en geeft ook mij nog altijd energie.

Rob, your honest interest in this project is still beyond me, but has always made me feel incredibly supported. You have actively thought about sphere designs and read most (if not all) of the work that is presented here. More importantly though, you have succeeded in the daunting task of dragging me away from my laptop from time to time. Thank you for continuously challenging me, and for spurring me on whenever I feel like giving in. You make me feel like I can do anything and I'm excited for the future.

BIOGRAPHY

Myrthel Dogge was born on the 26th of July, 1989 in the Netherlands. She obtained her bachelor's degree in psychology at Utrecht University in 2010 (cum laude). After a short detour that resulted in master's degree in social psychology (cum laude), Myrthel decided that she wanted to pursue a career in research. Accordingly, she started with a research master's in social and health psychology, which she completed in 2013 (cum laude). In that same year she started working as a research assistant under the supervision of Prof. Van Haren at the Brain Center Rudolf Magnus in the UMCU. Here, she was involved in research on experienced control in patients suffering from schizophrenia, and also worked on a research grant proposal under supervision of prof. Aarts. This grant was funded by the NWO in 2014, allowing Myrthel to start the PhD project that is described in this dissertation under supervision of Prof. Aarts, Dr. Custers and Dr. Hofman. She is currently working in Dr. Peelen's Visual Cognitive Neuroscience lab at the Donders Institute for Brain, Cognition and Behaviour in Nijmegen.