Endogenicity and awareness in voluntary action

Elisabeth Parés Pujolràs

Institute of Cognitive Neuroscience

University College London

Supervisor: Prof. Patrick Haggard

PhD thesis

2019

Declaration

I, Elisabeth Parés Pujolràs, confirm that the work presented in this thesis is my own. Where information has been derived from other sources, I confirm that this has been indicated in the thesis.

Elisabeth Parés Pujolràs

30/09/2019

Abstract

The idea that we can trigger and control our actions at will is central to our experience as agents. Here, we investigated different cognitive mechanisms involved in voluntary action control.

In the first part of the thesis, we investigated the relationship between motor preparation and awareness of intention. To do so, we used spontaneous action paradigms and combined them with novel random and real-time EEG probing techniques. We investigated two main questions. First, do people know that they are about to do something before they do it? Second, to what extent are delayed intention judgements informed by prospective motor preparation rather than retrospective reconstruction? Our findings suggest that people have some feeling of motor intention before acting and can use it to voluntarily control action initiation in real-time. However, their recall-based intention judgements are strongly influenced by overt events happening after the time of probing. Because most daily-life voluntary actions occur in interaction with the environment, in the second part of the thesis we embedded self-paced actions in a decision-making context. We investigated two ways in which endogenous factors can contribute to action selection. First, as a symmetry-breaking mechanism in contexts of external ambiguity. Second, by top-down modulating decision-making processes. We identified the neural correlates of an internal decision-variable that tracks perceptual decisions and also indexes dynamic changes in endogenous goals. Further, we show that the readiness potential can be found not only preceding spontaneous actions, but also in contexts where actions are informed by evidence but preserve a self-paced nature.

In sum, this thesis provides new insights into the cognitive mechanisms underlying conscious experience of intention and provides new tools to investigate voluntary control over action initiation and selection processes.

Impact statement

The idea that human beings are capable of conscious control of behaviour is fundamental to the structure of most contemporary societies, which are largely based on legal systems that rely on individual liability for actions and attribute agents with moral responsibility. Further, the feeling that we can control our actions is essential to our experience as agents. Yet, the cognitive mechanisms underlying experiences of voluntary control and their relation to decision-making and action generation mechanisms are poorly understood.

The first part of the thesis provides novel insight into the mechanisms underlying the *feelings* and *judgements* of motor intention. Our results suggest that agents know what they are about to do in real-time, but what they actually end up doing has a stronger influence on their recall-based intention *judgements*. Our contribution to a deeper understanding of the mechanisms linking conscious thought and action control paves the way for the development of clinical applications for disorders of volition such as Gilles de la Tourette syndrome, utilization behaviour and anarchic hand syndrome, which show characteristic impairments in the ability to consciously inhibit unwilled movement initiation. Further, our findings are of potential relevance to the fields of moral philosophy and legal responsibility, since they shed new light into the cognitive processes underlying experienced and reported intention in voluntary action.

The second part of the thesis combines methods from voluntary action research and perceptual decision-making literature to investigate voluntary control over action selection processes in more naturalistic contexts. We identified a signal that tracks endogenous control processes, and thus constitutes a valuable target signal to investigate voluntary action control and its impairments in a number of disorders. Further, our approach is theoretically novel: the study of voluntary control of action in evidence-informed contexts bridges the gap between the fields of voluntary action and perceptual decision-making research, and broadens the scope and tools available to voluntary action research.

Finally, all the experimental paradigms described in this thesis were developed for the purposes of this research and are designed to overcome some limitations traditionally encountered by the field. Thus, they provide new, original tools for future research in the field of motor intention and voluntary action in healthy and clinical populations.

In sum, the results and methods described in this thesis significantly contribute to our understanding of voluntary action and provide valuable tools for clinical applications and future research in the field.

Acknowledgements

First and foremost, I would like to thank my supervisor Patrick Haggard for his guidance, support and encouragement throughout my PhD. I am extremely grateful for the freedom that he has given me to explore ideas and methods, for the many opportunities to attend stimulating meetings and conferences, and for allowing me to develop as an increasingly independent researcher. I am also very thankful for his support as a mentor and many insightful career chats, which have been a source of inspiration and motivation.

I would also like to thank all past and present members of the Action and Body lab, who have been a constant source of help and have also been, above all, kind and fun colleagues that have made my PhD experience a most enjoyable time. Special thanks go to Anne Löffler for introducing me to the wonderful world of coding, and Eoin Travers for patiently answering countless questions and providing infinite amounts of coffee. Double thanks to both of them for their helpful comments on previous versions of this thesis. Further, I would like to thank Yoana Petrova and Jeremy Hatchuel, who helped me collect the data in chapters 4 and 5 during their BSc and MSc projects and were most diligent students.

Further, I would like to thank John-Dylan Haynes for giving me the opportunity to visit his lab at the Humboldt Universität in Berlin for three months. The real-time EEG study reported in chapter 3 would not have been possible without his support. I would also like to specially thank Matthias Schultze-Kraft, with whom I directly collaborated throughout the project from the initial experimental design stages to the final analysis, and Karla Matic, who helped us design and run the experiment in Berlin. They were great colleagues who made the experience fruitful and most enjoyable. I would also like to thank Yong-Wook Kim and Chang-Hwan Im for their help with the canonical-correlation analysis reported in chapter 1.

Finally, I would also like to specially thank my family and friends for always being there and being a constant reminder of the most important things in life. In particular, I would like to thank my mum, who has always been and continues to be an example of strength and kindness, and a source of unconditional support.

Table of contents

List of figures and tables

Chapter 1

Chapter 2

Chapter 3

Chapter 4

Chapter 5

Appendix A

Appendix B

Appendix C

Chapter 1

Rethinking voluntary action

Synopsis

The idea that human agents are capable of voluntary control of their actions is fundamental to both our experience as agents and to the organisation of our societies. The neuroscientific study of the cognitive mechanisms related to voluntary action control has traditionally been based on reductionistic operationalisations of volition that fail to account for most instances of naturalistic voluntary behaviour. The aim of this chapter is to provide an updated, cognitively-grounded theoretical model for the neuroscientific study of voluntary action. We propose a hierarchical framework including *global* and *local* mechanisms of control. *Global* control mechanisms involve intentional top-down modulations of low-level decision-making processes driving action selection, while *local* control refers to the ability to allow or inhibit the execution of a specific movement. We believe that the proposed framework overcomes some of the traditional limitations of the field and offers new opportunities for the development of a neuroscience of voluntary action control.

1.1.What are voluntary actions?

As agents, we feel that we can control our actions. When we make a decision, we normally feel we are free to choose what to do. We experience ourselves as deliberating, and we feel more or less certain about our choices. Once we have made up our minds, we feel that we can initiate movements at will, and we typically feel that we are in control of our actions and also of their outcomes (Haggard and Tsakiris 2009). All of these experiences are central to our self-representation as agents. Further, the idea that human beings are capable of conscious control of behaviour is fundamental to the structure of most contemporary societies, largely based on legal systems that rely on individual liability for actions and attribute agents with moral responsibility. But what are voluntary actions?

1.1.1. Traditional definitions

In the neuroscientific literature, specific features typically ascribed to voluntary action have been used as standard definitions (see Haggard (2018) for a recent exhaustive review). Broadly, these definitions can be grouped according to three of the classic Aristotelian causes (Barnes 1984). That is, they respond to the question about "why are some actions voluntary" by referring to either their efficient causes (i.e. the events that triggered the action), final causes (i.e. the goals of the action), or formal causes (i.e. the representation driving the action).

A first family of definitions of voluntary action focuses on the nature of the event that triggered it – its *efficient* cause. In particular, voluntary actions are typically defined as endogenously generated or self-initiated (Libet et al. 1983; Jahanshahi and Frith 1998; Libet 1999). That is, the source of the movement or the choice lies within the agent, and not in a stimulus in the immediate environment. These types of actions are typically defined in contrast to exogenously-driven actions (e.g. braking in response to a red traffic light), and also from reflex actions (e.g. withdrawing the hand from a burning surface), which are a direct result of an environmental stimulus. Paradigmatic examples of voluntary action under this definition are spontaneous movements that do not obey any external imperative (e.g. tapping the fingers on a table just because one feels like it) or free choices between equivalent options (e.g. picking the apple on the left or the one on the right).

A second type of characterizations emphasize the goal-directedness that is characteristic of some voluntary actions (Haggard 2005, 2018), the *final* cause. That is, that voluntary actions have a purpose. Voluntary action has been defined as "the realization of the anticipated content of the goal-presentation"(Ach 2006). This teleological characterization puts the focus on the ends towards which actions are directed, in the attainment of a goal. It has been suggested that actions are selected by means of an inverse-modelling process (Wolpert and Kawato 1998), by which an action plan is formed backwards: starting from goal of the action, all motor programs that could produce it are evaluated. Under this definition, an action such as reaching for a glass of water is voluntary.

A final group of *formal* characterizations define voluntary actions as those actions that follow from some idea, intention or representation. People can give reasons for their voluntary actions (Davidson 1963). Actions come about "in virtue of" some previous mental state, an intention, that people typically refer to and identify as the cause of their actions (Uithol et al. 2014). The concept that actions result from some kind of intention represented in the brain can be traced back to some aspects of Carpenter (1852) and James's (1890) ideomotor theory, which postulated that voluntary actions come about in virtue of an *idea.* More recently, it has been suggested that there may be different, hierarchically organized types of intentions (Pacherie 2008). These would include *distal* intentions (e.g. eating an apple), *proximal* intentions (e.g. actions that could serve the purposes of achieving a distal intention – e.g. going to buy an apple) and *motor* intentions (i.e. sending a command to specific muscles). However, the status of intentions as mental representations is controversial. Some have argued that the more abstract *distal* intentions in particular may not correspond to a specific representation, and that the concept may not be useful for cognitive research (Uithol et al. 2014). This family of *formal* definitions is closely related to the *final* ones described above: intentions typically *are* representations of a goal, at various levels of abstraction. However, goal-directed action can happen in the absence of intention. This is the case, for example, in alien hand syndrome patients (e.g. Debray and Demeestere 2018), who perform goal-directed actions but describe them as involuntary because they did not *intend* to execute them. That is, the goals of their actions do not correspond to a prior plan or intention representation that patients relate to.

Investigating voluntary action under any of these definitions is problematic, because they are often contradictory. For example, many actions that follow from intentions are externally triggered. One can decide to compete in a race, and have an intention to do so. However, athletes typically wait for a gunshot to start running. The action to start running is voluntary under *formal* definitions (i.e. it responds to the intention to run the race), but not under *efficient* ones (i.e. the action was effectively triggered by an external event). Also, some spontaneous actions are not goal-directed. One can decide to spontaneously lift one's arm, yet such a spontaneous and endogenously initiated action does not have any goal, it is capricious, and thus not voluntary under the *final* definitions of action. Thus, by focussing on some of its most salient features or by postulating theory-laden constraints, traditional definitions fail to capture the complexity of the phenomenon and paint a caricaturesque portrait of voluntary action and its cognitive underpinnings.

1.1.2. Towards a naturalistic definition

Here, we want to take a step back and look at the phenomenon that gives rise to all of these seemingly incompatible definitions. Most of us have an immediate, intuitive understanding of which of our actions are voluntary and which are not (Haggard 2018), and these are the kinds of *naturalistic* voluntary actions that we aim to study.

All naturalistic voluntary actions have two basic elements in common. First, that they are actions – they involve doing things, moving in certain ways. Second, that we experience those actions *as if* we are able to *control* them. Under this perspective, the above-reviewed traditional definitions highlight different ways in which we can feel in control, rather than describing essential characteristics of *all* voluntary actions. For example, we can feel in control if our actions produce certain outcomes (*final* aspects), if we know why we are doing something or feel our actions follow from our intentions (*formal* aspects) or by deciding autonomously, independently from exogenous factors (*efficient* aspects).

We do not always feel we control actions in all these relevant ways, but for all actions that we would intuitively class as voluntary we do feel *some* kind of control, and different types of control seem to relate to different aspects of our perception of moral agents. In particular, being able to act independently of the environment is relevant for our perception of actions as free, whereas the ability to deliberate and produce actions that are in agreement with our intentions plays a more relevant role in our evaluation of moral responsibility (Vierkant et al. 2019). Whether we actually *are* in control in any of these ways, beyond *feeling* that we are, is a question for empirical investigation.

1.1.3. Outline

In section 1.2. we review the contemporary neuroscientific tradition on voluntary action, which has focussed on investigating the endogenous mechanisms responsible for spontaneous action initiation. We discuss the limitations of its theoretical assumptions and experimental paradigms, and we propose a metacognitive framework that overcomes some of the conceptual limitations of the traditional approaches and provides useful tools for further experimental research.

Then, in section 1.3 we describe the decision-making processes involved in action selection, and we argue that explaining voluntary action control at that level requires going beyond the traditional study of endogenous contributions to “free” decisions. Rather than conceiving endogenous states in opposition to exogenous ones, we will argue that it is necessary to study their interaction to understand how specific actions can be selected to match higher-order goals. In particular, we will suggest that voluntary control over action selection may involve distributed top-down effects from high-order endogenous states to low-level decision-making processes.

Finally, section 1.4. integrates the reviewed literature in a unitary model of voluntary action. The goal of the model is to explain the cognitive mechanisms involved in action generation, and how our phenomenal feelings of control relate (if at all) to actual control mechanisms. Hopefully, this cognitively-grounded model will provide a framework for the study of voluntary action that avoids the caveats of sticking to constrained traditional definitions and incorporates traditionally neglected but relevant cognitive processes into the picture.

1.2. Forward we go! Voluntary control over action initiation

The contemporary tradition of cognitive neuroscience of voluntary action has been particularly interested in the cognitive mechanisms and phenomenology underlying action initiation. We feel our conscious decisions *cause* subsequent movements and actions, and this feeling is central to our experience as agents. Thus, the ability to initiate actions “at will” has often been conceived as a hallmark of voluntary action (Libet 1999; Haggard 2005). Understanding how internal (endogenous) events of this kind (i.e. a “conscious decision”) can trigger bodily movements in the absence of any external (exogenous) imperative has been the focus of the contemporary neuroscientific tradition of voluntary action. This tradition thus focussed on the study of the *efficient* aspects of voluntary actions. That is, on the idea that the ability to spontaneously initiate action is characteristically voluntary and essentially different from exogenously-triggered action.

1.2.1. The neural basis of action initiation

A classic theory of the neural mechanisms involved in action initiation supported the conceptual exogenous-endogenous dichotomy (Passingham 1987).

Early lesion studies in monkeys showed that lesions in the Supplementary Motor Cortex, which includes the pre- and Supplementary Motor Areas (pre-SMA and SMA; Picard and Strick 1996), resulted in animals being unable to initiate self-paced actions, while they retained the capacity to execute cued movements (Passingham 1987). In turn, lesions in the lateral premotor cortex (PMC) caused impairments in the ability to initiate cued movements, while self-initiated actions remained unaltered (Passingham 1987). These findings suggested that externally-triggered actions and internally-generated ones may lead to action execution through two different pathways. The model suggested that self-paced (i.e. internally triggered) actions originate in the medial SMA, whereas externally-triggered movements reach the M1 through the lateral premotor cortex (PMC). Single-cell recordings in monkeys (Okano and Tanji 1987; Romo and Schultz 1987; Kurata and Wise 1988) further showed preferential (although not exclusive) activation in the SMA preceding self-paced movements and in the PMC preceding cued movements. Additionally, human neuroimaging (Deiber et al. 1991, 1999; Jahanshahi et al. 1995), human single-cell recordings (Fried et al. 2011) and clinical studies in lesion patients (Laplane et al. 1977) have reported findings compatible with the two-pathway hypothesis.

The idea that two different pathways underlie action generation triggered by exogenous and endogenous factors, respectively, is further supported by analysis of the anatomy and connectivity between brain areas in primates. It has been proposed that motor areas can be grouped in two broad categories: areas that primarily receive input from parietal cortex, and areas that receive it from prefrontal cortex (Rizzolatti and Luppino 2001). The parietal cortex receives rich sensory input through the dorsal pathway, while the prefrontal cortex is involved in higher-order processes such as long-term planning (Tanji and Hoshi 2001) and free choice (e.g. Frith et al. 1991; Hyder et al. 1997). This evidence suggests that while the parietal-dependent areas (including F2, equivalent to the human PMC) are involved in quick sensorimotor transformations, prefrontal-dependent areas (including F6, equivalent to the human pre-SMA) are related to execution of slower, long-term, higher-order plans (Rizzolatti and Luppino 2001). Further, recent computational analysis have suggested that primate F6 and F7 (human SMA) tend to cluster separately from the motor areas receiving primarily parietal input and are suggested to be an interface between prefrontal areas and primary motor cortex (Caminiti et al. 2017).

Although the idea that there is a clear two-pathway modular organization in the motor system is controversial (Nachev et al. 2008; Nachev and Husain 2010; Hughes et al. 2011), the exogenous-endogenous dichotomy lies at the core of the contemporary neuroscience of voluntary action and inspired the early paradigms – which continue to be used nowadays.

1.2.2. The Libet Experiment: a case against conscious control over action initiation

The first experimental attempts at linking the conscious experience of “willing”, or “moving now” and the neural mechanisms underlying self-initiated action were made by Benjamin Libet (Libet et al. 1983), whose seminal experiments set the ground for contemporary research in voluntary action.

In his studies, Libet recorded the EEG activity while participants were performing a simple task. They were instructed to make a quick movement (e.g. flex the wrist) whenever they felt like it while staring at a clock with a rotating hand. After they had executed the movement, they were asked to report the position of the rotating hand of the clock when they “felt the urge to move”. These judgements were effectively a report of “motor intention” awareness, and are known as “W-judgements” (i.e. time of “Will” judgements). Libet showed that participants reported an intention to move only about 200 ms before actually moving, but that a neural signal – the Readiness Potential (RP) - was consistently present up to 1 s before electromyogram (EMG) onset.

The *Bereitschaftspotential* (BP, German for “Readiness Potential”, RP) was first described as a slowly increasing negativity over the motor cortex preceding self-paced movements by Kornhuber and Deecke (1965). The RP is typically divided in two components: an early slowly increasing component with a symmetric and maximal distribution around the midline and a late, steeper one which is asymmetric and maximal on the contralateral side of the movement (Shibasaki and Hallett 2006). Both its presence preceding self-paced actions and the time of intention reports that Libet measured have been consistently replicated (Keller and Heckhausen 1990; Haggard and Eimer 1999; Trevena and Miller 2002; although see Schlegel et al. 2013). Importantly, the RP is typically found before spontaneous actions, but not externally-driven movements (Papa et al. 1991; Jahanshahi et al. 1995). Because of this specificity, the RP has traditionally been interpreted as a neural marker of endogenous action initiation, as opposed to stimulus-driven movement generation. Further, the sources of the early RP lie in the endogenous action generation system proposed by Passingham (1987). The early slow component has been shown to originate from the pre-SMA and SMA with EEG (Praamstra et al. 1996; Toma et al. 2002) and MEG dipole source localization techniques (Erdler et al. 2000), as well as high-resolution EEG analysis (Cui and Deecke 1999). In turn, the late, lateralised fast component originates from the motor cortex contralateral to the movement (Nagamine et al. 1996).

The fact that the time at which participants reported being aware of an *urge to move* in Libet’s experiment was much later than the onset time of the RP was interpreted as follows: the conscious experience of being about to move (or “willing” to move) does not *cause* the endogenous movement preparation processes. Rather, the unconscious mechanisms driving the RP signal cause the experience of being about to move. The conscious decision to move is thus conceived as a percept resulting from an unconsciously initiated process of action initiation (Haggard 2005; Mirabella 2007), rather than the other way around. However, Libet suggested that while conscious decisions did not trigger movement initiation, his results were compatible with a different type of voluntary control. He argued that the short time-window separating the conscious experience from the action onset still allowed the possibility to consciously veto the unconsciously initiated spontaneous movements (Libet et al. 1983; Libet 1999).

Yet, Libet's classic interpretation of these results remains controversial (Schlegel et al. 2013; Verbaarschot et al. 2015; Brass et al. 2019). Several objections have been raised, targeting both conceptual and methodological issues regarding the RP, the W-judgements and the theoretical framework.

Limitations of the classic interpretation of the RP

The idea that the RP is a *specific* marker of voluntary motor preparation has been challenged (e.g. Schurger et al. 2012). Classical RP studies lock EEG data to movement onsets and hence study neural activity immediately prior to action - but ignore it at other times. This biased sampling means that the RP pattern might not be specifically related to voluntary action execution, nor to awareness of intention (Mele 2011). In particular, RP-like fluctuations might also occur in the absence of actions and conscious intentions, but those RPs would be invisible to the action-locked methods used in classical RP paradigms. Recent computational models have shown that simply averaging an ongoing stochastic signal time-locked to a threshold crossing event can reproduce the form of the RP, both in humans (Schurger et al. 2012) and mice (Murakami et al. 2014, 2017). According to this interpretation, accumulation of spontaneous neural activity determines *when* to perform a voluntary action. Since the underlying fluctuations in such models are continuous, RP-like signals should also occur even in the absence of action or an intention to act. Thus, rather than reflecting the consistent neural correlates of an intentional process of motor preparation, the RP may reflect stochastic processes.

A related concern applies to Libet’s interpretation of the causal relationship between the RP and action. In particular, Libet’s claim that movements in his task were unconsciously initiated assumes that the RP is a deterministic marker of movement preparation. In other words, that the neural activity on any given *individual* trial contributing to the *averaged* RP signal would unavoidably lead to a movement, unless consciously vetoed. However, this inference is not justified. Libet’s results only show that whenever there is a self-paced action, an RP is found after averaging over trials. They do not show that every time that there is an RP a movement follows. If the above criticism is right, and RP-like activity can also be present in the absence of overt movement, there is no reason to assume that the neural activity that produces the RP *alone* suffices to trigger action.

Recent studies have shown that it is possible to increase the probability of interrupting a self-paced action using an RP-based real-time algorithm (Schultze-Kraft et al. 2016) on a single-trial level. The study used a veto instruction: whenever an RP was detected, a red light was shown on-screen and participants were required to refrain from moving. Participants were often unable to do so, indicating that the RP is a valid marker of motor preparation (i.e. the presence of an RP was often followed by a subsequent movement that could not be inhibited). However, the results leave the question open as to *how* predictive it is. If *all* RPs lead to movement execution (unless actively vetoed), then surely the neural activity producing the RP is sufficient for triggering movements. However, in the Schultze-Kraft et al. (2016) study it remains unclear whether that is the case. The trials where participants did not move after being probed may reflect instances where they would not have moved *anyway,* regardless of the probe (i.e. the RP would not have been followed by a movement), rather than instances where participants successfully inhibited an incoming movement (i.e. the RP indeed indicated motor preparation, but participants vetoed action execution).

Thus, the causal relationship between the RP, motor preparation and conscious intention to move remain unclear. While the presence of an RP seems to make spontaneous movements more likely, it remains logically possible that some additional condition is required for this enhanced level of readiness to actually trigger an action. Such a condition could be a certain background state of the system (e.g. the agent's disposition to actually move), or a conscious 'Go' decision that effectively triggers – or perhaps *allows* - the movement to happen.

Finally, it remains unclear whether the RP is related to motor preparation only, or to intention awareness as well. Libet's original studies and most recent literature have studied action and intention together in movement-based paradigms (see Matsuhashi and Hallett (2008) for a rare exception). The reasons for that are clear: it is difficult to study brain activity when no objective behavioural event (e.g. movement) is present. However, the fact that motor preparation and intention awareness are always perfectly correlated in these experiments means that the RP may be related to voluntary motor preparation but not intention awareness, or vice versa. To explain how the cognitive mechanisms involved in voluntary action triggering relate to its phenomenal aspects, it is crucial to investigate the extent to which the neural markers of voluntary action are specific to its cognitive (motor preparation) or phenomenal (experience of an "urge") aspects or rather are common to both. Yet, the classic Libet paradigm critically conflates both.

Limitations to the classic interpretation of W-judgements

Libet’s interpretation of the W-judgements relies on two assumptions. First, that they refer to a genuine phenomenological experience. Second, that they are accurate. *Figure 1.1.* illustrates different possible interpretations of the results.

The classic interpretation of the data assumes a *prospectivist* notion of intention. That is, it assumes that a conscious experience of an *urge* genuinely happened before action execution. However, it has been suggested that there may not be any genuine experience of “being about to move” at all before an action is executed, but rather the reported times are a postdictive reconstruction (Banks and Isham 2009, Wegner 2002) or a distorted experience of time (Dennett and Kinsbourne 1992). These *retrospectivist* interpretations suggest that participants were asked to report a time after they executed an action, and so they could have provided a guess. However, as far as the experimenters know, this report need not correspond to any actual phenomenal experience. In fact, it has been shown that people can be induced to confabulate intentions about actions they did not cause (Aarts et al. 2005), or that they presumably failed to inhibit (Kühn and Brass 2009).

Relatedly, the fact that participants always executed an action *before* reporting an intention leaves open the question as to whether “pure prospection” exists. That is: to what extent is the feeling of an “urge to move” intrinsically linked to action execution? In other words, are there instances of motor intention without action execution? Retrospectivist theories would predict that there is no such thing: intentions are only inferred after actions. Thus, no intention should be expected in the absence of an action. Yet, Libet-type paradigms cannot answer the question.

Other criticisms to the W-judgements focus on the methods employed to obtain the intention estimate, granting the assumption that "urge" experiences genuinely exist before action execution. Several authors have speculated that the visual processing of the clock might not be independent of either the actions made or of the experiences reported (Miller et al. 2011), and the fact that awareness reports were only provided *offline* (i.e. always post-hoc, *after* action execution) allows for processes happening during and after action execution to affect the report. In fact, there is evidence that perturbations of neural activity *after* action execution modulate W-reports (Lau et al. 2007). These types of objections have resulted in efforts to develop *online* report methods which aim to obtain immediate intention reports (e.g. Matsuhashi and Hallett 2008) that do not encounter these difficulties.

In general, these observations highlight an important distinction: the one between *feelings* as opposed to *judgements*, which has been discussed in the context of other agentic experiences such as the sense of agency (Synofzik et al. 2008). Feelings are the conscious experiences *per se*, as perceived by an individual. Instead, judgements are second-order reports of those experiences by means of some kind of explicit description. While Libet assumed that the W-judgement accurately reflected a genuine feeling (the "urge") preceding the action, the fact that it was delayed with respect to the presumed genuine experience introduces the possibility for other processes (e.g. visual processing of the clock, motor-related information, recall-related processes) to influence the report.

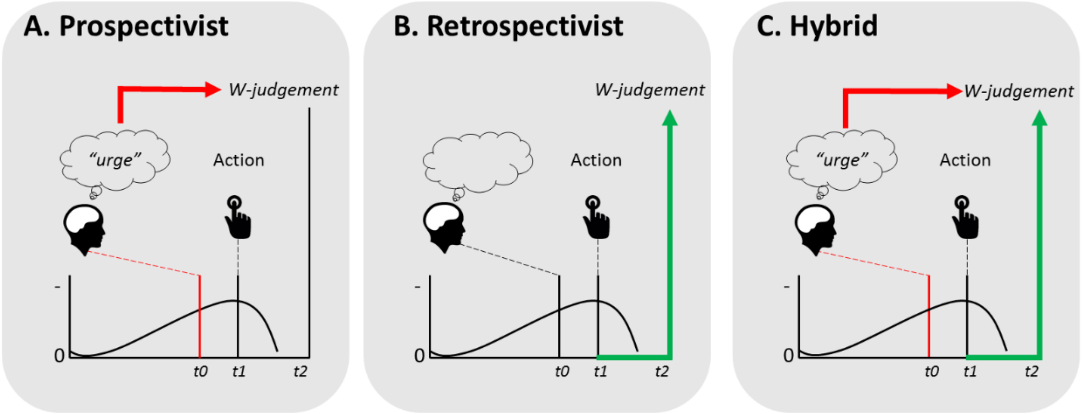

Figure 1.1. Schematic of the Libet experiment and its possible interpretations. Participants executed an action (*t1*) and subsequently provided a W-judgement (*t2*), reporting when they felt an "urge" to move (*t0*). **(A)** The traditional interpretation is a prospectivist account, and assumes a genuine phenomenal experience of motor intention (e.g. an "urge") is accessible to agents before an action is executed, and accurately reported. Conversely, retrospectivist **(B)** accounts suggest that there is no phenomenal experience of intention before action execution, and rather W-judgements report retrospectively inferred times. Finally, hybrid **(C)** accounts of motor intention suggest that the percept of motor intention results from a combination of both prospective information about motor preparation and retrospective integration of action-related and post-action information. Prospective (*red*) and retrospective (*green*) contributions to W-judgements are colour-coded throughout.

Impact

The analysis of the Libet experiments and its limitations pose a number of challenges to voluntary action research.

First, the above review highlights the striking lack of diversity in paradigms for the study of proximal control of action and motor intention. Libet's rotating clock method was developed in the 1980's and has been widely used since then. Yet, very few alternative methods that overcome its inherent difficulties have been developed, and hence the interpretation of the results remains controversial.

Second, despite its limitations and ambiguities, the Libet tradition has had a great impact on the field of voluntary action. The idea that spontaneous movements are not consciously initiated is widely accepted, and some have claimed that the idea that the conscious experience of willing triggers action is an illusion (Wegner 2002). These theories suggest that while we feel conscious experiences of "willing" play a directive role in motor control (i.e. we feel they play a causal role in determining following actions), they are actually, if anything, a description of unconscious motor preparation processes – a mere epiphenomenon. It has been pointed out that a widespread acceptance of these conclusions and the embracement of "willusionism" (i.e. the idea that conscious "will" is an illusion) may be detrimental at a societal level (Vohs and Schooler 2008; Nahmias 2011), and it may be detrimental at a scientific level too. A corollary of these views is that the study of awareness of motor intention is pointless if one aims to investigate voluntary action control. But if the interpretation is wrong, and volitional experiences *do* play a role of some kind in action control, then such a research program is worthy. In the following section, we argue that there is a case to be made for the idea that agentic experiences of "willing" may play a role at the level of action triggering.

1.2.3. Beyond Libet: a case for conscious control over action initiation

Evidence for the role of volitional experiences in behavioural control comes from disorders of volition. The most common type of disorders of volition involve experiencing self-initiated movements as involuntary.

In Gilles de la Tourette Syndrome (GTS) patients experience tics that sometimes look like normal voluntary actions, but can range from small facial muscle contractions to full-body movements, including verbalization (Leckman 2002; Ganos et al. 2014).

Their movements do not typically respond to any immediate environmental cue, nor to any spontaneous decision to move. It has been suggested that the action generation process might be driven by noise in the motor system (Misirlisoy et al. 2015) and insufficient inhibitory control (Georgiou et al. 1995, although see Ganos et al. 2014), and there is evidence that the neural sources of the movements are located in the medial system typically associated with endogenous, self-paced actions (Bohlhalter et al. 2006; Wang et al. 2011). In turn, utilization behaviour (Lhermitte 1983; Iaccarino et al. 2014), alien (Geschwind and Mega 1996) and anarchic hand syndrome (Della Sala et al. 1991) patients produce movements that are perfectly adequate to the environment (e.g. picking up glasses, grasping the arm of an experimenter). In this case, it is not noise, but immediate exogenous stimuli that drive the action generation process. The actions available in the environment suffice to trigger action.

Interestingly, the phenomenal experience and coping strategies accompanying the involuntary movements varies greatly between some of the disorders.

In GTS, patients sometimes feel a “premonitory urge” before the tics (Wang et al. 2011; Ganos et al. 2015). That is, they feel they are “about to move”. Although the urge is often described as involuntary, many patients perceive the execution of the movement as voluntary, much as feeling the need to scratch an itch (involuntary) and then scratching it (voluntary) (Kranick and Hallett 2013). There is some evidence showing that patients who experience the “urges to move” earlier are better able to voluntarily suppress the tics (Ganos et al. 2015), suggesting that while the motor preparation processes underlying the tic may be unconscious and involuntarily, the cognitive mechanisms underlying the phenomenal experience of an “urge” may be relevant for cognitive inhibition of the incoming movement. Patients cannot control the processes generating the movement, and the feeling of an “urge” is not perceived as the “cause” of the movement, but they are able to monitor them to some extent and cognitively supress the tics at later stages (to a certain extent).

In turn, patients with alien or anarchic hand syndrome do not have any phenomenal experiences related to their anarchic limbs. Rather, they report the hand as “having a will of its own” (Banks et al. 1989), and they typically try to prevent the hand from accomplishing its planned movement by grabbing it with the other, non-anarchic/alien hand (e.g. Debray and Demeestere (2018)) rather than by suppressing the movements cognitively like GTS patients often do.

In sum, these two types of patients have different subjective experiences over their symptoms, and these correlate with different behavioural control strategies. Thus, volitional aspects of action seem to be closely related to action control mechanisms. When “urges” are present, patients are able to cognitively control those involuntary movements. Instead, when such urges are not present patients try to control the involuntary movements reactively by external means. This association suggests that volitional experiences may be related to the possibility of cognitive control over action execution. But how is that implemented neurocognitively?

A metacognitive framework for “volition”

Libet conceived “the conscious will” as a somehow uncaused phenomenon, not deterministically following from previous neural activity (Libet 1999). Thus, the claim that a conscious phenomenon could cause a physical event was not problematic under his conception of volition. However, such dualistic conceptions of consciousness are largely rejected in the contemporary neuroscientific research (Haggard 2005; Shibasaki and Hallett 2006; Mirabella 2007). Most neuroscientists work on the physicalist assumption that conscious experiences are not something essentially different from physical states in the brain. Thus, the philosophical question of interaction arises. How are those phenomenal states that we take as playing a directive role related to the cognitive mechanisms underlying action? If conscious states somehow result from brain activity, the question can better be posed as: how do the neurocognitive processes underlying “urge” experiences relate to the neurocognitive processes responsible for motor control? Here, it is suggested that the relationship is a metacognitive one.

Metacognition is broadly defined as “cognition about cognition” (Fleming et al. 2012), and it refers to the ability to monitor and evaluate our own cognitive processes. It operates in multiple domains including memory (Nelson and Narens 1990), multiple sensory modalities (e.g. Beck et al. 2019) and also action (Metcalfe and Greene 2007; McCurdy et al. 2013). The core notion common to all these domains is that first-order processes take place (i.e. memory, perceptual judgements, etc), and that a second-order process *about* that first-order process accompanies them (Fleming et al. 2012). It is typically assumed that the role of metacognitive abilities may be to help control behaviour (Fernandez-Duque et al. 2000; Shea et al. 2014), and it has been suggested that “metacognitive loops” may allow for online monitoring and executive control by connecting areas related to first-order processing with higher-order structures (Fleming and Dolan 2012). Further, several theories of consciousness postulate a strong link between higher-order metacognitive processes and consciousness (Lau and Rosenthal 2011; Fleming 2019). That is, they suggest different ways in which second-order processes may be fundamental for conscious experiences.

In this framework, one can think of phenomenal aspects of volition (i.e. the feeling of an “urge”, “being about to move” or “willing to move”) as the product of a second-order read-out of first-order motor preparation processes. Because it disentangles the first-order, generative processes (i.e. motor preparation, action initiation) from the second-order, monitoring ones (i.e. the ones resulting in “urges” to move), this metacognitive conception of volition may help explain how phenomenal and action initiation aspects of voluntary action can dissociate in pathologic circumstances. The fact that motor execution and phenomenal experience are doubly dissociable (as we shall review below) suggests that the neural basis of the cognitive and phenomenal aspects of voluntary action are not completely overlapping. However, the fact that they are typically connected in healthy people and that alterations in their coupling seem to affect motor control suggests that they are coextensive to *some* degree.

A metacognitive circuit for cognitive control over action initiation

It has previously been suggested that a pathway involving bidirectional projections between parietal and motor areas may be involved in motor awareness, whereas another parallel pathway involving pre-SMA and SMA projections to M1 would be involved in movement preparation and execution (Douglas et al. 2015).

Consistent with this proposal, individuals with lesions in parietal areas including the Angular Gyrus (AG) report delayed intention awareness (Sirigu et al. 2004) and tDCS anodal stimulation of both the AG and M1 results in early W-judgements, while the same stimulation over SMA does not affect intention reports (Douglas et al. 2015). Also, intracranial stimulation of parietal areas results in the feeling of an urge to move (Desmurget et al. 2009), without the production of an accompanying movement. Furthermore, there is evidence from single-cell recordings in monkeys that parietal areas have specialized neural populations that represent ongoing and intended movements (Seal et al. 1982; Crammond and Kalaska 1989; Snyder et al. 1997; Buneo et al. 2002). However, TMS over pre-SMA after action affects W-judgements (Lau et al. 2007) and intracranial stimulation of the pre-SMA has also been reported to result in reported “urges to move” and feelings of “moving” in the absence of overt movement (Fried et al., 1991), and single-cell activity in the human SMA predicts both the time of movement and conscious intention (Fried et al. 2011). This suggests that the pre-/SMA complex is also relevant for intention awareness.

Evidence from disorders of volition provides support for the involvement of both the SMA and parietal areas in circuit involved in the monitoring and control of voluntary action. Alien hand syndrome has been reported after lesions in the SMA (Brainin et al. 2008), parietal (Martí-Fàbregas et al. 2000; Assal et al. 2007) and medial fontal areas (Della Sala et al. 1991). Further, parietal areas and M1 have been reported to be more active during tics than during voluntary movements in GTS patients, and the intensity of that activity has been shown to correlate with tic severity (Wang et al. 2011). Further, the mechanisms of cognitive control seem to involve frontal areas. Utilization behaviour patients, who fail to inhibit actions triggered by environmental stimuli, characteristically have frontal damage (Lhermitte 1983; Eslinger 2002; Ghosh and Dutt 2010). The role of frontal areas in voluntary motor inhibition is well-known in healthy populations (Munakata et al. 2011), and it has been suggested that the symptoms in utilization behaviour may be due to the lack of normally occurring frontal inhibition of motor output mediated by the SMA (Eslinger 2002).

Thus, it is suggested here that the pre-/SMA complex, parietal and frontal areas may be part of a circuit involved both in metacognitive monitoring of action initiation processes and online motor control. In particular, we suggest that motor preparation processes in the SMA are monitored in superior parietal areas, which in turn feed back into frontal areas that enable cognitive control over the motor output. *Figure 1. 2* illustrates a schematic of the proposed circuitry of metacognitive monitoring and control.

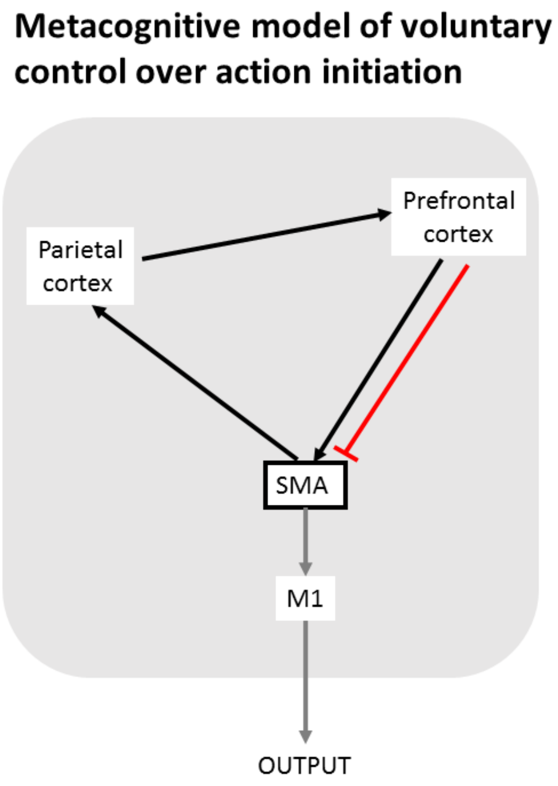

Figure 1. 2. A metacognitive model of voluntary control over action initiation. In healthy populations, motor output is normally suppressed through inhibitory projections from frontal areas to the SMA (red line). Motor initiation states from the SMA are monitored by posterior parietal areas, which send the information to the frontal cortex. In turn, frontal areas feed back into the SMA to either allow movement initiation or inhibit it (black arrows). Disruptions in the monitoring (involving SMA and the parietal cortex) or feedback mechanisms (involving parietal and frontal areas) can lead to missing or abnormal experiences of action initiation and lead to impaired cognitive control mechanisms.

This cognitively-grounded metacognitive model provides a framework for the study of the phenomenology and control mechanisms associated with action initiation processes that moves away from the traditional dualist vs epiphenomenal interpretations. While the Libet tradition was strongly grounded in a dualistic concept of “the will”, its Wegnerian and “willusionist” counterparts relegated the conscious experiences of “willing” to an epiphenomenon (i.e. a feeling with no causal impact), a kind of by-product of cognitive processes that makes no difference in the state of the world. A metacognitive approach takes the best of both worlds: it explains how some kind of causal link between conscious experiences of volition and action control can exist, while avoiding metaphysically implausible assumptions.

1.2.4. Challenges: from segregation to integration

We have so far argued that control over action initiation is an important aspect of voluntary action, and we have provided a metacognitive framework for the study of the volitional aspects of action initiation. Yet, the discussion has relied on two assumptions. First, that endogenous actions are intrinsically different form exogenous ones. Second, that action initiation can be studied as a stand-alone process, separable from action selection. While the reviewed literature focussed on action initiation processes and our ability to control those, naturalistic actions are not typically initiated out of the blue, for no reason. Actions are selected amongst the available alternatives according to current goals and intentions, and execution is only the final step. Thus, action selection and initiation are part of a continuum, and they are embedded in a world with exogenous stimuli.

Libet’s experiments aimed to investigate how actions could be triggered in the absence of exogenous imperatives. Yet, it is not obvious how the type of actions participants performed in Libet’s experiments relate to naturalistic voluntary actions, and thus the naturally occurring action initiation processes (Bayne 2011). It has been argued that the kind of spontaneous movements performed are at best degenerate rather than paradigmatic instances of voluntary action (Roskies 2011). Certainly, the degrees of freedom agents have over their actions in the Libet paradigm are severely restricted compared to naturalistic voluntary actions. Successful control of behaviour involves not only being able to trigger or inhibit pre-specified actions, but also selecting the ones that are better suited to serve one's intentions, goals and desires. Such ability requires processing external information in the light of current goals and desires to guide action.

In the following section, we provide an overview of the cognitive mechanisms involved in action selection and we discuss how they may relate to voluntary action control in naturalistic scenarios.

1.3. Being-in-the-world: voluntary control over action selection

The action repertoire available to humans is vast. At any given point in time, there is an infinite number of actions that could be executed. Thus, the ones that are eventually produced are the result of an action selection process. Perceptual decision-making processes have been widely studied, but rarely related to voluntary action. Yet, the ability to select the actions that best fit our purposes is essential to voluntary control of action. For an action to be goal-directed or reasons-responsive, exogenous stimuli informing about the state of the world need to be integrated with endogenous states specifying intentions and reasons so that the best course of action can be selected. We know that agents are able to tune their decision-making processes to meet their goals, and that they are able to flexibly adapt to changing environments to pursue stable intentions (Deecke 2012). However, little is known about the mechanisms involved in voluntary control over those tuning processes.

1.3.1. The neural basis of action selection

An influential model suggests that multiple possible actions are represented in parallel in frontal and parietal areas (Cisek 2007). Sensory information coming from the dorsal stream (Mishkin et al. 1983; Goodale and Milner 1992) enters a fronto-parietal loop and simultaneously activates different neural populations that represent the currently available *affordances* (Gibson 1966). That is, the actions that are possible in a given context, such as grasping a cup, or picking up a pen. These simultaneous representations of multiple possibilities allow efficient interaction with the environment (Rizzolatti and Luppino 2001; Cisek 2007).

Human functional imaging evidence suggests that these available actions are represented hierarchically (Grafton and Hamilton 2007). The features of the action are specified by different neural populations with varying levels of abstraction, ranging from high-order goal representations to low-level kinematic aspects (Grafton and Hamilton 2007). This hierarchical representation hypothesis is further supported by evidence from single-cell studies in primates, which have shown that individual cells encode general actions (e.g. holding, grasping), specific motor commands (e.g. position of fingers) and temporal aspects of action (Rizzolatti et al. 1998). Further, these hierarchical representations seem to be segregated in multiple effector-specific parallel loops. For example, while the lateral intraparietal area (LIP) is related to gaze control (Snyder et al. 1997), the medial-parietal area (MIP) encodes potential arm movements (Kalaska and Crammond 1995).

Thus, in any given context, a range of possible actions is represented in a hierarchical manner. But how do we choose which one to execute?

Action selection has been framed as an evidence accumulation process during which external and internal factors are integrated and jointly contribute to decision-making (Cisek 2007). However, most previous research has studied the cognitive mechanisms related to exogenous and endogenous action selection separately from each other.

In the following section, we will briefly describe some relevant models of perceptual decision-making that have been put forward to describe how exogenous stimuli determine decisions. Then, we will discuss the way in which endogenous contributions to decision-making have traditionally been studied (i.e. as symmetry-breakers in decisions underdetermined by the evidence). Finally, we will suggest a different approach that may be better suited to explain the way in which endogenous factors may be relevant for voluntary action control.

1.3.2. Evidence accumulation during perceptual decision-making

Research has widely exploited perceptual paradigms in which decisions need to be made on the basis of sensory input, and several models of decision-making have been proposed to explain how exogenous stimuli drive decisions. They share the core concept that decisions are made by continuously sampling and accumulating evidence until a given threshold is reached (Bogacz et al. 2006).

Different types of sequential sampling models exist (*Figure 1.3*). Race models (Brown and Heathcote 2008) suggest that evidence in favour of each alternative is evaluated independently by separate "accumulators". The decision process is terminated when one of the accumulators reaches a threshold. While some race models assume strict independence of the accumulators, others allow for mutual inhibition between them resulting in a winner-take-all competition (Usher and McClelland 2001). In turn, drift diffusion models (DDM) suggest that there is only one single "accumulator" that represents the *relative* evidence in favour of a given choice option (Ratcliff 1978; Ratcliff et al. 2016). Different choice options are represented by different decision bounds, and hence, the final choice depends on which one of the alternative bounds is reached first. These models have gained widespread popularity due to their ability to accurately predict choices and reaction times in a variety of tasks (Forstmann et al. 2016).

A recent influential model incorporates aspects of both race and diffusion models. The "affordance competition hypothesis" (Cisek 2007) suggests that available actions are represented in separate neural populations, but all actions share a same threshold for selection and they compete with each other to reach it (Cisek 2012). Both exogenous and endogenous factors are constantly sampled, accumulated over time and used to increase or decrease the probabilities of one affordance to be selected over another (Cisek 2007).

Each of these models postulates a number of parameters to describe how evidence is processed and accumulated over time on individual trials. For example, in a two-alternative forced choice (2-AFC) task, diffusion models explain decision-making processes based on four main parameters: the drift rate (i.e. the average slope or speed of evidence accumulation), the boundary separation (i.e. how far apart two options are), the starting point (i.e. the bias towards choosing one or the other option) and the non-decision time (i.e. time required for non-decisional processes such as stimulus encoding and motor response preparation) (Forstmann et al. 2016). The decision depends on a 'composite' variable that corresponds to the difference between action alternatives. In contrast, race models assume a certain threshold for all actions and separate accumulators that can each have an independent baseline state and build-up rates. The decision thus depends on a parallel rather than integrated evaluation of the available evidence.

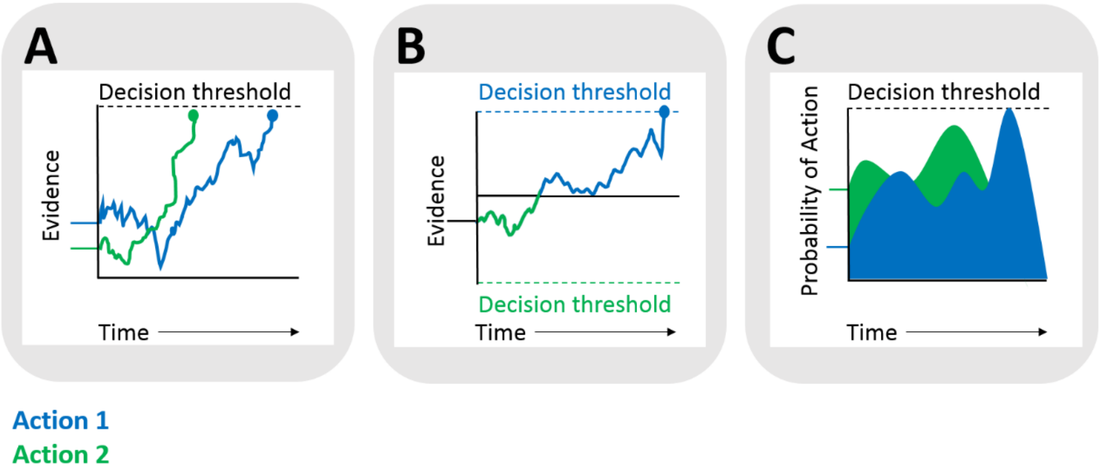

Figure 1. 3. Decision-making models. Schematic illustration of race (A), diffusion (B) and affordance competition (C) models of decision-making. A two-alternative choice example is described. In all models, evidence is sampled and accumulated over time. Evidence accumulation in favour of each action starts at a given baseline (*horizontal lines*) and grows at a certain pace over time (*x axis*). An action is selected as soon as a threshold is reached. In race models (A) independent accumulators process evidence for different actions, and the first one to reach the threshold determines the decision. In diffusion (B) models, a single accumulator integrates evidence for all available actions and drifts towards one out of the multiple available action thresholds, specific to each action. Finally, the affordance competition model (C) assumes that specific neural populations encode the evidence favouring individual actions relative to all available options, and the decision is determined by competition between those different alternatives.

Accumulator models have proven successful in predicting behavioural performance in simple perceptual tasks such as visual random-dot motion tasks (e.g. Shadlen and Newsome, 1996). Furthermore, there is evidence that single-cells in a number of sensorimotor areas participating in decision computations exhibit a build-up activity pattern compatible with a cumulative processing strategy (e.g. Frontal Eye Fields (FEF): Gold and Shadlen, 2000; Medial Temporal (MT): Ditterich et al., 2003;

Prefrontal Cortex (PFC): Kim and Shadlen, 1999, Lateral Intraparietal (LIP): Shadlen and Newsome 1996, 2001; Hanks et al. 2006; Shushruth et al. 2018). Thus, it is generally accepted that decision-making involving sensory information involves some kind of accumulation process.

1.3.3. Endogenous factors as symmetry-breakers

While the contributions of exogenous information to decision-making processes have been systematically studied by directly manipulating the exogenous information itself, the contributions of endogenous factors to action selection have typically been studied by exclusion. That is, by making external evidence either irrelevant or ambiguous. In "picking" scenarios, a free-choice must be made between two equivalent options (Ullmann-Margalit and Morgenbesser 1977; Furstenberg et al. 2015), and in "perceptual guessing" paradigms perceptual decisions need to be made on the basis of ambiguous evidence (Bode et al. 2012, 2013). It is typically assumed that when exogenous evidence is uninformative, endogenous factors are to be held responsible for decisions. Thus, endogenous factors have often been studied in their role as symmetry-breakers in ambiguous scenarios.

The ability to make decisions in ambiguous contexts is important from an evolutionary point of view. A paradigmatic example of this is the frequently quoted Buridan ass (Bode et al. 2013; Haggard 2018). When confronted with two identical piles of hay, the anecdote goes, the donkey was unable to decide from which one to eat, and starved to death. However, unlike Buridan's ass, people are able to make decisions in cases where actions are underdetermined by the environment, and endogenous mechanisms such as intrinsic variability may account for this ability. It has been argued that endogenous, random fluctuations in neural activity may account for variable responses given identical stimuli (Deco and Romo 2008), and thus such variability alone may enable action in the absence of strong exogenous imperatives. In fact, intrinsic variability of neural responses has been shown to contribute to cognitive processing by maximising the available information in mice (Padmanabhan and Urban 2010), and it has been suggested that the active introduction of variability in a deterministic system may provide evolutionary advantages by generating unpredictable behaviour (Maye et al. 2007).

It has been shown that the outcome of a dichotomous free choice can be predicted from the preceding frontoparietal neural activity in simple left vs. right motor tasks (Soon et al. 2008; Bode et al. 2011) and also more abstract choices, such as deciding whether to perform a mathematical addition or subtraction (Soon et al. 2013). Similar to exogenous action selection, the mechanisms underlying such endogenously-selected actions have also been framed as accumulations. It has been proposed that free choices between equivalent options share neurocognitive mechanisms with perceptual guesses about ambiguous stimuli (Bode et al. 2013), and that an 'implicit asymmetry' (Furstenberg et al. 2015) in neural activity – generated by the intrinsic variability in neural responses described above - could determine which option is selected.

However, in naturalistic situations, 'picking' and 'guessing' are typically only one amongst multiple possibilities. When facing a situation that is underdetermined by the environment, one can decide to gather more information, refrain from acting until something changes, or explore a different environment to seek alternatives. Successful control of voluntary action crucially involves the ability to choose the action strategy that best fits the current intentions and goals, and then executing it. Thus, while intrinsic variability can explain how decisions are made in under-determined circumstances, the symmetry-breaking role of endogenous stochastic variability does not explain the more complex issue of how endogenous factors can afford reasons-responsive, deliberate *control* of action.

1.3.4. Endogenous factors as volitional modulations of decision-making processes

Selecting the action that best meets the current requirements involves integrating high-order representations such as intentions and desires with the low-level decision-making processes that drive action. Here, it is suggested that endogenous factors play a determinant role in voluntary control over action selection by implementing higher-order preferences through the top-down modulation of lower-level decision-making processes. *Figure 1.4* depicts a schematic of the theoretical model. How would that work?

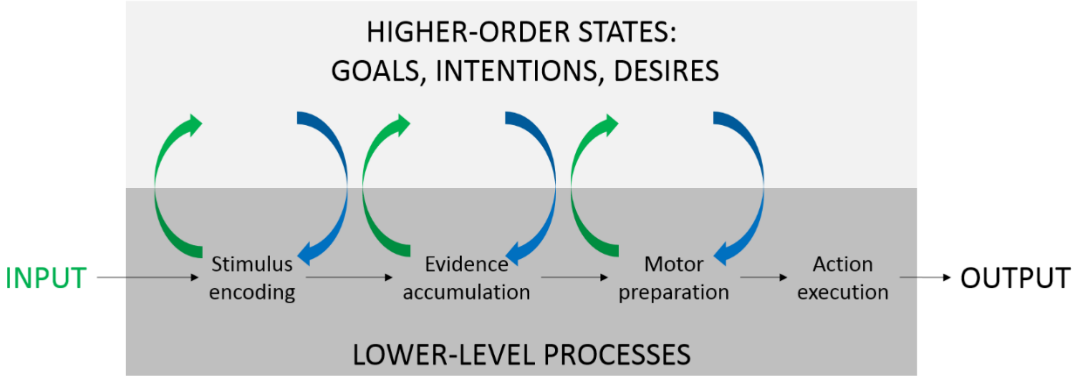

***Figure 1. 4. Background states as distributed effectors of voluntary action control*.** Sensory evidence is constantly sampled, encoded and accumulated in low-level decision-making processes (dark grey, black arrows), while being monitored by higher-order areas (green arrows). However, the very nature of this processing depends on the background state of the system (grey boxes), which is in turn dependent on top-down modulations (blue arrows) from higher-order representations (light grey). Note: the serial description of the decision-making processes is for illustrative processes only.

Theoretical framework

Any decision process takes place given an initial state of the deciding system (i.e. the agent). The initial (i.e. endogenous) state of the agent at any given time is determined by a complex combination of previous experiences, current desires, beliefs and intentions. In turn, such states determine how lower-level decision making processes take place. That is, they specify the 'background conditions' that define which processes can occur, and how. Such background conditions include, for example, the baseline activity of specific neural populations in early sensory areas, or the excitability of the motor cortex. These background conditions can be bottom-up modulated by external stimuli (e.g. Strafella and Paus 2000), but they can also be modulated by top-down influences from higher-order endogenous states (e.g. Dorris and Glimcher 2004; Galaro et al. 2019). For example, if one is searching for vegetarian places to eat tonight, information about seafood restaurants will most likely have a minimal influence in the decision outcome – it is irrelevant given the constraints imposed by the higher order intention that the restaurant be vegetarian. In turn, if one is very hungry and really wants to eat soon, some types of information such as waiting time or proximity will perhaps be weighted more than other factors such as the number of options on the menu, or how nice the venue is.

Achieving this type of control requires being able to adaptively tune how evidence is processed in the light of the current intentions, and these top-down regulations involve several types of well-known processes. In the previous example, selective attention can be argued to filter out irrelevant evidence, and endogenous modulations of the speed-accuracy trade-off (SAT, Wickelgren 1977) may be involved in achieving a fast response. Thus, volitional control of action selection engages several kinds of distributed processes exerting top-down influences on all levels of the sensorimotor hierarchy.

The idea that processes such as attention may operate passively in response to the environment but also be actively engaged by volitional processes has been suggested before (Deecke 2012), but it has received little empirical attention in the voluntary action literature. Yet, it is argued here that the ability to endogenously modulate processes of that kind constitutes a core feature of voluntary control of action.

Empirical approach

To investigate this type of top-down modulations, one can use the tools traditionally used in perceptual decision-making research. Rather than experimentally controlling the characteristics of external evidence, one can manipulate specific features of the higher-order states that guide decision-making processes, and investigate its neural signature. As illustrated in the example above, an intention has multiple dimensions. It has a "content" (e.g. finding a restaurant to eat out tonight), and also some "dynamic" characteristics that specify how that content is to be achieved (e.g. how much time is available for implementing that intention). These dynamic characteristics can be experimentally manipulated, and provide a useful tool for investigating how endogenous factors (or rather, specific features of endogenous intentions such as the requirement to decide quickly) can influence lower-level decision making processes. Thus, a simple way to investigate how endogenous goals shape decision-making is by studying some specific features of endogenous goals, such as SAT.

The effects of SAT manipulations have been described at different levels of the sensorimotor hierarchy. It has been shown that low-level encoding of sensory evidence is boosted in high-urgency paradigms (Steinemann et al. 2018). Further, modulations of parameters in accumulation models can explain how cognitive flexibility might be implemented at a computational level (Ratcliff et al. 2016). It has been proposed that fast responses could be achieved by lowering the action threshold (Cisek et al. 2009), increasing the baseline activation state (Forstmann et al. 2008; Steinemann et al. 2018) or increasing the rate of build-up to threshold (Twomey et al. 2015). A number of brain areas in action preparation and execution networks have been shown to have increased baseline activity in high-urgency tasks. Functional imaging studies have found evidence that the striatum, the pre-SMA (Forstmann et al. 2008) and the dorsolateral prefrontal cortex (dlPFC) (Veen et al. 2014) have higher baseline activities in tasks with higher urgency. Further, single-cell recordings in monkeys have found evidence for increased baseline activity in premotor areas and primary motor cortex (Thura and Cisek 2016) in urgent scenarios, and stimulation of the medial frontal cortex produces effects compatible with baseline-to-threshold distance reduction (Jha et al. 2016).

While effects of experimentally controlled SAT modulations have been extensively investigated, they have only been recently linked to volitional action control (Thura and Cisek 2014, 2016). Studies of perceptual decision-making in monkeys have shown that the firing patterns in motor areas track an internal decision-variable that is modulated by self-imposed SAT (Thura and Cisek 2014, 2016). Monkeys performed a perceptual decision-making task in two blocked conditions. In one environment, the faster monkeys executed actions, the more trials they could perform. Thus, hasty actions were more likely to maximise reward by increasing the number of available trials, despite decreasing the overall accuracy of responses. In another environment, slow actions were more likely to maximise the reward, since the execution of fast actions did not increase the number of trials they could perform. The studies found that monkeys voluntarily adapted the speed of their reactions to both environments, and the neural signature of this voluntary adaptation was found in motor areas. The firing rate of motor areas was found to correlate with changes in volitional adjustments of the SAT when monkeys decide to make fast vs. accurate decisions. In particular, they found that fast decisions correlated with higher baseline activations.

In humans, voluntary SAT modulations of this kind may respond not only to changing external environments, but also to internal deliberations. For example, one may be pondering about whether to catch a train or drive home, and suddenly remember that actually the car is broken. Together with the knowledge that the last train is about to depart, this may trigger a sense of urgency that may in turn affect the dynamics of subsequent decision processes. For the sake of the argument, let us assume that the car did not *just* break, and that the person was not reminded of the train schedule because they *just* saw it somewhere. The sense of urgency in this example did not result from any external *event*. Rather, internal endogenous states of the agent, in this case their beliefs about the car and the train, crucially affected the decision-making process leading to action. The ability to adapt behaviour (i.e. run to the train) to reach a goal (i.e. get home at a reasonable time) is a fundamental type of voluntary action control. The conceptualization of the adjustment of the SAT as an essential mechanism of voluntary action control can contribute to understanding how specific features of high-order representations such as intentions or distal goals can shape the way in which we perceive and interact with the world as autonomous agents.

1.4. A hierarchical model of voluntary action control

We started this review with the aim of describing the cognitive mechanisms relevant for a naturalistic account of voluntary action. Here, we propose a hierarchical model of voluntary action control (see *Figure 1.5*).

Voluntary actions result from the selection of an action amongst the multiple available affordances, and its eventual execution. Action selection and execution are thus part of a continuous decision-making process that takes place as an interaction between the endogenous states of the agent and the exogenous input it receives from the environment. The agent's internal states at any point in time result from a complex combination of previous states, current goals and desires, and contextual contingencies. We refer to this internal state as the *background* state. In turn, this *background* state can be modulated both bottom up, by the influence of external exogenous stimuli, and top-down, as a result of changes in higher-order endogenous processes. The background state of an agent determines how decision-making processes take place, from stimulus encoding in early sensory cortices to motor preparation.

Global voluntary control of action refers to the ability to modulate some of these top-down effects as a result of deliberative processes. For example, when one resolves to make a quick decision in urgent situations, processes such as sensory encoding, evidence accumulation or motor excitability are duly affected. While the ability to act under time pressure is widespread across the animal kingdom, other kinds of global modulations may be able to account for other types of uniquely human behaviours. For example, one could study how different moral or political stances influence decision-making processes by means of top-down modulations.

Hierarchically positioned below this *global* level of control, we suggest that *local* voluntary action control allows direct, conscious control over action triggering. Since action-triggering plays a pivotal role in behaviour, this type of control constitutes a kind of “safety-net” mechanism. The achievement of any goal ultimately depends on the execution or inhibition of a specific series of actions and, in fact, it has been argued that the possibility to choose between doing or not doing is the most central feature of voluntary action (Ebbesen and Brecht 2017). Once an action has been selected for execution (i.e. it has reached a threshold after a sampling process), a motor command needs to be sent. If agents have no conscious insight into the cognitive processes responsible for triggering an action, then their ability to consciously control the execution or inhibition of specific actions seems limited.

We have suggested that motor initiation processes encoded in the medial frontal cortex (pre-/SMA complex) are metacognitively monitored in parietal areas. In turn, that information is fed to frontal areas which close the metacognitive loop and enable cognitive control. This metacognitive monitoring system is the cognitive basis of phenomenal “urges” to move, and it affords the possibility to voluntarily control the execution or inhibition of a specific action through executive control – through frontal metacognitive feedback. This type of *local* voluntary action control is subject to *global* control effects. The whole action selection process is determined by top-down mechanisms, including motor initiation. Thus, understanding *local* action control requires studying it in the context of this more *global* cognitive control system. While the ability to consciously inhibit undesired movements seems essential for our experience of control, many voluntary actions are executed without a strong accompanying experience of “willing”. In fact, the most common marker of voluntary action might not be a strong conscious “fiat”, but rather a certain “awareness of consent" (Ach 2006). That is, a feeling that the action naturally follows from endogenous motivations or serves current purposes.

Joëlle Proust's (2008) definition of mental action nicely summarizes the way in which both *global* and *local* voluntary action control are conceptualized here:

> "A *willing* or a *trying* is a mental event through which an operation from the repertoire 1) is called on because of its instrumental relationship to a goal, and 2) is thereby made available to executive processes."

In our hierarchical framework, mental events such as deliberate decisions or "phenomenal urges" are conceived as cognitive mechanisms that enable specific operations of *global* or *local* top-down regulation, respectively, which in turn constitute the basis of voluntary action control.

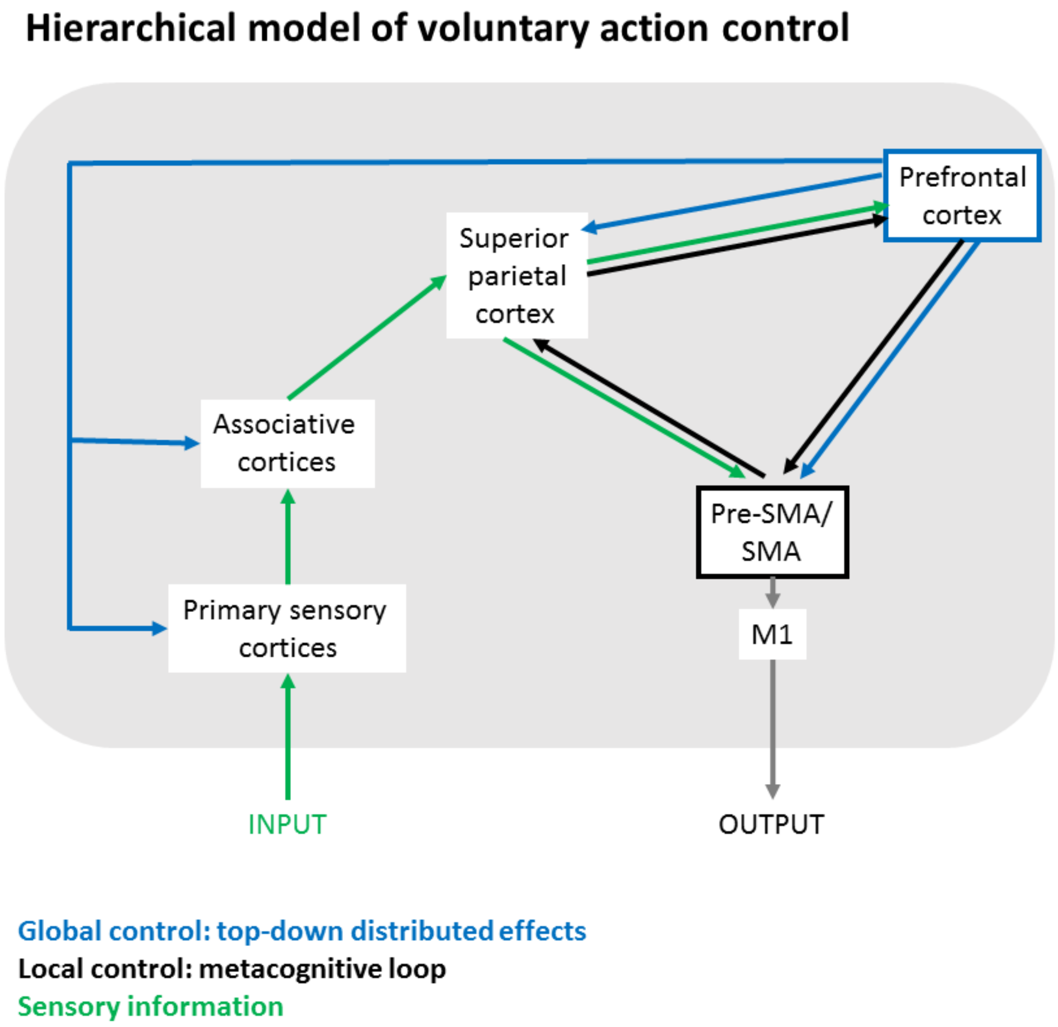

Figure 1. 5. A hierarchical model of voluntary action control. At any given time, information about the environment is received through sensory input (*green arrows*). The way in which this information is processed depends on the background state of the system (*grey box*). This state depends on continuous top-down modulations from frontal areas to all levels of the sensorimotor hierarchy (*blue arrows*). Global voluntary control of action involves the ability to modulate some of these top-down effects as a result of deliberative processes. In turn, local voluntary control of action involves a metacognitive monitoring and feedback loop (*black arrows*). Motor information about action initiation processes in the SMA is sent to superior parietal areas. In turn, this information is fed back into frontal areas, which can top-down modulate its inhibitory effect on the SMA.

1.5. Summary

In this review, we have presented an integrative framework for the study of voluntary action that overcomes some of the caveats of traditional approaches. We have argued that voluntary action control can happen at two hierarchical levels. *Global* voluntary action control spans all levels of the sensorimotor hierarchy, and it involves the ability to intentionally (i.e. as a result of a deliberative process) tune decision-making processes so that they are optimal to realise the agent's current intentions and reach the desired goals. In turn, *local* voluntary control of action allows control over action initiation and execution processes and primarily involves supplementary motor areas, parietal and frontal areas. In sum, we believe that this hierarchical model of voluntary action control provides a fertile cognitively and theoretically-principled framework to investigate naturalistic voluntary actions.

1.6. Thesis overview

Within this framework, we developed novel methods to investigate both *local* and *global* aspects of voluntary action control. The findings presented here contribute to advancing the scientific understanding of conscious control of voluntary action at the levels of action selection and initiation.

In **Part I**, we investigated the relationship of the RP with awareness of motor intention. In particular, we aimed to investigate the distinctive contributions of prospective motor preparation processes and retrospective reconstruction on intention awareness without conflating the experience of intention and action execution. In **Chapter 2**, we developed a series of novel paradigms based on random probing of awareness during a self-paced task. We acquired different types of awareness reports, and we investigated how the signal preceding the awareness probe correlated with the reported intention. In **Chapter 3**, we used a real-time EEG monitoring technique to deliver probe stimuli either in the presence or absence of the RP, and then require participants to report their awareness at the time of the probe. We combined this real-time probing method with a *Go/No-Go* task to investigate the interaction of prospective motor preparation signals with retrospective reconstruction.

In **Part II**, we investigated global aspects voluntary action control. In **Chapter 4**, we developed a paradigm in which people made self-paced actions on the basis of external input. The evidence was either strong (i.e. an easy decision) or neutral with respect to the decision (i.e. a kind of perceptual guess), and there was no time pressure to respond. This allowed us to study the neural correlates of low-level evidence accumulation and action initiation processes in a context where actions were informed by external evidence, but not directly triggered by it. Further, in **Chapter 5** we modified the paradigm developed in chapter 4 to incorporate manipulations of the speed-accuracy trade-off. We investigated how the correlates of decision-making dynamically tracked changes in endogenous goals (operationalised as "being fast" or "being accurate"), and how action initiation processes differed in various urgency conditions.

Part I. Forward now!

“It often happens, when the absence of imperative principle is perplexing and suspense distracting, that we find ourselves acting, as it were, automatically, and as if by a spontaneous discharge of our nerves, in the direction of one of the horns of the dilemma. But so exciting is this sense of motion after our intolerable pent-up state, that we eagerly throw ourselves into it. 'Forward now!' we inwardly cry, 'though the heavens fall.'” William James.

Chapter 2

From awareness to the brain: the neural correlates of motor intention awareness

Synopsis

The aim of this chapter is to investigate the relationship between awareness of intention and the readiness potential (RP). We developed a new paradigm to overcome some of the limitations of the classic Libet task. In particular, the problem of biased sampling of the EEG signal and the offline method of reporting awareness. In this study, we ran three different EEG experiments where participants performed a self-paced task and were randomly interrupted and asked about their motor intention at that time. In Experiment 1, participants reported motor intentions by executing an action. In Experiment 2, by inhibiting an action that they might have been preparing at the moment of the interruption by the probe. In Experiment 3, they provided a rating on a Likert scale to indicate how ready they felt they were to move at the time of probing. We found evidence that the RP is related to prospective awareness of intention (Experiments 1 and 3), although delayed judgements of intention awareness may be less sensitive to motor preparation processes at the time of probing (Experiment 2).

2.1. Introduction

An experience of intention to move accompanies execution of some voluntary actions. The classic Libet studies (Libet et al. 1983) suggested that the RP starts before intention is consciously accessed as measured by offline recall-based reports, yet the interpretation of the RP and its temporal relation to awareness of intention remain controversial.

Random probing and online reports

In the original Libet paradigm (Libet et al. 1983) the estimation of the conscious experience of intention is only reported 'off-line'. Participants noted the time at which they became aware of their intention to move, but reported this only after they actually executed the action. The traditional interpretation of the experiment assumes that the reported time corresponds to a genuine experience preceding action. However, the method famously introduces the possibility of a reconstructive or 'postdictive' process. For example, people might not have any genuine experience of intending to act at all but, having noticed that they had just acted, they might infer an intention, and project a corresponding experience backwards into the stream of consciousness to provide a candidate cause for their action (Wegner 2002).

In the three experiments in this chapter, our participants performed self-paced key presses while viewing a letter stream (cf. Soon et al. 2008). An 'interrupting' probe was inserted into the letter stream *at random*. Thus, the cue would sometimes intercept voluntary motor preparation before an action was executed, and sometimes not. Crucially, the sampling strategy of our paradigm differs from most RP studies. Classical RP studies lock EEG data to movements and hence show neural activity prior to action, but ignore neural activity at other times. This biased sampling means that the RP pattern might not be *specifically* related to voluntary action preparation, nor to awareness of intention (Mele 2011). In particular, RP-like fluctuations might also occur in the absence of actions and conscious intentions, but those RP-like fluctuations would be invisible to the action-locked methods used in classical RP paradigms. Recent computational models showed that simply averaging an ongoing stochastic signal time-locked to a threshold crossing event can reproduce the form of the RP (Schurger et al. 2012; see also Murakami et al. 2014). Since the underlying fluctuations in such models are continuous, RP-like forms should also occur even in the absence of action, and would not be related to a process of intentional action preparation as classically suggested.

The studies presented here search for neural precursors of voluntary action using a method that avoids this particular bias. We interrupted ongoing EEG with a random cue. If RP-like signals simply occur as part of ongoing stochastic fluctuation and are unrelated to awareness, they should be equally visible whether the cue interrupts an intention to act or does not. On the other hand, if the RP *is* specifically related to the participants' experience of intention, we should see some RP-like signal prior to an interrupting cues where participants reported a feeling of intention by reacting to the cue, but not otherwise.

Functional relevance of motor awareness

Libet's method contains another assumption that has received less attention. It implicitly assumes that people can only consciously access their preparation for the impending movement once this has reached a specific, fixed threshold. However, motor preparation might be accessible in some form even before this threshold is reached. It has been suggested that the experience of intention to move may develop progressively, rather than having a single categorical threshold and a fixed onset (Matsuhashi and Hallett 2008; Verbaarschot et al. 2016). We hypothesise that, although people may normally experience intention to move only at late stages (when some ‘spontaneous threshold’ is exceeded), the intention is *potentially accessible* at earlier stages (‘latent threshold’). That is, an intention *may be accessed* given certain conditions, before it becomes a foreground mental content.

The concept of variable thresholds for awareness of intention makes sense for many actions. In the restricted setting of the Libet experiment, the action is not motivated by any reason or constraint, nor does it bear any consequences in the external world beyond its own execution. However, where information about motor preparation is crucial for guiding behaviour, a stronger experience of preparation at earlier stages may be highly functional. For example, people are not normally aware of their intention to extend the leg before crossing a road. However, if a motorbike happened to be approaching rapidly, it would be useful to be aware of the impending movement so as to modify or inhibit the action and prevent any accident. This flexible-threshold model of intention awareness thus resembles models of perceptual attention: a close-to-threshold stimulus is more often consciously perceived when attention is directed to it.

In our paradigm, the interrupting cue acted as an intention probe, directing participants’ attention to their current motor preparation state. The participant’s experience of their own motor preparation determined how they should respond to the cue, thus being behaviourally relevant. In Experiment 1, participants were instructed to respond to the interrupting cue if they felt they had already begun preparation of their next movement, and to ignore it if they were not preparing to move (‘contingent action’, i.e. responding to the cue was contingent on participants’ experience of intention). In Experiment 2, participants were given the opposite instruction: they had to veto any incoming actions if they felt they were about to move on presentation of the orange letter cue, and to ignore the cue otherwise ('contingent veto', i.e. experience-dependent veto). Finally, in experiment 3, participants had to provide a rating indicating how ready they felt at the time of moving on a scale from 0 ('not preparing at all') to 7 ('about to move').

In line with Matsuhashi and Hallett (2008), we hypothesised that, if probed, participants are able to consciously access early stages of motor preparation and can use their experience of intention as a basis for action decisions.

2.2. Experiment 1: Contingent Action

In this experiment, participants performed a self-paced movement task while staring at a letter stream (similar to Soon et al. 2008). Occasionally, they were presented a visual cue as an intention probe. If they felt they were preparing their next movement at the time the probe appeared (i.e. if they were aware of a motor intention), they were instructed to respond to the cue by executing a movement. The awareness report method was, thus, a 'contingent action' (i.e. an action contingent on their awareness of motor intention). This type of report was designed to be as immediate as possible. To report an intention, participants had to provide a response by executing the same movement that they might have been preparing. Our online probing and report method allowed us to 1) interrupt motor preparation at various stages, 2) make real-time awareness of intention relevant for task performance and 3) eliminate the possibility for reconstruction by turning the action itself into a report. This is in contrast with Libet's classic offline report method, and aims to overcome some of its limitations.

Further, in this experiment we aimed to directly test one of the concerns related to the post-hoc Libet report. In particular, the relationship between perceptual attention to the clock and the time of reported awareness. Libet's participants reported the position of a rotating clock hand after every action. Several authors have speculated that the visual processing of the clock might not be independent of either the actions made or of the experiences reported, and that the RP might reflect overlapping potentials related to spatial attention (Libet 1985; Miller et al. 2011; Guggisberg and Mottaz 2013).

In this experiment, after executing an action, participants were sometimes asked to report the letter that was on the screen when they first felt an intention to move, a recall-based mental chronometry method analogous to Libet's one. The kind of periodic stimuli we used elicits a steady-state visual evoked potentials (SSVEP) at the stimulation frequency and its harmonics. SSVEP amplitude reflects the variations in visual pathway processing gain with visual attention. For example, switches of attention between two stimuli of different frequencies can be tracked by analysing fluctuations of SSVEP amplitude at the corresponding frequencies (Müller et al. 1998). Canonical correlation (CCA) between EEG and a target frequency can be used to study the allocation of visual attention (Lin et al. 2006). In typical BCI applications, multiple stimuli flicker at different frequencies. The frequency that has the highest CCA score is interpreted as selected in attention (e.g. Chen et al. 2015). In this study we used a single stimulation frequency, corresponding to the letter stream presentation, and we used CCA analysis to study how attention to the letter stream was modulated over time. In particular, we used canonical correlation to investigate the temporal relation between RP, conscious intention, and the dynamics of visual attention to the letter stream used to report intention.

2.2.1. Materials and methods

Participants

A previous similar study (Schurger et al. 2012) showed a large effect size ($d_z = 0.9$) for the contrast between the pre-stimulus EEG amplitude for responses to a tone that involved lower vs higher reaction times (Schurger A, personal communication). A power analysis for a paired-samples *t*-test on mean RP amplitudes contrasting the two conditions of main interest in the previous study indicated a required sample size of 19 participants for a power of $\beta = 0.95$ and $\alpha = 0.05$. Although our core inferences were based on cluster statistics (permutation tests) rather than parametric tests (see below), this estimate was used as a stopping rule for the current experiment.

Twenty-six subjects were initially recruited from the Institute of Cognitive Neuroscience Subject Database. All participants were healthy, right-handed, young adults with normal or corrected to normal vision, no known disabilities and no history of neurological or psychological disorder. The study was approved by the UCL Research Ethics Committee and written informed consent was obtained from all participants before beginning the experiment. Subjects were paid £7.50 per hour.

All participants were invited to a single EEG session. Two participants did not show any RP, based on analysis of an independent subset of self-paced epochs used for selection purposes only (see below) and were therefore excluded from further analysis. Visual inspection of the data revealed that five additional participants exhibited excessive noise throughout the whole EEG time course due to technical problems during the recording session, and were excluded from further analysis.

Nineteen participants (13 female) were therefore included in the final dataset (M_{age} = 22.5, *SD* = 3.13; range: 19-30 years). For the CCA analysis only, data from one further participant could not be used due to noisy occipital EEG channels (thus n = 18 for these analyses).

Stimuli and experimental design

Procedure: Participants sat in a quiet room and viewed the stimuli on a computer monitor at 50 cm distance. The instructions for the task were first displayed on the computer screen and then verbally repeated by the experimenter before the beginning of the experiment.

Participants performed a simple motor task. The task was programmed in Matlab R2014b and Psychophysics Toolbox v3 (Brainard 1997, Kleiner et al. 2007). Subjects made actions by pressing the space bar key of a standard computer keyboard with the right index finger.

Stimuli: The task was divided into 6 blocks of 20 letter streams each. The letter streams were a stream of random lower-case consonants (b, c, d, f, g, h, j, k, l, m, n, p, q, r, s, t, v, w, x, z). Any letter repetitions were separated by at least 7 other letters. Letters were either black or orange and were presented on a light grey background without any blank interval. Black letters were presented for 216 ms, and orange letters for 266 ms. Each stream contained multiple orange letters, which appeared at random times 3 s to 20 s after stream onset or after the preceding orange letter. The duration of each letter stream was variable and dependent on pseudorandomised constraints (see *Time of awareness report* section below).

Task

Self-paced condition: Participants were instructed to fixate the letter stream and to make self-paced keypresses whenever they felt like it. They were specifically told to not pre-decide to respond to a specific letter and to try to be as spontaneous as possible (e.g. not to make a keypresses after a fixed interval). Furthermore, participants were instructed to pay special attention to the moment they first felt an 'urge' to move and to remember the letter that was on the screen at that time as they would sometimes be asked to report it.

Latent awareness/No awareness condition: Occasionally, one letter was presented in orange rather than black (see *Stimuli* section above). Participants were told to respond to the orange letters only if they felt they were already preparing the next self-paced movement when the orange letter appeared. Thus, they reported conscious intention to perform a self-paced action by a motor response that was similar to the action they had been preparing. We chose this method as being more immediate than verbal reports or report via a different motor response. We reasoned these alternative methods would involve additional cognitive processes of switching between alternative actions. Any keypress occurring within 2 s of an orange letter was considered a report of conscious intention, as distinct from a self-paced action. These were labelled 'Latent awareness' epochs. Orange letters with no keypress within 2 s after the presentation were labelled 'No awareness' epochs, and considered a report of absence of conscious intention at the time of the orange letter. We deliberately chose a response window that was well in excess of normal simple reaction time because participants' response to the orange letters involved a two-step cognitive process (i.e. first, deciding about their motor preparation state and second, executing an action (or not) to report the outcome of that decision). Further, the experiment can be viewed as a mix of an endogenous (self-paced key presses) and an exogenous (response to the orange letter conditional to awareness) task, and previous studies have shown that switching between endogenous and exogenous modes of action results in costs in reaction time (Obhi and Haggard 2004).

Time of awareness report: After some actions selected at random, the letter stream was terminated and a prompt appeared asking: ‘Which letter was on the screen when you first felt the urge to move?’. Participants responded by pressing the corresponding key on the keyboard. The letter they pressed was viewed on the screen and they had to confirm that they reported the correct letter by pressing the space bar. Because we expected there would be more self-paced than latent awareness actions, the probability of terminating a letter stream after a self-paced action was set at 0.2, while it was set at 0.5 for latent awareness actions. This precaution was taken to maximise the number of time of awareness estimates for latent awareness actions, by asking the time of awareness question in a higher proportion of those actions. After responding to the awareness question, there was a 2 s interval before the following letter stream started.

Figure 2.1 provides a schematic illustration of the task.

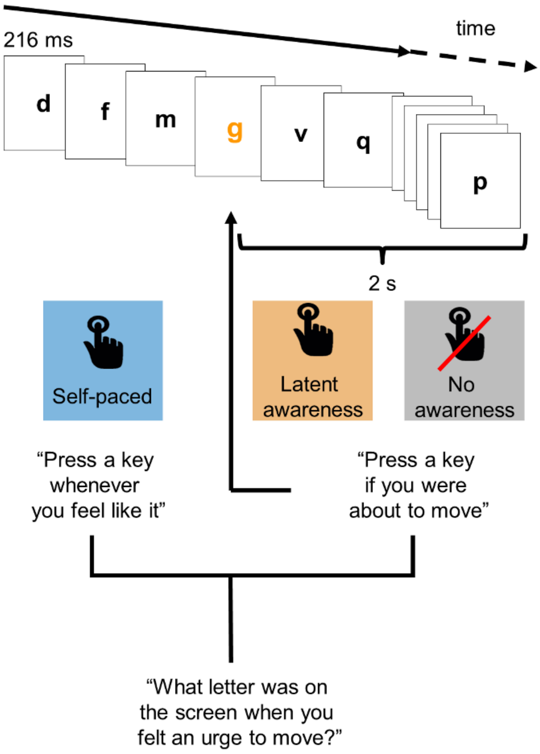

Figure 2. 1. Experiment 1 design. Participants were asked to press a key with their right hand whenever they felt like it (Self-paced). They were instructed to respond to orange letters (bold) only if they felt they were already preparing their next movement (Latent awareness), and to ignore them otherwise (No awareness). This was our online measure of intention awareness. Additionally, participants were asked to report the letter that was on the screen when they felt the urge to move after 20% of self-paced actions and after 50% of actions executed within 2 s after the orange letter. This was our offline measure of intention awareness.

EEG recording

EEG was recorded from 26 scalp sites (FZ, FCZ, CZ, CPZ, PZ, POZ, FC1, FC2, C1, C2, CP1, CP2, F3, F4, F7, F8, C3, C4, CP5, CP6, FC5, FC6, P3, P4, O1, O2) using active electrodes (g.LADYbird) fixed to an EEG cap (g.GAMMAcap) according to the extended international 10/20 system. EEG data were acquired using a g.GAMMAbox and g.USBamp with a sampling frequency of 256 Hz and 0.01 Hz high-pass and 100 Hz low-pass online filters. Signal was recorded using g.Recorder (G.tec, medical engineering GmbH, Austria). All electrodes were online referenced to the right ear lobe. Vertical and horizontal electroocular activity was recorded from electrodes above and below the right eye and on the outer canthi of both eyes.

Behavioural data analysis

The time of awareness was estimated based on the reported letters in both self-paced and latent awareness actions. On average, 75 reports per participant were analysed in the self-paced condition (*SD* = 12.70) and 39 in the latent awareness condition (*SD* = 13.92). A goodness-of-fit Chi-square test for uniformity was performed on the distribution of reaction times (RT) to orange letters for each individual participant using the *chisq.unif.test* function from the *spgs* R package. For this analysis, reaction times were divided into 9 bins, corresponding to the 9 letters presented from the orange letter onset to the end of the response window (orange letter included). The same test was run on the distribution of self-paced keypresses before the presentation of orange letter, which was an equivalent period of time. For this analysis, keypress times were divided in 10 bins, corresponding to the 10 letters presented before the orange letter onset. Ten bins instead of 9 were chosen for the self-paced analysis because one participant did not perform enough keypresses in the 9 letters before the orange probe for the analysis to be robust. For the rest of participants, inference results did not change when using only the keypresses during the last 9 letters.

EEG analysis

Preprocessing*:* EEG data were processed using Matlab R2014b (MathWorks), Matlab R2017b (MathWorks), SPM12 (Statistical Parametric Mapping software, version 12), EEGLAB version 13.5.4b (Delorme and Makeig 2004) and Signal Processing Toolbox R2017b.

First, scalp and eye electrodes were re-referenced to the average of two mastoid electrodes. Continuous EEG and EOG data were filtered with a 0.01 Hz high-pass filter. Then, data were downsampled to 200 Hz and filtered with a 30 Hz low-pass filter. These filters were applied off-line using a 5th order Butterworth filter with zero phase shift.

Second, an independent component analysis (ICA) was computed on the continuous data using the EEGLAB *runica* algorithm. Vertical eye movement components were visually identified and removed from the signal. Removal of horizontal eye movements in participants where they were identifiable (n = 10) did not change the main inferential results.

Next, EEG signals were locked to either a) orange letters and b) the next letter appearing after a keypress. We will refer to a) as stimulus-locked data, and b) as action-locked data (we locked to the letter immediately after the keypress so that EEG activities evoked by the letters themselves would have the same influence on both action- and stimulus-locked analysis). Epochs started 2.5 s before the event and finished 1 s after it. Baseline correction was performed using the 500 ms at the beginning of the epoch [-2.5 s to -2 s relative to event] for RP analysis. Finally, artefact rejection was performed by removing all epochs with >120μV fluctuations from baseline in any of the preselected channels (FCZ, CZ and C3 for RP analysis, POZ, O1 and O2 for CCA analysis – see below). Epochs in which there was a key press in the [-3 to 0 s] interval preceding the event of interest (i.e. orange letter or action) were rejected to prevent overlapping evoked-potentials. The number of epochs analysed was on average 55 per participant (*SD* = 28.43) for latent awareness and 65.36 for no awareness reports (*SD* = 53.68).

RP analysis for participant exclusion: Twenty percent of self-paced actions were selected randomly for each participant (*M* = 60.09, *SD* = 18.47) and visually inspected. If no RP was apparent (i.e. no increasing negativity towards the time of the action was visible), the participant was excluded. Only the remaining 80% of self-paced actions were used for statistical inferences drawn from non-excluded participants (*M* = 233.89, *SD* = 76.17).

Canonical correlation analysis: Canonical correlation coefficients between a 4.63 Hz (i.e., 1/216 ms) reference sine wave and the signal of each of the three preselected occipital electrodes (POZ, O1 and O2) were calculated (Lin et al. 2006) for both self-paced and latent awareness actions in action-locked data, and for latent awareness and no awareness epochs in stimulus-locked data. The electrodes were selected based on previous studies using CCA for SSVEP analysis (e.g. Chen et al. 2015). The maximum CCA coefficient was calculated for each sine wave – electrode pair on single-epoch data using a sliding window of 324 ms window length with 95 % overlap. The coefficients of the 3 electrodes were then averaged for each participant.

Change-point analysis: We estimated the onset of the RP in action-locked EEG data and the changing-point in CCA scores with a regression-based method (Mordkoff and Gianaros 2000). We used the Signal Processing Toolbox *findchangepts* function in Matlab in order to identify the best-fitting change point for both data types (RP and CCA scores) in each participant's average trace. For the RP onset analysis, we applied the function to the period preceding action [-2.5 0 s] in order to specifically estimate the onset of the RP rather than other abrupt changes happening after action execution. In contrast, we had no prior hypothesis about the time of visual attention modulations, so no time assumptions were made for the CCA change-point analysis: we searched for the optimal change-point through the entire epoch [-2.5 to 1 s relative to action].

Statistical analysis:

Behavioural data was analysed using Statistical Package for the Social Sciences, version 22 (SPSS Inc, Chicago, IL, USA).

Statistical tests on averaged EEG data were run using FieldTrip toolbox (Oostenveld et al. 2011) cluster-based permutation analysis (Maris and Oostenveld 2007). The main contrast of interest involved a stimulus-locked analysis, comparing potentials preceding orange letters that interrupted a conscious intention, and thus elicited a response, and those that did not. An additional response-locked analysis compared the amplitudes of the readiness potentials preceding self-paced versus latent awareness actions. Three electrodes over the SMA-preSMA and contralateral motor cortex (FCZ, CZ and C3) were preselected for analysis as being most relevant for motor preparation (Khalighinejad et al. 2018). The cluster-based tests were performed on electrodes using the following parameters: one-tailed dependent samples t-test, time interval = [-2 0 s relative to the event of interest], at least two neighbouring electrodes contributing, number of draws from the permutation distribution = 10000.

We also analysed single-trial EEG using linear mixed effects modelling. We fitted a logistic regression to predict the probability of reporting awareness based on one continuous predictor: the mean EEG amplitude at the time of probing [-0.1 to 0 s relative to orange letter onset], averaged across the three electrodes of interest (FCZ, CZ and C3). Each participant was treated as a random variable and random effects for the intercept and the slope of the continuous predictor were estimated to take into account the within-subject nature of the experimental design. The model was generated using the *bglmer* function in the R package *blme*. We used a weakly informative Wishart distribution for the covariance matrix prior (df = 4, scale = 10, see Chung, Gelman, Rabe-Hesketh, Liu & Dorie, In press) to obtain a non-degenerate covariance matrix and the default flat prior for fixed effects. The goal of this analysis was to test whether the probability of reporting awareness of intention depends on the EEG amplitude at the time the orange probe was presented.

2.2.2. Results

2.2.2.1.Behavioural results

Preliminary analysis

Participants executed a self-paced action on average every 11.29 seconds (*SD* = 2.53). In the latent awareness task, participants reported an intention to move after 50% of orange probes (*SD* = 17.97). The average waiting time between self-paced key presses and the frequency of latent awareness reports were negatively correlated (ρ = -0.51, *p* = 0.02). Participants who pressed more frequently in the self-paced task were also more likely to report awareness of intention in response to an orange letter probe.

In a very few actions, participants reported a letter that was not shown during the 2 s before the action (*M* = 3.14% of self-paced actions, *SD* = 3.97 %; *M* = 3.69% of latent awareness actions, *SD* = 3.29 %). These reports were assumed to be errors or lapses of attention or memory, and were excluded from behavioural analyses. Responses to orange letters in latent awareness epochs were made on average 1 second after the presentation of the orange letter (*M* = 1.06 s, *SD* = 0.54 s).

To investigate whether participants were indeed responding/not responding to the orange letters as a function of their experience of intention, we analysed the distribution of reaction times (RT). If participants were ignoring the letters and our instructions, and focussed only on the self-paced action task, we would expect a uniform distribution of keypresses after orange letters, since these were presented at random times. First, we performed a goodness of fit Chi-squared test to the distributions of self-paced key-presses over ten letters presented *before* orange letter, across all trials within each participant. This analysis aimed to test whether the assumption that random probing would result in a uniform distribution of self-paced keypresses. Out of the 19 tested participants, only one showed a distribution of keypresses that significantly differed from the expected uniform one (*Figure A1*). Second, we ran the same analysis on the distribution of keypresses over the 9 letters *after* the presentation of an orange probe (i.e. the RT distribution). The RT distributions of each participant significantly ($p<0.05$, n = 11) or marginally ($p<0.1$, n = 2) differed from uniform in most participants (total n = 13 out of 19) (*Figure A1*). This suggests that the orange letters were indeed processed and influenced participants' behaviour in accordance with the instructions.

Timing of awareness

We estimated the timing of awareness by subtracting the time of the presentation of the letter during which participants responded from the time of presentation of the reported letter. The times of spontaneous awareness were -0.43 s (*SD* = 0.19 s) for self-paced actions, and -0.50 s (*SD* = 0.2 s) for latent awareness actions ($t_{(18)} = 1.82$, $p = 0.08$).

In latent awareness reports, individual actions were sorted according to whether the reported letter was presented *before* or *after* the orange letter (if the orange letter itself was reported, the epoch was included in the *after* category). A significant difference in the percentage of epochs in these two categories was found ($t_{(18)} = 25.16$, $p < 0.001$). On average, conscious intention was reported to be *after* the orange letter (*M* = 90.86%, *SD* = 7.08%) more often than it was reported to be *before* (*M* = 9.13%, *SD* =

7.08%). This pattern is consistent with the idea that participants' intention was consciously access*ible* to them, though it had not yet been spontaneously access*ed* before the presentation of the orange letter.

2.2.2.2.EEG results

We analysed the EEG data in both a stimulus-locked and an action-locked manner. For stimulus-locked analysis, we compared potentials preceding orange letters in epochs where participants reported an intention to act (by pressing a key in the following 2 seconds) with epochs where they did not. Second, we further fitted a logistic regression to test whether the probability of reporting awareness could be predicted from the EEG amplitude at the time of probing. For action-locked analysis, we compared potentials preceding self-paced actions to the potentials preceding actions that occurred within 2 s of an orange letter (latent awareness), and were therefore related to an intention.

Awareness of intention is sufficient to find an RP in averaged EEG epochs

We compared the activity over premotor and motor areas preceding orange letters with and without a response. In the *a priori* selected region of interest (see *Methods* section), the cluster-based permutation tests revealed a significant difference between the latent awareness and no awareness epochs. There is a large variability in the EEG activity preceding self-paced actions at the single-trial level (e.g. VaezMousavi and Barry 1993). The RP thus reflects the fact that, *on average*, brain activity preceding self-paced actions shows a negative deflection more often than a positive deflection. Therefore, we had a clear *a priori* hypothesis that, on average, RP-like activity should be *more* strongly associated with awareness of intention than with lack of awareness, and we accordingly performed a one-tailed test. Latent awareness epochs indeed showed stronger negativity preceding orange letters than no awareness epochs (p =

0.039). The observed potentials had the form of a partial RP (*Figure 2.2*). A clear EEG periodicity was also visible, due to entrainment of the EEG by the letter stream.

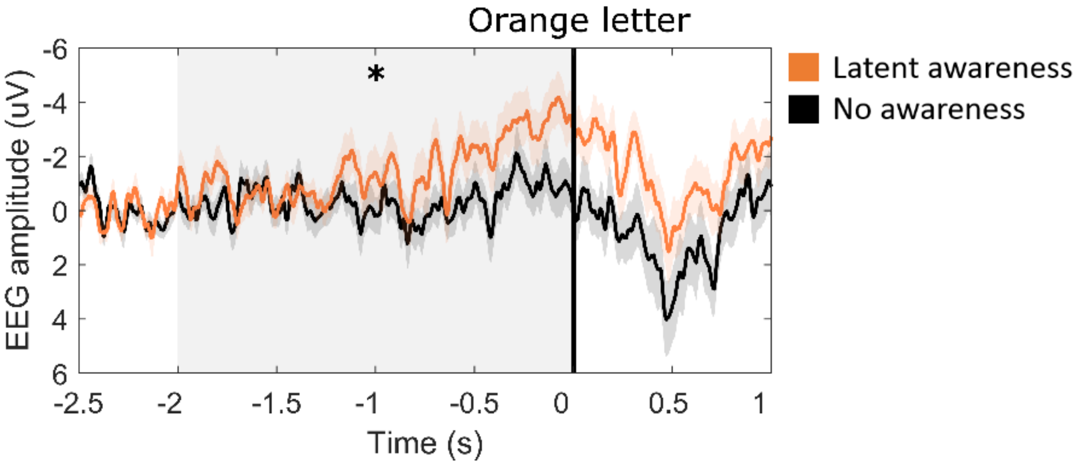

Figure 2.2. EEG activity preceding awareness probes. Grand-averaged EEG amplitude over motor areas (FCZ, CZ, C3) locked to orange letters and sorted by presence (Latent awareness, *orange line*) or absence (No awareness, *black line*) of a keypress during the 2 s following the orange letter. Shaded area indicates cluster analysis period. *, p < 0.05.

We further explored individual differences in this effect by calculating the difference between mean EEG signals averaged over FCZ, CZ and C3 from -2 s to 0 s in latent awareness and no awareness conditions for each participant. We found that this difference did not correlate significantly with the frequency of participants' latent awareness reports ($\rho = 0.18$, $p = 0.45$) or with the average reaction time to the orange letter in latent awareness actions ($\rho = -0.23$, $p = 0.32$), suggesting that the effect was independent of response strategy (*Figure A2*).

The results show that the presence of awareness is sufficient to find an RP in averaged EEG epochs, independently of the individual variability in response strategy.

EEG fluctuations predict awareness at a single-trial level

To study whether the EEG signal at a single-trial level is sufficient to predict awareness reports, we also fitted a logistic mixed model to predict the probability of reporting awareness based on the average signal of the channels of interest (FCZ, CZ and C3) at the time of probing [-0.1 to 0 s relative to orange letter]. We found that the average negativity at the time of probing significantly predicted the probability of reporting awareness (β_1 = -0.004, *SE* = 0.002, *p* = 0.02, see *Figure 2.3.*). That is, the more negative the EEG at the time of probing, the more likely participants were to report awareness (see *Figure A3* showing single subject EEG amplitude distributions at the time of probing).

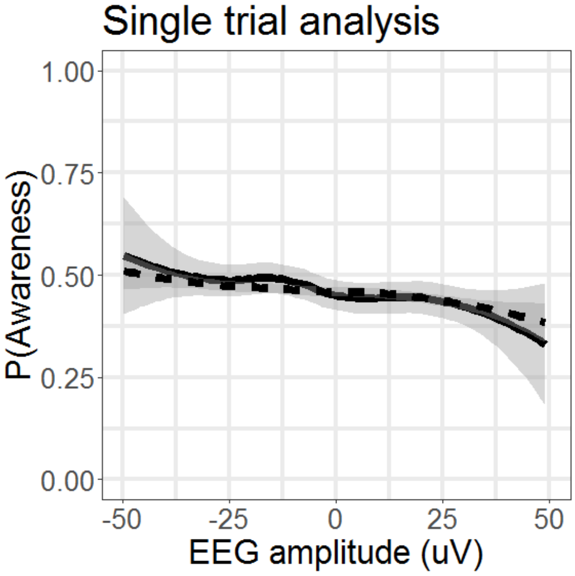

Figure 2.3. Single-trial EEG activity predicts awareness reports. Observed (*solid line*) and predicted (*dashed line*) probability of reporting awareness given the single-trial EEG amplitude at the time of probing [-0.1 0s], averaged over the channels of interest (FCZ, CZ, C3). Participants were more likely to report awareness of an intention the more negative the signal was at the time of probing. Note that there is no particular meaning for P(Awareness) = 0.5, since there is no experimentally controlled stimulus, and no 'correct' response.

Latent awareness actions have smaller readiness potentials

We also compared the RPs preceding self-paced actions to key presses following orange letters that were used to report awareness. We hypothesised that the RP would be *more* fully developed, and thus have higher amplitude, in self-paced actions than in latent awareness actions, because the orange letter would highlight a latent intention, accelerating the action itself, and causing it to occur at a lower level of preparation than normal. We therefore performed a one-tailed test. In the *a priori* selected region of interest (see *Methods*), the cluster-based permutation tests revealed a significant difference between the self-paced and latent awareness conditions ($p = 0.044$). Latent awareness actions showed smaller RP amplitudes than self-paced actions (*Figure 2.4*).

We further explored individual differences in this effect by calculating the difference between mean EEG signals averaged over FCZ, CZ and C3 from -2 s to 0 s in self-paced and latent awareness conditions for each participant.

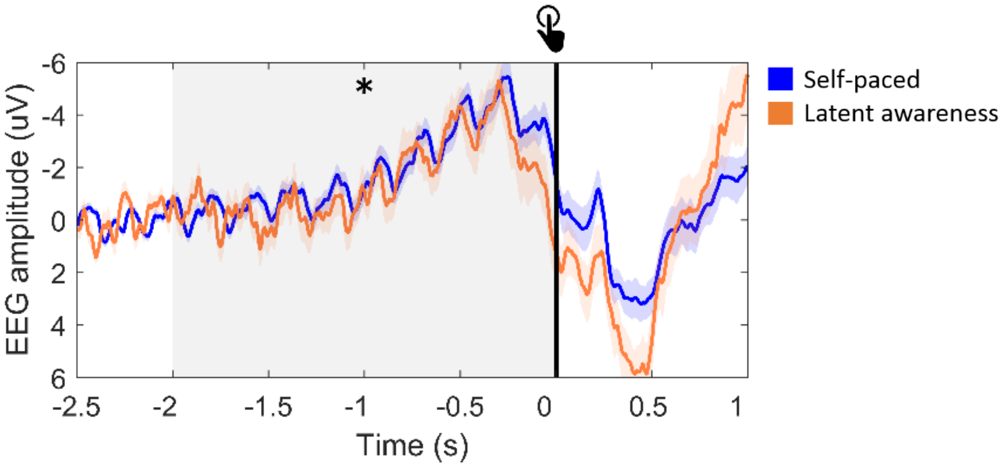

Figure 2.4. The RP in self-paced and latent awareness actions. Grand-averaged EEG amplitude over motor areas (FCZ, CZ, C3) preceding self-paced key presses (Self-paced, blue line) and key presses after an orange letter (Latent awareness, orange line). Shaded area indicates cluster analysis period. *, $p < 0.05$. Note: EEG data were locked to the letter presented immediately after the keypress, not to the keypress itself.

We found a positive correlation between this difference and the frequency of latent awareness reports (ρ = 0.63, p = 0.003), indicating that the amplitude difference between self-paced and latent awareness actions was larger in participants who more frequently reported latent awareness (*Figure A4).* We further observed a significant negative correlation between the magnitude of the effect and the average reaction time (ρ = -0.47, p = 0.04), showing that participants who reported latent awareness more often also responded faster to the orange letter. Finally, the average difference in this action-locked analysis did not correlate with the average difference between conditions found in the stimulus-locked analysis (ρ = -0.28, p = 0.24).

Intention precedes attention

We investigated the relation between motor preparation, dynamic visual attention to the letter stream, and awareness of intention, using CCA coefficients from occipital electrodes (POZ, O1 and O2)

The CCA coefficients increased prior to action (see *Figure 2.5*). In order to estimate the onset of the increase, we calculated the optimal change point for the CCA coefficients for each individual participant (see *Methods*). We estimated the onset of the RP using the same algorithm.

In self-paced actions (*Figure 2.5A*), the estimated onset of the RP was approximately 1s before action (*M:* - 0.96 s, *SD* = 0.36 s), whereas the onset of the increase in visual attention started later (*M* = - 0.34 s, *SD* = 0.34 s). The onset difference was significant ($t_{(17)}$ = 7.00, p < 0.001).

Latent awareness actions showed the same pattern (*Figure 2.5B*). The RP onset was also estimated to happen around 1 s before action (*M* = - 1.03 s, *SD* = 0.66 s), whereas the increase in visual attention started later (M = - 0.48 s, SD = 0.42 s). The difference was also found to be significant ($t_{(17)} = 2.98$, $p = 0.008$).

Interestingly, the onset of visual attention increase and the reported time of awareness were not significantly different in either the self-paced ($t_{(17)} = 0.94$, $p = 0.36$) nor the latent awareness ($t_{(17)} = 0.18$, $p = 0.86$) actions. This is consistent with the possibility that conventional mental chronometry measures of awareness are influenced by cross-modal synchronisation to the chronometer display itself. However, a null result should not be taken as evidence of absence of a difference. For these results, we therefore also calculated JZS Bayes factors (Rouder et al. 2009). These showed modest evidence that increased attention occurred at the same time as reported awareness, in both self-paced ($BF_{01} = 2.79$) and latent awareness trials ($BF_{01} = 4.05$).

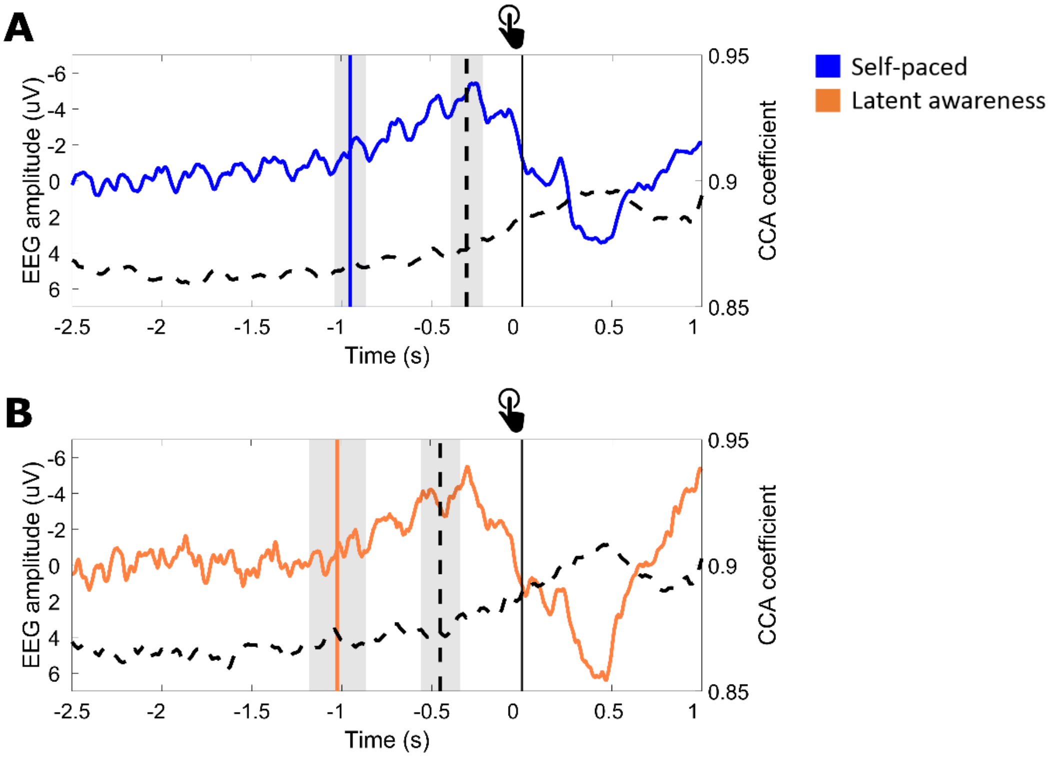

Figure 2.5. Visual attention modulations occur after motor preparation initiation. Grand-averaged EEG amplitude over motor areas (FCZ, CZ, C3, continuous line) and grand-averaged canonical correlation analysis (CCA) coefficients over occipital electrodes (POZ, O1, O2, dashed line) preceding self-paced **(A)** and latent awareness **(B)** key presses. Vertical lines and shaded areas indicate the mean ± SEM estimates of slope onsets for RP and CCA.

2.2.3. Discussion Experiment 1

This experiment suggests that the RP is specifically associated with awareness of intention. Participants' response to an orange letter was contingent on whether they felt they were already preparing their next movement or not. Crucially, these responses were not speeded forced-choice reactions to stimuli (Schurger et al. 2012), but rather reports regarding the subjective experience of intention at the time of the orange letter. We found a stronger RP-like increasing negativity preceding orange letters that interrupted a reported intention to act (and thus elicited a response) compared to that before orange letters that were not associated with an intention to act (*Figure 2.2*). This result has two main implications.

First, it provides direct evidence that the RP is linked to awareness of intention prior to voluntary action. When awareness of intention was present, a strong RP-like signal was present. Reportable conscious intention is a *sufficient* condition for an RP-like EEG pattern, on average.

Second, our analysis using mixed-models (*Figure 2.3*) allows us to make one additional claim: the probability of reporting awareness is linked to the underlying RP-like fluctuations. Our random probing design allowed us to study the probability of reporting awareness when probed in various brain states, and we found that participants were more likely to respond to an orange letter probe when EEG from our frontal motor cluster was more negative. It is important to note that RP is a negative-going signal on average, yet positive-going EEG signals are often seen prior to action on individual trials (see *Figure A5*). The averaged RP thus reflects the fact that brain activity preceding self-paced actions is more often negative than positive. Our regression analysis shows that awareness of intention follows the same pattern: participants are more likely to report awareness the more negative the EEG signal is, although awareness of intention can occur given a wide range of EEG states (*Figure A3*). In sum, our results show that participant's reports of intention awareness are sensitive to EEG fluctuations. Stronger negative-going EEG makes awareness reports more likely, and this is visible on average in the shape of an RP-like signal preceding the orange letters that were followed by a latent awareness report.

We further show that latent awareness actions had smaller RP amplitudes than regular self-paced actions (*Figure 2.4*). We initially hypothesised that action on latent awareness actions might be executed at earlier stages than normal self-paced actions. Specifically, the orange letter probe might boost participants' latent awareness of their motor preparation, pushing it into conscious experience at an unusually early stage. Voluntary action might now occur with a significantly lower degree of preparatory neural activity, implying a reduced RP amplitude for responses to orange letters, compared to self-paced actions.

However, our analysis of individual differences suggests an alternative interpretation. The frequency of latent awareness reports and the average reaction time to the orange letter probes correlated with the magnitude of the action-locked effect (i.e. the difference in the RP amplitude between self-paced actions and latent awareness ones, see *Figure A4*) but not with the stimulus-locked effect (i.e. the difference in EEG signal between latent awareness and no awareness epochs, see *Figure A2*). Participants who frequently reported awareness were generally faster at responding to the orange letter probe and showed a larger difference in the RP amplitude between self-paced and latent awareness trials. While the stimulus-locked effect was independent of the frequency and speed of their responses to orange probes (*Figure A2*), action-locked effects may be related to individual differences in response strategy rather than a systematic influence of neural activity at the moment when the probe interrupted motor preparation. Those participants who responded frequently and rapidly to the orange letter probes may indeed have executed action at an earlier stage, as initially hypothesized. They would therefore show smaller RP amplitudes in latent awareness than in self-paced actions. Conversely, participants who responded infrequently and slowly seemed to reach the same motor preparation level for both latent awareness actions and normal self-paced actions. Further studies are required to investigate the relationship between conscious accessibility of motor preparation signals and the potential variability of the threshold for action execution.

Our periodic visual letter stimulation additionally allowed us to study changes in visual attention to the letter stream during the motor preparation period preceding an action. The CCA traces show a similar pattern for self-paced and latent awareness actions (see *Figure 2.5*), with values increasing gradually towards the time of the action, and peaking just after action execution. Interestingly, the onset of the RP occurred significantly earlier than the increase in CCA scores, suggesting that motor preparation preceded increases in visual attention – at least those detectable by means of CCA. Our results cannot rule out the possibility that *some* modulation of visual attention occurred before the CCA increase. However, our results do show that at least *some* attentional modulation, as measured by CCA increase, coincides in time with the offline reports of awareness of intention. This supports the idea that offline, recall-based estimates of the time of intention awareness may be influenced by visual attention processes (Banks and Pockett 2007). Because the vast majority of awareness estimates (>90%) in latent awareness actions indicate times *after* the orange letter was presented, we infer that people can access their intentions, and use them to guide behaviour, *before* visual attention is modulated. We suggest that participants registered the orange letter, detected their awareness of intention, and then noted the letter that was displayed at that point. They did not appear to postdict their awareness of intention to some other time-marker. The orange letter was presented, on average, 1 s before action, but visual attention increased approximately 0.5 s before action execution (see behavioural results). In other words, visual attention would follow access to intention, rather than precede it. This result appears to rule out the possibility that conscious intentions are mere artefacts of the cross-modal distribution of attention demanded by the Libet task. Hence, online, recall-independent measures of intention awareness, as in the present study and in previous ones (Matsuhashi and Hallett 2008), might provide less attention-dependent estimates of awareness of intention.

Our CCA data seem to provide evidence against concerns that the RP might be an artefact of visual attention. Our data showed a clear attentional modulation during motor preparation, but this began *after* the initiation of motor preparation. We assume that the same result would hold in the classic Libet paradigm, and in other paradigms where people use an external visual chronometer as a cross-modal index of an internal event. Our results therefore appear to rule out the possibility that some modulation of visual attention causes initiation of voluntary actions.

Summary of results

The results of this first experiment suggest that 1) awareness of motor intention as reported by means of action execution is sufficient to find an RP on average, 2) the awareness of motor intention reports are driven by RP-like fluctuations on a single trial level, and 3) offline intention reports correlate with increases in visual attention, but the online detection of the RP precedes those modulations. In the next experiment, we used a very similar paradigm with the converse motor intention report to test whether reporting a motor intention by inhibiting an action would produce the same results.

2.3. Experiment 2 – Contingent Veto

In this experiment, we aimed to replicate the findings from Experiment 1 by means of a different report method. Here, participants performed a self-paced movement task while staring at a letter stream identical to that in Experiment 1 and they were occasionally presented a visual cue as an intention probe. However, participants were now instructed to *not* respond to the cue if they felt they were about to move when it was presented, or to ignore it otherwise. That is, they now reported awareness of an intention by vetoing the incoming movement. The awareness report method employed in this experiment was, thus, a ‘contingent veto’. After a 3 s delay, participants were required to provide an additional explicit report of intention. This delayed report was necessary: if participants had not provided an explicit report, there would be no overt marker to distinguish trials where participants actively vetoed an action from trials where they happened not to move.

2.3.1. Methods

Participants

The experiment was designed based on the same power calculation as Experiment 1.

Twenty-five subjects were initially recruited from the ICN Subject Database. All participants were healthy, right-handed, young adults with normal or corrected to normal vision, no known disabilities and no history of neurological or psychological disorder. The study was approved by the UCL Research Ethics Committee and written informed consent was obtained from all participants before beginning the experiment. Subjects were paid £7.50 per hour.

All participants were invited to a single EEG session. Two participants were excluded because they did not understand the instructions. Two participants were excluded due to excessive EEG noise across all EEG channels. Finally, two further participants did not show a RP in the ‘selection’ subset of self-paced trials (see below) and were excluded from the analysis.

Eventually, 19 participants (14 female) were included in the study (M_{age} = 22.5, *SD* = 2.66; range: 18-29 years).

Stimuli and experimental design

***Procedure*:** Participants sat in a quiet room and were presented visual stimuli on a computer monitor. The instructions for the task were first displayed on the computer screen and then verbally repeated by the experimenter before the beginning of the experiment.

Participants performed a simple motor task. The task was programmed in Matlab R2014b and the Psychophysics Toolbox v3 (Brainard 1997). Subjects responded by pressing the space bar key of a standard computer keyboard with the right hand.

The task was divided into 8 blocks of 10 minutes each. Participants could take breaks for as long as they wished between blocks. *Figure 2. 6* schematically describes the structure of the task.

Stimuli: A continuous letter stream was presented onscreen. All letters were lower-case consonants and the same letter was never repeated within 8-letter sequences. Letters were either black or orange and were presented on a light grey background without any blank interval between consecutive letters. All letters were presented for 250 ms. Orange letters appeared at a random time 6 s to 24 s after the preceding orange letter. If participants did not press a key during the 3 s following an orange letter, a tone was played. See below for further details.

Task

Self-paced condition: Participants were instructed to fixate the letter stream and to make self-paced keypresses whenever they felt like it. They were specifically told not to pre-decide to respond to a specific letter and to try to be as spontaneous as possible (e.g. not to count the seconds between keypresses). This precaution was taken to ensure that actions were truly self-paced, and generated anew each time, rather than responses to some internal rule.

In response to the orange letters, participants were told to evaluate whether they felt they were already preparing their next movement.

Veto condition: If they felt they were preparing their next movement, they were instructed *not* to move (i.e. to *veto* their impending action) and to wait for 3 seconds until a tone played. They were instructed to respond to the tone to indicate that they vetoed an action (i.e. that they were aware of a motor intention at the time the preceding orange letter was presented).

Ignored/Post-orange action condition: If participants did not feel they were preparing to move when an orange letter was presented, they were instructed to ignore it and continue with the task. They were explicitly told they were not required to wait for the tone and were free to press a key whenever they felt like it. We intentionally made this explicit to prevent the orange letter from being interpreted as a mere "No-Go" instruction. If participants happened to press a key within the 3 s following the orange letter, the tone was not played. Note that we will refer to these trials as ignored when talking about them in the context of stimulus-locked data (i.e. the orange letter was ignored because participants did not feel they were about to move) and post-orange when talking about action-locked data.

No Veto condition: If participants did not feel they were preparing to move and did not happen to press a key during the 3 seconds elapsed between the orange letter and the tone, they were instructed not to respond to the tone to indicate that their non-response to the orange letter did not represent a veto, but merely the absence of an intention.

Both the ‘ignored/post-orange’ and the ‘no-veto’ conditions represented instances where participants did not feel they were preparing to move. In contrast, the Veto condition represented instances where participants were aware of an intention to move at the time of the probe. Thus, for the main contrast of interest, the ‘ignored/post-orange’ and the ‘no-veto’ conditions were averaged (‘No Awareness’) and compared to veto (or ‘Awareness’) condition.

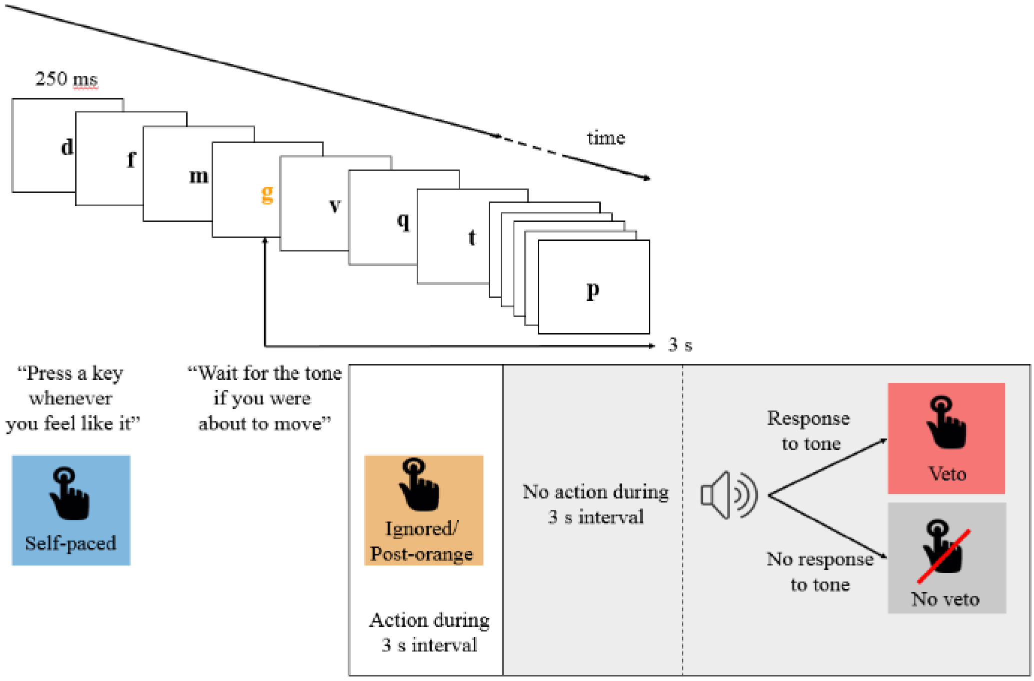

Figure 2. 6. Experiment 2 design. Participants were asked to press a key with their right hand whenever they felt like it (Self-paced). In response to the orange letters, they were asked to inhibit any movement and respond to the tone, which was presented 3 s later, to indicate they vetoed an action (Veto). If they did not feel they were about to move they were told to ignore the letter and were not required to wait for the tone. If they happened to press a key during the 3 s following an orange letter, the tone would not be presented (Ignored/Post-orange). If they did not happen to move during that interval, they were instructed not to respond to the tone to indicate that they did not veto an action when the last orange letter was presented (No veto).

EEG recording

The EEG signal was recorded using the same hardware and parameters as in Experiment 1.

Behavioural data analysis

Statistical analysis: The descriptive and statistical analyses of behavioural data were performed using Statistical Package for the Social Sciences, version 22 (SPSS Inc, Chicago, IL, USA). Greenhouse-Geisser correction was used when the sphericity assumption was violated and Bonferroni correction was applied to post-hoc multiple comparisons. Behavioural analyses were performed before the trial exclusion criteria employed for EEG analysis.

EEG analysis

Preprocessing*:* All pre-processing was performed with the same parameters as in Experiment 1. For EEG analysis, we excluded all trials where a keypress had been executed within 3s before the event of interest (i.e. orange letter or action execution) to avoid overlapping potentials. After this selection and for the included participants only (see below), the number of trials in each of the orange-locked conditions was on average 45.57 (*SD* = 19.67) for vetoed actions, 37.10 for non-vetoed actions (*SD* = 14.09) and 27.42 for ignored orange letters (*SD* = 26.89; the same trials were included for action-locked analysis).

RP analysis for participant exclusion*:* As in Experiment 1, in order to exclude participants who did not show RPs, 20% of self-paced trials were selected randomly for each participant (*M* = 65.17, *SD* = 27.55) and visually inspected. If no RP was visible, the participant was excluded. Two participants were excluded on these grounds. Only the remaining 80% of self-paced trials were used for statistical inference in the final participant sample (*M* = 303.63, *SD* = 101.72).

***Statistical analysis**:* Similar to Experiment 1, statistical tests were run using Fieldtrip toolbox (Oostenveld et al. 2011) cluster-based permutation analysis. Three electrodes over the motor cortex (FCZ, CZ and C3) were preselected for analysis. There two main contrasts of interest. First, the stimulus-locked Veto versus No awareness (i.e. Ignored + No veto) trials. Second, the amplitudes of the readiness potentials preceding self-paced versus post-orange (i.e. post-orange actions, where orange letters were ignored) actions.

The cluster-based tests were performed on the selected electrode using the following parameters: one-tailed dependent samples t-test, time interval = [-2 0 s relative to the event of interest], minimum number of neighbouring electrodes required = 2, number of draws from the permutation distribution = 10000.

2.3.2. Results

2.3.2.1.Behavioural results

Responses to orange letters

First, we analysed the responses to orange letters to test whether participants showed some preferential response to the orange letters (veto, no veto or ignored). The percentage of trials of each type of was not significantly different ($F_{(2,36)} = 1.07$, $p = 0.353$). On average, participants reported vetoing an action on 97 trials (*SD* = 66.47), not vetoing it on 75.10 (*SD* = 33.03) and ignored the orange letter on 85.21 trials (*SD* = 39.74). However, pooling the no-veto and ignored conditions revealed that participants reported *not* being aware of prior intention (No awareness (No-Veto + Ignored): *M* = 160.31, *SD* = 25.33) significantly more often ($t_{(18)} = 8.97$, $p < 0.001$)

than they reported awareness of an intention (Veto). Thus, participants reported vetoing an action only after 29.31% (*SD* = 10.25%) of orange letters, whereas they reported no intention to move after most probes (Ignored: *M* = 37.02 %, *SD* = 16.16%; No-Veto: *M* = 33.65 %, *SD* = 15.79%).

Before the orange letters

Second, we studied the distribution of key presses preceding orange letters in each one of the conditions. We calculated the percentage of trials in which there had been a self-paced key press within the 3 s preceding the presentation of the orange letter. A repeated measures ANOVA revealed there was a significant difference in this proportion ($F_{(2,36)} = 31.78$, $p < 0.001$, $\eta_p^2 = 0.64$). Post-hoc paired-samples t-tests revealed that this difference was driven by the No veto condition, in which there were self-paced key presses before the orange letter in a higher percentage of trials (*M* = 66.20%, *SD* = 18.38%) compared to veto (*M* = 37.39%, *SD* = 16.40%) and ignored (*M* = 35.67%, *SD* = 23.05%) conditions (all $p < 0.001$). This indicates that when orange letters were presented shortly after a self-paced keypress, participants were more likely to not move at all (No-Veto) than to than to indicate vetoing a movement (Veto) or perform one in a self-paced manner (Ignored condition) (see *Figure A8*).

For EEG analysis, we excluded all trials where a self-paced keypress was present within 3 s before the orange letter probe to avoid overlapping potentials. After this exclusion, we checked whether the time at which the last self-paced keypress was executed differed between conditions. The last self-paced keypress before presentation of an orange probe happened, on average, 7.69 s (*SD* = 3.46) before Ignored trials, 8.07 s (*SD* = 3.14) before No-Veto trials, and 6.89 s (*SD* = 1.88) before probes where participants reported a prior intention. These times were not significantly different ($F_{(2,36)}$= 1.21, p = 0.309, η_p^2 = 0.063).

2.3.2.2.EEG results

No significant differences in potentials preceding orange letters

We compared the activity over premotor and motor areas preceding orange letters. We had hypothesized that, when participants reported vetoing an action, an RP-like negativity should precede the orange letter probes. In the *a priori* selected region of interest, the cluster-based permutation tests found no significant clusters between the Awareness and No-awareness conditions (see *Figure 2.* **7**), suggesting that the EEG activity preceding orange letters was not systematically different in various awareness conditions.

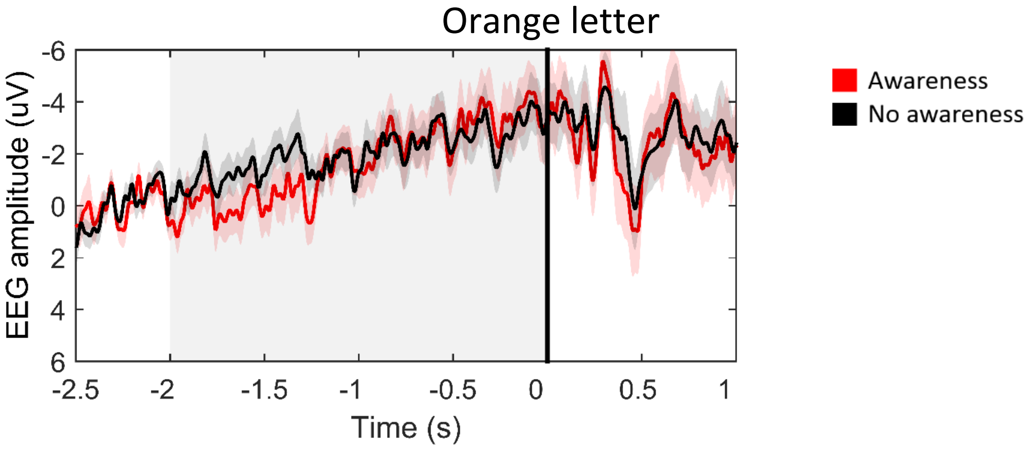

Figure 2. 7. EEG activity preceding awareness probes. Grand-averaged EEG amplitude over motor areas (FCZ, CZ, C3) locked to orange letters and sorted by response. Red line indicates participants responded to a tone presented after 3 s after the orange letter to indicate they vetoed an action (Awareness). Black line indicates trials where participants did not report an intention to move at the time of orange letter presentation (No awareness). This conditions includes trials where participants either pressed a key within 3 s from the orange letter presentation (i.e. ignored the cue) or did not move during that interval but they also did not respond to the tone, indicating they did not inhibit an action at the time of the orange letter. Shaded area indicates cluster analysis period.

No significant differences in action-locked data

For consistency with Experiment 1, we also compared the RPs preceding self-paced actions to the ones preceding orange probes (key presses following orange letters). In the *a priori* selected region of interest, the cluster-based permutation found no significant difference between the self-paced and post-orange conditions (*Figure 2. 8*).

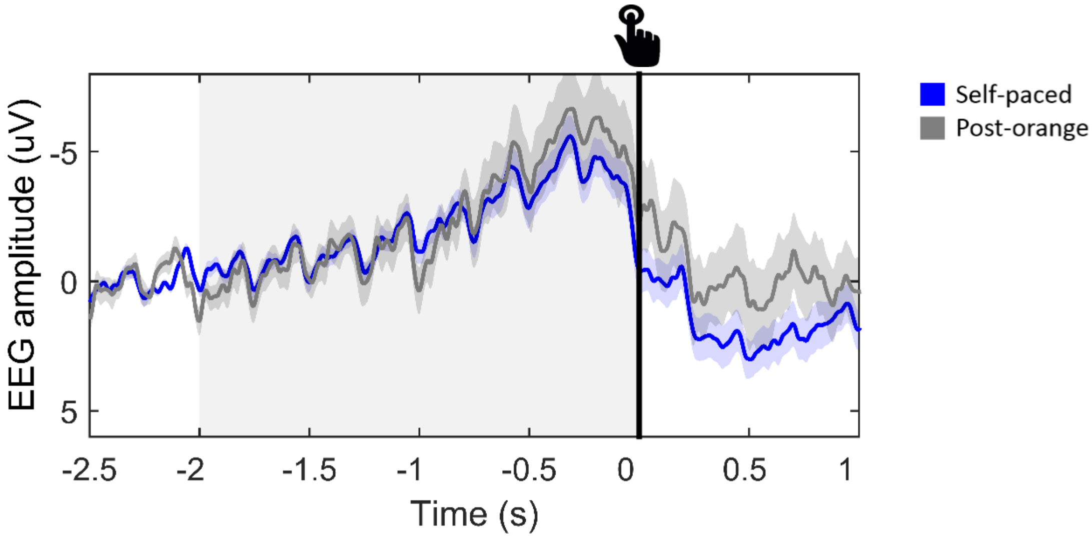

Figure 2. 8. The RP in self-paced and post-orange actions. Grand-averaged EEG amplitude over motor areas (FCZ, CZ, C3) preceding self-paced key presses (blue) and key presses within 3 s after the presentation of an orange letter (orange). Shaded area indicates cluster analysis period.

2.2.4. Discussion Experiment 2

The ‘Contingent Veto’ experiment did not replicate the stimulus-locked findings of our first ‘Contingent Action’ experiment. We expected that the response vs. no response difference in Experiment 1 would be replicated in the analogous veto vs. no awareness conditions. However, there are both methodologic and cognitive reasons that might explain the results.

A potential methodological limitation in Experiment 2 is that explicit reports were provided only after a 3 s delay. If participants felt they were preparing to move when an orange letter was presented, they had to inhibit an action and wait for a tone to indicate they vetoed an action. Instead, in Experiment 1 they had to respond immediately if they felt they were about to move. This ‘delayed’ report of awareness in Experiment 2 could have introduced some noise in participants’ reports of intention. In that sense, the contingent veto reports used here were *less* immediate and online than the contingent action reports used in Experiment 1. That is, they were more akin to *judgements* than to immediate reports of an instantaneous *feeling* of intention.

In fact, the task in Experiment 2 was comparatively more complex than the task in Experiment 1. If participants felt they were preparing their next movement at the time the orange letter appeared, participants had to 1) inhibit the movement and 2) report that they vetoed a movement after the beep, 3 seconds later. Such a two-step procedure was necessary for the intended stimulus-locked analysis. If participants had not provided the explicit veto reports after the beep, we would have had no overt indicator to determine which of the orange letters elicited an experience of “being about to move” and hence a lack of response as per the “veto” instruction, and which of the orange letters people just happened to ignore (a limitation that previous studies encountered, see Matsuhashi and Hallett, 2008). However, as a result of our design, the explicit veto reports we used to investigate pre-stimulus activity may not have been as immediate as they were in Experiment 1. While the *feeling* of being about to move presumably drove the decision to move or not move in response to each orange letter, the delayed *judgement* of awareness at the time of the beep might have been influenced by a number of factors. Looking at the RT distribution (*Figure A8*), one can observe a dip in the number of actions performed after an orange letter was presented. This suggests that participants might have detected instances where motor preparation was present (i.e. they did feel they were about to move when prompted) and subsequently inhibited an action *but* failed to report it afterwards.

A potential cognitive explanation for such a dissociation relates to the factors that influence intention judgements. Some theories about motor intention suggest that judgements of intention are influenced both by prospective motor preparation processes and retrospective information related to action execution (Douglas et al. 2015). In particular, it has been suggested that intentions may be typically reconstructed only after an action is executed (Banks and Isham 2009, Wegner 2002). A corollary of this theory is that the probability of reporting a prior intention is decreased in the absence of an action. Another potential explanation involves memory. It may be the case that rather than an active reconstruction based on the fact that they did not act, participants feeling of intention decays over time and is sometimes no longer accessible after the 3 s delay, so participants fail to recall that they indeed intended to move at the time of the probe. A final possibility is that the reconstructive and memory effects interact. People may be less able to post-hoc access (and hence less likely to report) motor preparation states if no action is executed. In other words, if participants execute an action, such execution might consolidate the feeling of intention for a longer time and make it accessible for offline reports.

All of these explanations would account for the fact that participants provided a comparatively low percentage of awareness reports in this experiment. While in Experiment 1 participants reported an intention to move after roughly 50% of the orange letter probes, only 30% of orange letters in this task were followed by veto reports. Further, the stimulus-locked data *Figure 2.* **7**) suggest that although an RP-like signal was often present, participants failed to detect it as a motor preparation state (or a prior intention to move) in many occasions. When participants reported having vetoed an action, an RP-like signal was present. This is compatible with the findings in Experiment 1. Yet, a similar signal was present when participants did *not* report a motor intention. Together with the low frequency of awareness reports, this suggests that participants' reports may have included a high proportion of false negatives.

A methodological corollary is that a contingent veto task may be intrinsically under-powered to detect motor preparation by using awareness reports as a proxy.

2.4. Experiment 3 – Readiness reports

Whether conscious experiences have a binary nature (i.e. either one is conscious or one is not) or rather can be gradual is an open question that has been investigated in some areas of consciousness research (Overgaard et al. 2006) but remains largely unexplored in the field of motor intention. In the previous experiments, we used dichotomous report methods, were participants reported either being aware of an intention or not. Here, we aimed to investigate whether the amplitude of averaged RPs would correlate with subjective graded reports of "readiness". In this experiment, participants performed a self-paced movement task identical to that in the previous experiments. Occasionally, they were presented a visual cue as an intention probe. Participants were then shown a Likert scale and were asked to indicate how ready they felt to move at the time of the probe, from 0 ("not ready at all") to 7 ("about to move").

2.4.1. Methods

Participants

Twenty-one subjects were initially recruited from the Institute of Cognitive Neuroscience Subject Database. All participants were healthy, right-handed, young adults with normal or corrected to normal vision, no known disabilities and no history of neurological or psychological disorder. The study was approved by the UCL

Research Ethics Committee and written informed consent was obtained from all participants before beginning the experiment. Subjects were paid £7.50 per hour.

All participants were invited to a single EEG session. Visual inspection of the data revealed that four participants exhibited excessive noise throughout the whole EEG time course, across all channels, due to technical problems during the recording session, and were excluded from further analysis.

Seventeen participants (13 female) were therefore included in the final dataset (M_{age} = 23.76, *SD* = 3.45).

Stimuli and experimental design

***Procedure*:** Participants sat in a quiet room and viewed the stimuli on a computer monitor at 50 cm distance. The instructions for the task were first displayed on the computer screen and then verbally repeated by the experimenter before the beginning of the experiment.

Participants performed a simple motor task. The task was programmed in Matlab R2014b and Psychophysics Toolbox v3 (Brainard 1997, Kleiner et al. 2007). Subjects made actions by pressing the space bar key of a standard computer keyboard with the right index finger, and reported “readiness” judgements by pressing the numbers on the keyboard numeric pad.

***Stimuli*:** The task was divided into 8 blocks of 36 letter streams each. The letter streams were a stream of random lower-case consonants (b, c, d, f, g, h, j, k, l, m, n, p, q, r, s, t, v, w, x, z). Any letter repetitions were separated by at least 7 other letters. Letters were either black or orange and were presented on a light grey background without any blank interval. All letters were presented for 250 ms. Each stream contained one single orange letter, which appeared at a random time between 3 s to 20 s after stream onset. The duration of each letter stream was variable and dependent on pseudorandomised constraints.

Task

Self-paced task: Participants were instructed to fixate the letter stream and to make self-paced keypresses whenever they felt like it. They were specifically told to not pre-decide to respond to a specific letter and to try to be as spontaneous as possible (e.g. not to make keypresses after a fixed interval).

Readiness rating: Occasionally, one letter was presented in orange rather than black (see *Stimuli* above). Participants were not given any specific instructions regarding whether to move or not to move in response to the orange letter. After the orange letter, the letter stream stopped and a Likert scale with numbers 0 to 7 was presented on screen with the question “How ready were you to move?”. Participants had to indicate how ready they felt they were to execute the next self-paced movement by pressing the corresponding number keys on the keyboard. There was no time limit for the response rating. There was a 2 s inter-trial interval, after which the next letter stream started. *Figure 2. 9* schematically describes the structure of the task.

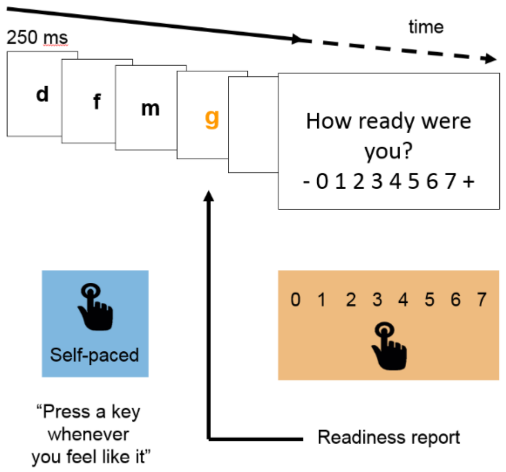

Figure 2. 9. Experiment 3 design. Participants were asked to press a key with their right hand whenever they felt like it (Self-paced). Occasionally, an orange letter would appear on screen and the letter stream would be interrupted. Afterwards, they would be asked to rate how ready they were to move when the orange letter appeared (Readiness rating) on a scale from 0 (not ready at all) to 7 (about to move).

EEG recording

EEG was recorded from 26 scalp sites (FZ, FCZ, CZ, CPZ, PZ, POZ, FC1, FC2, C1, C2, CP1, CP2, F3, F4, F7, F8, C3, C4, CP5, CP6, FC5, FC6, P3, P4, O1, O2) using active electrodes (g.LADYbird) fixed to an EEG cap (g.GAMMAcap) according to the extended international 10/20 system. EEG data were acquired using a g.GAMMAbox and g.USBamp with a sampling frequency of 256 Hz. Signal was recorded using g.Recorder (G.tec, medical engineering GmbH, Austria). All electrodes were online referenced to the right ear lobe. Vertical and horizontal electroocular activity was recorded from electrodes above and below the right eye and on the outer canthi of both eyes.

Data analysis

Behavioural analysis

All participants provided 288 readiness ratings (8 blocks, 36 trials each). For each rating, we calculated the time elapsed since the last self-paced keypress and the orange letter probe.

Participants were not given any instructions regarding whether to move or not after the orange letter was presented and before the rating. In some trials where people moved, it may have been the case that movements were mere reactions to the orange letter, or that the orange letter happened to appear onscreen at the time when they were about to execute a self-paced movement. Regardless of the drivers of post-probe actions, trials where participants moved differ substantially from those where they did not move. Thus, we excluded all those trials in which they moved after the probe to remove the possibility that retrospective reconstruction and reaction time influenced the judgement of intention in those trials.

EEG analysis

***Preprocessing**:* EEG data were processed using Matlab R2014b (MathWorks), Matlab R2017b (MathWorks) and EEGLAB version 13.5.4b (Delorme and Makeig 2004).

First, scalp and eye electrodes were re-referenced to the average of two mastoid electrodes. Continuous EEG and EOG data were filtered with a 0.01 Hz high-pass filter. Then, data were downsampled to 200 Hz and filtered with a 30 Hz low-pass filter. These filters were applied off-line using a 5th order Butterworth filter with zero phase shift.

Second, an independent component analysis (ICA) was computed on the continuous data using the EEGLAB *runica* algorithm. Vertical eye movement components were visually identified and removed from the signal.

Next, EEG signals were locked to either orange letters. Epochs started 1.5 s before the event and finished 1 s after it. Baseline correction was performed using the 500 ms at the beginning of the epoch [-1.5 s to -1 s relative to event] for RP analysis. The shorter 1 s epoch for RP analysis was chosen to maximise the number of available trials in all conditions. Finally, artefact rejection was performed by removing all epochs with >120µV fluctuations from baseline in the preselected channel (CZ). Epochs in which there was a key press in the [-2 to 0 s] interval preceding the event of interest (i.e. orange letter or action) were rejected to prevent overlapping evoked-potentials. After these rejection procedures, the number of epochs analysed was on average 128 (*SD* = 37.57).

For each orange letter probe, we extracted the maximum amplitude at CZ in the last 100 ms before the probe. This was used as a measure of motor preparation at the time of probing.

With the selected subset of trials, we analysed the data with linear mixed effects ordinal regressions using the MCMCglmm function the R package MCMCglmm (Hadfield 2010). The aim of the analysis was to test whether participants' readiness reports were predictable by the amplitude of the RP at the time of probing. We included the time elapsed between the probe and their previous keypress as a predictor, because the time may correlate with the brain oscillations underlying the RP, and also because participants may be basing their ratings on the time rather than on the internal RP brain signal.

2.4.2. Results

Behavioural results

Participants' use of the rating scale varied considerably (see *Figure A9).* Some mostly reported intermediate values of the scale, producing a normal distribution of readiness ratings. Others reported mostly the extremes, resulting in a bimodal distribution.

EEG results

We run an ordinal regression on a single-trial level to test whether RP amplitudes significantly predicted the readiness ratings while controlling for the time elapsed from the last keypress.

We found no significant relationship between the EEG amplitude at Cz preceding the orange probe on the subsequent rating of readiness ($pMCMC = 0.834$). That is, continuous changes in negativity at Cz at the time of probing did not ordinally correlate with the reported readiness. However, we observed that the highest readiness ratings (7: "About to move") showed the strongest RP-like signal, whereas the lowest readiness ratings (0: "Not ready at all") showed no RP-like signal at all (see *Figure 2.* **10**). A post-hoc t-test between the subset of participants who provided both 0 and 7 ratings at some point (N = 14) revealed that this difference was significant ($t_{(1,13)} = 2.28$, $p = 0.039$).

Further, we found that the time of the last keypress had a significant effect on the readiness ratings ($pMCMC < 0.001$). The longer the interval elapsed between the last keypress and the probe, the greater the readiness rating. This suggests that participants were using the time elapsed from their last keypress as a proxy to judge their readiness to act.

No significant interaction between the RP amplitude and the time of the last keypress was found (*pMCMC* = 0.128).

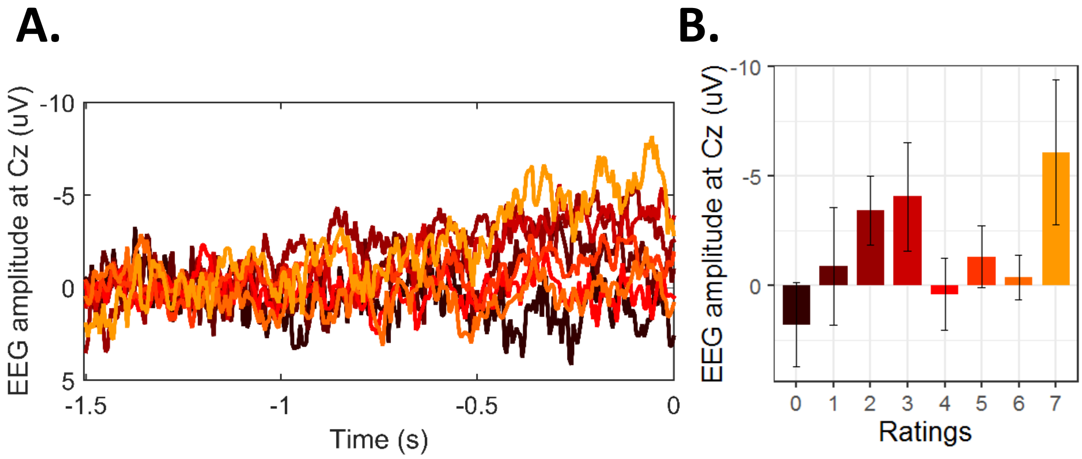

Figure 2. 10. EEG activity preceding readiness reports. **A)** Grand-averaged EEG amplitude over Cz locked to orange letters (time 0) and sorted by readiness rating (0 to 7, dark brown to light yellow). **B)** Grand-averaged amplitude (±SEM) over the last 100 ms before orange letter presentation. No significant correlation was found between the average amplitudes of the EEG signal preceding the probe and the readiness rating given by participants.

2.4.3. Discussion Experiment 3

In this experiment, we used a graded scale rather than a dichotomous report of intention awareness. We found that participants' ratings were heavily influenced by the time elapsed between the last self-paced keypress and the orange probe. This is in line with the time differences observed in Experiment 2, where participants were less likely to report an intention to move when probed shortly after a self-paced keypress. Further, we found no ordinal correlation between the RP amplitude and the readiness ratings. However, a post-hoc analysis revealed a significant difference between the RP amplitudes at the distribution extremes: there was a higher negativity when participants reported the highest readiness level compared to when they reported not being ready to move at all. This suggests that while motor intention may not be graded (or at least not explicitly reportable), the extremes of the scale replicate the findings in

Experiment 1. That is: when participants report being about to move, an RP is present. Instead, no RP is present when they report not being ready at all. It is important to note, however, that since the pairwise comparison supporting this claim was a post-hoc analysis and the task was not designed for this specific comparison, the interpretation of this finding is suggestive, rather than conclusive.

2.5. General discussion

In this chapter, we investigated the relationship between the RP and motor intention awareness by means of different types of awareness reports. In Experiment 1, we used a “contingent action” method to obtain awareness reports. Participants had to execute an action to indicate an intention to move at the time of probing. In Experiment 2, we used a complementary “contingent veto” method, in which participants had to refrain from moving if they were about to move at the time of probing, and then provide a delayed, explicit report of intention at the time of probing. Finally, in Experiment 3 we asked participants to provide graded ratings of readiness. We have discussed the experimental results in detail in the discussion section of each experiment individually. Here, we provide an overview of the implications of the three experiments.

Summary of results

Experiment 1 provided some evidence that participants have some conscious access into motor preparation processes. Participants were able to respond depending on their motor preparation state, as indexed by the RP (*Figure 2.2*). When they indicated an existing motor intention at the time of the probe, an RP preceded such probe. Instead, in Experiment 2 no difference in the potentials preceding the orange letter probes was found when participants reported an intention compared to when they did not (*Figure 2. 7*). In Experiment 3, we found no evidence for gradual insight into motor preparation processes. However, a post-hoc analysis revealed that there was indeed a difference in the RP amplitude between the two extreme ratings (*Figure 2. 10*). When participants reported being about to move, an RP was present. When they reported not being ready at all, no RP was present.

Implications

These results have both cognitive and methodological implications.

Cognitively, these results suggest that 1) the RP is a relevant signal for the *feeling* of motor awareness and can be used to guide action behaviour (Experiment 1), 2) delayed *judgements* of intention reports may not be informed by the RP - its influence on them may be reduced due to retrospective reconstruction or memory effects (Experiment 2) and 3) the RP does not seem to be gradually correlated with the feeling of motor "readiness". (Experiment 3).

Methodologically, the results of our experiments suggest that the shorter the delay between the presence of the RP and the intention report, the stronger the association between the two. In Experiment 1, participants made their report decision as soon as the orange letter appeared. In Experiment 3, the rating was provided immediately after the orange letter, and in the extreme ratings a pattern of results compatible with Experiment 1 was visible. In Experiment 2, instead, participants only explicitly reported their motor intention 3 s after the presentation of the orange letter. If the RP is indeed accessible to consciousness (i.e. some phenomenal *feeling* accompanies its occurrence) and can guide behaviour as Experiment 1 suggests, participants should be able to inhibit movements in response to the orange letter, just as they were able to trigger them. However, because the explicit report was only provided 3 s afterwards, participants' post-hoc *judgements* may fail to correspond to the *feeling* at the time of the orange letter.

Limitations

The three paradigms developed in this study constitute an initial step for the development of novel methods to investigate awareness of motor intention and its relationship to the RP. However, further research is needed to replicate these results and address the limitations of the current paradigms.

Regarding Experiment 1, one might be concerned that the awareness reports are actually self-paced actions that randomly happened around the time of the orange letter presentation. Although we do not think this is the case, as argued in the discussion section of Experiment 1, future experiments could directly control for that possibility by using a “contingent action” report where the action in self-paced movements is different to that in intention reports (e.g. if participants’ self-paced movements were executed with the right hand, the reports could be provided with the left hand).

In Experiment 2, the fact that explicit judgements of intention were only given after a 3 s delay may have reduced participants’ ability to recall motor intention. Thus, future experiments aiming to use ‘contingent veto’ instructions could reduce the response-time window to obtain the inhibition reports closer to the time of probing. However, it may still be the case that reconstruction effects intrinsically reduce the probability of people reporting awareness in ‘veto’ paradigms.

Further, in Experiment 3, the post-hoc test between the extreme ‘not ready at all’ and ‘about to move’ ratings can only be interpreted with caution. The experiment was designed to test gradual relationships between EEG activity and readiness ratings, and therefore it is underpowered for the restricted comparison made here. Further, participants did not execute any action beyond the one providing the rating. Thus, the potential reconstructive effects linked to action inhibition may have also reduced participants sensitivity to their motor preparation states in this experiment. However, at least descriptively, the results suggest that people's feelings of intention may not be gradual, but rather dichotomous. A paradigm with contingent action reports like the ones in Experiment 1 plus confidence ratings following actions might provide a stronger method to investigate potentially gradual relationships between the RP and motor intention at the levels of immediate action control (contingent report) and delayed metacognitive judgement (confidence rating). If accessibility to motor preparation is boosted (or perhaps simply enabled) by action execution, participants should be better able to discriminate their level of readiness after executing an action. While this task would probably introduce additional confounds (e.g. participants may rely on reaction times rather than internal motor preparation states to provide their readiness ratings), it would contribute to the understanding of the factors underlying motor intention judgements.

Finally, a common limitation to all of the paradigms presented in this chapter is that they rely on random probing of awareness and reverse correlation, i.e., analysing brain data contingent on a subjective response. Many empirical approaches to the identification of the neural correlates of consciousness (NCC) follow this pattern: given a set of experimentally controlled manipulations (in our case random probing), participants are required to provide a report of a conscious state, and then the brain states correlating with different reported experiences are then studied. If a conscious state C reliably correlates with a neural state N that is different from other neural states N' correlating with non-conscious perception, then N is identified as a neural correlate of C. But this correlation is not causal. It has only been shown that if $C \rightarrow N$, but this does not imply that if $N \rightarrow C$. Thus, in our case, we can claim that whenever an intention is present, an RP is present (in Experiments 1 and 3), but not that whenever an RP is present, a motor intention is consciously accessible. Yet, the latter type of assertion is relevant for interpreting the causal role of RP in voluntary action initiation and the possibility of conscious motor control. If the presence of an RP is not always followed by movement, or is not always accessible, the implicatiosn for voluntary action control are radically different. In chapter 3, we tackle this specific limitation.

Chapter 3

From the brain to awareness: real-time probing of motor awareness

Synopsis

In a real-time EEG experiment with a 2 by 2 design, we studied the contribution of prospective (presence/absence of motor preparation) and retrospective (presence/absence of an action) components to awareness of intention. Participants performed a self-paced pedal press task and occasionally interrupted by either green Go ("Press the pedal as fast as possible") or red No-Go ("Do not move") probes. Crucially, probes were triggered by either the presence or absence of a RP, using an algorithm trained previously. After the allowed 1.5s response time, participants were asked to report whether they were about to move at the moment the cue appeared by saying "Yes" or "No". Our results show participants more frequently reported conscious intention at the time of the probe in the Go compared to the No-Go condition, suggesting a strong effect of retrospection on awareness of intention. Further, we also found evidence for prospective insight into motor preparation processes. Participants were more likely to report an intention when the cue was triggered by an RP, compared to when no RP was present. This suggests that judgements of awareness of intention in voluntary action involve both prospective and retrospective information.

3.1. Introduction

When we perform spontaneous, voluntary movements, our subjective experience is that of a coherent flow of conscious events, from forming the intention to act to executing the movement. Similarly, neurophysiological data suggest a succession of neural events, with the execution of voluntary movements being preceded by brain signals that indicate motor preparation (Kornhuber and Deecke, 1965; Shibasaki and Hallett 2006). Despite extensive research on the relationship between awareness of intention and motor preparatory processes, it remains unclear when an intention to move becomes accessible to consciousness.

One line of research suggests that we have conscious access to our motor preparation processes and thus become aware of our intention to move *before* action initiation (Libet et al. 1983; Haggard and Eimer 1999; Matsuhashi and Hallett 2008). That is, we know what we are about to do before actually doing it (*prospective hypothesis*). In contrast, others have argued that intentions are postdictively inferred *after* action initiation (Wegner, 2002; Kühn and Brass 2009). In this view, it is the fact of having done something that generates the idea that we intended to do it (*retrospective hypothesis*). An intriguing third possibility is an integrative model that views awareness of intention as a process extended in time during which prospective and retrospective effects might be integrated (Lau et al. 2007; Douglas et al. 2015; Verbaarschot et al. 2016). Comparator models of action control suggest that efferent copies of motor commands and sensory feedback after action execution are integrated, and these efferent copies are postulated to carry different types of information about the movement characteristics, including *when* it will be executed (Wolpert and Kawato 1998). We suggest that a similar integration mechanism is in play for intention awareness that is over and above the prospective and retrospective effects. We refer to this as the *temporal integration hypothesis*.

The current study

In this study, we investigate how motor preparation processes and action execution interact over time to produce the experience of intention that accompanies spontaneous movements. To investigate this, we combined real-time EEG-monitoring of a self-paced task with a classical *Go/No-Go* task. We aimed to interrupt participants with Go/No-Go cues either while they were *preparing* to execute a self-paced movement, or at a time when they were not preparing at all. After each interruption, we asked them about their intention to move. Our real-time method allowed us to investigate in an unbiased way how motor-preparatory states influence conscious intentions.

One candidate motor preparation signal is the RP. As reviewed in chapters 1 and 2, the RP is an increasing negativity over the motor cortex that was first shown to precede self-paced movements by Kornhuber and Deecke (1965) and was later linked to awareness of intention (Libet et al. 1983; Haggard and Eimer 1999; Sirigu et al. 2004; Schlegel et al. 2013). While the RP can be observed from 2s before movement execution, people generally report becoming aware of an impending movement only about 200 ms before movement onset. While concerns have been raised about the validity of the RP as a neural signal uniquely linked to voluntary motor preparation (Schurger et al. 2012; Jo et al. 2013), recent evidence shows that it can be used to predict voluntary movement at a single trial level (Schultze-Kraft et al. 2016). This does not rule out the possibility that RP-like activity can be found in the absence of an action, but it does suggest that the presence of an RP reflects an increased probability of action. The RP remains one of the best candidate signals providing information about motor preparation, and is thus a potentially relevant signal for prospective intention awareness.

In the current study, participants performed a self-paced task during which they were instructed to press a footpedal at any time they wished after trial onset. Occasionally, they would be interrupted by either a green (*Go*) or a red (*No-Go*) cue, instructing them to press the pedal immediately or inhibit any movement, respectively. Importantly, cues were triggered either at a random time when no motor preparation was detected (*RP-*), or as soon as a readiness potential was detected (*RP+*) by a brain-computer interface (BCI) monitoring participants' electroencephalogram (EEG) in real-time. Each trial was randomly assigned a combination of motor preparation state (*RP+/RP-*) and action execution instruction (*Go/No-Go*). When interrupted by a cue, participants were given time to respond accordingly (execute/inhibit a movement) and were additionally asked to verbally report (*Yes/No*) whether they were preparing to move at the time the coloured cue was presented.

Our experimental design allowed us to directly test several predictions that follow from the hypotheses formulated above.

First, if people have conscious access to motor preparatory processes, awareness of intention should be more likely in the presence of a signal indicating motor preparation. Thus, the prospective hypothesis predicts a higher rate of awareness reports when probed in the presence of an RP than when probed in the absence of an RP. Second, the retrospective hypothesis states that it is the action execution itself that produces the experience of intention. Hence, it predicts that participants should be more likely to report an intention if an action is eventually executed (i.e. reports following a green *Go* cue) than when it is not (i.e. reports following a red, *No-Go* cue). Finally, the temporal integration hypothesis proposes that intention judgements depend on the integration of motor representations preceding action and on sensory feedback after action execution. We predict that prospective information about motor preparation based on transient signals such as the RP is only available for integration with (retrospective) sensory feedback for a short time. Thus, we expect that the effect of motor preparation processes on intention judgements should be stronger if a movement is executed shortly after the presence of an RP.

3.2. Methods

Participants

We investigated 23 healthy, naive participants (17 female, $M_{age} =$ 30, SD = 5.2 years). The experiment was approved by the local ethics board and was conducted in accordance with the Declaration of Helsinki. All participants gave their informed oral and written consent. Seven participants were excluded following the participant exclusion procedure (see below), and a final sample of N = 16 was included for the main analysis.

Experimental setup

Participants were seated in a chair facing a computer screen at a distance of approximately 1 m. They were asked to place their hands in their lap and their right foot 2 cm to the right of a 10 cm x 20 cm floor-mounted switch pedal (Marquardt Mechatronik GmbH, Rietheim-Weilheim, Germany). Throughout the experiment, EEG was recorded at 1 kHz with a 64-electrode Ag/AgCl cap (EasyCap, Brain Products GmbH, Gilching, Germany) mounted according to the 10-20 system, online referenced to FCz and re-referenced offline to a common reference. Given that the signal of interest, the readiness potential, is observed predominantly over central electrodes, EEG was recorded from the following 30 electrodes: F3, F1, Fz, F2, F4, FC5, FC3, FC1, FC2, FC4, FC6, C5, C3, C1, Cz, C2, C4, C6, CP5, CP3, CP1, CPz, CP2, CP4, CP6, P3, P1, Pz, P2, P4. In addition to the EEG, the right calf electromyogram (EMG) was recorded using surface Ag/AgCl electrodes in order to obtain the earliest measure of movement onset. The amplified signal (analog filters: 0.1, 250 Hz) was converted to digital (BrainAmp MR Plus and BrainAmp ExG, Brain Products GmbH, Gilching, Germany), saved for offline analysis, and simultaneously processed online by the Berlin Brain-Computer Interface (BBCI, github.com/bbci/bbci_public) Toolbox. The Pythonic Feedback Framework (Venthur et al. 2010) was used to generate visual feedback. Verbal reports in response to the prompting task (see below) were recorded by a microphone that was placed on the table and manually registered trial-by-trial after the experiment. Verbal reports were chosen over movement reports to disentangle the effects used in the main motor task (see below) from the intention reports.

Experimental design and task

The experiment was divided in three stages (*Figure 3.1*). In a preparatory experimental stage I, participants performed a simple self-paced task. The data collected in stage I were used to train a classifier to monitor EEG activity in real-time during stage II. In stage II, the main experiment, participants performed the same self-paced task and a prompting task. In a supplemental stage III, participants performed a cued reaction task.

Stage I: Collection of training data for the classifier

Participants performed a simple self-paced task. The start of a trial was signalled by a traffic light display appearing on the screen with all three coloured lights (green, yellow, red) turned off. Participants were instructed to wait for roughly 2 seconds, after which they could press the pedal at any time. They were asked to avoid preplanning the movement, avoid any obvious rhythm, and to press when they felt the spontaneous urge to move (Kornhuber and Deecke, 1965; Libet et al., 1983). When the pedal was pressed the yellow light was turned on for 1 second, after which the traffic light disappeared and was replaced by a fixation cross. This constituted the end of a trial. The fixation cross remained onscreen for a 3s intertrial period. Participants were asked to fixate and remain relaxed without moving. Each participant performed a total of 100 trials in stage I, with the possibility of taking a break after each 25 trials.

Stage II: Main experiment

Participants performed the same self-paced task indicated above, but additionally they would sometimes be interrupted by either the green (*Go*) or red (*No-Go*) traffic light turning on before they moved. Participants were instructed to press the pedal as fast as possible in response to the green light, and to withhold from moving or to abort any potentially planned pedal press in response to the red light. They were given 1.5 s to respond to this Go/No-Go task. A *Go* trial was considered correct if the pedal was pressed while the green light was on. When participants pressed the pedal, the yellow traffic light turned on for 1 second. A *No-Go* trial was considered correct if the pedal was not pressed while the red light was on. If a trial was not executed correctly, it ended with a fixation cross and was discarded from further analysis. After correct trials, the question “*Were you about to press?”* appeared on screen for 1.5 s.

Participants were instructed to verbally report (*”Yes”/”No”*) whether they were preparing to move at the time the coloured cue appeared on screen.

Each trial in stage II was randomly assigned to one of four conditions defined by a combination of two factors. The first factor was the Action execution instruction (*Go/No-Go*), while the second one was the motor preparation state that would be used to trigger the instruction (*RP+/RP-*). Thus, while the former determined *which* light would be turned on in the trial (green in Go trials and red in *No-Go* trials), the motor preparation state determined *when* the light would be turned on. Note that this assignment of trials was putative rather than absolute, because participants sometimes performed self-paced movements, as in stage I, before they were interrupted by any cue. This was a consequence of how the timing of the cues was realized in *RP+* and *RP-* trials, which we describe in detail below in the section *Real-time BCI predictor*. Stage II had a total duration of 60 minutes, with the possibility to take a break every 15 minutes.

Stage III: Supplementary task

Participants performed a simple, cued reaction task. The start of a trial was signalled by the traffic light display appearing on the screen with all three coloured lights turned off. In each trial, the green light would turn on for 1.5 seconds after a random time, chosen from a uniform distribution between 2 and 5 seconds. Participants were instructed to respond as fast as possible with a pedal press as soon as the green light appeared. When they pressed the pedal, the yellow light was turned on for 1 second after which the traffic light disappeared and was replaced by a fixation cross. This constituted the end of a trial. The fixation cross remained onscreen for a 3s intertrial period. Each participant performed this task for a total time of 8 minutes, with the possibility to take a break after 4 minutes. The aim of this stage was to obtain measures of speeded reaction times in the absence of a self-paced task, and to compare them to the reaction times to *Go* cues obtained in stage II.

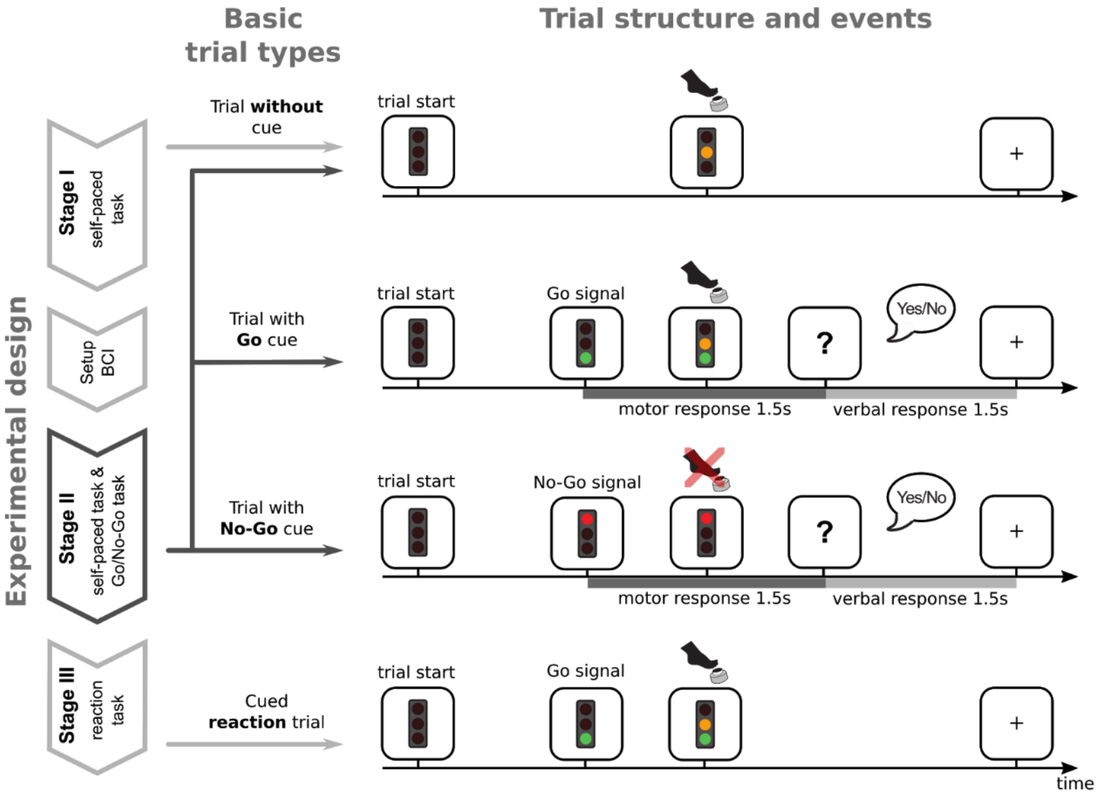

Figure 3.1. Experimental design. **Stage I:** participants performed spontaneous, self-paced pedal presses in single trials. No green or red cues were elicited and a trial always ended with a pedal press. **Setup BCI:** After stage I, a BCI was trained on EEG data recorded during stage I. **Stage II:** main experiment. Participants again performed spontaneous, self-paced pedal presses, but were occasionally interrupted by either the green or the red light turning on. If they pressed the pedal before either light was turned, the trial ended as in stage I. If the green or the red light was turned on, participants had 1.5 seconds to execute the corresponding motor response, i.e. to press immediately after a green cue or not press / inhibit after a red cue. Subsequently, the question "Were you about to press?" appeared on screen for 1.5 seconds (indicated by a question mark) during which participants were asked to verbally respond. **Stage III:** participants performed a simple reaction task: In each trial, the green light would turn on after a random time (chosen from a uniform distribution between 2 and 5 seconds), after which participants were instructed to press the pedal as quickly as possible.

Training of real-time BCI predictor

For the BCI predictor used in stage II, a linear classifier was trained on EEG data from the 100 pedal presses recorded during stage I.

EMG onset detection: For each trial, we assessed the movement onset. In order to be more precise we defined the onset of the movement based on the EMG rather than based on the final completion of the movement with the pedal press. To obtain EMG onset we high-pass filtered the EMG signal at 20 Hz. Then the standard deviation of the signal during the first 1000 ms after each trial start cue was determined as an "idle" baseline. For each trial individually, the standard deviation of subsequent, overlapping 50 ms windows was computed and EMG onset identified as the end of the first 50 ms window where the standard deviation exceeded the idle baseline by a factor of 3.5.

Class specification: Based on these movement onsets, two periods were defined as "Move" and "Idle" for the training of the classifier: The *Move* periods were 1200 ms long segments preceding EMG onset, while the *Idle* periods were 1200 ms long segments preceding the trial start cue (i.e. the onset of the traffic light).

Feature extraction and classifier training: EEG data from the two classes were baseline corrected to the mean signal in the time interval between −50 and 0 ms w.r.t. EMG onset or trial start cue, respectively. These were then averaged over time windows defined by the time points 1200, 900, 650, 450, 300, 200, 100 and 50 ms w.r.t. EMG onset or trial start cue, respectively. The choice of the baseline correction interval being locked to the end of the EEG segment (as opposed to the traditional choice of being locked to the beginning of the segment) and the choice of unequal time intervals were both based on a piloting analysis on previous data (Schultze-Kraft et al. 2016) that showed improved classification accuracy with these parameters. The resulting values were concatenated and used as features to train a regularized Linear Discriminant Analysis (LDA) classifier with automatic shrinkage (Blankertz et al. 2011). During stage II, the so-trained classifier was used to monitor the ongoing EEG in real-time. Therefore, every 10 ms a feature vector was constructed from the immediately preceding 1200 ms of EEG data, as outlined above, and used as input to the classifier, generating a classifier output value every 10 ms. This output variable was a continuous signal that probabilistically classified the current EEG segment either to the *Idle* or to the *Move* class.

Threshold setting: The classifier threshold was set for each participant individually. Because the classifier output signal was likely to mirror the stochastic nature of the EEG, a conservative threshold was defined in order to avoid many cues to be prematurely triggered by noise. For this, we performed a 10-fold cross-validation on the 100 trials from stage I and – mimicking a real-time predictor with a sliding window – computed for each trial the time of first threshold crossing of the classifier output for different threshold values. Let us refer to the time of first threshold crossing in a trial as a "prediction" event. Now, again for stage I, we define predictions occurring somewhere between trials start and up to 600 ms before movement onset as false alarms (FA), predictions occurring between 600 ms before movement onset and the time of movement onset as Hits, and predictions occurring after movement onset or not occurring at all as Misses. From this, the F-measure (Powers 2011) $F_\beta(\theta) = \frac{(1+\beta^2)Hit(\theta)}{(1+\beta^2)Hit(\theta)+\beta^2 Miss(\theta)+FA(\theta)}$ was computed for different threshold values θ. The largest F thus corresponds to the threshold θ were the Hit rate is maximal, while at the same time the FA and Miss rates are minimal. Moreover, by choosing β=0.5, we aimed at giving the minimization of FAs more weight than minimizing Miss rate. We prioritized minimizing the number of false alarms, at the cost of potentially missing some actions.

The resulting F values from stage I were smoothed and the threshold with the highest F value chosen. Classification accuracy for cross-validated data was > 75% for all participants.

Time of probing during stage II

Timing of *RP*+ cues: During stage II, if a trial was assigned as a *RP+* trial, the BCI was inactive during the first 1500ms after trial start. This ensured that an *RP+* cue was not elicited during the minimum self-paced waiting time of 2 sec as instructed to participants. After 1500ms, the BCI was activated and either the green or the red light were turned on as soon as the classifier reached the specified threshold.

Timing of *RP*- cues: During stage II, if a trial was assigned as a *RP*- trial, a cue was elicited after a predefined time that was chosen at random for each trial individually before trial start. To ensure that the cue was displayed at a plausible time given behavioural characteristics of the participant, a random time was selected from a uniform distribution between the 15 and 85 percentiles of the waiting times (time from trial start to EMG onset) of the 100 pedal presses in stage I. We further ensured that there was no EEG evidence for movement preparation at the randomly selected time points by eliciting *RP*- cues only if the classifier output was below threshold, indicating that the signal belonged to the “Idle” class.

Participant exclusion

The readiness potential is the target brain signal that we used to manipulate prospective information about motor preparation. While the RP is a potentially informative feature used by the classifier, it is not guaranteed *a priori* that the EEG features that were extracted by the classifier in order to separate both classes were based on the presence and absence of an RP over central channels. In turn, this does not guarantee that *RP+*

and *RP-* cues in stage II were triggered by the presence and absence of RPs, respectively. Therefore, we investigated whether such BCI-based manipulation was effective in each individual participant. We examined the recorded EEG data of each participant individually and found that for 7 out of the 23 participants this prerequisite was not met. These participants were thus excluded from all analyses. Details about the analysis procedure are provided below.

A qualitative assessment of EEG data from stage I used to train the classifier (*Figure 3.2*A) shows that for most participants the signals look as expected, with EEG signals preceding self-paced movements displaying the typical negative trend of a readiness potential (“Move” class), while EEG signals preceding trial start cues do not show any particular trend (“Idle” class).

A visual inspection of stage II data (*Figure 3.2*B) suggests that for most participants *RP+* cues were effectively preceded by an RP-like negativity, while *RP-* cues were not (with some conspicuous exceptions). This means that the classifier did often use the RP signal to trigger the cues. Note that here we only consider *RP+* and *RP-* cues that were elicited *before* any movement onset, thus excluding EEG data that would otherwise be contaminated with signals related to movement execution. In order to formally test whether we could rely on the BCI-triggered cues during stage II to discriminate *RP+* from *RP-* activity in each individual participant, we performed the following analysis. Channel Cz was chosen for analysis because readiness potentials preceding foot movements are typically most distinct over that channel (Schultze-Kraft et al., 2016). For each trial individually, we subtracted the average EEG signal in the time interval -200ms to 0ms from the average EEG signal in the time interval -1200ms to -1000ms, with respect to the time of cue presentation. These values represent the relative change in amplitude in channel Cz during the 1.2 seconds before cue presentation. If the BCI relied on the readiness potential for classification, EEG signals over central channels preceding *RP+* cues should be on average more negative than signals preceding *RP-* cues. To test this hypothesis, we ran an unpaired, one-sided t-test, for each participant separately. The box plots in *Figure 3.2*C show, for each participant and for *RP+* and *RP-* cues individually, the distributions of amplitude changes, with participants ordered by the t-statistics of the t-test from largest (left) to smallest (right). For the first 16 participants the t-test showed that signals preceding *RP+* cues became significantly more negative during the 1.2 sec interval than signals preceding *RP-* cues. For the remaining 7 participants the t-test did not give evidence of stronger negativities preceding RP+ cues compared to those preceding RP- cues. That is, the classifier had not made predictions based on the presence or absence of RP-like events in the EEG. Consequently, these participants are excluded from all subsequent analyses. *Figure 3.2*D shows individual and grand average EEG signals preceding *RP+* and *RP-* cues of the 16 participants selected for the final sample.

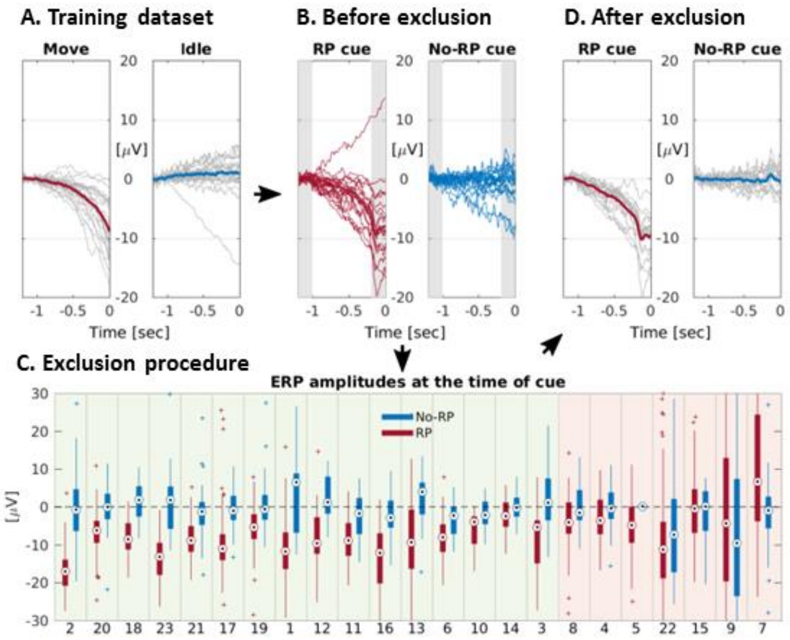

Figure 3.2. Participant selection procedure. Selection of suitable participants based on EEG signals of channel Cz. **A)** Event-related potentials (ERPs) of EEG signals recorded during stage I, time-locked to self-paced EMG onsets (left, “Move”) and trial start cues (right, “Idle”). ERPs are baseline-corrected in the interval -1200 to -1000 ms, and shown for individual participants (grey) and as grand-average (coloured). **B)** ERPs of EEG signals recorded in stage II *before* participant exclusion, time-locked to RP+ cues (left) and RP- cues (right). ERPs are baseline corrected in the interval -1200 to -1000 ms and shown for individual participants. **C)** For each participant and for RP+ and RP- cues individually, the box plots show the distribution of EEG signal amplitude changes between the time interval -1200 to -1000 ms and -200 to 0 ms with respect to cue onset. Participants are ordered in ascending order by the t-statistic of a two-sample one-sided t-test that tests whether the mean change in RP+ trials was more negative than in RP- trials. Participants for which p<.01 are highlighted in green, otherwise in red. **D)** ERPs of EEG signals recorded in stage II only for the selected N=16 participants (grey) and the corresponding grand average (coloured).

Statistical analysis

We ran two primary logistic mixed-effects analysis using the *glmer* function in the *lme4* R package (Bates, D, Mächler, M, Bolker, B.M., Walker 2015). In both analyses, we used a model comparison approach to select the optimal random effect structure, as suggested in (Matuschek et al. 2017). Full details of the step-by-step random effect selection process in the Appendix B (*Table B1* and *Table B2*).Below, we provide a summary of the analysis.

Prospective vs. retrospective contributions: The aim of the first analysis was to investigate whether intention reports are influenced by prospective (motor preparation) and retrospective (action execution) components. For this, we used all selected *Go* and *No-Go* trials and ran a logistic regression to predict the probability of reporting awareness given the presence or absence of an RP (*RP+*/*RP-*), the execution of an action (*Go/No-Go*) and their interaction (*RP x Action*).

Dynamic integration of prospective and retrospective cues: The second analysis aimed to study whether retrospective reconstruction and motor preparation interact in a time-dependent manner. For this, we studied *Go* trials only, because integration of prospective motor preparation and retrospective sensory feedback was only possible when participants executed an action (i.e. when sensory feedback was present). In *Go* trials, participants could execute an action at different times after cue presentation. We refer to this time delay as reaction time (RT). For each participant, we excluded trials where the RT was above or below 3*SD* from the individual mean to exclude outlier trials where participants were probably distracted. Because we were interested in the effect of absolute time passed between cue presentation and movement initiation rather than the deviation from mean RT, we did not centre the RT predictor variable. Finally, we ran a logistic regression to predict the probability of reporting awareness given the presence or absence of an RP (*RP+*/*RP-*), the continuous reaction time (*RT*) and their interaction (*RP* x *RT*).

3.3. Results

Data description & preliminary analysis

The number of trials in which participants were presented a cue, as well as the exact times when cues were presented, could not be precisely experimentally controlled. In case of *RP+* trials, this is because the BCI was calibrated so as to elicit cues preferably during the interval just before a movement, based on the detection of a readiness potential. However, the detection of transient events in the EEG in real-time by means of an asynchronous BCI is only possible with a limited accuracy, bound by the noisy nature of EEG signals. In turn, the timing of cues in *RP-* trials was predetermined at random at trial start. In the following, we characterize all recorded trials by whether and when a movement and/or a cue occurred during each trial, and the trial selection procedure. *Figure 3. 3* illustrates the types of trials that occurred during the task and highlights the ones that were included for analysis.

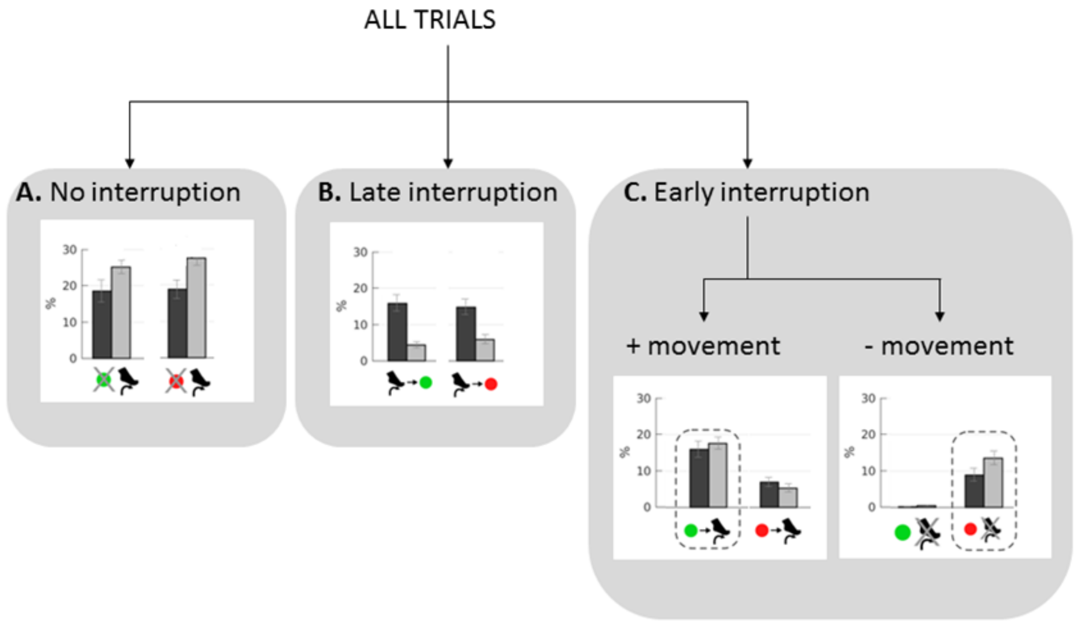

Figure 3. 3. Types of observed trials and trial selection procedure. Bar graphs represent the grand-averaged percentage (+ SEM) of trials within each category, for Go (*green*) and No-Go (*red*) trials separately, in the RP+ (dark grey) and RP- condition (light grey). The pictograms below the bar graphs indicate the temporal relation between cue presentation (*green/red circle*) and movement onset (*foot*). **A)** No interruption trials. In some trials, participants executed a movement and no cue was presented at all. In these trials, no awareness report was collected, and no further analysis were conducted. **B)** Late interruption trials. In some trials, cues came "too late", shortly after participants had already started moving. All these trials were discarded from further analysis. **C)** Early interruption trials. Cues were shown before any EMG onset was detected. In the *Go* condition (*green circle*), only trials where participants moved after the green cue presentation were included for analysis (*dashed rectangle*). In the *No-Go* condition (*red circle*), only trials where no EMG onset was detected after the red cue presentation were included for analysis (*dashed rectangle*).

In some cases, participants pressed the pedal *without* a cue being elicited (*Figure 3. 3A*). In *RP+* trials, these represent instances where the BCI failed to detect a readiness potential (*RP+/Go*: *M* = 18.5%, *SEM* = 3.1%; *RP+/No-Go: M* = 18.8%, *SEM* = 2.6%). In *RP-* trials, these represent instances where participants pressed the pedal before the random predetermined time of the cue (*RP-/Go: M* = 25.3%, *SEM* = 1.9%; *RP-/No-*

Go: M = 27.5%, *SEM* = 2.0%). In all these cases, since no cue was presented, no awareness report was collected. Thus these trials are excluded from further analysis.

In another subset of trials, both a cue was presented *after* EMG onset (*Figure 3. 3B*). In some *RP+* trials, a readiness potential was presumably correctly detected by the BCI, but cue was presented after participants had already started moving (*RP+/Go: M* = 15.9%, *SEM* = 2.3; *RP+/No-Go: M* = 14.8%, *SEM* = 2.1). In turn, the *RP-* trials where a cue was presented after participants' movement reflect rare instances where the predetermined probing time by chance coincided with the self-paced time of movement (*RP-/Go*: 4.3%, *SEM* = 0.9; *RP-/No-Go: M* = 5.9%, *SEM* = 1.2). For our purposes, these cues came too late and the corresponding awareness reports are thus excluded from further analysis.

In another subset of trials, the cue was presented *before* EMG onset, and participants moved after the cue. In the *Go* condition (*Figure 3. 3C*), these trials fulfil our prerequisite that *Go* cues must be followed by a movement, and thus the corresponding awareness reports are used in the main analysis (*RP+/Go: M* = 16.0%, *SEM* = 2.2; *RP-/Go: M* = 17.7%, *SEM* = 1.7). In the *No-Go* condition (*Figure 3. 3*C), participants sometimes *initiated* a movement after a cue was presented (*RP+/No-Go: M* = 6.9%, *SEM* = 1.3; *RP-/ No-Go:* 5.3%, *SEM* = 1.1). Although they were often able to abort a movement before fully pressing the pedal in some of these trials, the very initiation of a movement might suffice for participants to reconstruct an awareness of intention in the awareness probes that followed those cues. Thus, these trials were excluded from further analysis.

Finally, in some trials a cue was elicited before any EMG onset but no movement was produced after it (*Figure 3. 3D*). In the *Go* condition, these very rare occurrences reflect trials where participants failed to respond with a pedal press to a green cue (*RP+/Go*: *M* = 0.15%, *SEM* 0.06; *RP-/Go*: *M* = 0.4%, *SEM* = 0.2). In contrast, as expected, in the *No-Go* condition this occurred more frequently (*RP+/ No-Go: M* = 8.9%, *SEM* = 1.8; *RP-/ No-Go: M* = 13.6%, *SEM* = 1.8). In these trials, participants successfully followed the instruction to withhold any movement after a red cue. Because they fulfilled our prerequisite that *No-Go* cues must not be followed by a movement, the corresponding awareness reports were used in the main analysis.

Methodological considerations

A closer look into the distributions of EMG onsets with respect to the *Go/No-Go* probes provides better insight into the way in which our experimental design resulted in the observed proportions of trials (*Figure 3. 4*).

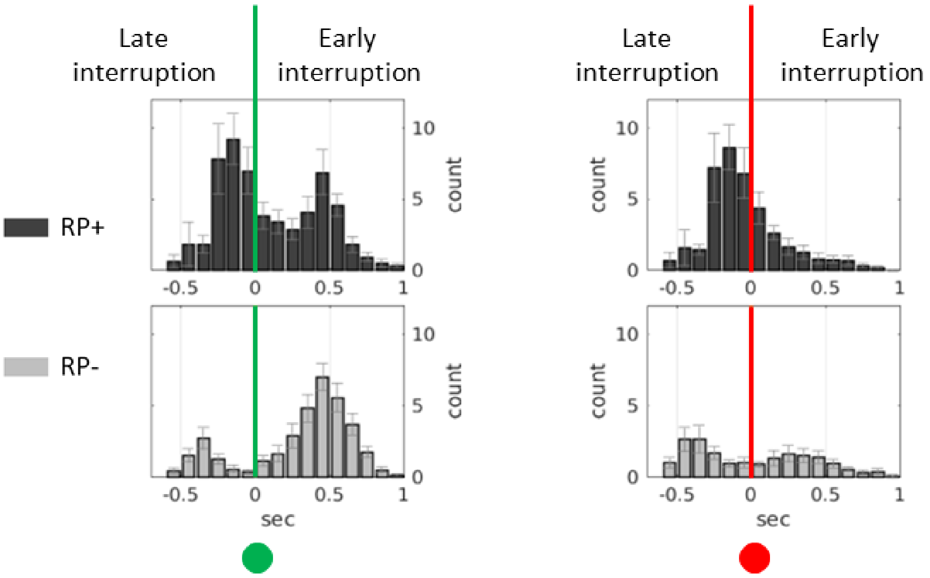

Figure 3. 4. Distribution of EMG onsets around the time of EEG-triggered awareness probes. Distribution of EMG onsets with respect to the time of probe presentation in Go (*green*, left) and No-Go (*red*, right) trials, in the RP + and RP - condition. Negative times correspond to the distribution of EMG onsets in Late interruption trials, in which the cue was presented after participants started moving. Positive times correspond to the distribution of EMG onsets in Early interruption trials (i.e. classic reaction times), where a cue was presented and a movement was initiated shortly afterwards.

Go and *No-Go* signals were often triggered *after* participants had started moving in the *RP+* conditions, while this was rarely the case in *RP-* condition. These "Late interruption" trials correspond to the distribution centred before cue presentation in the RP+ condition (*Figure 3.4*). These are instances of motor preparation states that were successfully detected by the BCI, but too late. However, the trials falling on the right-most tail of this distribution were instead instances where motor preparation was successfully interrupted early by the BCI. These trials can be interpreted as interruptions after the point of no-return (Schultze-Kraft et al., 2016). That is, trials in which participants would have moved *anyway* if a cue had not been presented. In fact, in a number of *No-Go* trials participants failed to inhibit a movement and an EMG onset was detected after the red cue. In turn, in the *Go* condition, the effect of these intercepted self-paced actions is visible in the higher count of trials with very fast responses (RT < 200 ms) in the *RP+/Go* condition compared to the *RP-/Go* condition. In sum, in the *RP+/Go* condition, very fast trials (<200 ms) include both self-paced movements that happened to occur just after the green *Go* signal (right tail of the "late interruption" distribution), and also reactions to the *Go* signal (left tail of the "early interruptions" distribution; see *Figure 3.4*). Instead, in the *RP-/Go* condition, movements produced very fast after the cue presentation were only reactions.

We checked that these very fast responses in the left tail of the "early interruption" distribution could physiologically be fast reactions rather than self-paced actions that the classifier did not predict, by looking at the RT distribution to *Go* cues on stage III (*Figure 3.5*). Reaction times show a skewed distribution that is typical for simple cued reaction time tasks. In that task, where no self-paced actions were being performed and participants were only reacting to *Go* cues presented at random times, we also observed some very fast reaction times (<200 ms) comparable to the ones found in the *RP*- condition in stage II.

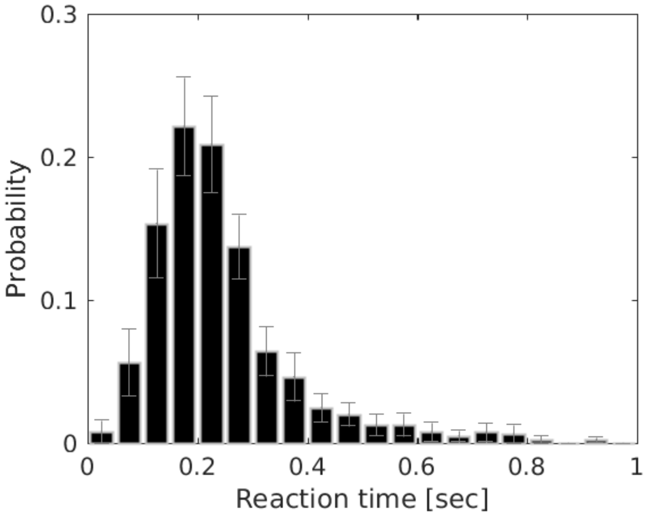

Figure 3.5. Stage III cued reaction time distribution. Grand-averaged distribution of reaction times from stage III (± SEM). The histogram shows in discrete 50 ms bins the probability of observing EMG onsets after presentation of the Go cue.

Interim considerations

The selected subset of trials for the main analysis included only those in which the cue was presented before any movement onset and was either followed by a movement in the *Go* condition or by no EMG onset at all in the *No-Go* condition (*Figure 3. 3C*). The analysed trials thus included an average of 31 (*SEM* = 4) *RP+*/*Go* trials, 27 (*SEM* = 3) *RP-/Go* trials, 18 (*SEM* = 4) *RP+/ No-Go* trials and 22 (*SEM* = 3) *RP-/ No-Go* trials, per participant.

Main analysis

To study how motor preparation processes and action execution processes influence the experience of intention, we fitted two logistic regression models to participants' responses. In each model we used two independent variables as predictors to test specific predictions made by the retrospective, prospective and temporal integration hypotheses (see *Methods* section).

Prospective and retrospective contributions to motor intention awareness

The prospective hypothesis suggests that participants have conscious access to their motor preparatory processes before movement initiation, and it thus predicts affirmative intention judgments to be more likely in the *RP+* than in the *RP-* condition. In turn, the retrospective hypothesis predicts that the execution of a movement will yield more awareness reports than the absence of an action.

As shown in *Figure 3.6A*, participants were significantly more likely ($X^2_{(1)} = 21.27$, $p < 0.001$) to report awareness in the *Go* (*M* = 35.2%, *SEM* = 6.4%) than in the *No-Go* condition (*M* = 16.0%, *SEM* = 4.7%). This suggests a strong effect of retrospection: executing an instructed action strongly increased the probability of participants' reporting an intention to move at the time of probing, compared to trials where no overt movement was present. Furthermore, participants were also significantly more likely ($X^2_{(1)} = 5.04$, $p = 0.024$) to report awareness of an intention to move in the *RP+* (*M* = 32.7%, *SEM* = 6.1%) than in the *RP-* condition (*M* = 24.0%, *SEM* = 6.0%), as shown in *Figure 3.6B*. That is, if they were preparing to move at the time the probe appeared, they were more likely to report an intention than if they were not preparing. No significant interaction was found ($X^2_{(1)} = 1.68$, $p = 0.193$), although descriptively the effect of prospection in the *No-Go* condition seemed smaller.

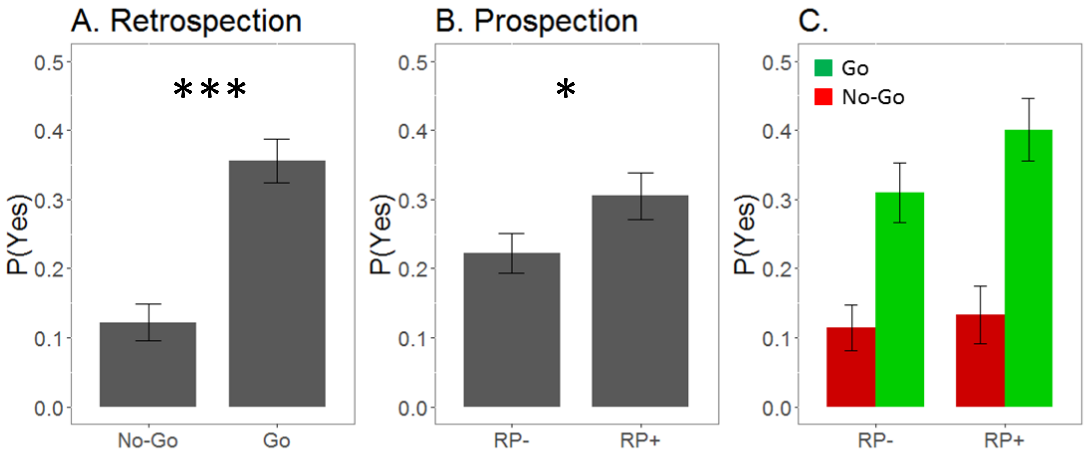

Figure 3.6. Prospection and retrospection effects in motor intention judgements. Pooled proportion of trials where participants responded 'Yes' in *Go* (*green*) and *No-Go* (*red*) trials. Error bars indicate 95% confidence interval. **A)** Participants were more likely to report an intention after executing an action (*Go*) than when no movement was present (*No-Go*), suggesting a strong effect of retrospective reconstruction. **B)** Participants were also more likely to report an intention when they were probed in the presence of motor preparation (*RP+*) than when no motor preparation was present (*RP-*), suggesting that intention judgements are influenced by prospective motor preparation states at the time of probing. **C)** Proportion of 'Yes' answers in the four conditions of interest. No interaction between Action (*Go/No-Go*) and RP (*RP+/RP-*).

Dynamic integration of prospective and retrospective information

We next investigated whether prospective and retrospective cues are dynamically integrated to form intention awareness judgements. The temporal integration hypothesis suggests that the experience of intention depends on the dynamic integration of motor preparation states and the sensory feedback following action execution. Because *No-Go* trials lacked any movement and thus a measurable reaction time, we restricted this analysis to *Go* trials only. First, this allowed us to test whether the RT modulates the retrospective reconstruction of intention. Second, it allowed us to investigate whether the effect of the motor preparation state triggering the cue

(*RP+*/*RP-*) on awareness of intention is dependent on the execution of the movement, i.e. the RT. We predicted that awareness of intention follows a mechanism similar to comparator models of motor control (Wolpert and Kawato 1998). In particular, we predicted that prospective information about motor preparation is only available for integration with (retrospective) sensory feedback for a short time. Thus, we expected intention judgments to be modulated by the time delay between the time elapsed between the motor preparation state triggering the cue and the execution of a movement.

As shown in *Figure 3.7,* the probability of reporting awareness decreased over time ($X^2_{(1)} =$ 70.74, $p < 0.001$). Participants were very likely to report awareness of intention if they initiated a movement shortly after a cue, but very unlikely to report awareness if they were slow. This was the case both in the *RP+* and the *RP-* condition. Furthermore, the interaction between the RT and the RP ($X^2_{(1)} =$ 11.04, $p < 0.001$) indicated that the presence of an RP significantly increased the probability of reporting awareness, but *only if* an action was executed within approximately 250 ms after cue presentation. At a hypothetical RT of 0 s, for example, the probability of reporting awareness predicted by the model is 0.944 in the *RP+* condition, while it is only 0.788 in the *RP-* condition.

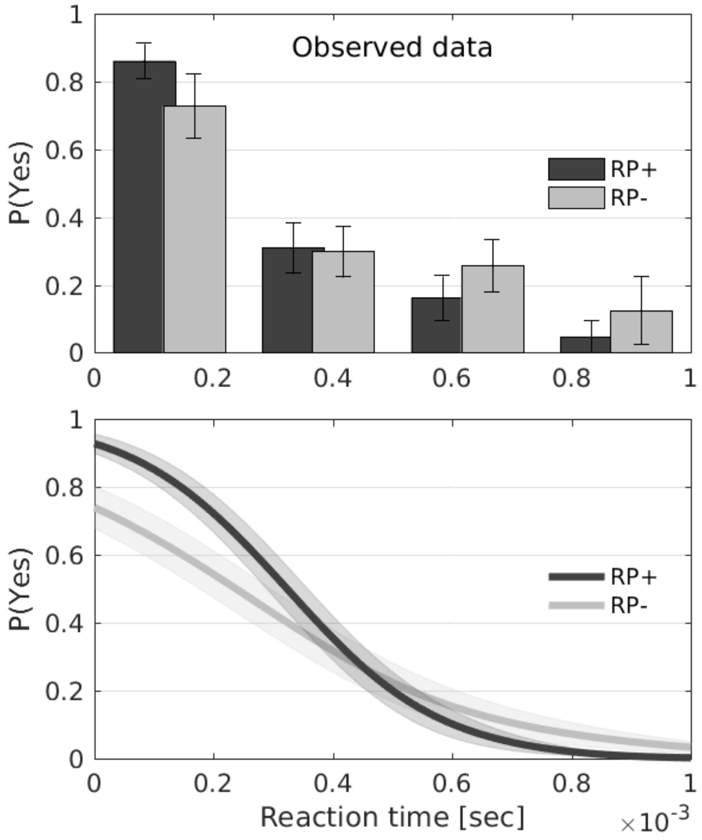

Figure 3.7. Dynamic integration of prospective and retrospective information in motor intention judgements. Grand-averaged probability (± SEM) of responding Yes in *Go* trials as observed (top) and predicted by the model (bottom) for *RP+* and *RP-* trials individually and for different reaction times (x-axis). In the observed data, average probabilities are reported according to reaction times in 4 discrete 250 ms bins. Model predictions are generated for a continuous reaction time variable.

3.4. Discussion

We conducted an EEG study of intention awareness using a BCI technique which allowed us to monitor motor preparation processes in real-time. Participants performed a self-paced movement task and were occasionally interrupted by a cue which instructed them to either execute or inhibit an action. They were then asked to report whether they were intending to move *at the time* the cue appeared. The time of presentation of the green and red cues was determined by a BCI trained to detect the presence or absence of an RP. This experimental design allowed us to investigate how awareness of intention is influenced 1) by the presence/absence of motor preparation (prospective hypothesis), 2) by the presence/absence of an action (retrospective hypothesis), and 3) by the time lag between the two (temporal integration hypothesis). Our results provide new insight into the elements contributing to the experience of intention.

Our first analysis showed a strong effect of retrospection (see *Figure 3.6*). Participants were overall more likely to report an intention to move at the time of probing when an action was executed (*Go* condition) compared to when no overt movement was present (*No-Go* condition). This is in line with retrospectivist theories suggesting that intention judgements are (at least partially) reconstructed after action execution (Banks and Isham 2009; Kühn and Brass 2009). Further, the estimated probability of reporting awareness was significantly higher for probes preceded by an RP (*RP+*) compared to those where no RP (*RP-*) was present, thus suggesting that motor preparation processes happening *before* action execution also influence motor intention perception, both when an action is executed and when it is not. A word of caution is needed here: although the interaction between prospective and retrospective factors was not found to be significant, the data descriptively showed a smaller effect of prospection in the *No-Go* condition (see *Figure 3.6C).* Thus, given the limited number of trials available for our analysis, it may be the case that no interaction was detected due to a lack of power rather than due to an absence of effect.

With our second model, we explored the dynamics of the prospective-retrospective interaction by taking into account the time at which actions were executed with respect to *Go* signals. First, this analysis showed that the retrospective reconstruction of intention found in the first model is time-dependent. Participants were very likely to report awareness when they responded fast to a *Go* signal, but this probability decreased sharply the larger the time difference was between the *Go* signal and movement initiation (*Figure 3.7*). These temporal dynamics were observed in both *RP+* and *RP-* trials. However, while even in the absence of movement preparation (*RP-*) a mere fast reaction (RT < 200 ms) was sufficient to produce a high percentage of awareness reports (*M*=72.8%, *SEM*=9.5%), equally fast trials preceded by movement preparation (*RP+*) resulted in an even higher rate of subsequent intention judgements (*M*=86.1%, *SEM*=5.1%). Thus, the model revealed a time-dependent effect of prospection: when a movement was executed, the presence of an RP increased the probability of awareness reports if the movement was executed within approximately 250 ms after a *Go* cue. Because the cues were locked to the presence or absence of an RP, the result can also be phrased as follows: in the *Go* condition the RP only made awareness reports more likely when an action took place shortly after movement preparation was detected. This suggests that information about prospective motor preparation is available for a limited period of time, and that it may be overwritten by retrospective reconstruction if potentially more salient information (e.g. sensory feedback) is present (*Go* condition). This time-constraint is consistent with everyday experience. Normal self-paced actions are typically executed immediately after an RP is present, and not otherwise. Therefore, it makes sense that only events (i.e. actions) happening at a physiologically plausible time after presence of a motor preparation brain signal are integrated with that brain state and perceived as its consequence.

A few considerations are worth noting before the final conclusions. First, this experiment used delayed reports of intention awareness. Because we were interested in studying both the prospective and retrospective contributions to motor intention awareness, participants provided their intention judgement after the allowed motor response time. Thus, although we asked participants about their intention at the time of the cue, our results cannot be interpreted as reflecting the *feeling* of intention participants had *at the time* the probe was presented. Rather, they reflect the influence of the brain states at the time of probing on subsequent judgements. The RP is a transient signal, and it decays over time. Thus, while the observed effect of the RP on intention awareness observed in the first analysis was small, the possibility remains that the effect of prospection could be higher if reports were obtained closer to the time of action.

Second, awareness of intention was not essential for successful task completion. Participant's main goal was to respond to the *Go* and *No-Go* cues correctly, and intention was only reported afterwards. This is in contrast with other awareness report methods, where action execution is contingent on instantaneous conscious experience (e.g. Matsuhashi and Hallett, 2008, and 'contingent action' experiment in chapter 2) and thus focus on the prospective component of motor awareness. Further research is required to evaluate to what extent purely prospective cues are accessible to guide action using paradigms where task performance is dependent on awareness.

Finally, our results show strong evidence for retrospective reconstruction of intention: participants were more likely to report an intention to move at the time of a cue if they executed a movement following such cue. This effect is time-dependent: they are more likely to report an intention if they respond fast to the cue, and unlikely to report it if their reaction to the cue is slow. This is consistent with retrospectivist theories suggesting that judgements of intention are (at least partially) reconstructed after action execution (Banks and Isham 2009; Kühn and Brass 2009, Wegner 2002). However, one question remains open. Assuming that participants have a certain "baseline" probability of reporting awareness (i.e. based on prospective cues or individual biases), two mechanisms might explain the observed retrospection results. First, it could be the case that execution of an action *increases* the probability of reporting an intention, as the aforementioned theories suggest. This would be a type of "positive reconstruction", where people "make up" intentions based on external events. However, it is also possible that, instead, inhibition of an action *decreases* the probability of reporting awareness. That is, there could be a kind of "negative reconstruction" in which participant's intention reports become less likely (i.e. due to loss of sensitivity to prospective cues, or increased reliance on external cues – or the lack thereof). Both mechanisms might coexist. While the present data show that some kind of retrospective effect is present, it cannot disentangle these two possibilities because no "baseline" awareness rate is available. Future research might attempt to investigate this question by developing a method that could provide a "baseline" estimate of the likelihood of reporting awareness (e.g. verbal reports) and comparing them, within subjects, to reports provided after action execution or action inhibition.

In sum, in this study we show that both prospective and retrospective cues influence delayed motor intention judgements, and that their integration is time-dependent. The presence of a motor preparation signal increases the probability of reporting an intention (prospective effect) and reports of intention are more likely after executing an action than in the absence of overt movement (retrospective effect). Further, the retrospective effect is modulated by response time (i.e. the probability of reporting awareness decreases the slower the RT), and there is a critical time window of approximately 250 ms during which prospective information is integrated with action execution feedback and influences awareness reports. This research sheds new light into the prospective and retrospective components of awareness of intention, and provides new methods to investigate the neural correlates of voluntary action execution and the related experiences of intention and agency.

Part II. Being-in-the-world

It is not the case that man 'is' and then has, by way of an extra, a relationship-of-Being towards the 'world'—a world with which he provides himself occasionally. Martin Heidegger.

Chapter 4

Breaking the symmetry: the neural correlates of evidence accumulation and action initiation in perceptual guesses

---

Synopsis

Most previous studies have defined voluntary actions as spontaneous and independent of the environment. However, naturalistic voluntary actions are typically reasons-responsive and informed by contextual information. We designed a decision-making task where actions were not spontaneous, but also not mere reactions to individual stimuli. Participants had to monitor evidence over time and decide *whether* to make an action on the basis of that evidence. The evidence was either strong (i.e. an easy perceptual decision) or neutral with respect to the decision (i.e. a kind of perceptual 'guess'), and there was no time pressure to respond, so actions were self-paced. We found that the p3 event-related potential tracks decision-making processes both when evidence is strong and when it is ambiguous, and that the RP can be found in easy decisions in which actions are strongly driven by the evidence. This suggests that the RP signal does not simply reflect the extent to which action selection is endogenous.

4.1.Introduction

Traditional definitions of voluntary action have relied on a strong endogenous-exogenous dichotomy. In particular, these definitions suggest that voluntary actions are endogenously generated, and they have typically been studied with methods aiming to minimise the role of exogenous input in the decision-making process. However, agents are intrinsically interactive. Humans are not isolated creatures, and successful behaviour requires the ability to meaningfully interact with the environment. Thus, action selection processes have not evolved to process endogenous and exogenous sources of information separately, but rather to optimise their integration.

The exogenous-endogenous continuum

Rather than constituting two separate categories of actions, the endogenous-exogenous dichotomy can be thought as two ends of a continuum. The relative contributions of endogenous and exogenous factors at the level of action initiation and action selection respectively define different types of actions. *Figure 4. 1* illustrates a schematic representation of the continuum.

At the endogenous end of the scale, research has investigated action initiation and action selection using tasks where exogenous factors are either absent or irrelevant. In ‘spontaneous action’ paradigms, participants are asked to execute actions “whenever they feel like it”, and no exogenous information is provided to inform when specific actions have to be initiated throughout the task (Libet et al. 1983; Matsuhashi and Hallett 2008; Schurger et al. 2012; Khalighinejad et al. 2018, and also **chapter**s **2** and **3** of this thesis). In turn, in ‘free choice’ or ‘picking’ paradigms (Frith et al. 1991; Soon et al. 2008; Bode et al. 2011, 2013) the available options are either extremely noisy or explicitly equivalent, and thus provide no strong evidence for the decision.

On the other extreme of the scale, the literature on perceptual decision-making has focussed on studying how external evidence is encoded, accumulated and used to drive action, typically in speeded tasks. When actions are based on strong external evidence, decisions are fast (Kelly and O'Connell 2013; Twomey et al. 2015), and people are confident about their choices (Kiani, Corthell, et al. 2014). The dynamics of easy perceptual decisions can be precisely modelled by taking into account the nature of the exogenous evidence (Ratcliff et al. 2016), while endogenous factors are typically only considered to account for residual variability in e.g. reaction times.

However, neither of these two extremes captures most instances of naturalistic, meaningful interaction with the environment. Most voluntary actions are not capricious and spontaneous movement, devoid of any connection with the immediate environment, nor are they mere reactions to strong perceptual stimuli. The aim of this chapter is to investigate actions that are more akin to the ones executed in daily-life, and thus lay on the middle of the endogenous-exogenous continuum.

People's actions are typically not simply *driven* by the environment, but rather *informed* by it. That is, actions are not simple reactions to stimuli, but rather follow from deliberate decisions that take into account the state of the external world. To achieve one's endogenous goals, it is necessary to be aware of the available means, which are partially dictated by the environment. However, environmental information is often noisy, and it is necessary to disambiguate the content of perceptual input (Parise and Ernst 2017). Many actions are thus akin to 'perceptual guesses' (Bode et al. 2012), where sensory information is not explicitly equivalent but rather too noisy to be conclusive.

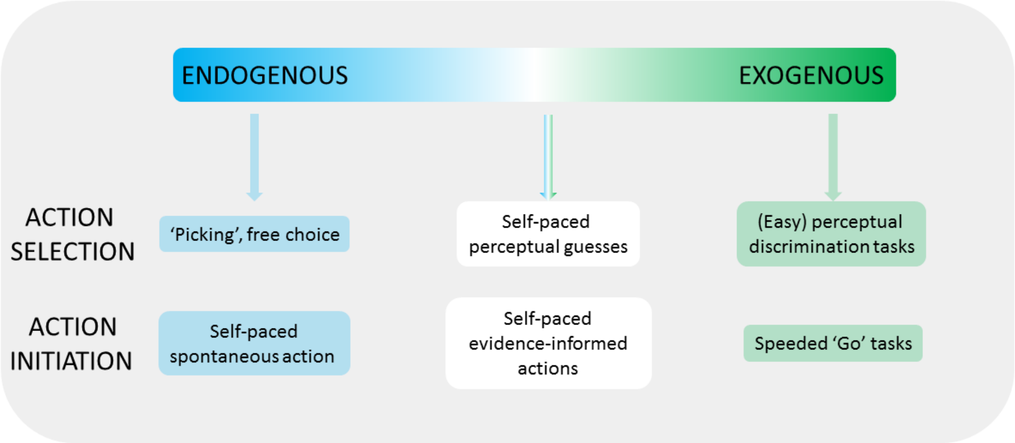

Figure 4. 1. The endogenous-exogenous action continuum. The boxes below indicate paradigmatic examples of actions in which either the outcome of the action selection process or the time of action initiation are determined by endogenous (blue), exogenous (green) or a combination of both sources of information (white). Note that the division between action selection and action initiation is not meant to depict these two as sequential, independent stages of a process. Rather, they illustrate the aspects that are the focus of the highlighted experimental paradigms.

Symmetry-breaking mechanisms

In ambiguous scenarios, one can think about two possible mechanisms to break deadlock situations and enable action. On the one hand, it has been suggested that noise in the sensory system can act as a symmetry-breaker, extracting signal from noise and thus resolving ambiguity (Shadlen et al. 1996; Deco and Romo 2008). On the other hand, it is possible that it is noise in the motor system that enables action triggering in the absence of sensory signal. In fact, recent studies have suggested that the RP may reflect a symmetry-breaking mechanism operating at the motor level in arbitrary actions, but absent in deliberate ones (Maoz et al. 2017). Indeed, ambiguous scenarios share similarities with spontaneous action initiation: both types of situations require the ability to act in the absence of compelling exogenous triggers. However, the neural correlates of these two potentially similar mechanisms have mostly been studied separately.

In our paradigm, participants had to decide *whether* to make an action based on either strong or ambiguous evidence. Actions were thus informed by evidence, but sometimes underdetermined by it. Further, decisions were made over long time scales (trial duration = 25 s), and actions were self-paced to a great extent. Thus, participants had a choice about whether to do something or not do anything at all, and their decisions were informed by either strong or weak evidence. Our task preserved the self-paced and optional nature that are typically ascribed to voluntary action, while embedding those actions in a decision-making context where evidence is relevant. The aim of our study is to investigate symmetry-breaking mechanisms in the context of an evidence-informed voluntary action task, where both might be operating either synergistically or as alternative means of generating action.

With this paradigm, we investigated the neural mechanisms underlying action decisions in ambiguous contexts at two levels. First, at the level of evidence accumulation. Second, during action initiation.

The neural correlates of noisy evidence accumulation

In this study, we used the p3 as a target neural signature of decision-making. The p3 is a classic event-related potential (ERP) that can be recorded in central and parietal electrodes beginning around 300 ms after stimulus onset and is often divided into early p3 (p3a) and late p3 (p3b) subcomponents (Polich 2009). The early p3 is closely linked to stimulus perception and 'novelty' detection, while the late p3 is strongly task-related and is thought to reflect stimulus processing rather than detection (Polich 2009). Recently, the p3 has been shown to encode a build-up to threshold decision variable that tracks decision-making processes in real-time, rather than merely reflecting their outcome (Twomey et al. 2015). It has been suggested that the p3 has the same neural sources of the centro-parietal positivity (CPP), a signal that tracks a continuous supramodal decision-variable that is sensitive to evidence strength and predicts decision time (O’Connell et al. 2012; Kelly and O’Connell 2013).

In our task, participants made decisions on the basis of long streams of discrete stimuli which were either neutral or informative, and the informative ones were either in favour of acting (positive evidence, +Ev), or against it (negative evidence, -Ev). This type of stimuli are in sharp contrast with other paradigms such as random-dot motion (Kelly and O’Connell 2013) or changing contrast (O’Connell et al. 2012), where different types of evidence and noise are presented continuously and simultaneously. Our design thus allowed us to investigate two main features of evidence accumulation that may underlie symmetry-breaking mechanisms.

First, the dynamic evolution of the p3 throughout the trial. Rather than investigating the dynamics of single ERP for each decision, we investigated the evolution of a decision-making process by analysing several sequential p3 components. If accumulation of evidence drives actions in all of our conditions, the evolution of the p3 amplitude throughout the trial should reflect the evidence accumulation process. However, the dynamics of the decision may vary across conditions and may reflect the symmetry-breaking mechanisms specifically operating in the context of perceptual guesses.

Second, whether different types of evidence are categorically encoded. The fact that different types of evidence evoked separate p3 components (+Ev/-Ev) allowed us to investigate whether asymmetries in the processing of different types of information underlie categorical decisions in ambiguous scenarios.

The neural correlates of action initiation in an evidence-informed task

Voluntary action research has traditionally investigated the neural correlates of action initiation in 'spontaneous action' tasks, where action execution is not the result of an external evidence accumulation process, but rather "endogenously generated". These tasks have identified the RP as a reliable neural precursor of spontaneous action (Libet et al. 1983; Keller and Heckhausen 1990; Trevena and Miller 2002; Schultze-Kraft et al. 2016; Khalighinejad et al. 2018), and because such signal is typically absent in cued-reaction tasks (Papa et al. 1991), it has traditionally been interpreted as a marker of endogenous action preparation and initiation. However, the role of the RP has long been controversial.

An alternative interpretation of the RP suggests that it merely reflects accumulation of random noise. It has been shown that RP shapes can be obtained by averaging of stochastic fluctuations in the motor system in humans (Schurger et al. 2012), and spontaneous foraging behaviour in mice (Murakami et al. 2014, 2017) and crayfish (Kagaya and Takahata 2010) is preceded by a build-up activity in motor areas similar to the one observed in the human SMA. Thus, recent studies have suggested that, rather than reflecting how endogenous or intentional an action is, the RP may enable action triggering in arbitrary actions, but be absent in deliberate ones (Maoz et al. 2017). This interpretation of the RP thus bears a strong resemblance to the idea that neural noise in perceptual systems drives decisions in ambiguous contexts (Shadlen et al. 1996).

Because action initiation is rarely spontaneous and independent from the environment, in our paradigm we embedded self-paced action initiation in an evidence-informed decision-making task. The different conditions in our task allowed us to discriminate different possible interpretations of the RP.

A traditional interpretation of voluntary action would suggest that when participants decided to act in ambiguous conditions they did so in a "more endogenous" manner than when they acted in response to strong evidence, and would therefore predict that an RP should be present prior to those more endogenous actions, but not prior to actions executed after strong exogenous input. Similarly, a symmetry-breaking interpretation of the RP (Maoz et al. 2017) would also predict that the RP should be present preceding arbitrary decisions (i.e. decisions made where the evidence did not strongly favour any alternative), but not in the easy ones where actions are strongly-driven by exogenous evidence. Alternatively, if the RP is related to planning and slow preparation of self-paced actions rather than to the amount of external evidence used to make a decision, we would expect to find an RP in all conditions in our task – since all of them are self-paced to a great extent, and participants can prepare them over extended periods of time.

Summary

In sum, in this study we designed a novel task that embeds voluntary action in a decision-making context where actions are evidence-informed. In our paradigm, actions are not spontaneous, but also not mere reactions to individual stimuli, and they preserve the self-paced nature characteristic of voluntary actions. This allowed us to investigate both the perceptual decision-making processes involved in action selection, and the action initiation processes in an evidence-informed action paradigm. In particular, we were interested in investigating the neural correlates of evidence accumulation in perceptual guesses and potentially in identifying the underlying asymmetries driving action. Further, our task operationalised voluntary action initiation in a context where actions were not mere reactions, but also not spontaneous and independent from the environment. Thus, it allowed us to shed new light into different interpretations of the readiness potential.

4.2.Methods

Participants

All participants were recruited from the ICN Subject Database. All participants were healthy, right-handed, young adults with normal or corrected to normal vision, no known disabilities and no history of neurological or psychological disorder. The study was approved by the UCL Research Ethics Committee and written informed consent was obtained from all participants before beginning the experiment. Subjects were paid £7.50 per hour, plus a performance-dependent reward (see below).

Nineteen subjects were initially invited to a single EEG session. Two participants were excluded because they did not understand the instructions. Three participants were excluded due to excessive EEG noise in the electrodes of interest (Cz & Pz). Eventually, 16 participants (13 female) were included in the study (*M* age = 22.66, *SD* = 3.19; range: 19-30 years).

Stimuli & experimental design

Procedure: Participants sat in a quiet room and were presented visual stimuli on a computer monitor. The instructions for the task were first displayed on the computer screen and then verbally repeated by the experimenter before the beginning of the experiment. Before the experiment, participants performed a practice version of the task (5 trials) to get familiar with the experiment.

Participants performed a decision-making task and provided responses with left and right-hand key presses. The experiment was divided into 3 blocks of 52 trials each. Every trial included a letter stream presented for 25 s followed by either one or two questions (see below). Trials were separated by a 2 s interval and participants could take a break half-way through and between blocks. In total, there were 39 trials per condition, and all conditions were randomised within blocks (13 trials per condition, per block). *Figure 4.2* depicts a schematic of the task, described in detail below.

Stimuli: A continuous stream of pseudo-randomised letters was presented at a 4Hz rate (duration of each letter 250 ms) on either a grey, pink or blue background (see below). Every trial lasted 25 s. The experiment was programmed in Matlab R2015a and Psychophysics Toolbox v3 (Brainard 1997). Participants made actions by pressing either the ‘j’ (pink) or the ‘f’ (blue) key of a standard computer keyboard with their left or right index finger respectively and answered the post-trial questions (see below) on a visual analog scale (VAS) by sliding a scroll-bar with a standard computer mouse.

All letters were lower-case black consonants presented without any blank interval between consecutive letters. There were two sets of task-relevant (target) letters, to which a task-relevant colour was assigned (in parenthesis): ‘b’ and ‘d’ (*bd,* blue), and ‘p’ and ‘q’ (*pq,* pink). The precaution was taken not to include letters ‘h’, ‘j’, ‘k’, ‘l’, ‘t’ in the letter stream to ensure that morphological similarities between these and the target letters would not confound the results.

Task: Participants were instructed to monitor the letter stream and decide which set of targets was appearing more frequently (i.e. *bd* or *pq*). Each trial started with either a blue or a pink background. Participants’ task was to make sure that the background colour matched the more frequent group of letters (i.e. if the most frequent set of targets was *bd*, the screen should be blue. If the most frequent set of targets was *pq*, the screen should be pink). If the default colour of the screen on a given trial matched the set of targets participants perceived as most frequent, they were not required to do anything (*Stay* trials). If the default colour of the screen on that trial did not match the most frequent set of targets, they had to press the key corresponding to the colour they wanted to switch it to (i.e. either pink or blue) with the left or right hand respectively (*Switch* trials). They were allowed to switch the colour of the screen only once per trial, which they were aware that lasted 25 s regardless of their behaviour. There was no incentive to respond fast, since trials were not terminated at the time of response and stimuli continued to appear on the screen until the end of the 25 s. Thus, in some trials participants would execute an action to change the colour of the screen (*Switch* trials), whereas in others they would not do so (*Stay* trials).

Self-reported measures: At the end of each trial, participants were asked a few follow-up questions. Participants were asked to report how confident they were that switching (or not switching) was the right thing to do on a Visual Analog Scale (VAS). If they had switched, they were also asked to estimate at what point during the trial they switched on a VAS.

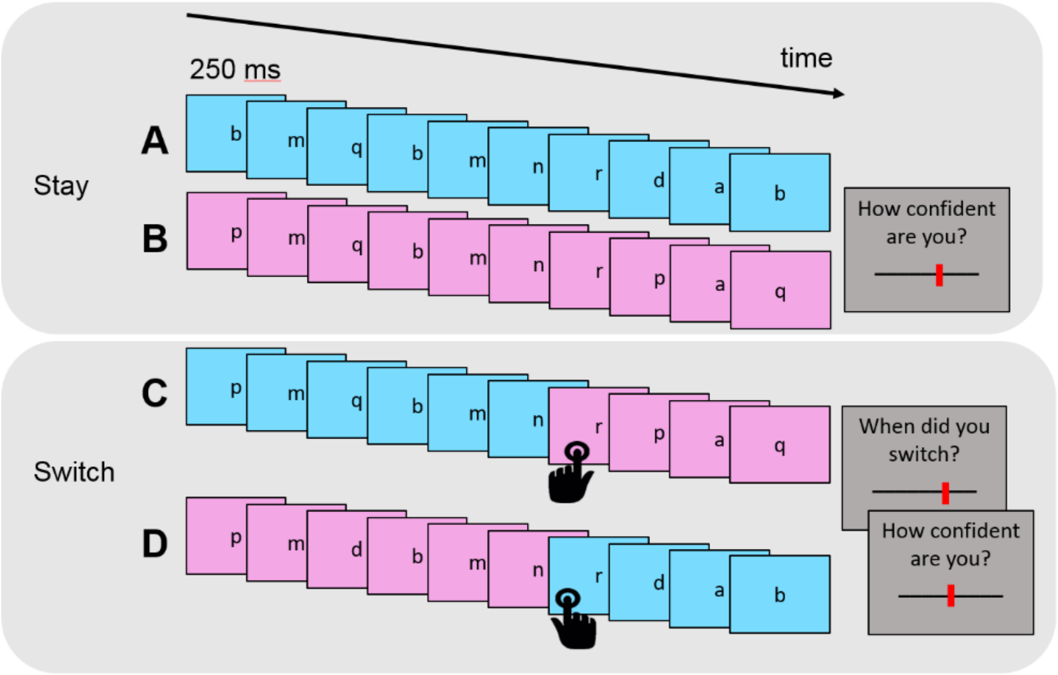

Figure 4.2. Experimental design. Each trial started with either a blue or a pink background. A letter stream was presented for 25 seconds, and participants had to decide which of the two target letter sets was most frequent (blue targets: ‘b’ and ‘d’, or pink targets: ‘p’ and ‘q’). Their task was to make sure that the colour of the screen matched the most frequent set of targets (e.g. if ‘p’ and ‘q’ were most frequent, the screen should be pink). Sometimes, the starting colour would already match the set of targets participants perceived as most frequent. For example, if ‘b’ and ‘d’ were most frequent, and the screen was already blue (**A**), or ‘p’ and ‘q’ were most frequent, and the screen was already pink (**B**). In those cases, they were instructed to not execute any action (**Stay**). In other trials, the most frequent targets did not match the given colour of the screen (**Switch**). In those cases, they could change the colour of the screen by pressing either the ‘f’ key with their left hand to turn the screen pink (**C**), or the ‘j’ key with their right hand to turn it blue (**D**). They were only allowed to change the colour of the screen once every trial. After 25 s, the letter stream was terminated and participants were asked confidence ratings and estimates of the time at which they made a decision in Switch trials only. Note that the example trials are for illustrative purposes only. In the experiment, relevant letters were always separated by at least 2 distractors, and 100 letters were presented in each trial.

Experimental conditions: The ratio between the frequencies of the two sets of targets together with the colour of the screen at trial onset defined four experimental conditions.

In the **Informative** conditions, the most frequent set of letters would appear twice as frequently (i.e. once every 1 to 4 s) as the other one (i.e. once every 1 to 8 s), and provided strong evidence either for or against switching. In the **Informative pro-Stay** (pSt) condition, the colour of the screen at trial onset and the most frequent set of targets matched (i.e. the frequency of $pq > bd$ and the screen was pink, or frequency of $pq < bd$ and the screen was blue). Hence, participants should not press a key to change the colour of the screen. In the **Informative pro-Switch** (pSw) condition, the colour of the screen and the most frequent set of targets did not match (i.e. the frequency of $pq > bd$ and screen was blue, or the frequency of $pq < bd$ and the screen was pink). Therefore, participants should press a key to switch the colour of the screen.

In **Neutral** conditions, both sets of letters were presented at the same frequency. In the **Neutral High** (NH) condition, evidence was abundant (i.e. one piece of +Ev and one piece of –Ev were presented at a random time every 1 to 4 s). In the **Neutral Low** (NL) condition, evidence was scarce (i.e. one piece of +Ev and one piece of –Ev were presented at a random time every 1 to 8 seconds). In these two neutral conditions, there was no strong evidence for or against switching in any given trial. Note, however, that participants were *not* told that in some conditions there was no net evidence.

In all conditions, letters from both sets of targets were never presented within intervals shorter than 0.5 s.

Coding scheme: In all conditions, half of trials started with a Blue screen, and half with a Pink screen. For analysis purposes, we averaged trials with different starting colours within conditions. Evidence was coded with respect to the colour that the screen would become if they decided to Switch. That is, in trials where the screen was Blue at trial onset, *pq* were coded as positive evidence (+Ev, evidence *for* switching), and *bd* as negative evidence (-Ev, evidence *against* staying). In turn, trials where the screen was initially Pink, *bd* were coded as +Ev and *pq* as –Ev. The coding scheme is illustrated in the figure captions of the results section.

Reward: In both Informative conditions, participants were rewarded with +1 penny for correct decisions (i.e. switching in the pro-switch trials and not switching in the pro-stay trials) and penalised for incorrect decisions (i.e. switching in the pro-stay trials and not switching in the pro-switch trials) by deducting 1 penny. In the neutral conditions, there were no correct or incorrect decisions. Hence, participants were assigned +1 penny reward or no reward at random on every trial. Participants were informed about the accumulated reward on the breaks between blocks.

EEG recording

EEG was recorded from 26 scalp sites (FZ, FCZ, CZ, CPZ, PZ, POZ, FC1, FC2, C1, C2, CP1, CP2, F3, F4, F7, F8, C3, C4, CP5, CP6, FC5, FC6, P3, P4, O1, O2) using active electrodes (g.LADYbird) fixed to an EEG cap (g.GAMMAcap) according to the extended international 10/20 system. EEG data were acquired using a g.GAMMAbox and g.USBamp with a sampling frequency of 256 Hz. Signal was recorded using g.Recorder (G.tec, medical engineering GmbH, Austria). All electrodes were online referenced to the right ear lobe. Vertical and horizontal electroocular activity was recorded from electrodes above and below the right eye and on the outer canthi of both eyes.

Behavioural data analysis

Perceived time of switch decision: After trials where a switch decision was made, participants were required to provide an estimate of the time at which they decided to switch. In order to calculate the bias in their estimation with respect to the actual time of the decision to switch we first linearly converted the 100-point VAS on which participants provided the response dividing it by 4, thus providing a 25-point scale equivalent to the 25 s duration of the trial. To calculate the bias (B) in their perceived time of decision, we then subtracted the actual time at which they executed an action (t_a) form the perceived time at which they switched (t_p), such that $B = t_p - t_a$.

Statistical analysis: All analysis of behavioural data were performed in R using mixed-models regression with the *glmer* function in the *lme4* package (Bates, D, Mächler, M, Bolker, B.M., Walker 2015). We fit models to compare the percentage of Switch decisions, their time, the bias in the perceived time of action between conditions. Further, we compared confidence in Switch and Stay decisions between all conditions. Conditions (Informative pro-Stay/Informative pro-Switch/NH/NL) - and decisions (Switch/Stay) in the confidence analysis - were included as fixed effects, and a random intercept was included in all models to control for individual variability. Post-hoc pairwise comparisons run as follow-up to the primary regression analysis were run using the *emmeans* function in the emmeans package (Lenth et al. 2018), and all pairwise t-tests were corrected for multiple comparisons using the Tukey method and Kenward-Roger method was used to estimate the degrees of freedom.

EEG analysis

***Preprocessing**:* EEG data were processed using Matlab R2014b (MathWorks), and EEGLAB version 13.5.4b (Delorme and Makeig 2004). First, scalp and eye electrodes were re-referenced to the average of two mastoid electrodes. Continuous EEG and EOG data were band-pass filtered between 0.01 Hz and 30Hz using an 8th order Butterworth filter. Then, data were down-sampled to 200 Hz. Second, EEG signals were epoched in two ways. For p300 analysis, EEG signal was locked from -0.5 before to 1 s after the presentation of relevant letters (‘p’, ‘q’, ‘b’ and ‘d’). For RP analysis, EEG signal was locked from -2.5 s to 0.5 s after the presentation of the letter following a keypress. Next, baseline correction was performed using the 500 ms at the beginning of the epoch [-0.5 to 0 s in p300 epochs, or -2.5 s to -2 s in RP epochs]. Then, an independent component analysis (ICA) was computed on the epoched data using the EEGLAB *runica* algorithm. Vertical eye movement components were visually identified and removed from the signal. Finally, artefact rejection was performed by removing all epochs with >120μV fluctuations from baseline in the preselected channels of interest (CZ and PZ).

p300 analysis: To study the dynamics of evidence accumulation as encoded in the p300, we extracted the p300 at PZ in response to every instance of relevant evidence (‘b’, ‘d’, ‘p’ and ‘q’), and we obtained the average amplitude of the whole duration of the component [0.3 to 0.8s post stimulus], as determined by visual inspection of the grand-averaged traces. We used these raw values for the dynamic analysis of the evolution of the p3 amplitude over time in a stimulus-locked (i.e. in absolute time from trial onset) and action-locked manner (i.e. in relative time to execution of action).

For single-trial average ERP analysis, we averaged the p3 components in response to +Ev and –Ev separately for each single trial, and we used these values for statistical inference. We ran a multilevel linear regression using the *blmer* function from the *blme* package (Dorie et al. 2016) using a weakly informative Wishart prior for the covariance matrix and a null prior for the fixed-effects. We fit a mixed model to test whether the single-trial average amplitude of the p3 significantly varied between different kinds of evidence (+Ev/-Ev), conditions (Informative/NH/NL) and Action decisions (Switch/Stay).

For initial and final amplitude analysis, we investigated whether the amplitudes of the *first* and *last* pieces of evidence presented from trial onset until action execution or until trial end (in Switch and Stay trials respectively) varied between different kinds of evidence (+Ev/-Ev), conditions (Informative/NH/NL) and Action decisions (Switch/Stay).

In all models, a random intercept was included to control for individual variability.

RP analysis: Statistical tests on averaged EEG data were run using FieldTrip toolbox (Oostenveld et al. 2011) cluster-based permutation analysis (Maris and Oostenveld 2007). The main contrast of interest was the comparison between neural activity preceding Switch actions between conditions (pro-Switch, NH, NL). One electrode over the medial frontal areas (CZ) was preselected for analysis. The cluster-based tests were performed on the individual participant averages using the following parameters: two-tailed dependent samples t-test, time interval = [-2 0 s relative to the event of interest], number of draws from the permutation distribution = 10000. No correction for multiple comparisons was performed.

4.3.Results

4.3.1. Behavioural results

Switch decisions: Participants performed the task well. On average, they switched on almost all of the pro-switch trials (*M* = 89.58%, *SEM* = 1.51%) and almost never in the pro-stay ones (*M* = 3.04%, *SEM* = 0.53%). All errors were excluded from further analysis. In the NH condition, participants switched on 40.54 % of trials (*SEM* = 2.72%), which was significantly more often than they did in the Neutral Low condition (*M* = 34.13% , *SEM* = 2.06%; $X^2_{(1)} = 6.67$, $p = 0.009$; see *Figure 4. 3A*).

Time of action: The time at which participants made the decision to switch varied across conditions ($X^2_{(2)} = 79.95$, $p < 0.001$). In particular, participants were significantly slower in the NL condition (*M* = 18.86 s, *SEM* = 0.39 s) than in the NH (*M* = 17.50 s, *SEM* = 0.55 s), which in turn was slower than the pro-switch (*M* = 16.09 s, *SEM* = 0.56 s) conditions (all pairwise comparisons $p < 0.001$, Tukey corrected). That is, they were faster when the evidence was strong than when it was weak, and they were the slowest when evidence was ambiguous, and scarce (*Figure 4. 3B*).

Perceived time of action: In all conditions, the average bias values were significantly below zero ($X^2_{(1)} = 53.80$, $p < 0.001$), indicating that the reported time preceded the actual time of the Switch. We found a main effect of condition ($X^2_{(2)} = 8.82$, $p = 0.012$), indicating that the time bias in the NH condition (*M* = -2.25 s, *SEM* = 0.47 s) was significantly lower than in both the pro-Switch (*M* = -2.77 s, *SEM* = 0.37 s; $t_{(1008)} = 2.74$, $p = 0.01$) and NL (*M* = -2.88 s, *SEM* = 0.37; $t_{(1008)} = 2.47$, $p = 0.03$) conditions.

Confidence: We analysed confidence ratings depending on condition and action decision. For the regression, we collapsed all accurate trials in the pro-Stay and pro-Switch conditions into one category, labelled ‘Informative’, and we included all trials in the NH and NL trials. We found a main effect of condition ($X^2_{(2)}$ = 483.44, $p <$ 0.001), indicating that participants were more confident about their decisions in the informative (M = 67.55, SEM = 2.87) than in the NH (M = 48.37, SEM = 3.12; $t_{(2391)}$ = 18.33, $p <$ 0.001) and NL (M = 50.15, SEM = 3.02; $t_{(2391)}$ = 15.96, $p <$ 0.001) conditions, which did not differ between them ($t_{(2391)}$ = 1.71, p = 0.20).

Further, we found an interaction between Action and Condition ($X^2_{(2)}$ = 20.15, $p <$ 0.001). In the NL condition, there was a main effect of action ($X^2_{(1)}$ = 14.29, $p <$ 0.001), indicating that participants were more confident of their choices when they Switched (M = 54.05, SEM = 4.35) compared to when they Stayed (M = 46.25, SE = 4.10). The difference was not significant in the NH ($X^2_{(1)}$ = 3.61, p = 0.057; Switched: M = 50.41, SEM = 4.38, Stayed: M = 46.34, SEM = 4.54), and the informative conditions ($X^2_{(1)}$ = 3.45, p = 0.063; Switched: M = 66.58, SEM= 4.12, Stayed: M = 68.53, SEM = 3.99; see *Figure 4. 3C*).

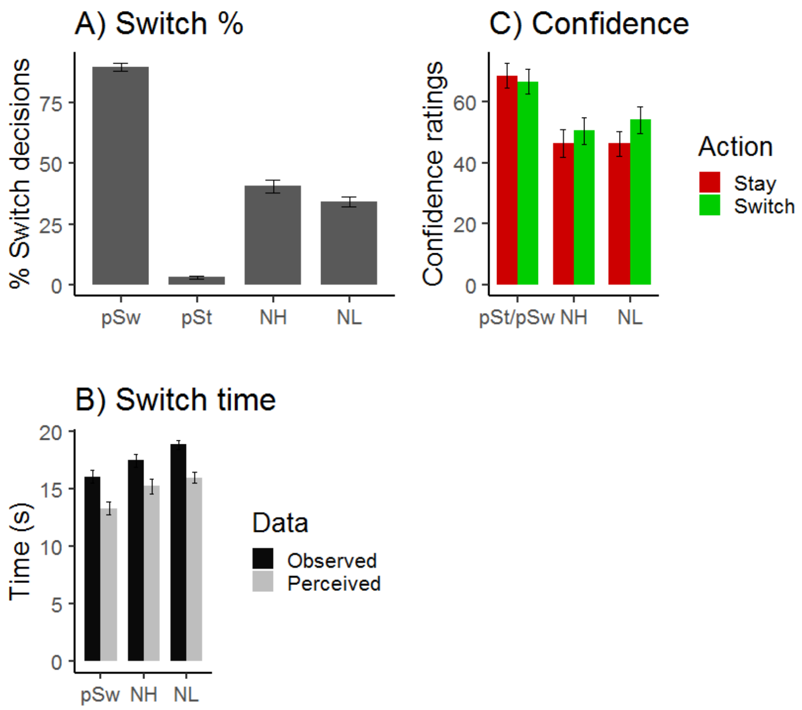

Figure 4. 3. Behavioural results. Grand-averaged (+SEM) behavioural results for all conditions (pro-Switch (pSw), pro-Stay (pSt), Neutral Nigh (NH) and Neutral Low (NL). **A)** Percentage of switch decisions in each condition. **B)** The perceived time of switch decisions (*grey*) was systematically earlier than the observed, actual time of action (*black*). **C)** Confidence ratings in Stay (*red*) and Switch (*green*) trials.

4.3.2. EEG results

The p3 encodes a categorical decision variable

We first investigated whether participants' decisions were categorically encoded by the p3. To do so, we ran a linear regression the average amplitude of the p3 up to the time of action in Switch trials or until the end of the trial in Stay trials differed given the decision (Switch/Stay), the type of evidence (+Ev/-Ev), and the condition (pSw/NH/NL). *Figure 4.4.* illustrates the main results of the analysis.

We found an interaction between the decision and the type of evidence ($X^2_{(1)}$= 7.92, $p = 0.004$). Hence, we run a regression for Switch and Stay trials separately. When participants decided to Switch, the amplitude of the p3 was significantly higher ($X^2_{(1)} = 5.68$, $p = 0.017$) in response to +Ev ($M = 2.59$ μV, $SEM = 0.43$ μV) than to –Ev ($M = 1.08$ μV, $SEM = 0.47$ μV). There was no significant difference ($X^2_{(1)}$ =2.43, $p = 0.11$) in the amplitudes of the p3 in response to +Ev and -Ev where participants decided to Stay, although descriptively the average results showed the opposite pattern than Switch trials. In Stay trials, the amplitude of the p3 in response to +Ev ($M = 1.28$ μV, $SEM = 0.39$ μV) was, on average, lower than that in response to –Ev ($M = 2.11$ μV, $SEM = 0.36$ μV).

Further, we found a main effect of condition ($X^2_{(2)}$= 9.98, $p = 0.006$). Descriptively, a clear gradient was observable in the average p3 amplitudes, where the average amplitudes of the p3 in the NL ($M = 2.73$ μV, $SEM = 0.35$ μV) were greater than in the pSw ($M = 1.69$ μV, $SEM = 0.31$ μV), which in turn showed greater amplitudes than the NH ($M = 0.94$ μV, $SEM = 0.39$ μV), although only the difference between the NL and the NH was significant ($p = 0.003$).

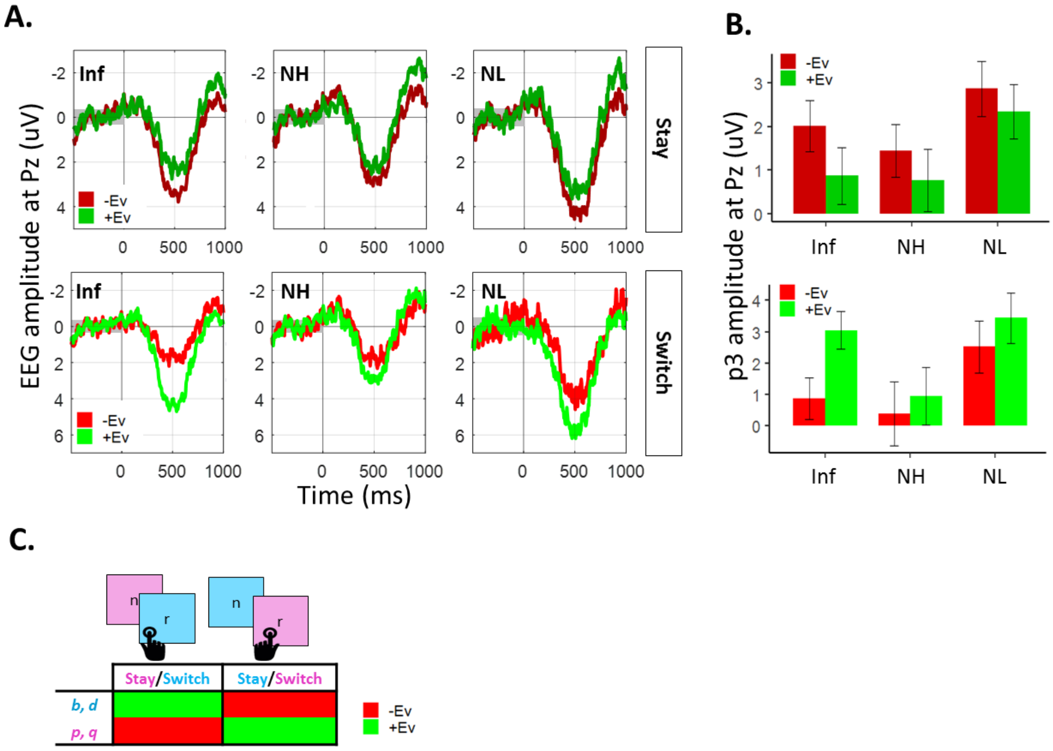

Figure 4.4. The p3 encodes an internal decision variable. **A)** Grand-averaged ERP traces at Pz in Stay (*top*) and Switch (*bottom*) trials in response to +Ev (*green*) and –Ev (*red*) in Informative (*Inf*, *left*), Neutral High (NH, *middle*) and Neutral Low (NL, *right*) conditions, locked to letter onset. **B)** Grand-averaged (±SEM) p3 amplitude between 300 and 800 ms, used for statistical analysis in Stay (*top*) and Switch (*bottom*) trials. **C)** Coding scheme. Evidence is coded with respect to the colour the screen would become if participants decided to Switch. (i.e. in trials where the screen was Blue at trial onset, *pq* are +Ev, and *bd* are -Ev).

These results show that the decision variable encoded by the p3 categorically encoded internal decision-making processes both in easy, informative conditions and difficult, neutral ones.

The p3 dynamics over time

Stimulus locked

We then investigated how the dynamic evolution of the p3 amplitude over the course of the trial underlay the observed average differences. For this analysis, we fitted a linear multilevel regression to predict the amplitude of the p3 evoked by every single relevant letter (*bd/pq*) based on the time elapsed from trial onset (in seconds) and its interaction with the type of letter that evoked it (+Ev/-Ev), participants decision (Switch/Stay), and the condition (Informative/NH/NL). In this analysis, we focussed on the main effect of time and its interaction with other factors. *Figure 4.5.* illustrates the results of this analysis.

We found a significant main effect of time ($X^2_{(1)}$= 18.94, p < 0.001), indicating that the amplitude of the p3 increased over the course of the trial in all conditions, all evidence types, and both when participants Switched and Stayed.

Additionally, we found a marginally significant interaction between time from trial onset and decision ($X^2_{(2)}$= 3.10, p = 0.07) which suggested that the increase in amplitude was descriptively smaller in Stay than in Switch trials, as visible in the plots (see *Figure 4. 5*). Further, another marginally significant interaction between time from trial onset, type of evidence and decision ($\mathrm{X}^2_{(2)}$ = 2.91, p = 0.08) suggested that, in Switch trials, the increase of amplitude was descriptively greater in response to +Ev than in response to –Ev (*Figure 4. 5*).

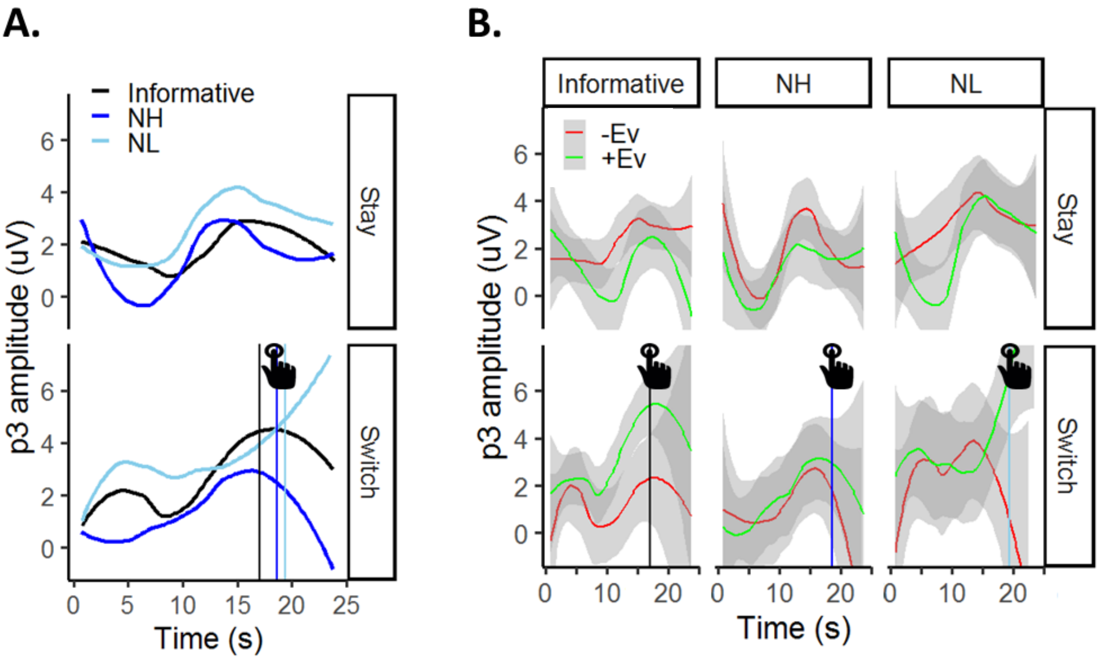

Figure 4. 5. Stimulus-locked evolution of the p3 amplitude. **A)** Evolution of p3 amplitude at Pz from trial onset, pooled across participants and types of evidence in Stay (*top*) and Switch (*bottom*) trials. The amplitude of the p3 increased over time in all conditions, although descriptively less in Stay trials. **B)** Evolution of the p3 amplitude (±SE) at Pz from trial onset, pooled across participants for each type of evidence separately in Stay (*top*) and Switch (*bottom*) trials. Note: vertical bars indicate mean time of Switch action in each condition.

Action-locked

Because actions happened at different times on each trial, stimulus-locked analysis necessarily have less power than action locked ones. Thus, we repeated the analysis for Switch trials only, locking the p3 data to the time of action. This allowed us to investigate the dynamics of the p3 with respect to an internal decision time rather than the external absolute time. This analysis revealed a significant interaction between the time to Switch action and type of evidence ($X^2_{(2)}$= 5.72, p = 0.01), indicating that the amplitude of the p3 significantly increased *only* in response to +Ev towards the time of action (*Figure 4. 6B).*

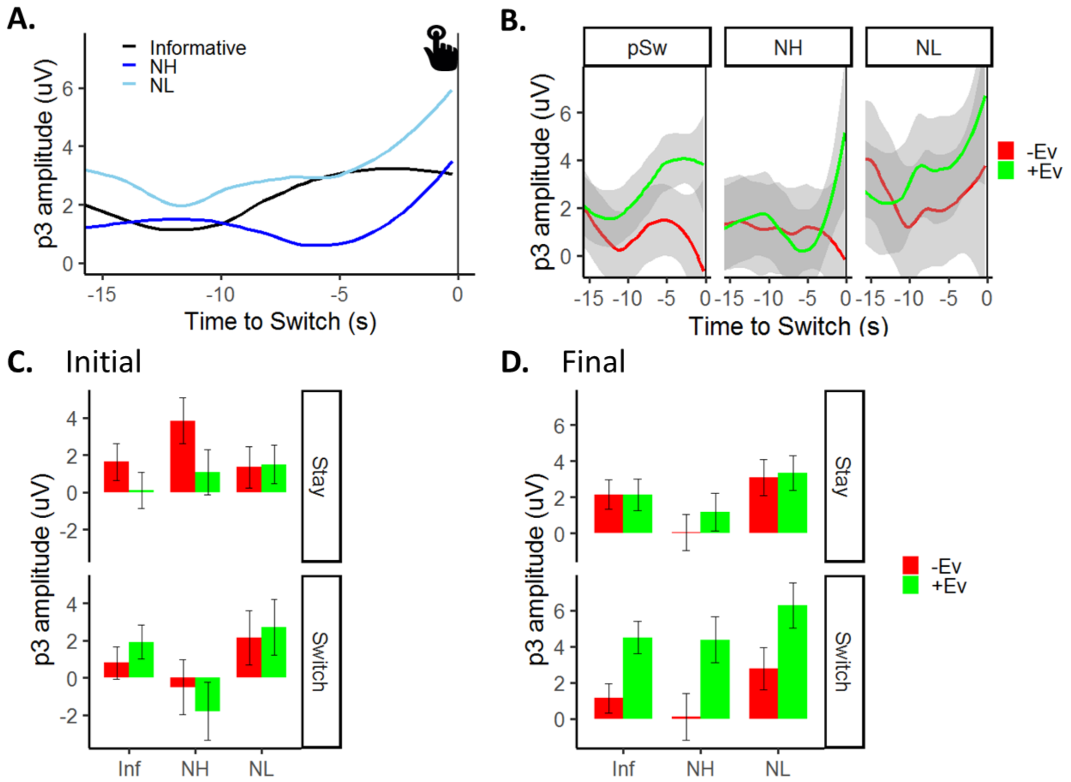

Figure 4. 6. Action-locked evolution of the p3 amplitude. **A)** Evolution of p3 amplitude at Pz locked to action in Switch trials, pooled across participants and types of evidence, temporally smoothed with the Loess method. The amplitude of the p3 increased over time in all conditions. **B)** Evolution of the p3 amplitude (±SE) at Pz locked to action in Switch trials, pooled across participants for each type of evidence separately. The amplitude of the p3 only increased significantly in response to +Ev. **C)** p3 amplitude in response to the *first* piece of +Ev (*green*) and –Ev (*red*) presented on Switch and Stay trials respectively. **D)** p3 amplitude in response to the *last* piece of +Ev and -Ev presented before action in Switch trials, or before the end of trial in Stay trials.

How is symmetry broken?

Because in the neutral conditions there was no net evidence available for accumulation, we further investigated whether asymmetries in the initial p3 component in response to +Ev and –Ev might have driven the evidence accumulation process. We define *initial* responses as responses to the *first* instances of +Ev and –Ev presented on a given trial. For this analysis, we run a mixed regression to investigate whether the amplitude of the first p3 differed between conditions (Inf/NH/NL), decisions (Switch/Stay) and types of evidence (+Ev/-Ev).

We found a significant interaction between decision and condition ($X^2_{(2)}= 8.37$, $p = 0.01$), indicating that in the Neutral High condition only, p3 responses at the beginning of the trial were greater in Stay than in Switch trials ($X^2_{(1)}= 6.86$, $p = 0.008$), as visible in *Figure 4. 6C.* However, the interaction between decision and type of evidence was not significant ($X^2_{(1)} = 1.87$, $p = 0.17$), suggesting that the baseline difference was generalised, rather than specific to a kind of evidence. Thus, initial asymmetries in the p3 response to +Ev and –Ev cannot account for participants' decision to Switch or Stay (*Figure 4. 6C*).

However, because a difference between the p3 in response to +Ev and –Ev could be seen in the averaged amplitudes (*Figure 4.4*) and the action locked-analysis revealed a selective increase in the amplitude of +Ev in Switch trials (*Figure 4. 6B*), we further compared the final amplitudes reached by the p3 in response to +Ev and –Ev. We define *final* amplitude as the p3 amplitude in response to the *last* instances of +Ev and –Ev presented before action execution in Switch trials, or before the end of the letter stream in Stay trials. For this analysis, we run a mixed regression to investigate whether the amplitude of the last p3 differed between conditions (Inf/NH/NL), decisions (Switch/Stay) and types of evidence (+Ev/-Ev). As expected, we found a significant interaction between the decision and the type of evidence ($X^2_{(1)} = 7.94$, $p = 0.004$) indicating that the amplitude of the p3 in response to the last piece of +Ev before action was significantly higher than that in response to –Ev ($X^2_{(1)} = 17.95$, $p < 0.001$), whereas no significant difference between types of evidence was found between the last pieces of evidence presented in Stay trials (*Figure 4. 6D*).

Finally, there was no interaction between type of evidence and condition ($X^2_{(2)} = 0.19$, $p = 0.905$), indicating that the asymmetry in the processing of +Ev and –Ev reflected by the p3 right before action in Switch trials was *not* significantly different when evidence was strong, and when it was weak.

The RP can be present in exogenously-driven actions

Finally, we investigated the neural activity preceding action execution to test whether neural correlates traditionally associated with voluntary action (i.e. the RP) would be present in any of the conditions.

We run pairwise comparisons between all conditions. We found no significant clusters in any of the comparisons (all $p > 0.05$). Descriptively, an RP-like shape was visible in the pro-Switch and NL conditions, while in the NH it was either absent or severely reduced.

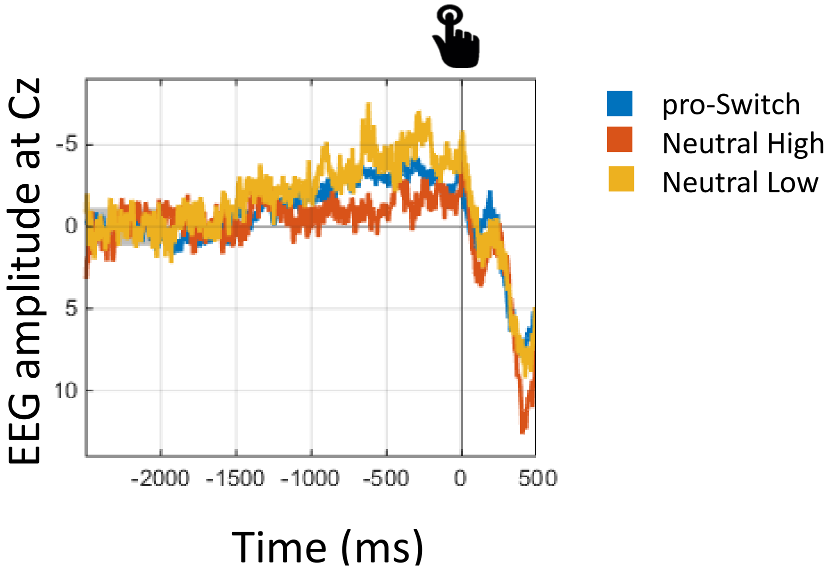

Figure 4. 7. The RP preceding evidence-informed voluntary actions. Grand-averaged EEG activity at Cz preceding Switch actions in the pro-Switch (blue), Neutral High (orange) and Neutral Low (yellow) conditions.

4.4.Discussion

In this experiment, we developed a paradigm to investigate evidence-informed voluntary action. In our task, decisions had to be made on the basis of perceptual input. However, actions preserved two aspects traditionally present in conventional voluntary action experiments: the self-paced nature of the movement, and, in the neutral conditions, the fact that decisions were underdetermined by exogenous evidence.

The p3 encodes an internal categorical decision variable

Our analysis of the p3 in response to single, discrete instances of evidence contrasts with the paradigms that recently described the p3 or CPP in decision-making processes, which used continuous, gradually changing stimuli such as contrast fluctuations (O'Connell et al. 2012). We show that evidence in favour of each available option is categorically encoded by the p3. The average amplitude of the p3 at the end of a trial tracks the decision people will make. In trials where participants decided to Switch, the amplitude of the p3 in response to +Ev was significantly greater than the p3 amplitude in response to –Ev. In turn, in Stay trials, the p3 response to –Ev was greater than that in response to +Ev, although only descriptively (see *Figure 4.4*).

Further, we found that this pattern of encoding was visible in *all* conditions – including neutral ones. In the pro-switch and pro-stay conditions, the most frequent set of targets was twice as frequent as the infrequent one. Thus, the overall evidence was very strong. Conversely, in NH and NL trials both types of letters appeared at the same frequency. Overall, there was no net evidence, and participants' actions were 'perceptual guesses'. Yet, the p3 also tracked the decision participants made. The results in all conditions are consistent with build-up to threshold models of decision-making: In all conditions, the amplitude of the p3 started low at the beginning of the trial, and reached a stable threshold before action execution in Switch trials (*Figure 4.6D*). This suggests that the p3 tracks an internal decision variable rather than the objective external evidence and suggests that evidence accumulation processes in neutral trials reflect endogenous asymmetries in the responses to +Ev and –Ev.

To investigate how these asymmetries originated, we investigated the evolution of the p3 over time. We found no systematic differences in the responses to initial presentations of +Ev and –Ev (*Figure 4. 6*C*)*, but found that the p3 amplitude selectively increased for +Ev when participants decided to Switch (*Figure 4. 6B),* the difference peaking right before action execution (*Figure 4. 6D)*,) and being visible in the averaged ERPs (*Figure 4.4*). This suggests that symmetry was not broken as a result of a baseline asymmetry in the processing of +Ev and –Ev, but rather was progressively built up throughout the course of the trial. The action-locked figures suggest a sudden, sharp increase in the +Ev p3 in the NH condition, whereas a slower ramping-up increase was visible in the NL one (*Figure 4. 6*). However, our current results cannot specify how or precisely when the asymmetry between +Ev and –Ev arises in the brain. One possibility is that trial-by-trial variations in the precise timing at which letters were presented sometimes made a few instances of +Ev appear close in time, making them seem more frequent than -Ev. Another possibility is that, given identical stimuli presentation times, stochastic fluctuations in the system resulted in a random boost in the response to one kind of evidence. Either scenario might have produced an asymmetry in neural responses that was further amplified during the trial and eventually resulted in an action.

Further analysis of this dataset may aim to tackle the question of how and when asymmetries arose by fitting the behavioural response time data to a diffusion model to investigate whether the variation in the timing of the stimuli suffices to explain the evolution of the decision variable tracked by the p3, or by analysing the effect of rhythmic fluctuations in parietal areas, which have been shown to influence the effect of new pieces of evidence in sequential evidence accumulation processes (Wyart et al. 2012).

The RP is related to the self-paced nature of movements, not their arbitrariness

Finally, we investigated the neural correlates of action initiation in a self-paced, evidence-informed perceptual decision-making task. We show that it is possible to find an RP even when actions are (partially) externally-driven. A clear RP was visible in the pro-Switch and in the NL conditions, and a smaller RP-like shape could also be observed in the NH (*Figure 4. 7*) condition. It is necessary to note that the lack of significant differences may be due to a lack of power. Since our aim was to preserve the characteristically unconstrained nature of voluntary actions, we did not want to use a 2-Alternative Forced Choice (2-AFC) task. However, because acting was optional, there is a limited number of trials where we could analyse behavioural and EEG activity preceding actions. Thus, the limited amount of trials unavoidably reduces our power to detect significant differences in the neural activity preceding actions between the three conditions, were there any. However, beyond cross-condition comparisons, the fact that an RP-like shape is visible in the informative conditions is interpretable.

The fact that an RP-like shape can be found preceding evidence-informed actions suggests that the RP does not reflect symmetry-breaking mechanisms in arbitrary decisions, as recently suggested (Maoz et al. 2017), nor is it specific to capricious, environment-independent actions. The actions in the informed condition of our paradigm were not arbitrary. Rather, they were informed by strong evidence. Our data show that the time at which participants decided to Switch corresponds to a stable peak amplitude reached by +Ev p3 responses, and hence there was clearly no symmetry to break. Thus, our results do not point to the RP and the p3 as alternative mechanisms of symmetry-breaking. Rather, a more plausible interpretation is that the p3 and the RP reflect different aspects of a synergic process where evidence is encoded, accumulated and fed to the motor system to enable action. Our results suggest that the RP may reflect gradual motor preparation processes that can result either from external evidence accumulation, such as in our evidence-informed self-paced task, or from "internal" stochastic noise accumulation in spontaneous actions. Thus, it seems that the time constraints in which movements are executed are more important determinants of the RP than the arbitrariness or exogenicity of the actions.

Chapter 5

Adaptation from within and from without: endogenous mechanisms of SAT control and perceptual changes of mind

Synopsis

The ability to adapt decision-making processes to varying exogenous environments and endogenous goals is essential for successful control of voluntary action. Here, we investigated how the p3 tracks endogenously generated variations in the speed-accuracy trade-off, and we investigated how these modulations may allow for changes of mind. For this purpose, we designed a two-step perceptual decision-making task where participants provided a first quick guess, and a second slow and deliberate decision. We found that the p3 reflects mechanisms compatible with a reduction of the baseline-to-threshold distance during urgent decision-making, and that sudden changes in the p3 amplitude track endogenous switches form an urgency to an accuracy decision regime. We further found that the p3 dynamically tracked changes of mind. Our results show that the p3 reflects a volitional process of top-down control that enables dynamic tuning of decision-making processes to suit varying endogenous goals.

5.1.Introduction

Human agents are in constant interaction with the environment, and most naturalistic actions result from the integration of endogenous and exogenous factors. In the previous chapter, we developed a paradigm to study actions that are neither purely endogenous, nor purely exogenous, and we described the neural correlates of decision-making and action initiation in contexts of strong and ambiguous external evidence. However, the ability to successfully control behaviour not only requires breaking deadlock situations, but especially the capacity to voluntarily tune decision-making processes to varying endogenous goals and adapt them to changing environments.

In this chapter, we investigate these two aspects of voluntary action control in the context of evidence-informed actions. On the one hand, we investigate how modulations of the speed-accuracy trade-off (SAT, Wickelgren 1977) are reflected in the neural correlates of evidence accumulation and action initiation processes, and we argue that the top-down modulations involved in the regulation of the SAT constitute a central mechanism for voluntary action control. On the other hand, we investigate how decisions are revised in the light of new evidence and how the decision variable encoded by the p3 dynamically tracks changes of mind in perceptual decisions.

Adjusting the speed-accuracy trade-off: a mechanism of voluntary action control

Given different environments, animals are able to adjust the time spent on decisions so that reward is optimised (Thura and Cisek 2014, 2016). Such adaptation happens spontaneously in naturalistic scenarios, and it is considered an important part of successful behavioural control. The speed-accuracy trade-off (SAT) has been defined as a ‘strategic adjustment in the decision process adapting to environmental demands’

(Heitz and Schall 2012), and it has been described in monkeys (Hanks et al. 2014; Thura and Cisek 2014, 2016), mice (Rinberg et al. 2006) and even bees (Ings and Chittka 2008). However, in humans, strategic adjustments of behaviour can result from more complex internal deliberations involving not only external rewards, but also subjective preferences and abstract goals. It can also result from simple arbitrary decisions. If told to act fast, people are able to do so, even if there is no objective benefit or reward. That is, the SAT can be modulated at will as a result of an internal decision-making process. Thus, while SAT adjustments are not always voluntarily control, they *can* be, to a certain extent, and they are therefore relevant for the study of general mechanisms of voluntary action control. Here, we investigate the neural correlates of voluntary SAT modulations in the context of evidence-informed decisions.

It is typically suggested that the mechanisms that allow fast actions involve a reduction in the distance between baseline and threshold. Recent studies in monkeys have described the effects of volitional, dynamic adjustments of the SAT at the motor level (Thura and Cisek 2014, 2016), showing that decisions made under urgency correlated with higher baseline activations in premotor and primary motor areas. In humans, several neuroimaging studies have shown increases in baseline activation of motor areas during urgent decisions (Forstmann et al. 2008; Ivanoff et al. 2008; Mansfield et al. 2011; Steinemann et al. 2018). Further, the effects of SAT modulations seem to span all levels in the sensorimotor hierarchy. In particular, urgent decisions have been shown to correlate with increased baseline firing rates in LIP neurons in monkeys (Hanks et al. 2014), and low-level encoding of sensory evidence has also been shown to be boosted in high-urgency paradigms (Steinemann et al. 2018).

In this study, we aim to gain further insight into the neural mechanisms underlying the voluntary implementation of various SATs at the level of evidence accumulation and motor preparation processes in the context of evidence-informed voluntary actions. That is, in actions that are not spontaneous and independent from the environment, but also not mere reactions to external stimuli. To this purpose, we modified the paradigm developed in chapter 4 to introduce variations in the speed-accuracy requirements of the task.

In the current experiment, participants performed the same task as in the previous chapter. That is, they monitored a letter stream and had to evaluate which out of two sets of targets (*bd* and *pq*) was most frequent. However, in the current version of the experiment, they were required to execute two actions on each trial. The first one needed to be fast, but also as accurate as possible in the light of the available evidence. The second one had to be accurate. These speed-accuracy requirements were encouraged by means of a reward scheme that will be described below, in the methods section.

Based on the previous literature, we made the following predictions. First, that the p3 responses to evidence presented during speeded decision-making should reflect patterns compatible with a reduction in the baseline-to-threshold distance, either by reducing the threshold or by increasing the baseline. Further, that the p3 amplitudes should reflect dynamic changes in the SAT. Second, we hypothesized that actions in urgent conditions would be more dependent on immediate evidence than slow actions made without time pressure. That is, that they would be more akin to reactions to external stimuli than to self-paced actions. Since the RP is typically not found preceding speeded reaction tasks (Papa et al. 1991), we hypothesized that no RP would be present preceding the urgent actions in our paradigm. In turn, following our results from chapter 4**,** we expected evidence-informed actions to be preceded by an RP.

Changing of mind: choice updating in the light of new evidence

A further important aspect of voluntary control of action is the ability to dynamically update decisions in the light of new evidence. The ability to change of mind has been described as essential to our survival and wellbeing (O'Connell and Murphy 2018) and a hallmark of cognitive flexibility (Fleming et al. 2018). Changes of mind typically improve accuracy (Resulaj et al. 2009; Kiani, Cueva, et al. 2014), and thus provide an important mechanism for behavioural optimization. Research on the mechanisms underlying changes of mind in perceptual decisions has received increasing interest in recent years (Fleming et al. 2009, 2018; Resulaj et al. 2009; Burk et al. 2014; Löffler et al. 2019). Computational models have been successful in predicting reaction times and changes of mind (Resulaj et al. 2009), and areas relevant for error monitoring have been identified as potential mediators of changes of mind in perceptual tasks (Fleming et al. 2018).

In this experiment, we aimed to investigate whether the decision variable encoded by the p3 dynamically tracked changes of mind in real-time.

In our task, participants were encouraged to make a fast guess early on in the trial. However, they had a chance to revise their decision with newly incoming evidence and change of mind in their second action. That is, sometimes, participants would choose 'Pink' in the first Action, and 'Blue' in the second one. In chapter 4, we found that the p3 categorically encodes discrete pieces of evidence. In particular, p3 amplitudes are smaller in response to evidence favouring the non-selected option than to evidence favouring the selected one. Thus, we predicted that the current paradigm should allow us to track changes of mind. In particular, we predicted that a change in the dominant p3 amplitudes (*pq<bd or bd>pq*) should index changes of mind. For example, in trials where participants chose ‘Pink’ in the first action and ‘Blue’ in the second one, we should see a greater p3 response to *pq* than to *bd* before the first action (‘Pink’), but the opposite pattern for targets presented between the first and second action (‘Blue’).

Summary

In sum, in this chapter we aim to investigate two aspects of cognitive flexibility which are important for voluntary control. First, the ability to flexibly adapt behaviour based on changing endogenous priorities (i.e. here operationalised as “acting fast” vs “being accurate”). Second, changes of mind in the context of perceptual decision-making. That is, the dynamic updating of decisions in the light of new exogenous information. We hypothesise that the p3 will be modulated by varying SAT, and that in urgent decisions the observed modulations will be compatible with a reduction in the baseline-to-threshold distance. Further, we expect that responses to different categorical alternatives will track changes of mind in real-time. Finally, we predict that the RP will be present in slow, self-paced actions, whereas no RP-like signal will precede the first, urgent actions at the beginning of the trial.

5.2. Methods

Stimuli & experimental design

Participants: Twenty-one subjects were initially invited to a single EEG session. Two participants were excluded because they did not understand the instructions. Two participants were excluded due to excessive EEG noise in the electrodes of interest (Cz & Pz). Eventually, 19 participants (10 female) were included in the study (M_{age} = 25.3, SD_{age} = 3.64; range: 21-34 years).

Procedure: Participants sat in a quiet room and were presented visual stimuli on a computer monitor. The instructions for the task were first displayed on the computer screen and then verbally repeated by the experimenter. Before the experiment, participants performed a practice version of the task (5 trials) to get familiar with the task. Participants performed a decision-making task and provided responses with left and right-hand key presses. The experiment was divided into 5 blocks, with a total of 114 trials. Every trial included a letter stream presented for 26.6 s and was followed by either one or two questions (see below). Trials were separated by a 2 s interval and participants could take a break between blocks. In total, there were 38 trials per condition, and all conditions were randomized across the experiment. The task was programmed in Python and the PsychoPy (Peirce et al. 2019) and Pandas (McKinney 2010) toolboxes.

Stimuli: A continuous stream of pseudo-randomised letters was presented at a 3.75Hz rate (duration of each letter 266 ms) on either a grey, pink or blue background. On every trial, 100 letters were presented (total duration = 26.6 s). All letters were lower-case black consonants presented without any blank interval between consecutive letters. There were two sets of task-relevant letters, to which a task-relevant colour was assigned (in parenthesis): ‘b’ and ‘d’ (blue), and ‘p’ and ‘q’ (pink). The precaution was taken not to include letters ‘h’, ‘j’, ‘k’, ‘l’, ‘t’ in the letter stream to ensure that morphological similarities between these and the target letters would not confound the results.

Participants made actions by pressing either the ‘j’ (pink) or the ‘f’ (blue) key of a standard computer keyboard with their right or left index finger respectively and answered the post-trial questions (see below) on a visual analog scale (VAS) by sliding a scroll-bar with the computer mouse.

Task: Participants were instructed to monitor the letter stream and decide which set of targets was appearing more frequently (i.e. *bd* or *pq*). Each trial started with a grey background. Participants were asked to execute two actions: a fast and a slow one. The first decision (Action 1) had to be as quick as possible, since their potential reward decreased up to the time of their first action. However, it had to be evidence-informed to a certain extent. The second decision (Action 2) had to be accurate, since if they made the wrong decision they would lose money (see *Reward* section below). Thus, they had to select the colour of the screen twice per trial. Trial duration was always the same, regardless of the timing of their second action. That is, they had to wait for all of the 100 letters to be presented before the trial ended. Trials in which participants chose the same colour twice (for the first and the second action) were labelled *congruent.* Instead, trials in which the selected colours differed between the first and the second action were labelled *incongruent* trials, and reflect changes of mind. *Figure 5. 1* shows a schematic of the experimental design.

Self-reported measures: At the end of each trial, participants were asked two follow-up questions. Participants were asked to estimate the time of their first and second actions on two VAS, and to report how confident they were on their final (i.e. second) decision on a VAS.

Experimental conditions: There were three experimental conditions randomised within blocks, each constituting one third (38) of the total amount of trials (114). Equivalent to the experiment in chapter 4, we had **Informative** conditions, and neutral conditions with either abundant instances of evidence (**Neutral High, NH**) or scarce evidence (**Neutral Low, NL**).

In the Informative condition, the most frequent set of letters would appear twice as frequently (12.5% of total letters in the stream) as the less frequent set (6.25% of letters). Thus, the perceptual decision was easy. In total, one third of the total amount of trials were Informative. Of these, half were Informative Blue ($p(bd) > p(pq)$) and the other half Informative Pink ($p(bd) < p(pq)$).

In the neutral conditions, both sets of letters were presented at the same frequency (18% of all letters were *bd* and *pq* in the NH condition, while only 14% of the total were *bd* and *pq* in the NL). In these two neutral conditions, the number of *bd* and *pq* was balanced for each trial, by the end of the trial. Thus, overall, there was no strong evidence for or against any option in any given trial. However, participants were not informed of the existence of neutral trials. Target letters were separated by a minimum of 0.5 s interval in all conditions.

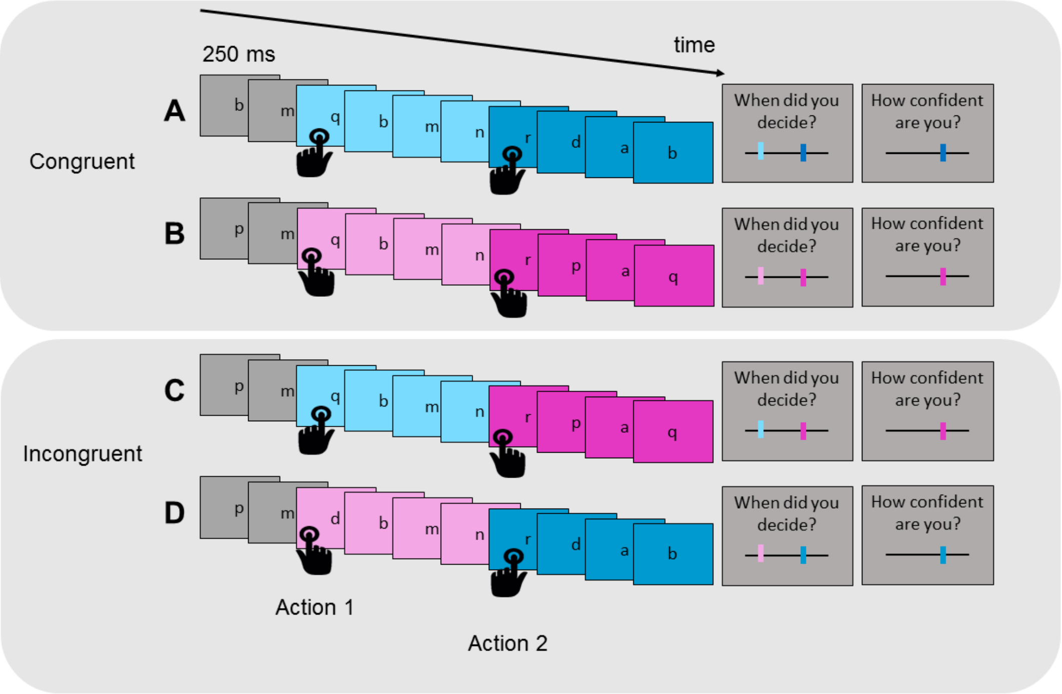

Figure 5. 1. Experimental design. Each trial started with a grey background. A letter stream of 100 letters was presented for 26.6 seconds, and participants had to decide which of the two target letter sets was most frequent (blue targets: 'b' and 'd', or pink targets: 'p' and 'q'). Their task was to make sure that the colour of the screen matched the most frequent set of targets (e.g. if 'p' and 'q' were most frequent, the screen should be pink). They were required to execute two actions. Action 1 had to be as fast as possible, since their maximal reward decreased until the time of the first action. Action 2 had to be as accurate as possible. If they made a mistake, they lost money. Participants sometimes chose the same colour both in Action 1 and Action 2 (*Congruent*), but sometimes they changed of mind (*Incongruent*). At the end of the trial, they were asked to provide an estimate of the time at which they executed each action, and a separate judgement about their confidence in their final decision (Action 2). Note that the example trials are for illustrative purposes only. In the experiment, relevant letters were always separated by at least 2 distractors, and 100 letters were presented in each trial.

Reward: The reward scheme was designed to encourage the first action (Action 1) to be fast but not completely random, and the second one (Action 2) to be as accurate as possible. *Figure 5. 2* illustrates the reward scheme with four possible types of trials. Participants could win a maximum of 25 pence in each trial. To encourage the first action to be fast, they were informed that they would lose 1 penny for each second they waited before executing the first action. However, they were also encouraged to pay attention to the evidence in that first action. Wrong actions were penalised by subtracting 12.5 p for each. Thus, in addition to the 1p/s penalisation up to the time of the first action, participants could be penalised with -12.5 p if either the first or the second action was wrong, or -25 p if both actions were wrong.

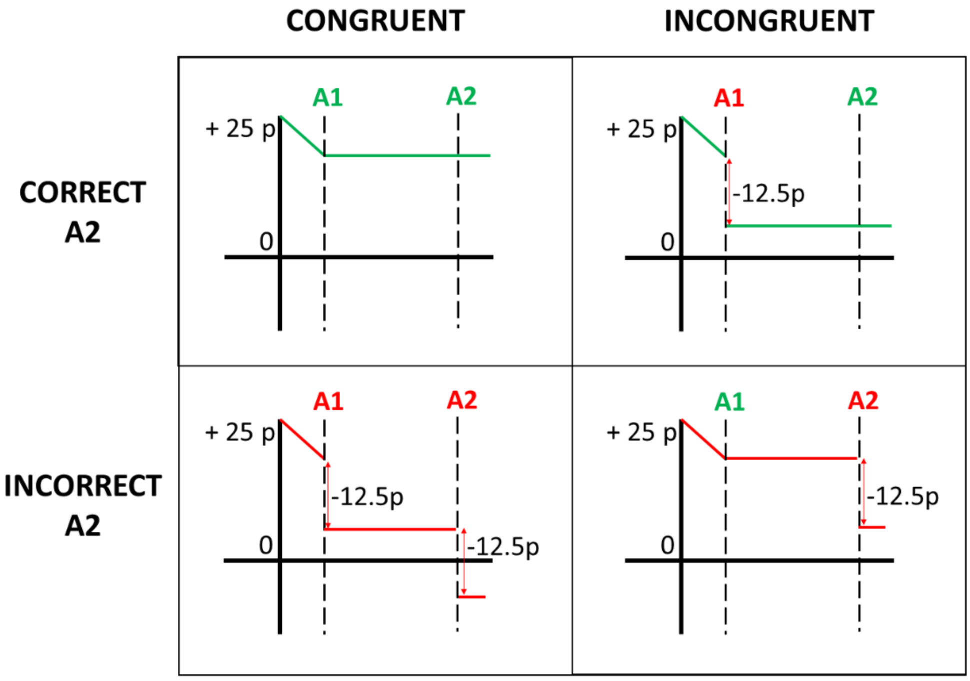

Figure 5. 2. Reward scheme. Participants could win a maximum of 25 p for each trial. One penny was subtracted from the maximum possible reward for each second they waited before Action 1 (A1). Additionally, 12.5 pence were subtracted from the reward at the time of A1 for any wrong choice. Given the same speed for Action 1 execution, Participants obtained the maximal reward if both of their actions were correct. In turn, their reward was minimal if they made the wrong choice consistently in Actions 1 and 2.

This reward scheme was only applied to informative conditions. In neutral trials, since there was no correct or incorrect choice, participants were not rewarded. If participants failed to execute two actions in any given trial, they did not obtain any reward. Participants were given averaged information about the reward obtained in each block at its end, during the breaks.

Coding scheme: We collapsed 'Pink' and 'Blue' decisions within condition, and evidence was coded with respect to participant's decisions. For example, in trials where participants chose 'Blue' *bd* are positive evidence (+Ev) and *pq* are negative evidence (–Ev). In turn, in trials where participants chose 'Pink' in both actions *pq* are +Ev and *bd* are –Ev. In incongruent trials, where participants made two different decisions in Action 1 and Action 2, the coding scheme for each specific analysis is illustrated in the results figures.

EEG recording: the same parameters as in chapter 4 were used for EEG recording.

Behavioural analysis

Primacy effect: We investigated whether the first piece of evidence participants saw in each trial biased their choices in neutral conditions. To do so, we ran a mixed logistic regression using the *glmer* function in the lme4 package (Bates, D, Mächler, M, Bolker, B.M., Walker 2015) to predict the probability of Action based on the first piece of evidence (+Ev/-Ev) on each single trial, the condition (NH/NL) and the interaction between both categorical variables. We included a random intercept in the model to account for variability between subjects.

Exogenicity analysis: We developed a measure to estimate the extent to which participant's actions were dependent on immediate external evidence. We assumed that, if participants were making a decision completely independently of the environment, the distribution of +Ev and -Ev preceding an action should uniformly distributed, as expected at random. However, if their actions are driven to some extent by immediate external evidence, the observed distribution just prior to Switch actions should deviate from the one expected at random.

To measure the magnitude of this deviation, we compared the observed distribution of –Ev (i.e. evidence *against* making a key press) and +Ev (i.e. *for* making a key press) to the expected uniform distribution during the 2.5 s preceding an action, divided in 266 ms time bins. For this comparison, we calculated a *Deviation Score* (DS) for +Ev and -Ev separately ($DS_{(+Ev)}$ and $DS_{(-Ev)}$). For each participant and condition, we subtracted the expected number of ±Ev ($\pm Ev_e$ = probability of any given letter being ±Ev * number of Switch trials) from the observed number of ±Ev ($\pm Ev_o$) in each time bin. We then divided the result by the expected number of targets to normalise the score (Generic formula: $DS = (\pm Ev_o - \pm Ev_e) / \pm Ev_e$). Positive values in the deviation score indicate that there were more targets than expected, while negative values indicate that there were less targets than expected.

Finally, we combined these two deviation scores in a single measure: the *Exogenicity Index (EI).* The Exogenicity Index was calculated by subtracting the $DS_{(-Ev)}$ from the $DS_{(+Ev)}$ (Generic formula: $EI = DS_{(+Ev)} - DS_{(-Ev)}$). The EI is thus a measure of how much the distribution of both +Ev and –Ev combined deviates from the one expected at random. The greater the EI, the greater the dependency of actions on immediate evidence. For statistical analysis, we averaged the EI in each time bin for the whole 2.5 s epoch preceding actions and run pairwise comparisons for each condition pair.

Perceived time of actions: For simplicity, we analysed the perceived temporal distance between Action 1 and Action 2 rather than the bias for each action individually. We first calculated the actual time difference between the two actions (Actual time difference = $\text{time}_{(\text{Action 2})}$ – $\text{time}_{(\text{Action 1})}$). Then, we calculated the difference between the perceived times of both actions (Perceived time difference = perceived time $_{(\text{Action 2})}$ – perceived time $_{(\text{Action 1})}$), after re-scaling the reported values on the 100 point rating scale to map the 26.6 s duration of the trial (rating/3.83). Finally, the bias measure was calculated by subtracting the perceived time difference from the actual time difference (Bias = Actual time difference – Perceived time difference). Positive values indicate that participants perceived actions as closer together in time than they actually were, whereas negative values indicate that actions were perceived as more distant in time.

EEG analysis

Preprocessing*:* EEG data were processed using Matlab R2014b (MathWorks), and EEGLAB version 13.5.4b (Delorme and Makeig 2004). First, scalp and eye electrodes were re-referenced to the average of two mastoid electrodes. Continuous EEG and EOG data were band-pass filtered between 0.01 Hz and 30Hz using an 8th order Butterworth filter. Then, data were down-sampled to 200 Hz. Second, EEG signals were epoched in two ways. For p300 analysis, EEG signal was locked from -0.5 before to 1 s after the presentation of relevant letters (‘p’, ‘q’, ‘b’ and ‘d’). For RP analysis, EEG signal was locked from -2.5 s to 0.5 s after the presentation of the letter following a keypress. Next, baseline correction was performed using the 500 ms at the beginning of the epoch [-0.5 to 0 s in p300 epochs, or -2.5 s to -2 s in RP epochs]. Then, an independent component analysis (ICA) was computed on the epoched data using the EEGLAB *runica* algorithm. Vertical eye movement components were visually identified and removed from the signal. Finally, artefact rejection was performed by removing all epochs with >120μV fluctuations from baseline in the preselected channels of interest (CZ and PZ).

p300 analysis: To study the dynamics of evidence accumulation as encoded in the p300, we extracted the p300 at PZ in response to every instance of relevant evidence (‘b’, ‘d’, ‘p’ and ‘q’), and we obtained the average amplitude of the whole duration of the component [0.3 to 0.8s post stimulus], as observed in the grand-averaged data. We used these raw values for the dynamic analysis of the evolution of the p3 amplitude over time in a stimulus-locked (i.e. in absolute time from trial onset) and action-locked manner (i.e. in relative time to execution of Action 1 and Action 2 respectively).

For average single-trial ERP analysis, we averaged the p3 component in response to +Ev and –Ev separately for each single trial, and we used these values for statistical inference. We ran a multilevel linear regression using the *blmer* function from the *blme* package (Dorie et al. 2016) using a weakly informative Wishart prior for the covariance matrix and a null prior for the fixed-effects. We fit the model to test whether the single-trial average amplitude of the p3 significantly varied between different kinds of evidence (+Ev/-Ev), conditions (Informative/NH/NL), Actions (A1/A2) and congruence conditions (Congruent/Incongruent).

We further investigated whether the amplitudes of the *first* (baseline) and *last* pieces of evidence presented from trial onset to action execution varied between different kinds of evidence (+Ev/-Ev), conditions (Informative/NH/NL), Actions (A1/A2) and congruence conditions (Congruent/Incongruent).

In all models, a random intercept was included to control for individual variability.

RP analysis: Statistical tests on averaged EEG data were run using FieldTrip toolbox (Oostenveld et al. 2011) cluster-based permutation analysis (Maris and Oostenveld

2007). The main contrast of interest involved comparing Action 1 to Action 2, for each condition separately. One electrode over medial frontal areas (CZ) was preselected for analysis. The cluster-based tests were performed on the individual participant averages using the following parameters: two-tailed dependent samples t-test, time interval = [-2 0 s relative to the event of interest], number of draws from the permutation distribution = 10000. No correction for multiple comparisons was performed.

5.3.Results

5.3.1. Behavioural results

Descriptive analysis

Accuracy: Participants performed the task well. On average, they executed two actions in >90% of trials in all conditions. In the informative conditions, their final choice was accurate on 93.55 % (*SEM* = 0.8%) of trials.

Congruency: The rate of incongruent trials differed significantly between conditions ($X^2_{(2)} = 39.36$, $p < 0.001$). Participants changed of mind significantly less often in the informative condition (*M* = 23.27%, *SEM* = 2.12%) than in the NH (*M* = 35.74%, *SEM* =3.14%; $t_{(2101)} = 5.18$, p < 0.001) and NL (*M* = 37.04%, *SEM* = 2.88%; $t_{(2101)} = 5.64$, $p<0.001$) ones, which did not significantly differ between them ($t_{(2101)} = 0.53$, $p = 0.856$; see *Figure 5. 3A*).

Response bias: In neutral conditions, participants had a small but significant bias towards selecting 'Pink' as their final choice (NH: *M* = 62.97%, *SEM* = 1.98 %*;* NL: *M* = 56.58%, *SEM* = 2.20%). We attribute this bias to the fact that the 'Pink' key was the right one in our experimental setup, and all of our participants were right-handed.

Confidence: There were significant differences between conditions ($X^2_{(2)} = 420.42$, $p < 0.001$). Participants were significantly more confident in informative (*M* = 63.79, *SEM* = 2.40) than in the NH (*M* = 48.86, *SEM* = 1.79; $t_{(2097)} = 17.10$, $p < 0.001$) and NL (*M* = 47.97, *SEM* = 1.77; $t_{(2097)} = 18.04$ = , $p < 0.001$) conditions, while the neutral conditions did not significantly differ in confidence reports ($t_{(2097)} = 1.2$, $p = 0.85$; *Figure 5. 3B*).

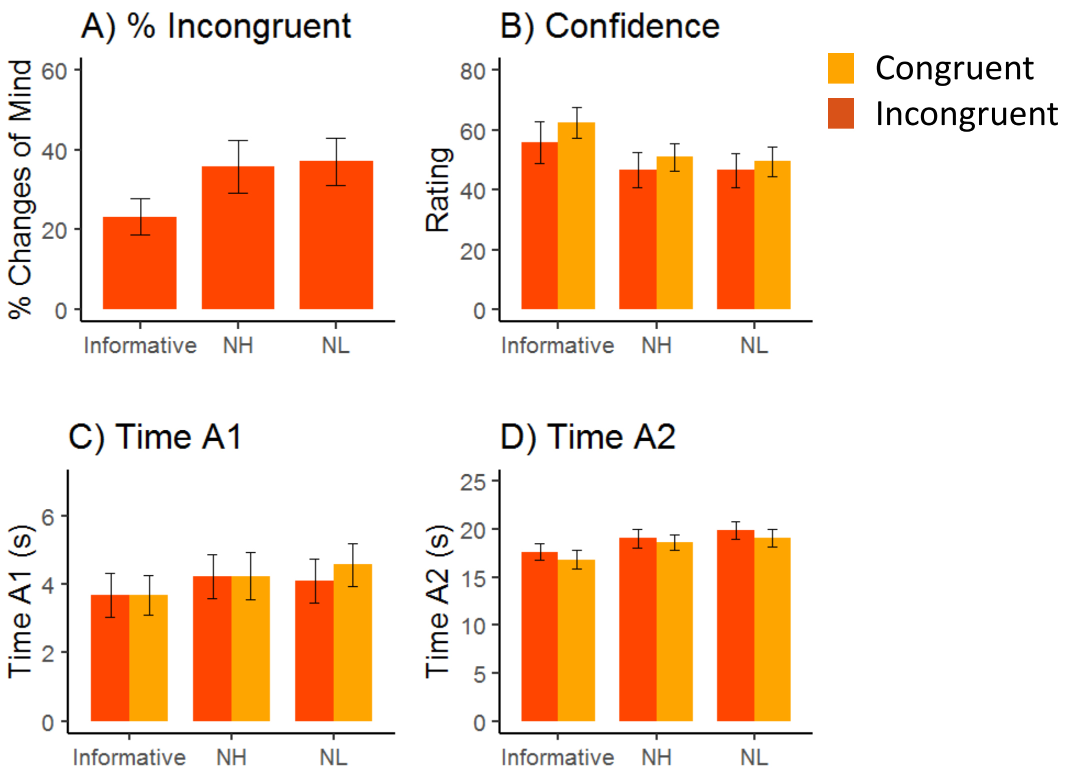

Figure 5. 3. Behavioural results. Bars indicate the grand-averaged mean (±SEM) across participants. **A)** Participants changed of mind significantly more often in the NH and NL conditions than in the informative ones. **B)** Participants were systematically more confident in Congruent than in Incongruent trials. Further, they were more confident in Informative trials than in neutral ones. **C)** In the first, urgent action (A1) there was a clear gradient of reaction times (Inf < NH < NL). There were no differences between Congruent and Incongruent trials. **D)** The second, slow action (A2) was systematically slower in Incongruent than in Congruent trials, for all conditions. As for Action 1, there was a clear gradient in RTs for Action 2 (Inf < NH < NL).

Time of action: *Figure 5. 3* illustrates the times of action. There was a significant effect of condition in the time of the first ($X^2_{(2)} = 99.38$, $p < 0.001$) and second ($X^2_{(2)} = 178.24$, $p < 0.001$) actions. Both actions in the informative condition (Action 1: $M = 3.91$ s, $SEM = 0.22$ s*;* Action 2: $M = 18.34$ s, $SEM = 0.33$ s) were faster than in the NH (Action 1: $M = 4.51$ s, $SEM = 0.24$ s*;* Action 2: $M = 20.03$ s, $SEM = 0.33$), which was in turn faster than NL (Action 1: *M* = 4.99 s*, SEM* = 0.27 s*;* Action 2: *M* = 20.76s, *SEM* = 0.33). All pairwise comparisons were significant (all $p < 0.01$, Tukey corrected).

Further, we found a significant main effect of congruency in the second action only ($X^2_{(1)} = 17.15$, $p < 0.001$), indicating that when participants changed of mind they were slower to execute the second action (*incongruent* trials: *M* = 20.08s , *SEM* = 0.30 s) than when their choice was the same as in the first action (*congruent* trials: *M* = 19.34s, *SEM* = 0.29 s).

Perceived time of action: For simplicity, we analysed the perceived temporal distance between Action 1 and Action 2 rather than the bias for each action individually. All values were negative, indicating that actions were generally perceived to be closer together in time than they actually were. We found a main effect of condition ($X^2_{(2)} = 11.23$, $p = 0.003$), indicating that the bias was significantly smaller in the Informative condition (*M* = -0.04 s, *SEM* = 0.32) than both in the NH (*M* = -0.66 s, *SEM* = 0.29; $t_{(2097)} = 3.43$, $p = 0.001$) and NL (*M* = -0.40 s, *SEM* = 0.32; $t_{(2097)} = 2.39$, $p = 0.043$) conditions, which did not differ between them ($t_{(2097)} = 1.00$, $p = 0.573$). We found no significant differences in the temporal bias in different congruence conditions ($X^2_{(1)} = 0.406$, $p = 0.524$).

Primacy effect analysis: We investigated whether the first piece of evidence participants saw in each trial biased their decisions for congruent and incongruent trials. We ran a logistic regression to predict Action 2 (Pink/Blue) based on the condition (NH/NL), congruence (Incongruent/Congruent) and first letter observed on that trial (*bd/pq*).

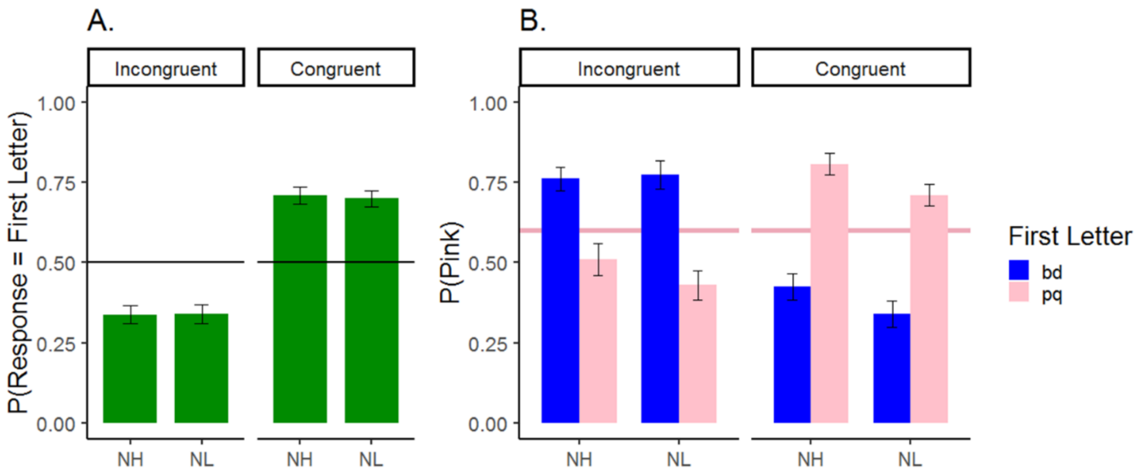

Figure 5. 4. Primacy effect. **A)** Grand-averaged (±SEM) percentage of Action 2 decisions that were consistent with the first instance of evidence presented on that trial. In congruent trials, participants were significantly more likely to choose the option corresponding to the first letter they saw both in Neutral High (NH) and Neutral Low (NL) conditions. In incongruent trials, participants initially chose the option suggested by the first piece of evidence, but then reversed their decision. Note: black horizontal line indicates proportion of responses congruent with the first instance of evidence expected at random. **B)** Probability of choosing 'Pink' in Action 2 given the first letter observed on that trial. In congruent trials, participants were significantly more likely than on average (*pink* line) to choose 'Pink' when the first letter they saw on that trial was a *p* or a *q* (*pink* bars), and significantly less likely to choose 'Pink' when the first letter they saw was a *b* or a *d* (*blue* bars). Note: the pink line reflects the percentage of trials where participants chose 'Pink' in Action 2.

We found a significant interaction between congruence and first letter ($X^2_{(1)} = 146.03$, $p < 0.001$). In congruent trials, participants' final choice was in agreement with the first letter that they saw on that trial significantly more often than expected at chance. Instead, in incongruent trials, participants initially chose an action compatible with the first piece of evidence that they saw, but they later reversed this decision and their choice at Action 2 was the opposite to the one suggested by the first piece of evidence in the trial (*Figure 5. 4*).

Exogenicity analysis: To compare the extent to which Action 1 and Action 2 were driven by immediate evidence, we compared the Exogenicity Index between Conditions (Inf/NH/NL) and Actions (A1/A2) in both congruence conditions (Congruent/Incongruent). *Figure 5. 5* summarizes the results. *Figure C1* and *Figure C2* provide the distribution histograms and deviations over time.

We found a main effect of condition ($X^2_{(2)} = 42.87$, $p < 0.001$), indicating that actions in the Informative condition ($M = 0.61$, $SEM = 0.12$) were less influenced by immediate evidence than those in both the NH ($M = 1.38$, $SEM = 0.13$; $t_{(93)} = 5.53$, $p < 0.001$) and NL ($M = 1.22$, $SEM = 0.13$; $t_{(93)} = 3.72$ $p = 0.001$) conditions. Neutral conditions were not significantly different from each other ($t_{(93)} = 1.48$, $p = 0.301$).

Further, we found an interaction between Action and Congruency ($X^2_{(1)} = 21.94$, $p < 0.001$). In congruent trials, the Exogenicity Index was significantly higher ($X^2_{(1)} = 72.91$, $p < 0.001$) in Action 1 ($M = 1.83$, $SEM = 0.13$) than Action 2 ($M = 0.64$, $SEM = 0.10$). That is, in congruent trials, urgent actions were more strongly driven by immediate external evidence than slow, accurate ones. Instead, in incongruent trials, the Exogenicity Index did not significantly differ ($X^2_{(1)} = 1.04$, $p = 0.306$) between Action 1 ($M = 0.90$, $SEM = 0.13$) and Action 2 ($M = 0.64$, $SEM = 0.14$). That is, there was no difference in the degree to which immediate evidence was responsible for triggering Action 1 and Action 2.

This difference between congruence conditions was driven by Action 1. The EI preceding Action 1 was significantly larger in congruent compared to incongruent trials ($X^2_{(1)} = 35.51$, $p < 0.001$), while the EI preceding Action 2 did not differ ($X^2_{(1)} = 0.02$, $p = 0.883$). Thus, in incongruent trials there was less evidence preceding Action

1 than there was in congruent ones. In other words: when participants changed of mind, their initial quick choice was driven by less immediate evidence.

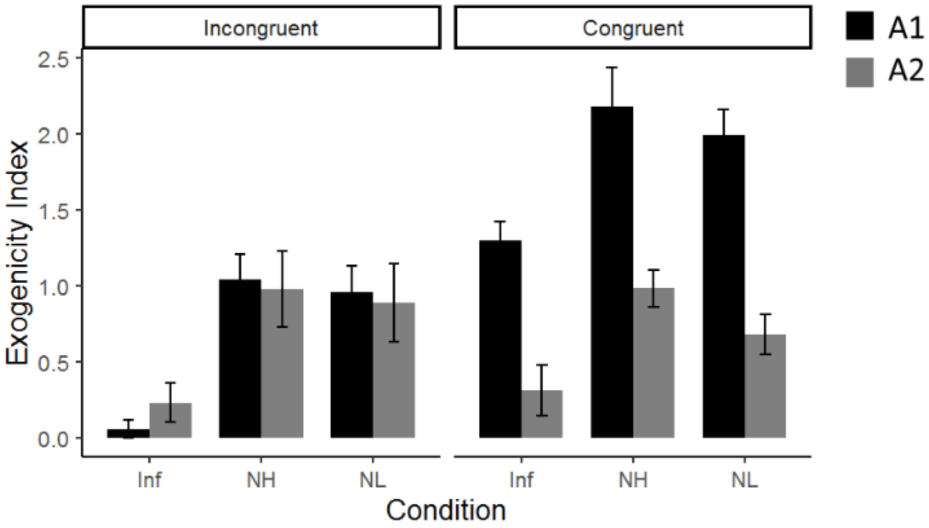

Figure 5. 5. Exogenicity Index. In Congruent trials, urgent actions (A1, *black*) were more influenced by immediate evidence than slow actions (A2, *grey*). In incongruent trials, there was no difference in the dependence of immediacy between the two types of actions. Urgent actions in congruent trials were *more* informed by immediate evidence than those in incongruent trials.

5.3.2. EEG results

We first investigated whether participants' decisions were categorically encoded by the p3 both before urgent Action 1 and slow Action 2. To investigate this, we ran a linear regression to test whether the average amplitude of the p3 differed between each type of evidence (+Ev/-Ev), condition (Inf/NH/NL), congruence (Congruent/Incongruent) on the averaged amplitude of each single trial, for each action respectively.

We found a three-way interaction between congruence, condition and action ($X^2_{(2)} = 6.00$, $p = 0.04$), and we followed it up with separate regressions for congruent and incongruent trials.

The p3 categorically encodes decisions in congruent trials

Figure 5. 6 illustrates the grand-averaged p3s preceding Action 1 and Action 2 in congruent trials. As expected, we found two main effects. First we found, a main effect of evidence type ($X^2_{(1)} = 9.09$, $p = 0.002$) showing that, on average, the p3 in response to +Ev ($M = 3.73$ µV, $SE = 0.17$ µV) had bigger amplitudes than the p3 in response to -Ev ($M = 2.72$ µV, $SE = 0.21$ µV) in both actions. This pattern was observable when looking at the Pink/Blue decisions separately as well (see *Figure C3*). Second, we found an effect of Action ($X^2_{(1)} = 33.35$, $p < 0.001$), indicating that the average amplitude of the p3s preceding the Action 1 were higher than those preceding Action 2. However, we found no interaction between Action and evidence type ($X^2_{(2)} = 0.62$, $p = 0.42$) suggesting that the p3 amplitude categorically encoded decisions both before the urgent Action 1 and the slow, deliberate Action 2.

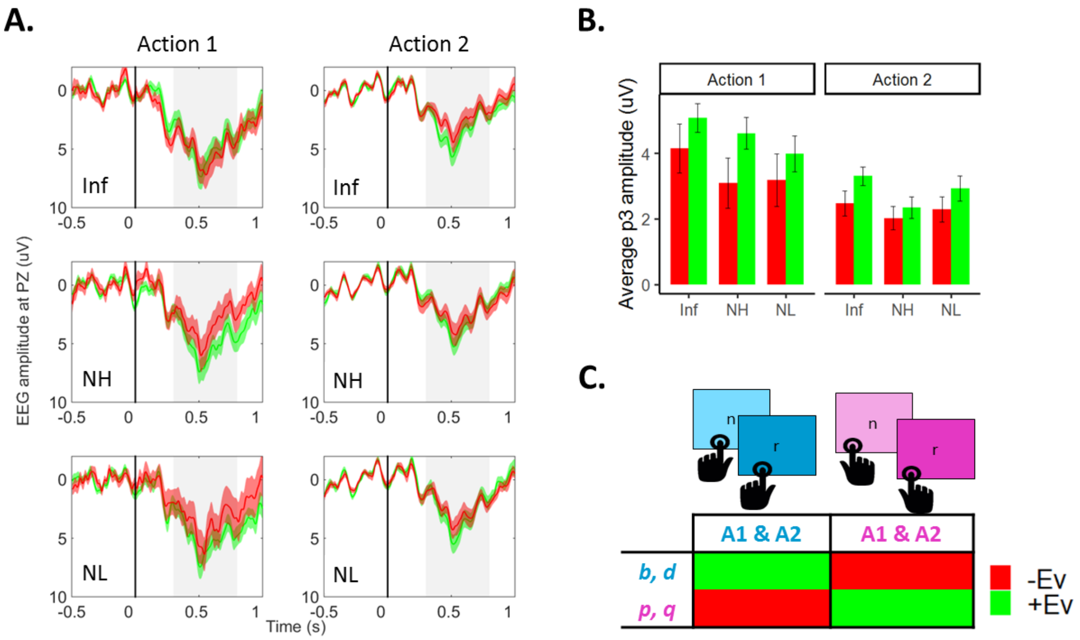

Figure 5. 6. The p3 categorically encodes decisions in congruent trials. **A)** Grand-averaged (±SEM) ERP traces at Pz in response to all +Ev (*green*) and –Ev (*red*) up to the time of Action 1 (*left*) and between Action 1 and Action 2 (*right*) in Informative (Inf, *top*), neutral high (NH, *middle*) and neutral low (NL, *bottom*) conditions, locked to letter onset. The shaded area indicates the time period used for statistical inference (300 to 800 ms from stimulus onset). **B)** Grand-average (±SEM) of single-trial mean p3 amplitude between 300 and 800 ms, used for statistical analysis. **C)** Coding scheme. Evidence is coded with respect to the decision participants made. In trials where participants chose 'Blue' in both Action 1 (A1) and Action 2 (A2), *bd* are +Ev and *pq* are –Ev. In turn, in trials where participants chose 'Pink' in both actions *pq* are +Ev and *bd* are –Ev.

The p3 in incongruent trials

Next, we investigated incongruent trials. In these, two different decisions were made in a single trial (e.g. Action 1: Blue, Action 2: Pink). To investigate whether the p3 tracked each of the two decisions we ran a linear regression to test whether the amplitude of the p3 differed given the type of letter shown (+Ev/-Ev), the condition (Inf/NH/NL), and the action (A1/A2) that it preceded. *Figure 5. 7* illustrates the grand-averaged p3 components preceding Action 1 and Action 2 in incongruent trials.

We found an interaction between Action and Condition ($X^2_{(2)} = 6.20$, $p = 0.045$). Thus, we investigated each condition with a separate regression. We found that, in the NH condition only, there was an interaction between the type of evidence and the action ($X^2_{(2)} = 4.38$, $p = 0.036$). This indicated that the p3 amplitude tracked the decisions made at A1 and A2. While the amplitude of the p3 in response to +Ev (i.e. here, evidence favouring the 2nd decision) increased from A1 to A2, the amplitude of the p3 in response to –Ev (i.e. here, evidence countering the 2nd decision) decreased (from A1 to A2). The difference between +Ev and –Ev was marginally significant before Action 1 ($X^2_{(1)} = 3.58$, $p = 0.05$).

In both the Informative ($X^2_{(1)} = 6.75$, $p = 0.009$) and the NL ($X^2_{(1)} = 8.05$, $p = 0.004$) conditions, we found an effect of Action indicating that the amplitudes of the p3 were significantly higher preceding Action 1 than they were, on average, between Action 1 and Action 2.

In contrast to the congruent trials, different types of evidence were not significantly different before Action 1 or between Action 1 and Action 2.

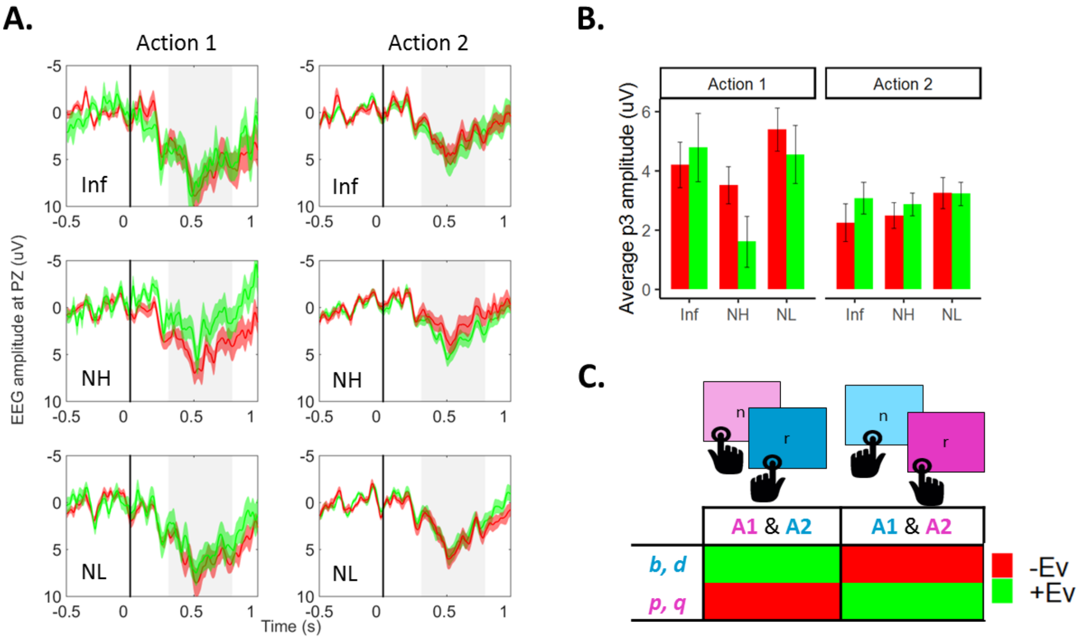

Figure 5. 7. The p3 in incongruent trials. **A)** Grand-averaged (±SEM) ERP traces at Pz in response to all +Ev (*green*) and –Ev (*red*) up to the time of Action 1 (*left*) and between Action 1 and Action 2 (*right*) in Informative (Inf, *left*), neutral high (NH, *middle*) and neutral low (NL, *right*) conditions, locked to letter onset. The shaded area indicates the time period used for statistical inference (300 to 800 ms from stimulus onset). **B)** Grand-average (±SEM) of single-trial mean p3 amplitude between 300 and 800 ms, used for statistical analysis. **C)** Coding scheme. Letters are coded with respect to the final decision, both in A1 and A2. In trials where participants chose ‘Pink’ in A1 and ‘Blue’ (*left*) in A2, *bd* are +Ev and *pq* are -Ev. In trials where participants chose ‘Blue in A1 and ‘Pink’ in A2 (*right*), *bd* are -Ev and *pq* are +Ev.

Stimulus-locked dynamic analysis of the p3

In order to investigate the dynamics of the variable producing the above-described average effects and the general effects of endogenous goals on the neural correlates of evidence accumulation, we investigated the evolution of the p3 amplitude over the course of the trial for both congruent and incongruent trials. Descriptively, we observed a non-linear evolution of the p3 amplitude in both types of trials from the time of trial onset. At the beginning of the trial, before Action 1, the amplitude started relatively high and then decreased over time. Towards the middle of the trial, the p3 amplitude started ramping up again up until the time of Action 2 (see *Figure 5. 8*).

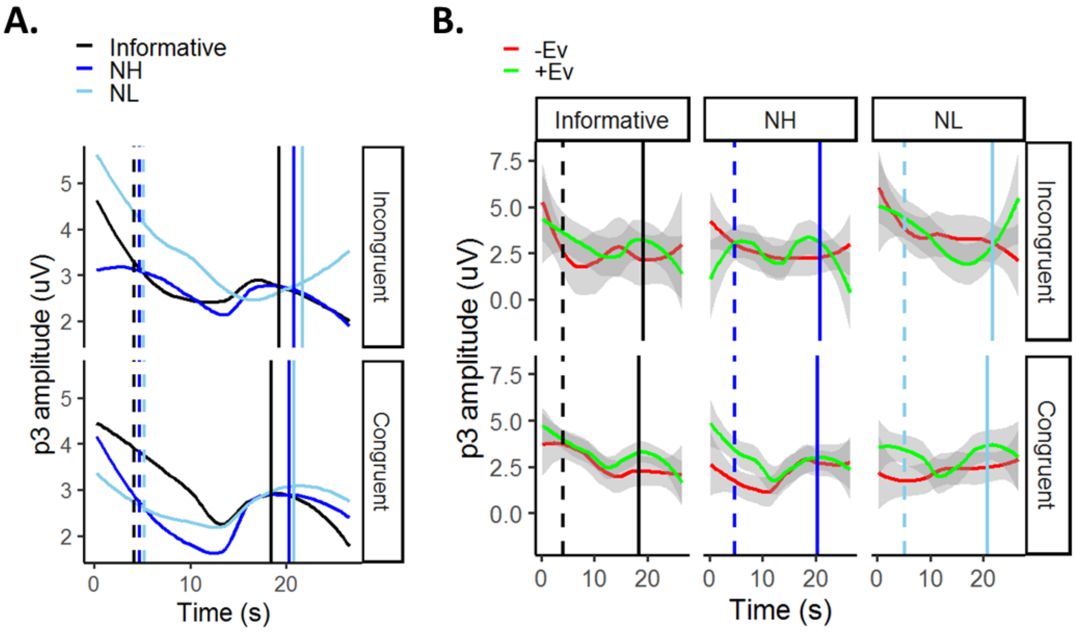

Figure 5. 8. Stimulus-locked evolution of the p3 amplitude throughout the trial. Vertical lines indicate the mean time of action (Action 1, *dashed;* Action 2, *solid*). The graded differences in decision times between conditions are clearly visible for both Actions (RT Inf<NH<NL). Further, RTs show that actions were executed when the internal decision variable reached a peak. **A)** Evolution of the p3 amplitude, pooled across participants and types of evidence. A non-linear evolution is visible in all conditions. **B)** Evolution of the p3 amplitude, pooled across participants, for each type of evidence.

To investigate these action-dependent dynamics, we investigated the evolution of the p3 relative to the first and second actions and ran a linear regression to predict the evolution of the p3 based on time to action (relative time), condition (Inf/NH/NL), evidence type (+Ev/-Ev) and congruence (Congruent/Incongruent), for Action 1 and Action 2 separately.

Evidence accumulation preceding fast, urgent Action 1

Figure 5. 9 illustrates the action-locked dynamics of the p3 preceding the first urgent action. We found a three-way interaction between condition, congruence and time to action ($X^2_{(2)} = 8.15$, $p = 0.016$), and another three-way interaction between condition, evidence type and time to action ($X^2_{(2)} = 10.58$, $p = 0.005$).

We only found effects of time in the Informative condition. In incongruent trials, a main effect of time indicated that the amplitude of both +Ev and –Ev increased towards the time of the first action ($X^2_{(1)} = 4.06$, $p = 0.04$). In Congruent trials, instead, the effect of time depended on the type of evidence ($X^2_{(1)} = 6.77$, $p = 0.009$), and only significantly increased for +Ev ($X^2_{(1)} = 4.57$, $p = 0.03$).

Since in the neutral conditions there was no objective evidence to be accumulated, and no increase in the amplitude of the p3 was observed towards the time of the first action, we reasoned that the averaged effect must be due to an asymmetry in the p3 responses to +Ev and –Ev. To investigate this, we compared the amplitudes of the initial responses to +Ev and –Ev (i.e. the p3 amplitude in response to the *first* instance of +Ev and –Ev presented on any given trial) based on the condition (Inf/NH/NL), congruence (Congruent/Incongruent) and type of evidence (+Ev/-Ev) with respect to the following action.

We found a main effect of evidence type ($X^2_{(1)} = 7.83$, $p = 0.005$), indicating that p3 responses to the first instance of +Ev were significantly higher than those in response to –Ev. We found no interactions with any other factor. We further investigated whether the amplitudes reached at the time of Action 1 continued to reflect this asymmetry by comparing the amplitudes of the *last* instance of +Ev and –Ev presented. That is, the final amplitude reached by each type of evidence. We found that the effect of evidence type was no longer significant just before Action 1 ($X^2_{(1)} = 3.51$, $p = 0.06$).

These results suggest that an initial asymmetry, rather than the final amplitude reached before Action 1 accounts for the observed average effects in the p3 responses preceding Action 1 and participants' decisions.

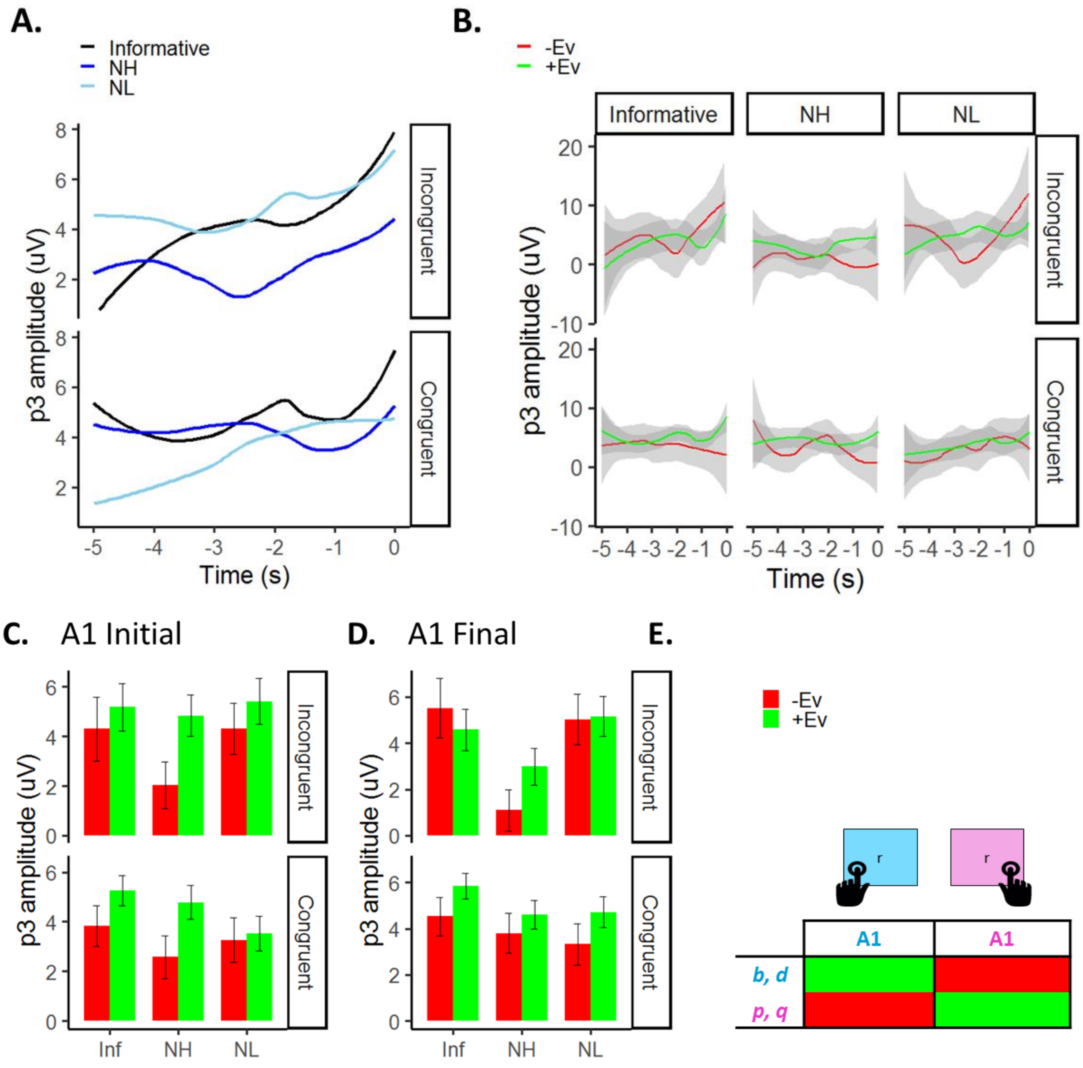

Figure 5. 9. Evolution of the p3 locked to Action 1. **A)** Evolution of the p3 amplitude at Pz, pooled across participants and types of evidence. A non-linear evolution is visible in all conditions. **B)** Evolution of the p3 amplitude at Pz (± SE), pooled across participants, for each type of evidence separately. A non-linear evolution is visible in all conditions. Note: for illustration purposes, only 5 s preceding Action 1 are plotted, but letters outside of those ranges were also included in the regression analysis. **C)** In all conditions, both for congruent and incongruent trials, the average p3 response to the first piece of +Ev was significantly higher than that to –Ev. Bars indicate p3 amplitude at Pz (± SE) in response to the first piece of +Ev and –Ev, pooled across participants. **D)** The differences between +Ev and –Ev observed at the beginning of the trial (C) were reduced or even inverted in some conditions. Bars indicate p3 amplitude at Pz (± SE) in response to the last piece of +Ev and –Ev before Action 1. **E)** Coding scheme. Blue and Pink conditions are pooled, and letters are coded with respect to Action 1.

Evidence accumulation preceding slow, deliberate Action 2

To investigate the evolution of the p3 leading to Action 2, we then analysed the p3 amplitude in response to letters appearing *between* Action 1 and Action 2, using a regression model with the same factors used in the previous analysis.

We found a significant effect of time in the evolution of the p3 preceding Action 2 ($X^2_{(1)} = 11.95$, $p < 0.001$), indicating that the amplitude of the p3 *increased* over time up to the time of the second action, as visible in *Figure 5. 10*. We found no significant interactions of time with any other factor.

To investigate how this evidence accumulation process unfolded over time, we further compared the evidence accumulation processes by investigating initial responses to +Ev and –Ev just after Action 1 (i.e. p3 amplitude in response to the *first* piece of +Ev and –Ev presented after Action 1) and the final amplitudes reached just before Action 2 (i.e. p3 amplitude in response to the *last* piece of +Ev and –Ev presented before Action 2).

In the analysis of initial amplitude, we found an interaction between congruence and condition ($X^2_{(2)} = 7.06$, $p = 0.029$). While in in incongruent trials there were no significant differences in the p3 amplitude just after Action 1 between conditions, in congruent trials the informative condition ($M = 3.07$ μV, $SE = 0.54$) showed higher initial amplitudes than the NH condition ($M = 1.46$ μV, $SE = 0.56$ *;* $p = 0.01$). We found no interactions with any other factor, nor a main effect of evidence type. This suggests that there were no instantaneous baseline asymmetries immediately after Action 1.

Instead, in the analysis of the final amplitudes, we found that the last piece of +Ev produced significantly greater p3 amplitudes than the last piece of -Ev ($X^2_{(1)} = 17.48$, $p < 0.001$).

In sum, contrary to the pattern observed in Action 1, there were no significant differences at the time the evidence accumulation for Action 2 started – assuming that it started just after Action 1. However, differences were significant just before execution of Action 2. This suggests that decisions in Action 2 were driven by the *final* differences present at the time of action, rather than *initial* asymmetries just after Action 1.

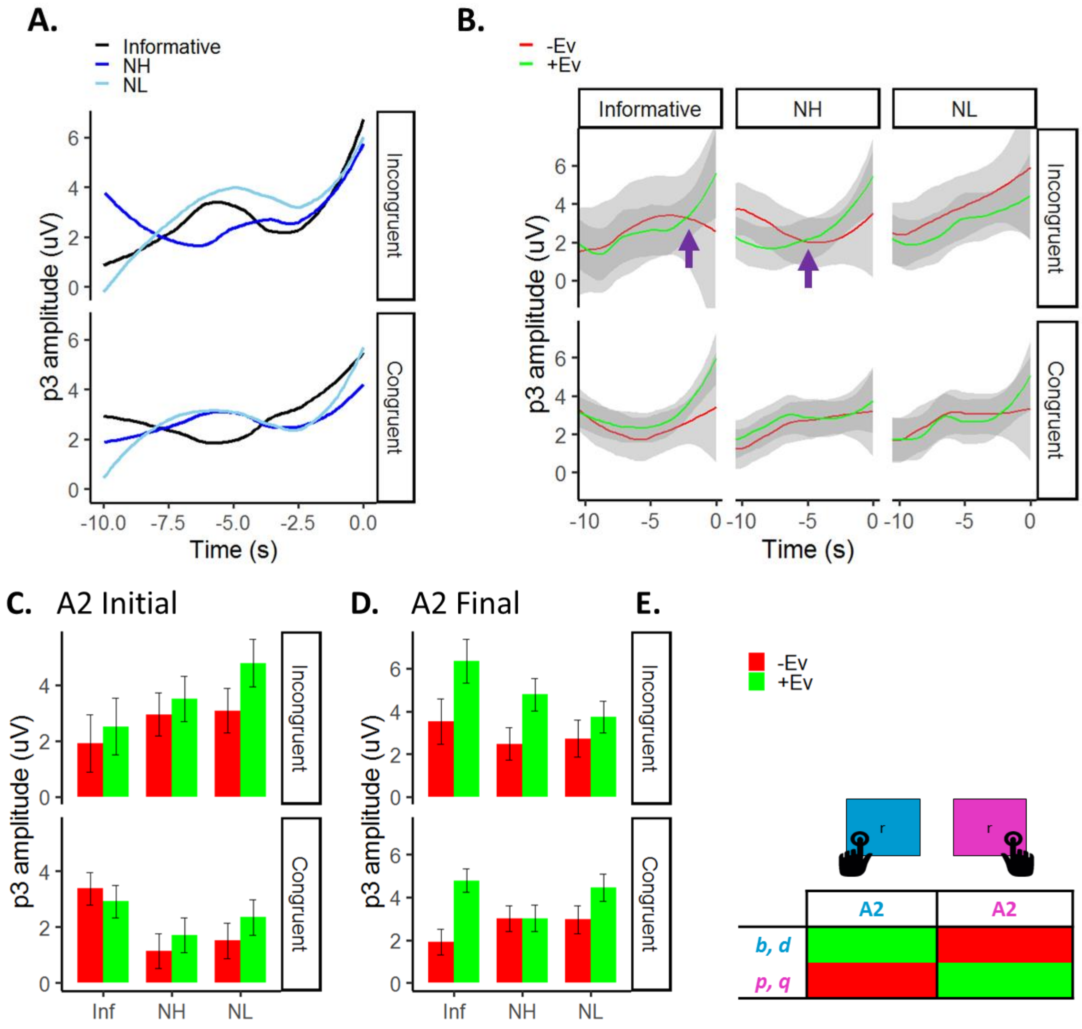

Figure 5. 10. Evolution of the p3 locked to Action 2. **A)** Evolution of the p3 amplitude at Pz, pooled across participants and types of evidence. **B)** Evolution of the p3 amplitude at Pz (± SE), pooled across participants, for each type of evidence separately. Purple arrows indicate a descriptive marker of changes of mind: the dominant stimulus (i.e. letters with the biggest p3 response) changed before the time of the second action. Note: for illustration purposes, only 10 s preceding Action 2 are plotted, but letters outside of those ranges were also included in the regression analysis. **C)** There were no significant differences in the initial responses to +Ev and –Ev, although descriptively +Ev elicited larger p3 responses than –Ev. Bars indicate p3 amplitude at Pz (± SE) in response to the first piece of +Ev and –Ev after Action 1. **D)** The differences between +Ev and –Ev generally increased, compared to baseline. Bars indicate p3 amplitude at Pz (± SE) in response to the last piece of +Ev and –Ev before Action 2. **E)** Coding scheme. Letters are coded with respect to Action 2.

Changes of mind

Based on the results from the previous chapter 4, we hypothesized that the p3 would encode a categorical variable and would dynamically track participants’ decisions.

We predicted that, if the p3 is the neural signature of an internal decision variable, a reversal in its “sign” should be visible during changes of mind. In incongruent trials, such a reversal is descriptively visible in the action-locked analysis in informative and NH conditions (see *Figure 5. 10B, purple arrows*). This indicates that the “dominant evidence” (i.e. the type of evidence that produces the largest p3 amplitudes) changes from the time of Action 1 to the time of Action 2.

Because modulations in +Ev/-Ev amplitude could occur at different times across trials, we directly compared the initial and final amplitudes of the p3 in response to different types of evidence (+Ev/-Ev), in all conditions (Inf/NH/NL). This corresponds to a direct comparison between the already reported baseline for Action 1 (*Figure 5. 9C)* and threshold for Action 2 (*Figure 5. 10D).* For illustrative purposes, we plot them together in a single figure below (*Figure 5. 11*). As expected, we found a significant interaction between the type of evidence and the time in the trial ($X^2_{(1)} = 9.21$, $p = 0.002$), indicating that the dominant amplitudes reversed (e.g. if $bd>pq$ at the beginning of the trial, $bd<pq$ at the end, before Action 2).

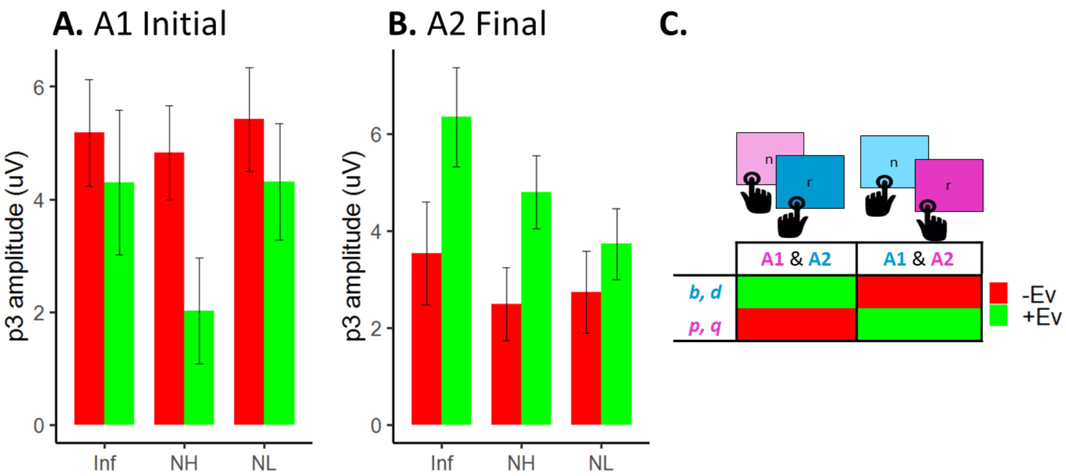

Figure 5. 11. The initial and final amplitudes of the p3 track changes of mind. Reversal of dominant p3 amplitude in changes of mind (i.e. incongruent trials). When participants changed of mind, the difference in the p3 components in response to different types of evidence reversed from trial onset (**A**) to the time of the second action (**B**). This is a neural marker of a perceptual change of mind. **C)** Illustrates the coding scheme. Evidence is coded with respect to the final action. Note that this the colour discrepancy between A and *Figure 5.10C*, which illustrates the same data with a different colour coding.

The p3 tracks dynamic modulations in the SAT conditions

Our results so far suggest that while urgent decisions are based on an asymmetry in the initial responses to +Ev and –Ev, slow deliberate actions are determined by the amplitude differences reached at the end of the evidence accumulation process. We investigated this by comparing baseline and threshold p3 responses within each decision process separately. However, we have not yet directly compared the two decision making processes.

Accumulation-to-bound theories of decision-making suggest that decisions made in urgent scenarios should exhibit modulations that result in a reduction in the baseline-to-threshold distance. This can be achieved either by increasing the baseline activation, lowering the threshold, or both. Here, we directly compare baseline and threshold activations between Action 1 and Action 2. We ran a mixed regression to predict the baseline and threshold amplitudes of the p3 based on condition (Inf/NH/NL), congruence (Congruent/Incongruent), Action (Action 1/Action 2) and type of evidence (+Ev/-Ev).

We found a three-way interaction between congruence, action and type of evidence ($X^2_{(2)} = 8.28$, $p = 0.01$), suggesting that the dynamics of congruent trials and trials containing changes of mind were different.

In congruent trials, we found that both initial and final p3 amplitudes were significantly higher in Action 1 than in Action 2 ($X^2_{(1)} = 21.86$, $p < 0.001$; initial p3 in Action 1: $M = 4.38$ µV, $SE = 0.31$; initial p3 in Action 2: $M = 2.54$ µV , $SE = 0.29$; $X^2_{(1)} = 8.77$, $p = 0.003$; final p3 in Action 1: $M = 4.66$ µV, $SE = 0.29$; final p3 in Action 2: $M = 3.36$ µV, $SE = 0.25$). This suggests that urgent decision-making processes started at higher baseline activations, and reached higher final amplitudes than slow, deliberate actions. No interactions were found, suggesting that the modulations happened across conditions, for all types of evidence (*Figure 5.* 12*)*.

In incongruent trials, there was an interaction between Action and type of evidence ($X^2_{(1)} = 6.53$, $p = 0.01$), indicating that while the baseline response to –Ev (i.e. evidence against the second decision) was higher in Action 1 ($M = 5.14$ µV, $SE = 0.52$ µV) than in Action 2 ($M = 2.76$ µV, $SE = 0.49$ µV), the baseline response to +Ev (i.e. evidence for the second decision) was not significantly different between those actions. In other words, while the response to –Ev decreased from trial onset to the time of Action 1, the response to +Ev remained the same.

Further, we found that the final amplitudes reached before each Action interacted with the condition ($X^2_{(2)} = 7.84$, p = 0.01) and the type of evidence ($X^2_{(2)} = 5.88$, p = 0.01).

Negative evidence (i.e. here, evidence for the first decision in Action 1, and against the second decision in Action 2) had higher amplitudes at the time of Action 1 (*M* = 4.18 μV , *SE* = 0.5 μV) than at the time of Action 2 (*M* = 2.83 μV, *SE* = 0.50 μV), on average, for all conditions. Instead, in NH trials only, the amplitude reached by +Ev at the time of Action 2 (*M* = 1.10 μV, *SE* = 0.90 μV) was significantly higher ($X^2_{(1)}$ = 10.24, *p* = 0.001) than that at the time of Action 1 (*M* = 4.79 μV, *SE* = 0.75 μV), while the +Ev at the time of Action 1 and 2 was the same in the NL and informative conditions. This interaction effect is illustrated by grey and black arrows in *Figure 5.* 12.

These results will be interpreted in detail in the discussion section.

From speed to accuracy: endogenous resetting of the SAT

Our design allowed us to further investigate the way in which the transition from an urgency to an accuracy decision-making regime is implemented at the neural level and reflected in the internal decision variable encoded by the p3.

Based on the pooled dynamics observed in the stimulus-locked analysis *Figure 5. 8,* we expected the amplitude of the p3 to decrease immediately after the execution of the first action. To test this, we compared the amplitude of the *last* p3 *before* execution of Action 1 to the amplitude of the first p3 *after* movement on every single trial. We run a linear regression including the position of the p3 with respect to the action (Before/After), the condition (Inf/NH/NL), the type of evidence (+Ev/-Ev) and the congruence of the trial (Congruent/Incongruent).

The interaction between congruence, action and condition was significant ($X^2_{(2)}$ = 8.59, *p* = 0.01).

In congruent trials, we found that the amplitude of the last p3 before Action 1 was significantly higher than that preceding Action 2 ($X^2_{(1)} = 39.23$, $p < 0.001$), suggesting that, as illustrated in *Figure* 5. 12 the amplitude of the p3 responses dropped immediately after execution of Action 1. We did not find any interactions with other factors. We refer to this effect as "baseline resetting": after an initial quick decision, baseline responses to evidence are tuned down and effectively re-set the initial conditions for the second, slow evidence accumulation process.

In incongruent trials, we found that the effect varied between conditions ($X^2_{(2)} = 9.35$, $p = 0.009$). In fact, the amplitude of the p3 only significantly decreased in the Informative condition ($X^2_{(1)} = 7.42$, $p = 0.006$). Instead, the p3 amplitudes in response to +Ev and –Ev showed opposing trends in the NH and NL condition, where either the –Ev p3 amplitude remained stable and +Ev one increased (in the NH condition), or the +Ev p3 remained stable but the –Ev p3 decreased (in the NL condition).

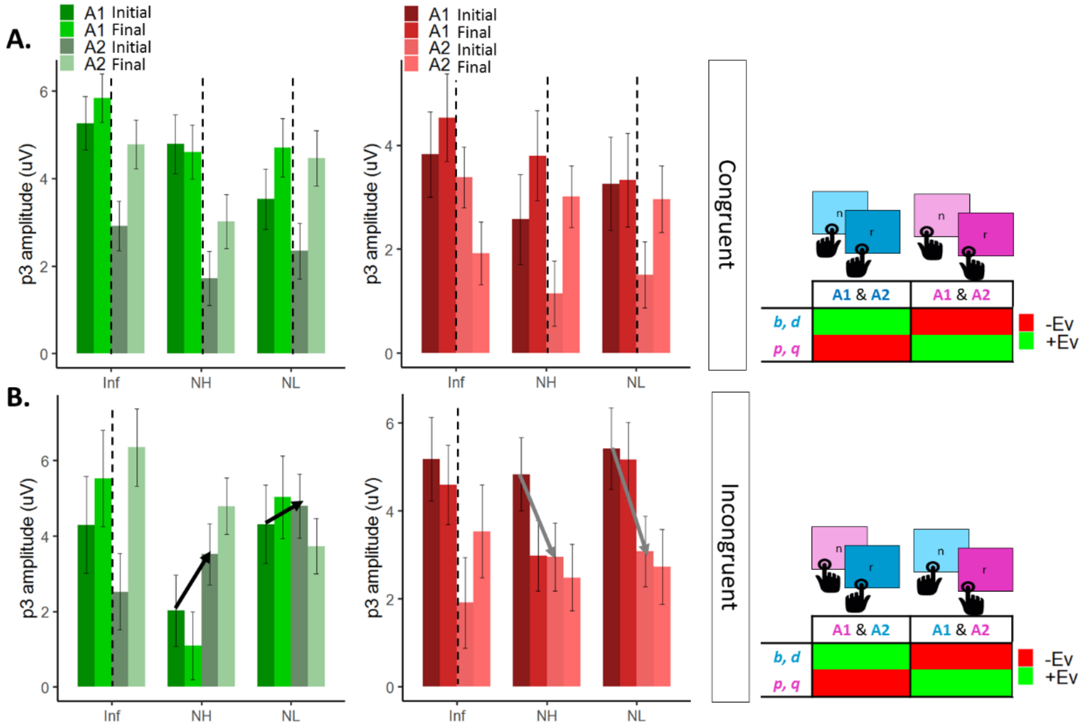

Figure 5. 12. The p3 tracks endogenous changes in the SAT conditions. The bars illustrate the initial and final amplitudes of the p3 in Actions 1 and 2 for each condition and type of evidence (+Ev, *green;* -Ev, *red*) separately. Darker shades correspond to Action 1, lighter shades to Action 2. **A)** In congruent trials (*top*), both the initial and final activities were significantly higher for the urgent Action 1 than for the slow, deliberate Action 2, in all conditions and for all types of evidence. The transition from an urgent to an accurate mode of action is reflected in a “baseline resetting” process, whereby the p3 amplitude drops just after the first Action was executed (*dashed line*). **B)** In incongruent trials, initial and final differences depend on condition and type of evidence. Baseline resetting occurred only in informative conditions (*dashed line*). In neutral conditions, a comparison between initial activities at Action 1 and Action 2 showed that while the amplitude of the p3 in response to -Ev decreased between trial onset and A2 baseline (*grey arrows*), the +Ev p3 remained stable (*black arrows*). Note: Evidence is coded with respect to the final action (A2). Initial refers to the *first* piece of evidence presented at trial onset or right after action 1 (for Action 1 and Action 2 respectively), and final refers to the *last* piece of evidence presented before action execution.

RP preceding urgent and deliberate evidence-informed actions

Finally, we investigated the EEG activity preceding action execution. Based on the previous experiment, our hypothesis were that an RP would be present before the informed Action 2 (at least in the informative and NL conditions), and that there would be no RP like activity before the urgent Action 1.

Figure 5. 13 shows the averaged RPs for the Action 1 (*red*) and Action 2 (*black*) for each condition separately, in congruent (*left*) and incongruent (*right*) trials. A clear RP was only apparent in the Action 2 in the NL condition, although negative-going shifts before the action were visible in the informative condition as well. In congruent trials, no significant clusters were identified in the informative and NH conditions, but a significant cluster (*p* = 0.047, uncorrected) was identified in the NL one, indicating that the second action showed larger negativities than the first action. No significant clusters were identified in any of the incongruent conditions.

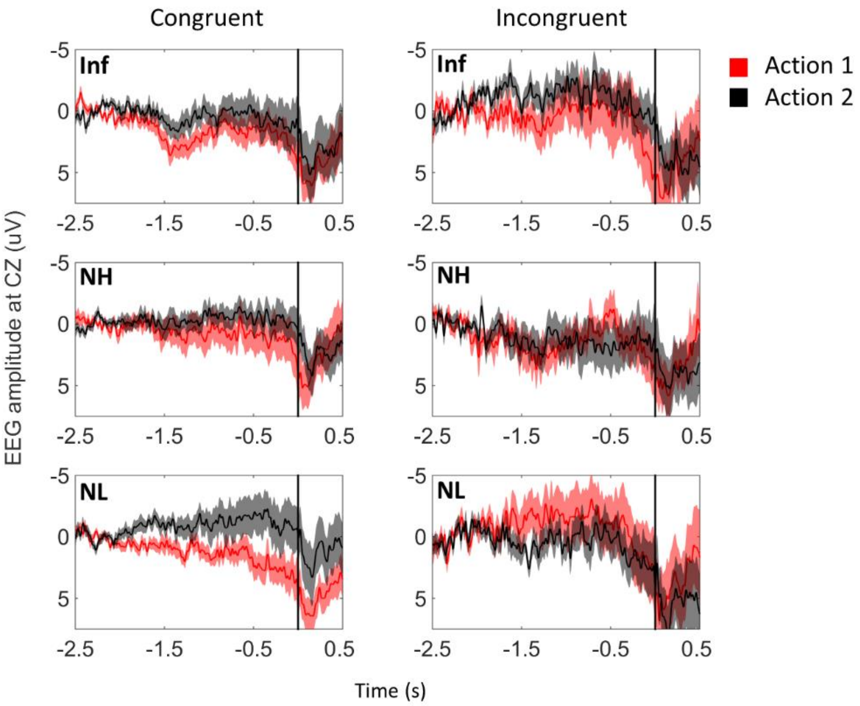

Figure 5. 13. The RP preceding urgent and deliberate evidence-informed actions. Grand-averaged (±SEM) RP traces preceding urgent Action 1 (*red*) and slow Action 2 (*black*) for each condition separately, in congruent and incongruent trials.

5.4.Discussion

In this experiment, we investigated 1) the effects of endogenous SAT modulations on the neural correlates of evidence accumulation and action initiation, and 2) the neural correlates of changes of mind in perceptual decision-making. We argued that adjustment of the SAT is a central aspect of voluntary control of action that can result from adaptations to the environment, but also from endogenous decisions. Further, we investigated whether the decision variable encoded in the p3 dynamically tracks the evolution of decision-making processes in changes of mind. Below, we provide an integrative summary and interpretation of the findings, illustrated in *Figure 5. 14*. We first discuss the results in congruent trials, and then move on to discuss how modulations in the described processes may underlie changes of mind.

In this study, our results from congruent trials replicated the main findings from chapter 4. We found that the p3 categorically encodes different types of evidence. In congruent trials, the average amplitude of the p3 in response to +Ev was greater than that in response to –Ev both before Action 1 and between Action 1 and Action 2, suggesting that decisions are categorically encoded by the p3 both when they are made under speed and when they are deliberate and slow (*Figure 5. 6)*. Again, we found this both for informative (easy) and neutral (difficult) trials, suggesting that the p3 reflects an internal decision variable rather than the objective strength of external evidence. We then analysed how the dynamic evolution of the p3 underlay those averaged results, and how they were modulated by variations in the SAT.

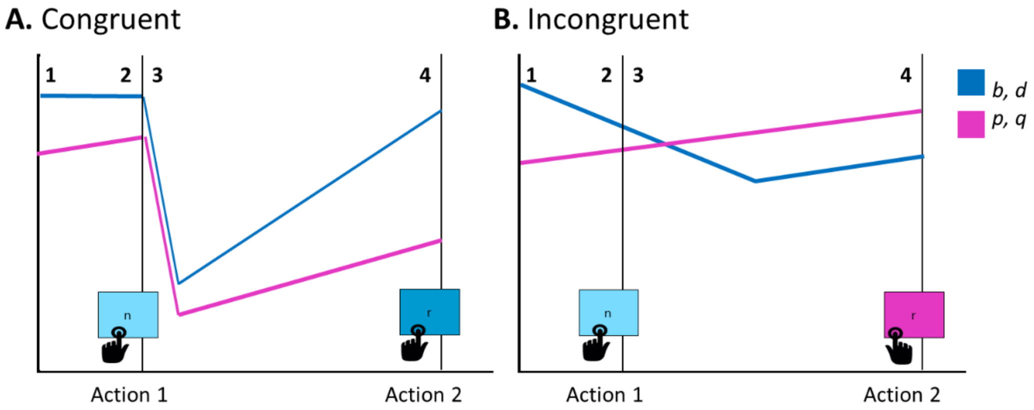

Figure 5. 14. Schematic representation of the p3 dynamics. The diagrams illustrate representative examples of the evolution of p3 responses in relation to actions, for different types of evidence in congruent (**A**) and incongruent trials (**B**). In both congruent and incongruent trials, a significant initial asymmetry (**1**) in the responses to different types of categorical evidence predicted the decision at the time of Action 1, and the amplitude of the p3 did not significantly increase up to the time of the first action (**2**). After the first action (**3**), the amplitude of the p3 dramatically decreased in congruent trials. We refer to this phenomenon as "baseline resetting". In incongruent trials, the amplitude of the p3 in response to different types of categorical evidence was differentially modulated. While evidence in favour the first decision (Action 1) decreased, evidence against that initial decision stayed stable, tending to increase. A clear difference in the p3 responses to different types of evidence was clearly visible before Action 2 both in congruent and incongruent trials (**4**). The comparison between the relative amplitudes of blue and pink evidence at times 1 and 4 shows that in congruent trials the dominant evidence remained stable (*bd*>*pq*), whereas in incongruent ones it flipped (*bd*>*pq* in 1, but *bd*<*pq* in 4). This change indexes changes of mind.

We predicted that if the p3 encodes an internal decision variable it should be sensitive to SAT changes, and in urgent decisions it should be modulated so that the baseline-to-threshold distance is reduced. In agreement with this prediction, we observed that, at the beginning of the experiment, p3 responses started at a relatively high amplitude. In fact, the amplitude of the p3 only significantly increased towards the time of Action 1 in informative trials. However, we observed significant differences in the averaged ERPs of all p3 components up to the time of Action 1 in congruent trials (*Figure 5. 6).* This suggests that while evidence was indeed accumulated in the informative condition, the mechanisms driving the observed difference in the averaged data preceding Action 1 in the neutral conditions must rely on implicit asymmetries rather than an active accumulation process. These implicit asymmetries were visible in the initial responses to the first instances of +Ev and –Ev, which were significantly different and predicted participants' decisions (*Figure 5. 9).* In sum, our analysis of the urgent decision-making process leading to Action 1 suggests that fast actions are achieved by increasing the baseline activation rather than reducing the threshold. This is compatible with previous results showing that urgency manipulations result in increases in baseline activations both at the level of sensory encoding and motor preparation (Forstmann et al. 2008; Mansfield et al. 2011; Steinemann et al. 2018).

We further showed that the transition process from an urgency to an accuracy regime is implemented through a "baseline resetting" mechanism whereby the p3 amplitude quickly and dramatically drops following Action 1. That is, the first p3 following Action 1 was significantly smaller than the last p3 before Action 2 (*Figure 5. 12).* This was observed in all conditions, for all types of evidence. Such a drop in the p3 response allows for a greater baseline-to-threshold distance, and thus enables evidence accumulation for longer time spans before a threshold is reached. Indeed, in contrast to Action 1, in the slow decision-making process leading to Action 2 we did find a significant increase in the p3 amplitude both for informative and neutral conditions (*Figure 5. 10*). Further, we directly compared the two decision processes from trial onset to Action 1, and from Action 1 to Action 2 by comparing the baseline activities at the beginning of evidence accumulation and the final amplitudes reached before each action. We found that the urgent decision-making process preceding Action 1 started with a significantly higher baseline response than that preceding the slow, deliberate Action 2. In congruent trials, this was the case for all conditions and all types of evidence, and was mediated by the baseline resetting mechanism described in the previous paragraph.

Changes of mind and the p3

When participants changed of mind, some of the processes described above showed differential characteristics.

First, while there was a difference in the initial p3 amplitudes in response to +Ev and -Ev at trial onset (*Figure 5. 9*), the averaged ERPs up to the time of Action 1 did not reflect consistent differences (*Figure 5. 7*). This is in contrast with congruent trials, where we found that the first fast decisions were already reflected in the averaged p3, in all conditions.

Further, the baseline resetting mechanism described above was significantly different in trials where participants eventually changed of mind. While in congruent trials there was a significant decrease in the p3 amplitude for all conditions and types of evidence after the first action, such decrease was only apparent in the informative condition in incongruent trials. In neutral conditions, p3 amplitudes stayed the same or increased after Action 1 (*Figure 5. 12*). This suggests that the endogenous modulation of the SAT in congruent and incongruent trials were substantially different, at least in the neutral conditions.

The comparison of the initial p3 amplitudes in response to different types of evidence at the beginning of the trial and just after Action 1 in incongruent trials provides some further insight into these differential dynamics. While the p3 amplitude in response to the type of evidence that dominated at the time of Action 1 (e.g. *bd* if Action 1 was Blue) decreased from trial onset to the time just after Action 1, the p3 amplitude in response to the non-dominant type of evidence (e.g. *pq* if Action 1 was Blue) stayed stable, tending to increase. Thus, while no systematic baseline resetting mechanism was observed in the neutral conditions, the mere lack of it produced a relative shift in favour of the option not chosen at Action 1, potentially enabling eventual changes of mind.

We initially predicted that if participants changed of mind, we should be able to see a reversal in the dominant amplitudes of the p3. That is, for example, that if participants chose blue in Action 1, the p3 amplitude in response to *bd* would be greater than in response to *pq* (p3(*bd*) > p3(*pq*)). If they then changed of mind and chose Pink in Action 2, the amplitude of the p3 should reverse and be now higher for *pq* than for *bd*. (Action 1, Blue: p3(*bd*) > p3(*pq*)*;* Action 2, Pink: p3(*bd*) < p3(*pq*)). While this reversal was not visible in the averaged ERP data (*Figure 5. 7*), we were able to detect it by comparing the initial amplitudes at trial onset and the final amplitudes before action execution (*Figure 5. 11*). Such a reversal occurred gradually, probably at different times in different trials. However, the fact that no clear "baseline resetting" was visible after the initial action constitutes a clear marker that, already at the time of Action 1, the decision-making process in congruent and incongruent trials was significantly different. But why would it be?

The magnitude of the difference between p3 responses to +Ev and -Ev at the time of Action 1 could be viewed as a proxy of participants' confidence in their initial, urgent decisions. Since this difference was absent or smaller in incongruent trials preceding Action 1, one might speculate that reduced confidence in initial decisions in those trials drove changes of mind. It has been suggested that the same neural populations involved in making decisions also encode the ‘confidence’ one should have in those decisions, and recent studies have shown that continuous monitoring of certainty on the environment and confidence on one’s own choices modulates explore/exploit decisions (Boldt et al. 2019). Further, several models have suggested that evidence presented after a decision has already been made is evaluated according to the perceived accuracy of the decisions (Fleming et al. 2018). Potentially reduced confidence at the time of Action 1 could also result from reduced support from immediate evidence. Our exogenicity analysis revealed that urgent actions in congruent trials were strongly driven by immediate external evidence, while in incongruent trials they were not (*Figure 5. 5*). That is, urgent actions that were followed by changes of mind were only weakly driven by immediate stimulus evidence.

Consistent with this possibility, our behavioural results showed that participants were indeed less confident in their decisions in incongruent than in congruent trials. Although we have no formal measure of confidence at the time of Action 1, one possible interpretation of this result is that they were less certain about the correct choice all throughout the trial, starting from their first decision – which was not strongly reliant on external evidence, and did not show a strong difference between +Ev and -Ev.

In sum, our results show that the p3 tracks a dynamic decision variable that continuously monitors evidence, categorically encodes decisions and responds to endogenous modulations of the SAT. Importantly, although SAT modulations in this task were imposed by the experimental instructions, the precise moment at which participants acted varied between trials, and was determined by the participant rather than by any external stimulus. A spontaneous change in SAT priorities could occur after purely internal deliberation in naturalistic situations: for example, one might notice one is unsure about a certain choice, and decide to take more time before making a final judgement. Thus, the description of the mechanisms by which it is implemented provides a valuable tool to investigate voluntary control of actions that meaningfully interact with the environment.

The RP in different SAT conditions

Finally, we investigated the motor activity preceding action execution. On the one hand, we hypothesized that urgent actions would not be preceded by an RP. Because participants were required to act fast, we predicted that they would rely more on immediate evidence, making their actions more akin to reactions than to self-paced ones. Our behavioural analysis of exogenicity confirmed that, as expected, participants were more dependent on immediate evidence prior to execution of the first action, compared to the second one (*Figure 5. 5*). Further, in agreement with our prediction, urgent actions were not preceded by an RP-like negativity in any of the conditions.

On the other hand, we assumed the second, slow actions in this paradigm were comparable to the evidence-informed actions in chapter 4. Thus, we hypothesized that we would be able to find an RP-like signal preceding slow, deliberate actions. However, we were only able to clearly see an RP in the Neutral Low condition, in congruent trials (*Figure 5. 13*).

Several reasons may explain these conflicting results. While in chapter 4 participants had a choice about *whether* to execute an action and *when* to do so, in this paradigm actions were compulsory. While in the previous chapter we argued that actions followed from a fluent evidence accumulation process where perceptual processes informed and eventually drove action, in this case there was a strong external imperative to act. This means that participants were compelled to act even if in some trials they were not necessarily confident about their choice. Actions in this experiment may thus contain a mixture of trials where participants reached a decision threshold spontaneously and fluently prepared an action, and other trials where they were decision-making processes did not reach any action threshold but participants pressed a key anyway due to experimental requirements. This type of response may involve different neural mechanisms and have different neural signatures than the fluent decision-making process studied in chapter 4. However, our design does not allow us to disentangle these two potential different types of trials.

Summary

In sum, in this chapter we have described how the p3 dynamically tracks modulations in the SAT and categorically encodes decision outcomes. We have further characterized an endogenous baseline resetting mechanism that allows the implementation of different speed-accuracy regimes, and we have identified how changes of mind may be afforded by distinctive modulations in that process. Finally, we have been unable to replicate the results from chapter 4 with respect to the RP preceding evidence-informed actions. However, the actions in this task were not exactly equivalent to the ones in the previous chapter (i.e. here, they were compulsory), and thus action initiation processes may have operated differently.

Chapter 6

General discussion

6.1.General discussion

In this thesis, we have investigated different aspects of voluntary action control. In **chapter 1,** we proposed a hierarchical model involving two main kinds of control: a *local* type at the level of action execution or inhibition, and a *general* type of control involving top-down regulatory mechanisms potentially spanning all levels of the sensorimotor hierarchy. The results of this thesis provide novel insight into both types of control and shed new light into the relationship between phenomenal aspects of voluntary action, the neural mechanisms involved in action selection and initiation, and their link to perceptual decision-making.

6.2.Summary of results

In **Part I** of this thesis, we investigated the relationship between motor preparation and awareness of intention in voluntary action. We were interested in two questions. First, do people *know* (i.e. have conscious access to) the fact that they are about to do something, and can they use this experience to control action in real time? In other words, do people's *feelings* of intention track some motor preparation signal, and are the neurocognitive underpinnings of these experiences relevant for behavioural control? Second, to what extent are offline, recall-based reports of intention informed by *prospective* motor preparation processes rather than *retrospective* reconstruction?

We approached these questions from a metacognitive perspective, as described in **chapter 1**. The RP has consistently been reported as a precursor neural correlate of self-initiated movements (Libet et al. 1983; Trevena and Miller 2002; Matsuhashi and Hallett 2008; Schultze-Kraft et al. 2016; Khalighinejad et al. 2018), and is thus taken as a marker of self-paced motor preparation. From this starting point, we hypothesized that the experience of awareness of intention is a percept resulting from the metacognitive monitoring of the first-order motor preparation processes indexed by the RP. Further, we hypothesized that such monitoring happens in real-time and allows online action control. That is, there is a prospective *feeling* of intention that can be used to proximally control action. Finally, we hypothesized that delayed, recall-based reports of intention result from a combination of prospective and retrospective cues. Chapters 2 and 3 investigate these questions from two complementary angles.

In **chapter 2**, we investigated whether awareness reports were sensitive to the presence of motor preparation signals. That is, the probability of finding a motor preparation signal given an awareness report (i.e. P(RP|Awareness)). To do so, we used a random probing method. While participants were performing a self-paced movement task, we introduced random awareness probes and obtained awareness reports with different methods (i.e. contingent action, contingent veto, readiness rating). Intention probes were presented *before* any action was executed to study prospective contributions to intention awareness exclusively.

In Experiment 1, we showed that when participants report an awareness to move by means of an action (i.e. "contingent action" instruction), an RP signal is present, whereas it is absent when they do *not* report an intention. A similar pattern was observed in the extremes of the scale in Experiment 3: when participants reported being about to move, an RP was present, whereas when they reported not being ready at all no RP was visible. Conversely, in Experiment 2, where people had to inhibit an action if they were aware of an impending movement (i.e. "contingent veto" instruction) we did not find any significant difference between awareness and no-awareness reports. As discussed in chapter 2, this may have been the case because the type of reports provided in Experiment 2 were more akin to judgements than to immediate reports of feelings.

In turn, in **chapter 3**, we investigated whether motor preparation signals predicted awareness reports. That is, the probability of people reporting awareness given a certain motor preparation state (i.e. P(Awareness|RP)). Further, we investigated the extent to which action execution retrospectively influences awareness of intention reports. To do so, we trained a linear classifier to discriminate self-initiated motor preparation states from non-preparation ones. We specifically targeted the RP as a motor preparation signal. Then, while participants were performing a self-paced task, we monitored their EEG signal in real time and triggered a Go/No-Go cue based on the presence or absence of an RP. After executing or inhibiting an action, participants were asked about their intention to move at the time the Go/No-Go cue was presented.

We found a strong effect of retrospective reconstruction: when participants executed an action in response to the Go cue, they were more likely to report an intention at the time of probing than when they refrained from moving in response to the No-Go cue. Further, we found an effect of motor preparation as well. Participants were more likely to report an intention to move when the Go/No-Go cues were triggered by the presence of an RP than when they were triggered in the absence of such signal. This suggests that motor preparation signals contribute not only to the *feelings*, but also to delayed *judgements* of motor intention.

In **Part II** of the thesis, we investigated voluntary action in a decision-making context. In particular, we investigated action selection and initiation processes in a context where actions were self-paced, but evidence informed. We were interested in two main ways in which endogenous factors can contribute to decision-making processes. First, by breaking the symmetry in ambiguous external scenarios mechanism. Second, by top-down modulating decision-making processes.

In **chapter 4**, we focussed on investigating potential symmetry-breaking mechanisms at the level of action selection and motor initiation. We designed a paradigm where participants decided *whether* to act based on either strong or ambiguous evidence. They had no time pressure to make a decision, so actions were largely self-paced. We showed that the p3 encodes an internal decision variable that categorically encodes different types of evidence, and predicts people's decisions even in the absence of net evidence. The evolution of the p3 amplitude both in easy and difficult conditions showed dynamics compatible with accumulation to bound theories of decision-making: the p3 in response to evidence in favour of the decision (+Ev) reached the same amplitude in all conditions before the time of action. We found no differences in the baseline responses to evidence in favour (+Ev) and against the decision (–Ev), and our dynamic analysis of the evolution of the p3 over trial time shows that the amplitude of the p3 progressively increased in response to evidence favouring the eventually selected option. This suggests that the symmetry between alternative responses was broken and gradually increased over time, eventually prompting a decision.

We were further interested in investigating the neural correlates of action initiation in a task where actions were not purely spontaneous, nor mere reactions to individual stimuli. We found that the RP can be detected both when actions were strongly informed by the evidence and when they were underdetermined by it. Thus, we concluded that the RP is not related to the arbitrariness of decisions (Maoz et al. 2017) nor to how "endogenous" the selection is. Rather, it seems to reflect the self-paced nature of actions that was preserved in all conditions of the task.

In **chapter 5**, we aimed to replicate the findings in chapter 4, and to further investigate a different way in which endogenous factors can influence action selection processes. While symmetry-breaking can be explained by noise in the neural system, voluntary control of action requires the ability to strategically tune decision processes to reach endogenous goals. As we argued in chapter 1, such manipulations may include a number of top-down processes, potentially spanning the whole sensorimotor hierarchy. One of these processes is the ability to voluntarily modulate the speed-accuracy trade-off. Here, we modified the paradigm in chapter 4 to investigate how different speed-accuracy regimes would be reflected in the p3 evolution. Participants now had to decide *what* do to, and they were required to execute a primarily fast action at the beginning of the trial, and then a second, accurate action at the end. The speed-accuracy trade-off was thus dynamically modulated within each trial. We found that the p3 dynamics show modulations compatible with a reduction in the baseline-to-distance when decisions are made under urgency, and that they dynamically track transitions from an urgency to an accuracy regime. Further, we replicated the findings in chapter 4, and we showed that the p3 dynamically encodes the evolution of decision-making processes, indexing both congruent decisions and changes of mind.

6.3.Theoretical implications

The contemporary neuroscientific tradition investigating the relationship between conscious experience and voluntary action control started with Libet's experiments (Libet et al. 1983), which relied on dualistic assumptions that are currently largely rejected (Haggard 2005; Mirabella 2007). However, both the methods and the conclusions derived from his experiments had a long-lasting impact on the following tradition. The theoretical framework, methods and results presented in this thesis aim to overcome some of the caveats that have hindered the development of the field and aim to provide a cognitively-grounded starting point for new developments in the neuroscience of voluntary action control.

6.3.1. Implications for theories of motor intention

In chapter 1, we provided a metacognitive framework for the study of the relationship between motor preparation processes and awareness of intention which does not rely on the dualistic assumptions that drove Libet's initial research. Thus, our framework inverted the direction of the traditional question. Rather than asking "how do conscious decisions cause motor preparation?" we asked "when is motor preparation made available to conscious experience?".

While the motor preparation processes indexed by the RP may result from accumulation of stochastic noise (Schurger et al. 2012), they may still constitute a relevant signal available for metacognitive monitoring, that may in turn allow top-down cognitive control of action. We argued that evidence from GTS, AHS and utilization behaviour patients suggests that a genuine phenomenology related to motor preparation exists before action execution and may result from neurocognitive mechanisms relevant for action control. However, Libet-like studies critically confounded *feelings* and *judgements* of intention, and were therefore unable to provide convincing evidence that the W-judgements reflected a genuine experience preceding action. In chapter 2, we tackled this question by directly probing the *feeling* of awareness in the absence of an action. Our results thus shed new light into *feelings* and *judgements* of motor intention. We provide evidence that people have some conscious insight into their motor preparation processes before action execution and can guide their behaviour in real-time according to these internal states. Further, in chapter 3 we showed that delayed *judgements* of intention result from a combination of prospective information about internal motor preparation states and retrospective reconstruction based on events (i.e. actions) happening between the time of probing and the time of intention report. Our findings suggest that healthy people have some *feeling* of motor intention before acting and can use it to voluntarily control action initiation. Yet, conscious access to discrete motor states is short-lived. Agents know what they are about to do in real-time, but what they *actually* end up doing has a stronger influence on their recall-based intention judgements.

6.3.2. Implications for the empirical approach to voluntary action research

We further questioned the idea that voluntary actions can be operationalised or defined as purely “endogenous”. We argued that agents are intrinsically interactive, and that our brains have evolved to meaningfully engage with the environment. Thus, while the study of spontaneous actions provides a valuable tool to specifically study spontaneous action initiation processes, it is not representative of most naturalistic voluntary actions, which take place in a rich environment where information is relevant for decision-making processes. Thus, in chapters 4 and 5 we advocated for the operationalisation of voluntary action in more naturalistic contexts, where external evidence is not excluded from the picture, but rather made relevant for action.

In particular, we were interested in empirically investigating the idea that an essential aspect of voluntary control of action is the ability to strategically control decision-making processes and optimise them for current goals. In the hierarchical model of action control proposed in chapter 1, it was argued that this ability requires top-down interaction between higher-order states and lower-level processes, mediated by fine-tuning the background states of the system where decision-making takes place. This tuning is proposed to be implemented in a distributed network potentially spanning the whole sensorimotor hierarchy – including action initiation- and involving several processes, such as selective attention and modulations of the speed-accuracy trade-off. As we argued in chapter 1, these processes need not always *be* under voluntary control, and indeed they are often modulated bottom-up by environmental changes. However, because they *can* be voluntarily modulated, they constitute an essential mechanism to explain how voluntary control over action selection and initiation processes is implemented.

In **chapter 5**, we investigated how an internal decision variable indexed by the p3 was modulated by volitional control processes. We described how the intention to act fast was implemented by modulations compatible with a reduction in the baseline-to-threshold distance, and that transitions from an urgent to an accurate action policy were implemented by means of an immediate baseline resetting mechanism. These results indicate that effective implementation of certain endogenous goals (or, at least, specific features of it, such as the required speed or accuracy) is mediated by an internal decision variable that can be objectively measured and tracked by means of the p3. This is an important development: it provides researchers with a metric to investigate specific features of endogenous intention implementation.

In our task, the SAT modulations were instructed. However, they did preserve a certain endogenicity in that the precise moment when the transition from an urgent to an accurate action policy happened depended on the participants' decision processes rather than on objective changes in the stimuli. Outside the laboratory, SAT modulations can occur in a spontaneous manner, resulting from internal deliberation processes. For example, one can spontaneously self-impose the requirement to act fast after realising that it is very late, and the last train is about to depart. The ability to implement these spontaneous intentions (i.e. run to the train station) and achieve one's goals (i.e. catching the last train) depends on being able to actively implement them by means of top-down modulations like the ones we have described. Investigating the implementation of those modulations in experimentally-controlled contexts provides us with tools and measurable variables to track naturally occurring volitional control processes. SAT adjustments and the modulations in the p3 that track them are only one specific kind of modulation – many other specific effects may relate to other features of intentional action control. Thus, this conceptual framework provides empirical methods to study voluntary control of actions in naturalistic scenarios, where agents meaningfully interact with the environment.

6.3.3. Implications for the interpretation of the readiness potential

Finally, our investigation of action initiation processes both in the conventional context of spontaneous actions (in chapters 2 and 3) and in the context of evidence-informed, self-paced actions (in chapters 4 and 5) allows us to sketch an updated interpretation of the RP. In chapters 4 and 5, we embedded self-paced action execution in an evidence-informed paradigm, where *whether* and *what* decisions had to be made on the basis of external evidence. This allowed us to investigate the neural correlates of action initiation to test whether the RP could be found in this non-reactive, but also not 'purely endogenous' context. We found that the RP could be found in some actions that were strongly driven by evidence, suggesting that it may be a valid signal of motor preparation in naturalistic scenarios. However, we did not find it in *all* our conditions. These results support a weak (or at least not strong) dependence relationship between voluntary action execution and the presence of the RP.

We suggest that a useful way to think about the potential causal relationship of the RP with voluntary action initiation involves a type of causality known as *INUS* causation. INUS causes are defined as Insufficient but Non-redundant parts of a condition which is itself Unnecessary but Sufficient set of conditions (Mackie 1965). The idea behind this complex acronym is that some event *A* is not *always* the cause of a certain effect

E (i.e. the presence of *A* is not *sufficient* to produce the effect *E*), but *A* may cause *E*, given an adequate set of circumstances *C* (i.e. if the event *A* happens given certain circumstances *C*, the effect *E* will be produced). However, the effect *E* could have occurred in the absence of *A*, if instead of *C* circumstances were *D*. An example given by Mackie in the original paper is a short-circuit. A short-circuit is not sufficient to set a house on fire, but together with another set of conditions (e.g. very dry atmosphere, wooden floors) it may be the event that actually sets the house on fire.

Similarly, it is suggested here that the RP may index processes that are not in themselves sufficient to trigger action (in normal conditions) but may do so in certain contexts (such as the Libet experiment). In this framework, the RP should not be singled out as *the* cause of spontaneous actions. Rather, the recorded actions happened by the joint contribution of higher-order processes driving people’s decisions to e.g. participate in the experiment, and the neural activity underlying the RP which triggered each specific movement. But that activity *alone* does not suffice to produce actions.

This interpretation of the RP suggests that, while stochastic fluctuations *can* trigger spontaneous self-paced movements, whether they actually *do* depends on further constraints. The RP may only suffice to trigger movements given a certain background state of the cognitive system. Were stochastic fluctuations over the motor cortex to happen in the absence of an intention to spontaneously move, or better, an active intention *not* to move, presumably those fluctuations would not suffice to trigger actions. Thus, while Libet concluded that action initiation processes are not initiated consciously because the RP onset precedes W-judgements, one could rather argue that those actions were only possible (and distally initiated) by the decision to take part in the experiment, which provided a suitable background for the RP to trigger actions in a spontaneous manner.

6.4. Practical implications

In addition to the theoretical implications, the methods developed in this thesis and the experimental results may have implications at various levels.

First, our results provide a better understanding of how *feelings* and *judgements* of motor intention come about. Some studies of patients with disorders of volition such as GTS (Moretto et al. 2011; Ganos et al. 2015) have used the Libet clock to obtain judgements of intention, and have investigated how those might reflect underlying cognitive impairments. We show that *judgements* of intention do reflect, to some extent, motor preparation processes preceding action. However, they are also heavily influenced by retrospective reconstruction. Here, we developed methods to directly investigate *feelings* of intention. Thus, those methods could provide a better measure to investigate impairments potentially emerging at the time of motor preparation, before actions are executed. Thus, research on clinical populations where metacognitive monitoring of endogenous action initiation processes is impaired may benefit from methods or ideas developed in these studies.

Second, our approach to the study of the RP provides valuable methods for future research in the field. We combined random probing paradigms with brain-based, targeted probing to investigate the causal relationship between the RP and motor intention feelings and judgements. While the classic approach to the study of the neural correlates of consciousness (NCC) typically involves a unidirectional approach (i.e., from awareness reports or measures to the brain activity correlating with those), causal claims about the role of a neural signal on the subsequent conscious experience require the converse direction as well (i.e. from brain activity to awareness reports). We believe that the combination of both methods could provide a very powerful strategy to precisely map the relationship between neural activity and conscious reports.

Finally, our results regarding prospective access to motor intention potentially have societal implications as well. The tradition emerging from Libet resulted in the wide-spread (and far-fetched) claim that people do not have free will. Previous research has suggested that such a belief has detrimental effects on individual behaviour (Vohs and Schooler 2008; Nahmias 2011), and indeed the idea that people are not capable of exerting control over their actions is problematic in societies where responsibility attribution depends on people's ability to "have done otherwise" (Kulakova et al. 2017). Here, we have argued that voluntary control of action involves a complex, distributed network of top-down control mechanisms to which all sensorimotor processes (including action initiation) are subjected. Our conceptualization of voluntary action control adds to many other criticisms of Libet's interpretation (Roskies 2011; Brass et al. 2019), and suggests that if the question of free will is to take into account evidence from neuroscience, it will need to consider much more than a single brain potential and a post-hoc subjective report.

6.5.Challenges & future directions

6.5.1. Practical considerations

The study of spontaneous voluntary action faces a series of unavoidable challenges. On the one hand, the self-paced nature of the object of study results in the unescapable fact that the experimental conditions and trial numbers are not experimentally controllable. Thus, the research presented in chapters 2 and 3 dealt with unbalanced designs where the number of trials per condition varied widely across participants. We tackled this limitation by using conservative and robust statistical methods (i.e. cluster analysis in chapter 2, and mixed linear regression in chapter 3). However, because the research presented in this thesis involves novel paradigms, replications of the reported findings will be required to establish them in the future.

A second limitation inherent to the study of conscious phenomena is the reliance on some kind of subjective reports. While no-report methods have been developed in studies of other types of conscious experiences (Tsuchiya et al. 2015), the golden-standard for validation of those measures always relies on a subjective report (Overgaard and Fazekas 2016). However, this is problematic. Given a scale, people may use it in different ways to report identical experiences. In binary choices, they may have different responses biases. Yet, we have no direct access to the experiences they report or their idiosyncratic differences. Thus, we need to use some of these imperfect measures. While each single experiment in chapters 2 and 3 o also suffered from this limitation, we aimed to partially overcome it by using not only one, but a range of reports. Each kind of report entails a specific set of limitations, but we reasoned that by combining different probing and report methods, the confounds arising from one method might be controlled by another one.

Finally, our investigation of motor preparation in chapters 2 and 3 relied on the assumption that the RP is a valid marker of spontaneous preparation. This is based on extensive literature describing its presence preceding voluntary actions (Libet et al. 1983; Keller and Heckhausen 1990; Trevena and Miller 2002; Desmurget et al. 2009; Schultze-Kraft et al. 2016; Khalighinejad et al. 2018). However, as discussed in the general introduction, the fact that whenever a self-paced action is present an RP is present does not entail that whenever an RP is present motor preparation (or a subsequent action) would be present. In these studies, we did not quantify the extent to which the presence of an RP-like activity predicts movement execution. However, the dataset and methods in chapter 3 potentially allow for such an analysis. For example, the same classifier used to interrupt self-paced movements could be used to offline calculate the extent to which the RP signal was predictive of spontaneous movements using a Receiver-Operating-Curve (ROC) type of analysis. While the interpretation of these analysis should be carefully interpreted in the light of the specific features and limitations of the classifier used, future analysis could aim to determine the extent to which the presence of an RP predicts movement initiation in a spontaneous action setting.

6.5.2. Theoretical considerations

In the general introduction we sketched a hierarchical model of voluntary action control, and we argued that the study of specific processes that are relevant for the implementation of some features of higher-order intentions such as modulations of the SAT and selective attention provided a valuable tool for the study of voluntary action research. In chapter 5, we specifically investigated how voluntary SAT modulations are dynamically tracked by the p3. However, in our task there were clear external instructions that guided the SAT changes throughout the task. Thus, the origin of the modulations can theoretically be traced back to an abstract representation of the instruction. We argued that such modulations also occur spontaneously and can result from internal deliberations. Yet, we have not tackled the question of how those deliberations might operate. Our model tackles the *implementation* of certain aspects of higher-order intentions, not their *generation*. Obviously, this is by no means a complete, exhaustive model of autonomous agency. Rather, it is a partial model that takes certain *given* intentions as a starting point to investigate how they are implemented. Future work on voluntary action should go upstream of those intentions, and investigate how they dynamically arise and produce changes in downstream processes.

6.6. Final conclusion

Most of us believe that we can control our actions. Yet, defining and testing the truth conditions of such belief is both a philosophically and empirically thorny problem. This thesis has tackled the question of voluntary control of action from a hierarchical, cognitively-grounded model as a starting point. We have suggested that voluntary control of action involves a distributed network of processes that enable the pursuit of high-level goals by means of top-down modulations of low-level decision-making processes. Further, we have argued that the conscious experience of motor intention may result from a second-order readout of first-order motor preparation processes, and that a metacognitive loop may allow cognitive control over action initiation in real time. This thesis provides new theoretical insights for the study of voluntary action, as well as relevant empirical observations and a range of novel methods that can be used in future research and potentially applied to several clinical populations.

References

Aarts H, Custers R, Wegner DM. 2005. On the inference of personal authorship: Enhancing experienced agency by priming effect information. Conscious Cogn. 14:439–458.

Ach N. 2006. On volition (T. Herz, Trans.). (Original work published 1910) [WWW Document]. URL http://www.uni-konstanz.de/kogpsych/ach.htm

Assal F, Schwartz S, Vuilleumier P. 2007. Moving with or without will: Functional neural correlates of alien hand syndrome. Ann Neurol. 62:301–306.

Banks WP, Isham EA. 2009. We infer rather than perceive the moment we decided to act. Psychol Sci. 20:17–21.

Banks WP, Pockett S. 2007. Benjamin Libet's Work on the Neuroscience of Free Will. Blackwell Companion to Conscious. 657–670.

Barnes J (Ed.). 1984. Complete works of Aristotle: The revised Oxford translation. Princeton University Press.

Bates, D, Mächler, M, Bolker, B.M., Walker SC. 2015. Fitting Linear Mixed-Effects Models using lme4. J Stat Softw. 251–264.

Bayne T. 2011. Libet and the case for free will scepticism. In: Swinburne R, editor. Free Will and Modern Science. Oxford:Oxford University Press. p. 1–22.

Beck B, Peña-Vivas V, Fleming S, Haggard P. 2019. Metacognition across sensory modalities: Vision, warmth, and nociceptive pain. Cognition. 186:32–41.

Blankertz B, Lemm S, Treder M, Haufe S, Müller KR. 2011. Single-trial analysis and classification of ERP components - A tutorial. Neuroimage.

Bode S, Bogler C, Haynes J-D. 2013. Similar neural mechanisms for perceptual guesses and free decisions. Neuroimage. 65:456–465.

Bode S, Bogler C, Soon CS, Haynes JD. 2012. The neural encoding of guesses in the human brain. Neuroimage. 59:1924–1931.

Bode S, He AH, Soon CS, Trampel R, Turner R, Haynes J-D. 2011. Tracking the unconscious generation of free decisions using ultra-high field fMRI. PLoS One. 6:e21612.

Bogacz R, Brown E, Moehlis J, Holmes P, Cohen JD. 2006. The physics of optimal decision making: A formal analysis of models of performance in two-alternative forced-choice tasks. Psychol Rev. 113:700–765.

Bohlhalter S, Goldfine A, Matteson S, Garraux G, Hanakawa T, Kansaku K, Wurzman R, Hallett M. 2006. Neural correlates of tic generation in Tourette syndrome: An event-related functional MRI study. Brain. 129:2029–2037.

Boldt A, Blundell C, De Martino B. 2019. Confidence modulates exploration and exploitation in value-based learning. Neurosci Conscious. 2019:1–12.

Brainard DH. 1997. The Psychophysics Toolbox. Spat Vis. 10:433–436.

Brainin M, Seiser A, Matz K. 2008. The mirror world of motor inhibition: The alien hand syndrome in chronic stroke. J Neurol Neurosurg Psychiatry. 79:246–252.

Brass M, Furstenberg A, Mele AR. 2019. Why neuroscience does not disprove free will. Neurosci Biobehav Rev. 251–263.

Brown SD, Heathcote A. 2008. The simplest complete model of choice response time: Linear ballistic accumulation. Cogn Psychol. 57:153–178.

Buneo CA, Jarvis MR, Batista AP, Andersen RA. 2002. Direct visuomotor transformations for reaching. Nature. 416:632–636.

Burk D, Ingram JN, Franklin DW, Shadlen MN, Wolpert DM. 2014. Motor effort alters changes of mind in sensorimotor decision making. PLoS One. 9:e92681.

Caminiti R, Borra E, Visco-Comandini F, Battaglia-Mayer A, Averbeck BB, Luppino G. 2017. Computational Architecture of the Parieto-Frontal Network Underlying Cognitive-Motor Control in Monkeys. Eneuro. 4:ENEURO.0306-16.2017.

Carpenter W. 1852. On the influence of suggestion in modifying and directing muscular movement independently of volition. R Inst Gt Britain.

Chen X, Wang Y, Nakanishi M, Gao X, Jung T-P, Gao S. 2015. High-speed spelling with a noninvasive brain–computer interface. Proc Natl Acad Sci. 112:E6058–E6067.

Chung, Y., Gelman, A., Rabe-Hesketh, S., Liu, J., Dorie, V., 2015. Weakly informative prior for point estimation of covariance matrices in hierarchical models. J. Educ. Behav. Stat. (in press).

Cisek P. 2007. Cortical mechanisms of action selection: the affordance competition hypothesis. Philos Trans R Soc Lond B Biol Sci. 362:1585–1599.

Cisek P. 2012. Making decisions through a distributed consensus. Curr Opin Neurobiol. 22:927–936.

Cisek P, Puskas G a., El-Murr S. 2009. Decisions in Changing Conditions: The Urgency-Gating Model. J Neurosci. 29:11560–11571.

Crammond DJ, Kalaska JF. 1989. Neuronal activity in primate parietal cortex area 5 varies with intended movement direction during an instructed-delay period. Exp Brain Res. 76:458–462.

Cui RQ, Deecke L. 1999. High resolution DC-EEG analysis of the Bereitschaftspotential and post movement onset potentials accompanying uni- or bilateral voluntary finger movements. Brain Topogr. 11:233–249.

Davidson D. 1963. Actions, Reasons, and Causes. J Philos. 60:685–700.

Debray SBE, Demeestere J. 2018. Alien hand syndrome. Neurology. 91:527–527.

Deco G, Romo R. 2008. The role of fluctuations in perception. Trends Neurosci. 31:591–598.

Deecke L. 2012. There are conscious and unconscious agendas in the brain and both are important-Our will can be conscious as well as unconscious. Brain Sci. 2:405–420.

Deiber M, Honda M, Ibañez V, Sadato N, Hallett M. 1999. Mesial Motor Areas in Self-Initiated Versus Externally Triggered Movements Examined With fMRI : Effect of Movement Type and Rate. J Neurophysiol. 81:3065–3077.

Deiber M, Passingham RE, Colebatch JG, Friston KJ, Nixon PD, Frackowiak RSJ. 1991. Cortical areas and the selection of movement: a study with positron emission tomography. Exp Brain Res. 393–402.

Della Sala S, Marchetti C, Spinnler H. 1991. Right-sided anarchic (alien) hand: A longitudinal study. Neuropsychologia. 29:1113–1127.

Delorme A, Makeig S. 2004. EEGLAB: An open source toolbox for analysis of single-trial EEG dynamics including independent component analysis. J Neurosci Methods. 134:9–21.

Dennett DC, Kinsbourne M. 1992. Time and the observer: The where and when of consciousness in the brain. Behav Brain Sci. 15:183–201.

Desmurget M, Reilly KT, Richard N, Szathmari A, Mottolese C, Sirigu A. 2009. Movement intention after parietal cortex stimulation in humans. Science. 324:811–813.

Ditterich J, Mazurek ME, Shadlen MN. 2003. Microstimulation of visual cortex affects the speed of perceptual decisions. Nat Neurosci. 6:891–898.

Dorie AV, Dorie MV, Bates D, Maechler M, Bolker B, Walker S. 2016. Package ‘ blme .’ 1–8.

Dorris MC, Glimcher PW. 2004. Activity in posterior parietal cortex is correlated with the relative subjective desirability of action. Neuron. 44:365–378.

Douglas ZH, Maniscalco B, Hallett M, Wassermann EM, He BJ. 2015. Modulating conscious movement intention by noninvasive brain stimulation and the underlying neural mechanisms. J Neurosci. 35:7239–7255.

Ebbesen CL, Brecht M. 2017. Motor cortex - To act or not to act? Nat Rev Neurosci. 18:694–705.

Erdler M, Beisteiner R, Mayer D, Kaindl T, Edward V, Windischberger C, Lindinger G, Deecke L. 2000. Supplementary motor area activation preceding voluntary movement is detectable with a whole-scalp magnetoencephalography system. Neuroimage. 11:697–707.

Eslinger PJ. 2002. The anatomic basis of utilisation behaviour: A shift from frontal-parietal to intra-frontal mechanisms. Cortex. 38:273–276.

Fernandez-Duque D, Baird JA, Posner MI. 2000. Executive Attention and Metacognitive Regulation. Conscious Cogn. 9:288–307.

Fleming SM. 2019. Awareness reports as inference in a higher-order state space. arXiv Prepr arXiv190600728.

Fleming SM, Dolan RJ. 2012. The neural basis of metacognitive ability. Philos Trans R Soc Lond B Biol Sci. 367:1338–1349.

Fleming SM, Dolan RJ, Frith CD. 2012. Metacognition: computation, biology and function. Philos Trans R Soc Lond B Biol Sci. 367:1280–1286.

Fleming SM, Mars RB, Gladwin TE, Haggard P. 2009. When the brain changes its mind: flexibility of action selection in instructed and free choices. Cereb Cortex. 19:2352–2360.

Fleming SM, Van Der Putten EJ, Daw ND. 2018. Neural mediators of changes of mind about perceptual decisions. Nat Neurosci. 21:617–624.

Forstmann BU, Dutilh G, Ridderinkhof KR, Brown S, von Cramon DY, Wagenmakers E-J, Neumann J. 2008. Striatum and pre-SMA facilitate decision-making under time pressure. Proc Natl Acad Sci. 105:17538–17542.

Forstmann BU, Ratcliff R, Wagenmakers E. 2016. Sequential Sampling Models in Cognitive Neuroscience : Advantages , Applications , and Extensions. Annu Rev Psychol. 67:641–666.

Fried, I., Amiram, K., McCarthy, G., Sass, K.J. Williamson P, Spencer, S., Spencer D. 1991. Functional Organization of Human Supplementary Studied by Electrical Stimulation Motor Cortex. 3656–3666.

Fried I, Mukamel R, Kreiman G. 2011. Internally generated preactivation of single neurons in human medial frontal cortex predicts volition. Neuron. 69:548–562.

Frith CD, Friston K, Liddle PF, Frackowiak RSJ. 1991. Willed action and the prefrontal cortex in man: A study with PET. Proc R Soc B Biol Sci. 244:241–246.

Furstenberg A, Breska A, Sompolinsky H, Deouell LY. 2015. Evidence of Change of Intention in Picking Situations. J Cogn Neurosci. 27:2133–2146.

Galaro JK, Celnik P, Chib VS. 2019. Motor cortex excitability reflects the subjective value of reward and mediates its effects on incentive-motivated performance. J Neurosci. 39:1236–1248.

Ganos C, Asmuss L, Bongert J, Brandt V, Münchau A, Haggard P. 2015. Volitional action as perceptual detection: predictors of conscious intention in adolescents with tic disorders. Cortex. 64:47–54.

Ganos C, Kühn S, Kahl U, Schunke O, Feldheim J, Gerloff C, Roessner V, Bäumer T, Thomalla G, Haggard P, Münchau A. 2014. Action inhibition in Tourette syndrome. Mov Disord. 29:1532–1538.

Georgiou N, Bradshaw JL, Phillips JG, Bradshaw JA, Chiu E. 1995. The simon effect and attention deficits in gilles de la tourette's syndrome and huntington's disease. Brain. 118:1305–1318.

Geschwind DH, Mega MS. 1996. Alien hand syndrome : Interhemispheric motor disconnection due to a lesion. 20187:802–808.

Ghosh A, Dutt A. 2010. Utilisation behaviour in frontotemporal dementia. J Neurol Neurosurg Psychiatry. 81:154–156.

Gibson JJ. 1966. The senses considered as Perceptual Systems. London: Allen and Unwin.

Gold JI, Shadlen MN. 2000. Representation of a perceptual decision in developing oculomotor commands. Nature. 404:390–394.

Goodale MA, Milner AD. 1992. Separate visual pathways for perception and action. Trends Neurosci. 15:20–25.

Grafton ST, Hamilton AFDC. 2007. Evidence for a distributed hierarchy of action representation in the brain. Hum Mov Sci. 26:590–616.

Guggisberg AG, Mottaz A. 2013. Timing and awareness of movement decisions: does consciousness really come too late? Front Hum Neurosci. 7:385.

Hadfield JD. 2010. MCMC methods for multi-response generalized linear mixed models: The MCMCglmm R package. J Stat Softw. 33:1–22.

Haggard P. 2005. Conscious intention and motor cognition. Trends Cogn Sci.

Haggard P. 2018. The Neurocognitive Bases of Human Volition. Annu Rev Psychol. 70:9–28.

Haggard P, Eimer M. 1999. On the relation between brain potentials and the awareness of voluntary movements. Exp Brain Res. 126:128–133.

Haggard P, Tsakiris M. 2009. The Experience of Agency. Curr Dir Psychol Sci. 18:242–246.

Hanks TD, Ditterich J, Shadlen MN. 2006. Microstimulation of macaque area LIP affects decision-making in a motion discrimination task. Nat Neurosci. 9:682–689.

Hanks TD, Kiani R, Shadlen MN. 2014. A neural mechanism of speed-accuracy tradeoff in macaque area LIP. Elife. 2014:1–17.

Heidegger, M. 1996. Being and time: A translation of Sein und Zeit. SUNY press.

Heitz RP, Schall JD. 2012. Neural Mechanisms of Speed-Accuracy Tradeoff. Neuron. 76:1–7.

Hughes G, Schütz-Bosbach S, Waszak F. 2011. One action system or two? Evidence for common central preparatory mechanisms in voluntary and stimulus-driven actions. J Neurosci. 31:16692–16699.

Hyder F, Phelps EA, Wiggins CJ, Labar KS, Blamire AM, Shulman RG. 1997. “ Willed action ”: A functional MRI study of the human prefrontal cortex during a sensorimotor task. 94:6989–6994.

Iaccarino L, Chieffi S, Iavarone A. 2014. Utilization behavior: What is known and what has to be known. Behav Neurol.

Ings TC, Chittka L. 2008. Speed-Accuracy Tradeoffs and False Alarms in Bee Responses to Cryptic Predators. Curr Biol.

Ivanoff J, Branning P, Marois R. 2008. fMRI evidence for a dual process account of the speed-accuracy tradeoff in decision-making. PLoS One. 3.

Jahanshahi M, Frith CD. 1998. Willed action and its impairments. Cogn Neuropsychol. 15:483–533.

Jahanshahi M, Jenkins H, Brown RG, Marsden CD, Passingham RE, Brooks DJ. 1995. Self-initiated versus externally triggered movements I . An investigation using measurement of regional cerebral blood flow with PET and movement-related potentials in normal and Parkinson ' s disease subjects. Brain. 118:913–933.

James W. 1890. The principles of psychology.

Jha A, Diehl B, Scott C, McEvoy AW, Nachev P. 2016. Reversed Procrastination by Focal Disruption of Medial Frontal Cortex. Curr Biol. 26:2893–2898.

Jo HG, Hinterberger T, Wittmann M, Borghardt TL, Schmidt S. 2013. Spontaneous EEG fluctuations determine the readiness potential: Is preconscious brain activation a preparation process to move? Exp Brain Res. 231:495–500.

Kagaya K, Takahata M. 2010. Readiness discharge for spontaneous initiation of walking in crayfish. J Neurosci. 30:1348–1362.

Kalaska JF, Crammond DJ. 1995. Deciding not to GO: Neuronal correlates of response selection in a GO/NOGO task in primate premotor and parietal cortex. Cereb Cortex. 5:410–428.

Keller I, Heckhausen H. 1990. Readiness potentials preceding spontaneous motor acts: voluntary vs. involuntary control. Electroencephalogr Clin Neurophysiol. 76:351–361.

Kelly SP, O'Connell RG. 2013. Internal and External Influences on the Rate of Sensory Evidence Accumulation in the Human Brain. J Neurosci. 33:19434–19441.

Khalighinejad N, Schurger A, Desantis A, Zmigrod L, Haggard P. 2018. Precursor processes of human self-initiated action. Neuroimage. 165:35–47.

Kiani R, Corthell L, Shadlen MN. 2014. Choice certainty is informed by both evidence and decision time. Neuron. 84:1329-1342.

Kiani R, Cueva CJ, Reppas JB, Newsome WT. 2014. Dynamics of neural population responses in prefrontal cortex indicate changes of mind on single trials. Curr Biol. 24:1542-1547

Kim JN, Shadlen MN. 1999. Neural correlates of a decision in the dorsolateral prefrontal cortex of the macaque. Nat Neurosci. 2:176–185.

Kleiner, M, Brainard, D, Pelli, D, Ingling, A, Murray, R, Broussard, C. 2007. What's new in Psychtoolbox-3? Perception 36 (14):1–16.

Kornhuber, HH, Deecke, L. 1965. Hirnpotentialänderungen bei Willkürbewegungen und passiven Bewegungen des Menschen: Bereitschaftspotential und reafferente Potentiale. Plfüger's Archiv für die gesamte Physiologie des Meschen und der Tiere, 284:1-17.

Kranick SM, Hallett M. 2013. Neurology of volition. Exp Brain Res. 229:313–327.

Kühn S, Brass M. 2009. Retrospective construction of the judgement of free choice. Conscious Cogn. 18:12–21.

Kulakova E, Khalighinejad N, Haggard P. 2017. I could have done otherwise : Availability of counterfactual comparisons informs the sense of agency. Conscious Cogn. 49:237–244.

Kurata K, Wise SP. 1988. Premotor and supplementary motor cortex in rhesus monkeys: neuronal activity during externally- and internally-instructed motor tasks. Exp Brain Res. 72:237–248.

Laplane D, Talairach J, Meininger V, Bancaud J, Orgogozo JM. 1977. Clinical consequences of corticectomies involving the supplementary motor area in man. J Neurol Sci. 34:301–314.

Lau H, Rosenthal D. 2011. Empirical support for higher-order theories of conscious awareness. Trends Cogn Sci. 15:365–373.

Lau HC, Rogers RD, Passingham RE. 2007. Manipulating the experienced onset of intention after action execution. J Cogn Neurosci. 19:81–90.

Leckman JF. 2002. Tourette's syndrome. Lancet. 360:1577–1586.

Lenth R, Singmann H, Love J, Buerkner P, Herve M. 2018. Package ‘ emmeans ’ R topics documented : R Packag version 115-15. 34:216–221.

Lhermitte F. 1983. “Utilization Behaviour” and Its Relation To Lesions of the Frontal Lobes. Brain. 106:237–255.

Libet B. 1985. Unconscious cerebral initiative and the role of conscious will in voluntary action. Behav Brain Sci. 8:529–539.

Libet B. 1999. Do We Have Free Will? J Conscious Stud. 47–57.

Libet B, Gleason CA, Wright EW, Pearl DK. 1983. Time of Conscious Intention To Act in Relation To Onset of Cerebral Activity (Readiness-Potential). Brain. 106:623–642.

Lin Z, Zhang C, Wu W, Gao X. 2006. Frequency Recognition Based on Canonical Correlation Analysis for SSVEP-Based BCIs. 53:2610–2614.

Löffler A, Haggard P, Bode S. 2019. Decoding Changes of Mind in Voluntary Action — Dynamics of Intentional Choice Representations. 1–14.

Mackie J. 1965. Causes and conditions. Am Philos Q. 2:245–264.

Mansfield EL, Karayanidis F, Jamadar S, Heathcote A, Forstmann BU. 2011. Adjustments of response threshold during task switching: A model-based functional magnetic resonance imaging study. J Neurosci. 31:14688–14692.

Maoz U, Yaffe G, Koch C, Mudrik L. 2017. Neural precursors of decisions that matter - an ERP study of deliberate and arbitrary choice. bioRxiv.

Maris E, Oostenveld R. 2007. Nonparametric statistical testing of EEG- and MEG-data. J Neurosci Methods. 164:177–190.

Martí-Fàbregas J, Kulisevsky J, Baró E, Mendoza G, Valencia C, Martí-Vilalta JL. 2000. Alien hand sign after a right parietal infarction. Cerebrovasc Dis. 10:70–72.

Matsuhashi M, Hallett M. 2008. The timing of the conscious intention to move. 28:2344–2351.

Matuschek H, Kliegl R, Vasishth S, Baayen H, Bates D. 2017. Balancing Type I error and power in linear mixed models. J Mem Lang. 94:305-315.

Maye A, Hsieh CH, Sugihara G, Brembs B. 2007. Order in spontaneous behavior. PLoS One. 2:e443.

McCurdy LY, Maniscalco B, Metcalfe J, Liu KY, de Lange FP, Lau H. 2013. Anatomical coupling between distinct metacognitive systems for memory and visual perception. J Neurosci. 33:1897–1906.

McKinney W. 2010. Data Structures for Statistical Computing in Python. Proc 9th Python Sci Conf. 51–55.

Mele AR. 2011. Libet on Free Will: Readiness Potentials, Decisions, and Awareness. In: Conscious Will and Responsibility: A Tribute to Benjamin Libet.

Metcalfe J, Greene MJ. 2007. Metacognition of agency. J Exp Psychol Gen. 136:184–199.

Miller J, Shepherdson P, Trevena J. 2011. Effects of clock monitoring on electroencephalographic activity: Is unconscious movement initiation an artifact of the clock? Psychol Sci. 22:103–109.

Mirabella G. 2007. Endogenous Inhibition and the Neural Basis of “Free Won’t.” J Neurosci. 27:13919–13920.

Mishkin M, Ungerleider LG, Macko KA. 1983. Object vision and spatial vision: two cortical pathways. Trends Neurosci. 6:414–417.

Misirlisoy E, Brandt V, Ganos C, Tübing J, Münchau A, Haggard P. 2015. The Relation Between Attention and Tic Generation in Tourette Syndrome. 29:658–665.

Mordkoff JT, Gianaros PJ. 2000. Detecting the onset of the lateralized readiness potential: a comparison of available methods and procedures. Psychophysiology. 37:347–360.

Moretto G, Schwingenschuh P, Katschnig P, Bhatia KP, Haggard P. 2011. Delayed experience of volition in Gilles de la Tourette syndrome. J Neurol Neurosurg Psychiatry. 82:1324–1327.

Müller MM, Teder-Sälejärvi W, Hillyard SA. 1998. The time course of cortical facilitation during cued shifts of spatial attention. Nat Neurosci. 1:631–634.

Munakata Y, Herd SA, Chatham CH, Depue BE, Banich MT, O'Reilly RC. 2011. A unified framework for inhibitory control. Trends Cogn Sci. 15:453–459.

Murakami M, Shteingart H, Loewenstein Y, Mainen ZF. 2017. Distinct Sources of Deterministic and Stochastic Components of Action Timing Decisions in Rodent Frontal Cortex. Neuron. 94:908–919.

Murakami M, Vicente MI, Costa GM, Mainen ZF. 2014. Neural antecedents of self-initiated actions in secondary motor cortex. Nat Neurosci. 17:1574–1582.

Nachev P, Husain M. 2010. Action and the fallacy of the 'internal': Comment on Passingham et al. Trends Cogn Sci. 14:192–193.

Nachev P, Kennard C, Husain M. 2008. Functional role of the supplementary and pre-supplementary motor areas. Nat Rev Neurosci. 9:856–869.

Nagamine T, Kajola M, Salmelin R, Shibasaki H, Hari R. 1996. Movement-related slow cortical magnetic fields and changes of spontaneous MEG- and EEG-brain rhythms. Electroencephalogr Clin Neurophysiol. 99:274–286.

Nahmias E. 2011. Why "willusionism" leads to "bad results": Comments on Baumeister, Crescioni, and Alquist. Neuroethics. 4:17–24.

Nelson TO, Narens L. 1990. Metamemory: A theoretical framework and new findings. Psychol Learn Motiv. 26:125–173.

O'Connell RG, Dockree PM, Kelly SP. 2012. A supramodal accumulation-to-bound signal that determines perceptual decisions in humans.

O'Connell RG, Murphy PR. 2018. U-turns in the brain. Nat Neurosci. 21:461–462.

Obhi SS, Haggard P. 2004. Internally generated and externally triggered actions are physically distinct and independently controlled. Exp brain Res. 156:518–523.

Okano K, Tanji J. 1987. Neuronal activities in the primate motor fields of the agranular frontal cortex preceding visually triggered and self-paced movement. Exp Brain Res. 421–426.

Oostenveld R, Fries P, Maris E, Schoffelen JM. 2011. FieldTrip: Open source software for advanced analysis of MEG, EEG, and invasive electrophysiological data. Comput Intell Neurosci.

Overgaard M, Fazekas P. 2016. Can No-Report Paradigms Extract True Correlates of Consciousness? Trends Cogn Sci. 20:241-241.

Overgaard M, Rote J, Mouridsen K, Zoe T. 2006. Is conscious perception gradual or dichotomous ? A comparison of report methodologies during a visual task. 15:700–708.

Pacherie E. 2008. The phenomenology of action: a conceptual framework. Cognition. 107:179–217.

Padmanabhan K, Urban NN. 2010. Intrinsic biophysical diversity decorrelates neuronal firing while increasing information content. Nat Neurosci. 13:1276–1282.

Papa SM, Artieda J, Obeso JA. 1991. Cortical activity preceding self-initiated and externally triggered voluntary movement. Mov Disord. 6:217–224.

Parise C V., Ernst MO. 2017. Noise, multisensory integration, and previous response in perceptual disambiguation. PLoS Comput Biol. 13:1–20.

Passingham RE. 1987. Two cortical systems for directing movement. Ciba Found Symp. 132:151–164.

Peirce J, Gray JR, Simpson S, MacAskill M, Höchenberger R, Sogo H, Kastman E, Lindeløv JK. 2019. PsychoPy2: Experiments in behavior made easy. Behav Res Methods. 51:195–203.

Picard N, Strick P. 1996. Motor Areas of the Medial Wall : A Review of Their Location and Functional Activation. Cereb Cortex. 342–353.

Polich J. 2009. Updating P300: An Integrative Theory of P3a and P3b. 118:2128–2148.

Powers DMW. 2011. Evaluation: From Precision, Recall and F-Measure To Roc, Informedness, Markedness & Correlation. J Mach Learn Technol. 2:37–63.

Praamstra P, Stegeman DF, Horstink MWIM, Cools AR. 1996. Dipole source analysis suggests selective modulation of the supplementary motor area contribution to the readiness potential. Electroencephalography and clinical. Electroencephalogr Clin Neurophysiol. 98:468–477.

Proust J. 2008. Epistemic agency and metacognition: An externalist view. Proc Aristot Soc. 108:241–268.

Ratcliff R. 1978. A theory of memory retrieval. Psychol Rev. 85:59–108.

Ratcliff R, Smith PL, Brown SD, McKoon G. 2016. Diffusion Decision Model: Current Issues and History. Trends Cogn Sci. 20:260–281.

Resulaj A, Kiani R, Wolpert DM, Shadlen MN. 2009. Changes of mind in decision-making. Nature. 461:263–266.

Rinberg D, Koulakov A, Gelperin A. 2006. Speed-Accuracy Tradeoff in Olfaction. Neuron.

Rizzolatti G, Luppino G. 2001. The cortical motor system. Neuron. 31:889–901.

Rizzolatti G, Luppino G, Matelli M. 1998. The organization of the cortical motor system: new concepts. Electroencephalogr Clin Neurophysiol.

Romo R, Schultz W. 1987. Neuronal activity preceding self-initiated or externally timed arm movements in area 6 of monkey cortex. Exp Brain Res. 67:656–662.

Roskies A. 2011. Why Libet's Studies Don't Pose a Threat to Free Will. In: Sinnott-Armstrong W,, Nadel L, editors. Conscious Will and Responsibility: A Tribute to Benjamin Libet. Oxford University Press. p. 45–66.

Rouder JN, Speckman PL, Sun D, Morey RD, Iverson G. 2009. Bayesian t tests for accepting and rejecting the null hypothesis. Psychon Bull Rev. 16:225–237.

Schlegel A, Alexander P, Sinnott-Armstrong W, Roskies A, Tse PU, Wheatley T. 2013. Barking up the wrong free: Readiness potentials reflect processes independent of conscious will. Exp Brain Res. 229:329–335.

Schultze-Kraft M, Birman D, Rusconi M, Allefeld C, Görgen K, Dähne S, Blankertz B, Haynes J-D. 2016. The point of no return in vetoing self-initiated movements. Proc Natl Acad Sci U S A. 113:1080–1085.

Schurger A, Sitt JD, Dehaene S. 2012. An accumulator model for spontaneous neural activity prior to self-initiated movement. Proc Natl Acad Sci U S A. 109:E2904-13.

Seal J, Gross C, Bioulac B. 1982. Activity of neurons in area 5 during a simple arm movement in monkeys before and after deafferentation of the trained limb. Brain Res. 250:229–243.

Shadlen MN, Newsome W. 1996. Motion perception: Seeing and deciding. 93:628–633.

Shadlen MN, Newsome WT. 2001. Neural basis of a perceptual decision in the parietal cortex (area LIP) of the rhesus monkey. J Neurophysiol. 86:1916–1936.

Shadlen N, Britten H, Newsome T, Movshon, Anthony J. 1996. A Computational Analysis of the Relationship between Neuronal and Behavioural Responses to Visual Motion. J Neurosci. 16:1486–1510.

Shea N, Boldt A, Bang D, Yeung N, Heyes C, Frith CD. 2014. Supra-personal cognitive control and metacognition. Trends Cogn Sci. 18:186–193.

Shibasaki H, Hallett M. 2006. What is the Bereitschaftspotential? Clin Neurophysiol. 117:2341–2356.

Shushruth S, Mazurek M, Shadlen MN. 2018. Comparison of decision-related signals in sensory and motor preparatory responses of neurons in area LIP. J Neurosci. 38:6350–6365.

Sirigu A, Daprati E, Ciancia S, Giraux P, Nighoghossian N, Posada A, Haggard P. 2004. Altered awareness of voluntary action after damage to the parietal cortex. Nat Neurosci. 7:80–84.

Snyder LH, Batista AP, Andersen RA. 1997. Coding of intention in the posterior parietal cortex. Nature. 386:167–170.

Soon CS, Brass M, Heinze H-J, Haynes J-D. 2008. Unconscious determinants of free decisions in the human brain. Nat Neurosci. 11:543–545.

Soon CS, Brass M, Heinze H, Haynes J. n.d. Unconscious determinants of free decisions in the human brain. 49:1–26.

Soon CS, He AH, Bode S, Haynes J-D. 2013. Predicting free choices for abstract intentions. Proc Natl Acad Sci U S A. 110:6217–6222.

Steinemann NA, O’Connell RG, Kelly SP. 2018. Decisions are expedited through multiple neural adjustments spanning the sensorimotor hierarchy. Nat Commun. 9:3627.

Strafella AP, Paus T. 2000. Modulation of cortical excitability during action observation: A transcranial magnetic stimulation study. Neuroreport. 11:2289–2292.

Synofzik M, Vosgerau G, Newen A. 2008. Beyond the comparator model: A multifactorial two-step account of agency. Conscious Cogn. 219–239.

Tanji J, Hoshi E. 2001. Behavioral planning in the prefrontal cortex. Curr Opin Neurobiol. 11:164–170.

Thura D, Cisek P. 2014. Deliberation and commitment in the premotor and primary motor cortex during dynamic decision making. Neuron. 81:1401–1416.

Thura D, Cisek P. 2016. Modulation of Premotor and Primary Motor Cortical Activity during Volitional Adjustments of Speed-Accuracy Trade-Offs. J Neurosci. 36:938–956.

Toma K, Matsuoka T, Immisch I, Mima T, Waldvogel D, Koshy B, Hanakawa T, Shill H, Hallett M. 2002. Generators of movement-related cortical potentials: fMRI-constrained EEG dipole source analysis. Neuroimage. 17:161–173.

Trevena JA, Miller J. 2002. Cortical movement preparation before and after a conscious decision to move. Conscious Cogn. 11:162–190.

Tsuchiya N, Wilke M, Frässle S, Lamme VAF. 2015. No-Report Paradigms: Extracting the True Neural Correlates of Consciousness. Trends Cogn Sci.

Twomey DM, Murphy PR, Kelly SP, O’Connell RG. 2015. The classic P300 encodes a build-to-threshold decision variable. Eur J Neurosci. 42:1636–1643.

Uithol S, Burnston DC, Haselager P. 2014. Why we may not find intentions in the brain. Neuropsychologia. 56:129–139.

Ullmann-Margalit E, Morgenbesser S. 1977. Picking and choosing. Soc Reseach. 44:757–785.

Usher M, McClelland JL. 2001. The Leaky, Competing Accumulator Model. Psychol Rev. 108:550–592.

VaezMousavi SM, Barry RJ. 1993. Positive and negative shifts of the readiness potential: preparatory effects. Int J Psychophysiol. 15:105-113.

Veen V Van, Krug MK, Carter CS. 2014. The Neural and Computational Basis of Controlled Speed-Accuracy Tradeoff The Neural and Computational Basis of Controlled Speed – Accuracy Tradeoff during Task Performance. 1952–1965.

Venthur B, Scholler S, Williamson J, Dähne S, Treder MS, Kramarek MT, Müller KR, Blankertz B. 2010. Pyff - A pythonic framework for feedback applications and stimulus presentation in neuroscience. Front Neuroinform. 4:1–17.

Verbaarschot C, Farquhar J, Haselager P. 2015. Lost in time...The search for intentions and Readiness Potentials. Conscious Cogn. 33:300–315.

Verbaarschot C, Haselager P, Farquhar J. 2016. Detecting traces of consciousness in the process of intending to act. Exp Brain Res. 234:1945–1956.

Vierkant T, Deutschländer R, Sinnott-Armstrong W, Haynes JD. 2019. Responsibility without freedom? Folk judgements about deliberate actions. Front Psychol. 10:1–6.

Vohs KD, Schooler JW. 2008. The Value of Believing in Free Will. Psychol Sci. 19:49–54.

Wang Z, Ph D, Maia T V, Ph D, Marsh R, Ph D, Gerber A, Ph D, Peterson BS. 2011. The Neural Circuits That Generate Tics in Tourette’s Syndrome. Am J Psychiatry. 168:1326–1337.

Wickelgren WA. 1977. Speed-accuracy tradeoff and information processing dynamics. Acta Psychol (Amst). 41:67–85.

Wolpert DM, Kawato M. 1998. Multiple paired forward and inverse models for motor control. Neural Networks. 11:1317–1329.

Wyart V, de Gardelle V, Scholl J, Summerfield C. 2012. Rhythmic Fluctuations in Evidence Accumulation during Decision Making in the Human Brain. Neuron. 76:847–858.

Appendix A – Chapter 2 supplementary material

Experiment 1

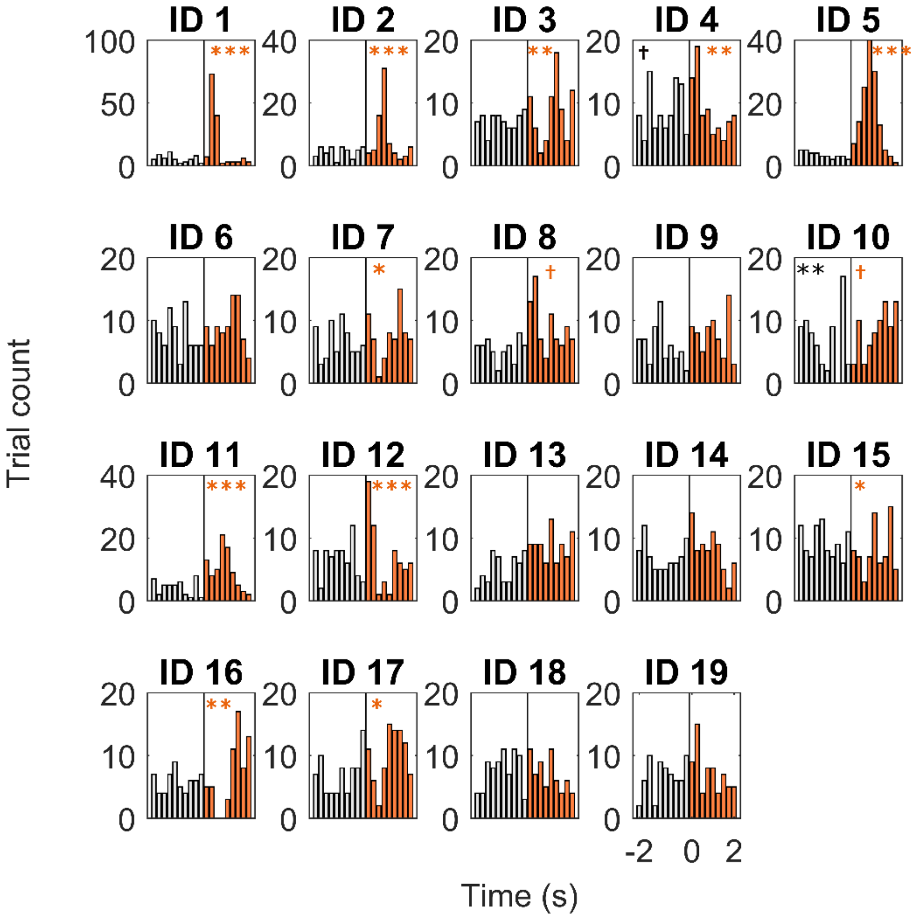

Figure A 1. Keypress distribution before and after orange letter presentation in Experiment 1. Distribution of self-paced keypresses (grey bars) and reaction times (RT, orange bars) to orange letters (presented at time 0) for each participant. A Chi-squared test was conducted to assess whether the distribution of RTs to the orange letters was uniform as would be expected if the letters were ignored. The test showed that the distribution was significantly (p<0.05, n = 11) or marginally (p < 0.1, n = 2) different form uniform in most participants (total n = 13 out of n = 19), suggesting that key presses following the letter were not simply self-paced movements that would have occurred at that time had the orange letter not been presented. In contrast, as expected, the distribution of self-paced keypresses before presentation of the orange letter (i.e. random probing) did only significantly differ from the expected uniform distribution in one participant. ***, p<0.001; **, p<0.01; *, p<0.05, †, p<0.1.

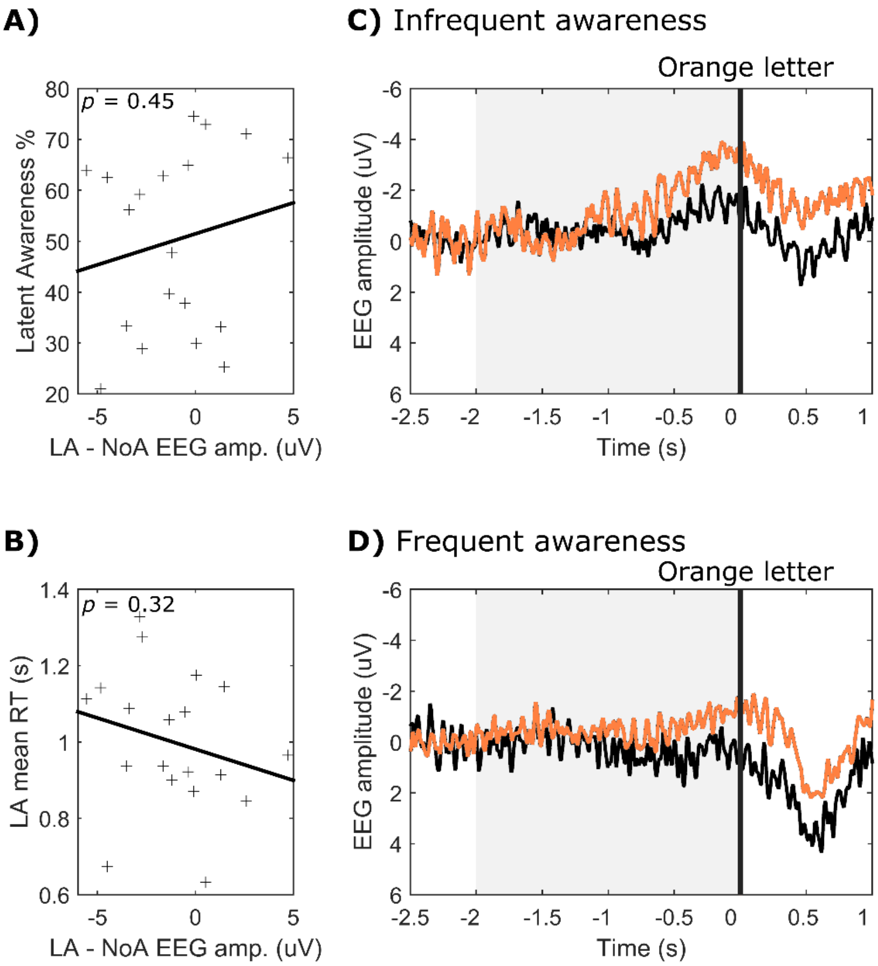

Figure A 2. Individual differences in pre-orange letter EEG activity. Stimulus-locked results are independent of awareness reports frequency and reaction times to orange probes in latent awareness epochs**.** The difference between latent awareness (LA, orange line) and no awareness (NoA, black line) stimulus-locked epochs was calculated by averaging the EEG amplitude across FCZ, CZ and C3, over [-2 to 0 s] prior to the orange letter probes. The magnitude of the difference did not correlate with the percentage of orange letters in which latent awareness was reported **(A)** nor with the mean response time to orange letters in latent awareness actions **(B)**. For descriptive purposes, we split participants in frequent (n = 10) and infrequent (n = 9) awareness reporters (i.e. reported awareness in more or less than 50% of orange letter probes respectively). The stimulus-locked grand-averaged EEG activity over motor areas (FCZ, CZ, C3) shows that the effect was visible both in infrequent **(C)** and frequent **(D)** awareness reporters**.** Note: the EEG signal was notch filtered at 4.63Hz for illustration purposes.

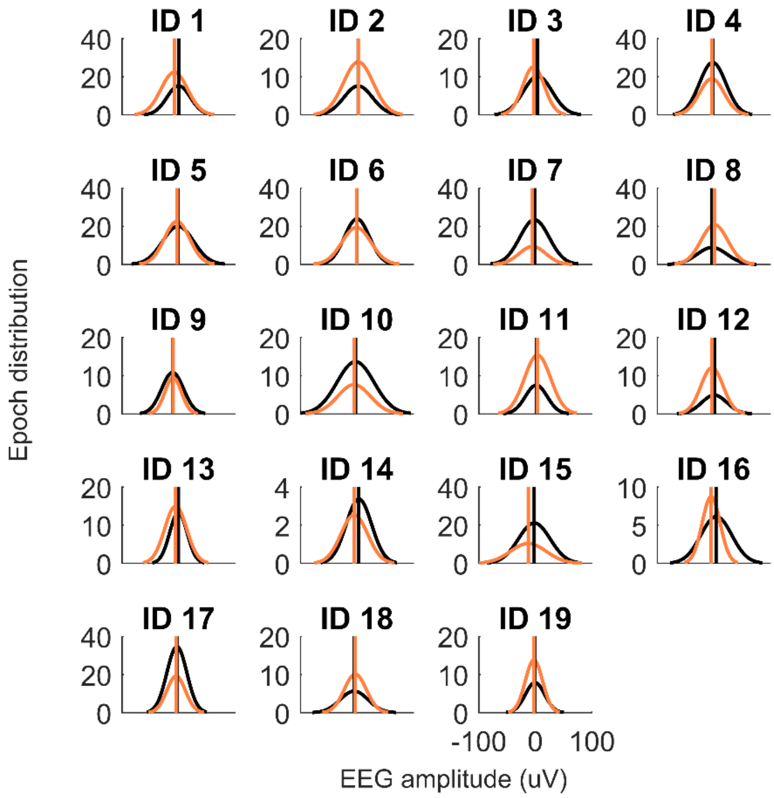

Figure A 3. Single-subject distribution of EEG amplitudes at the time of probing. EEG amplitude at Cz was averaged before orange letter presentation [-0.1 to 0s locked to orange letter] in the latent awareness (orange line) and No awareness (black line) conditions. In most participants, the mean of the latent awareness distribution was more negative than the mean of the no awareness condition. This indicates that, on average, the EEG signal was more negative when participants reported awareness than when they did not.

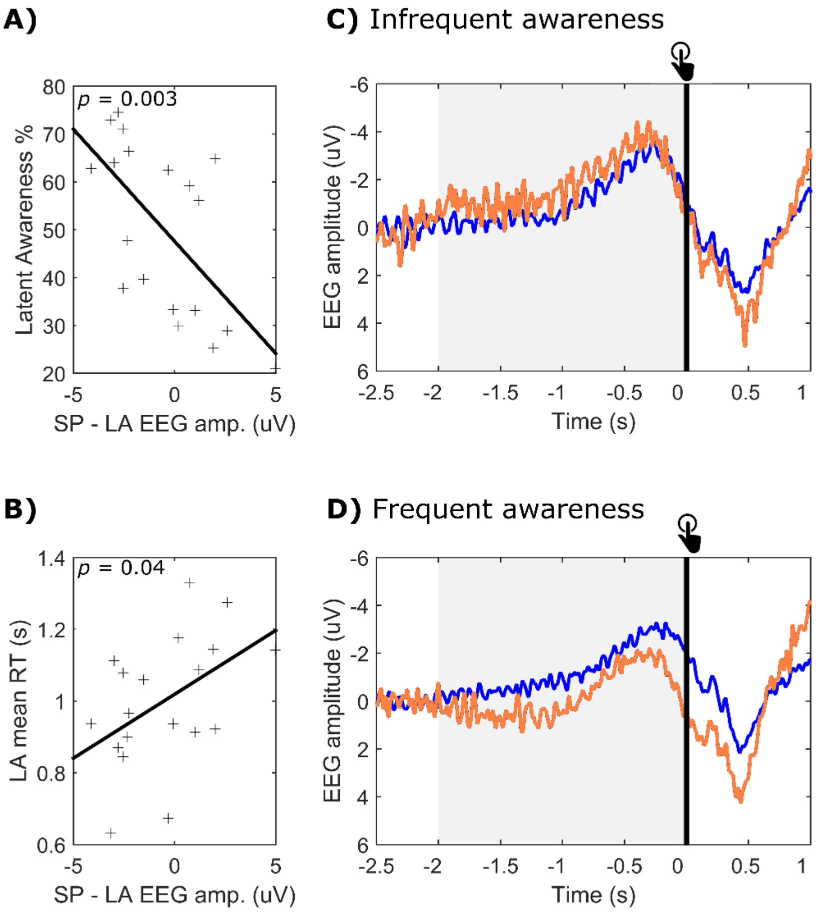

Figure A 4. Individual differences in the RP amplitude. Action-locked results are related to awareness reports frequency and reaction times to orange probes in latent awareness epochs. The difference between self-paced (SP, blue line) and latent awareness (LA, orange line) action-locked epochs was calculated by averaging the EEG amplitude across FCZ, CZ and C3, over [-2 to 0 s] prior to action. The magnitude of the difference significantly correlate with the percentage of orange letters in which latent awareness was reported **(A)** and with the mean response time to orange letters in latent awareness actions **(B)**. For descriptive purposes, we split participants in frequent (n = 10) and infrequent (n = 9) awareness reporters (i.e. reported awareness in more or less than 50% of orange letter probes respectively). The stimulus-locked grand-averaged EEG activity over motor areas (FCZ, CZ, C3) illustrates that no difference was visible preceding SP and LA actions in infrequent awareness reporters **(C)**, while the amplitude of the RP preceding SP trials was remarkably more negative than that preceding LA epochs in frequent reporters **(D).** Note: the EEG signal was notch filtered at 4.63Hz for illustration purposes.

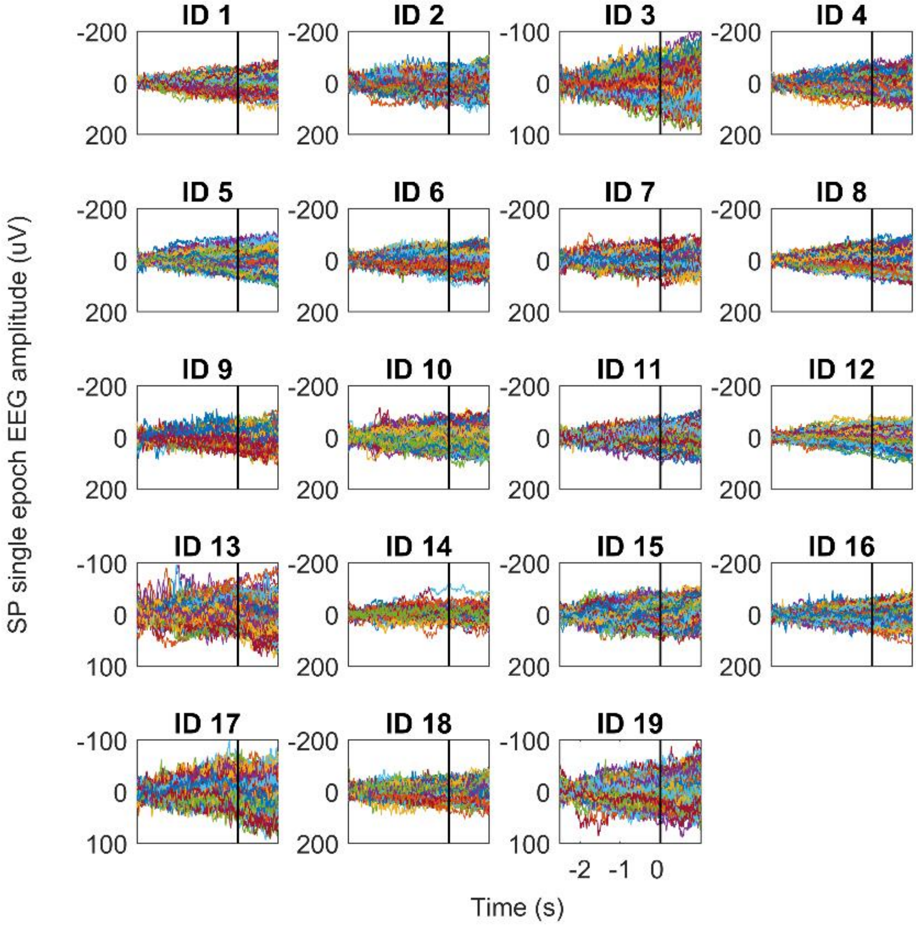

Figure A 5. Single-trial RP preceding self-paced trials. Self-paced actions can be preceded by a wide range of EEG amplitudes at single-trial level, including positive amplitudes. The averaged RP reflects the fact that brain activity preceding self-paced actions is more often negative than positive.

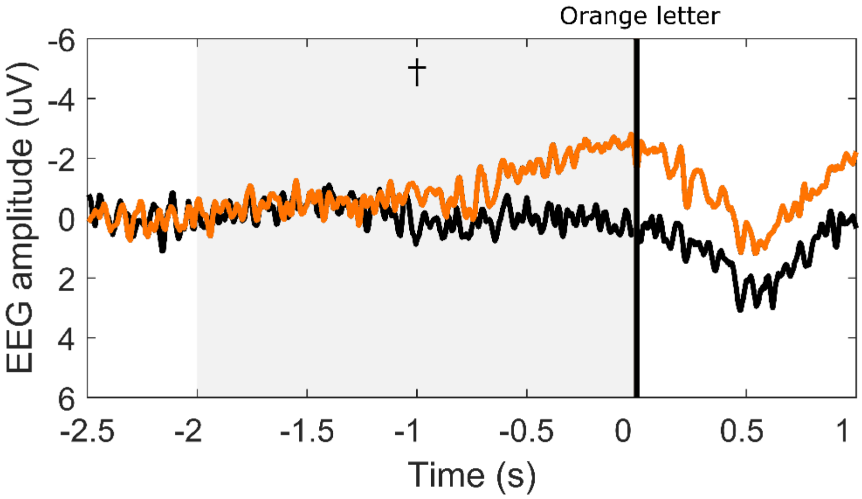

Figure A 6. Stimulus-locked analysis in a subset of participants. Stimulus-locked grand-averaged EEG activity over motor areas (FCZ, CZ, C3) in participants whose RT distribution was significantly not uniform (n = 11). The effect remained clearly visible, and still approached the conventional borders of statistical significance ($p = 0.06$). Latent awareness (orange line) epochs were preceded by greater RP-like negativities than No awareness (black line) epochs. Shaded area indicates cluster analysis period. †, $p < 0.1$.Note: the EEG signal was notch filtered at 4.63Hz for illustration purposes.

Experiment 2

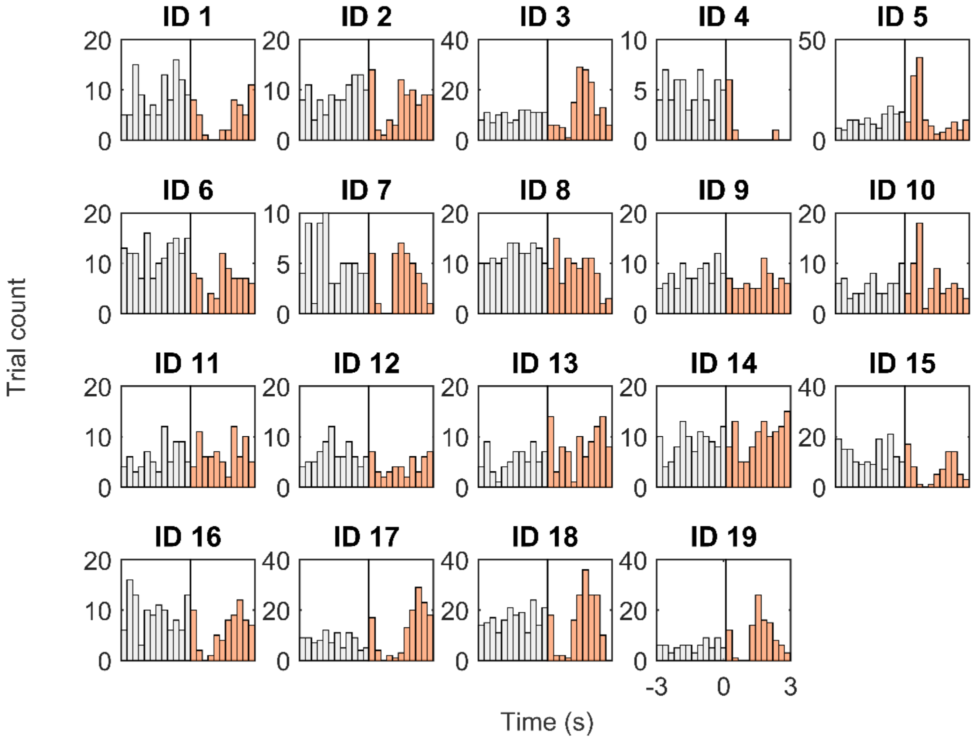

Figure A 7. Keypress distribution before and after orange letter presentation in Experiment 2. Distribution of self-paced keypresses (grey bars) and reaction times (RT, orange bars) preceding (-3 to 0) and following (0 to 3 s) orange letter presentation, for each participant in Experiment 2.

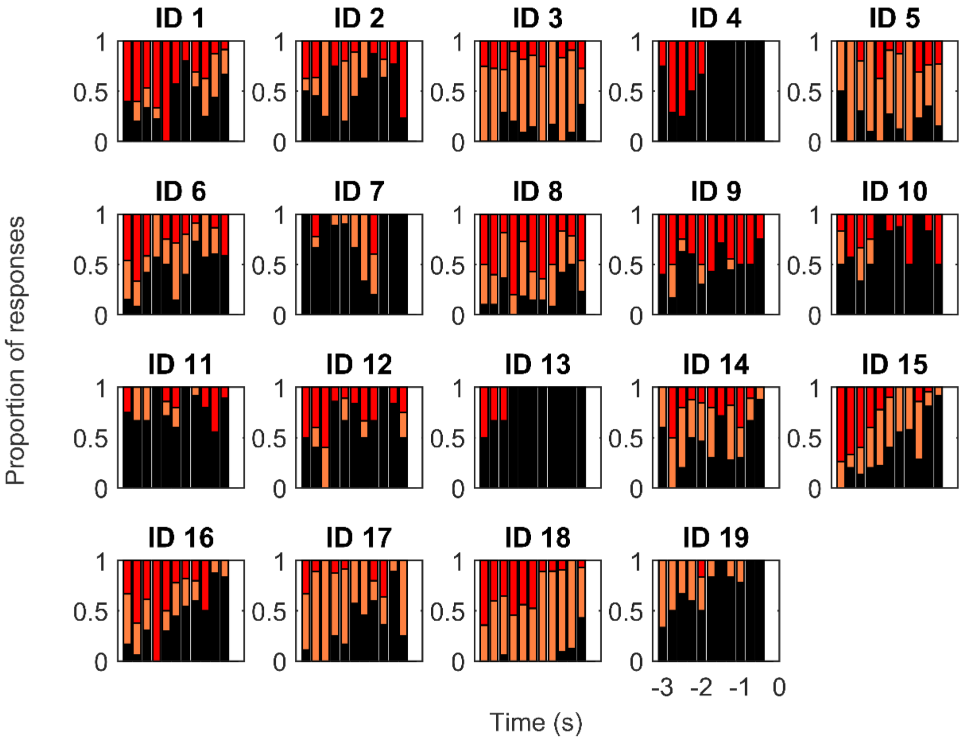

Figure A 8. Responses to orange letters given previous self-paced keypresses. Distribution of self-paced keypresses during the 3 s preceding presentation of the orange probe in Veto (*red*), No-veto (*black*) and Ignored (*orange*) conditions in Experiment 2. When orange letters were presented shortly after a self-paced keypress, participants mostly Ignored (*orange*) the probe and pressed a key within the 3 s interval that followed, or reported not vetoing (*black*) an action at the time of probe presentation. Instead, very few trials in which participants reported vetoing an action (*Red*) were preceded by a self-paced keypress shortly before the orange letter probe.

Experiment 3

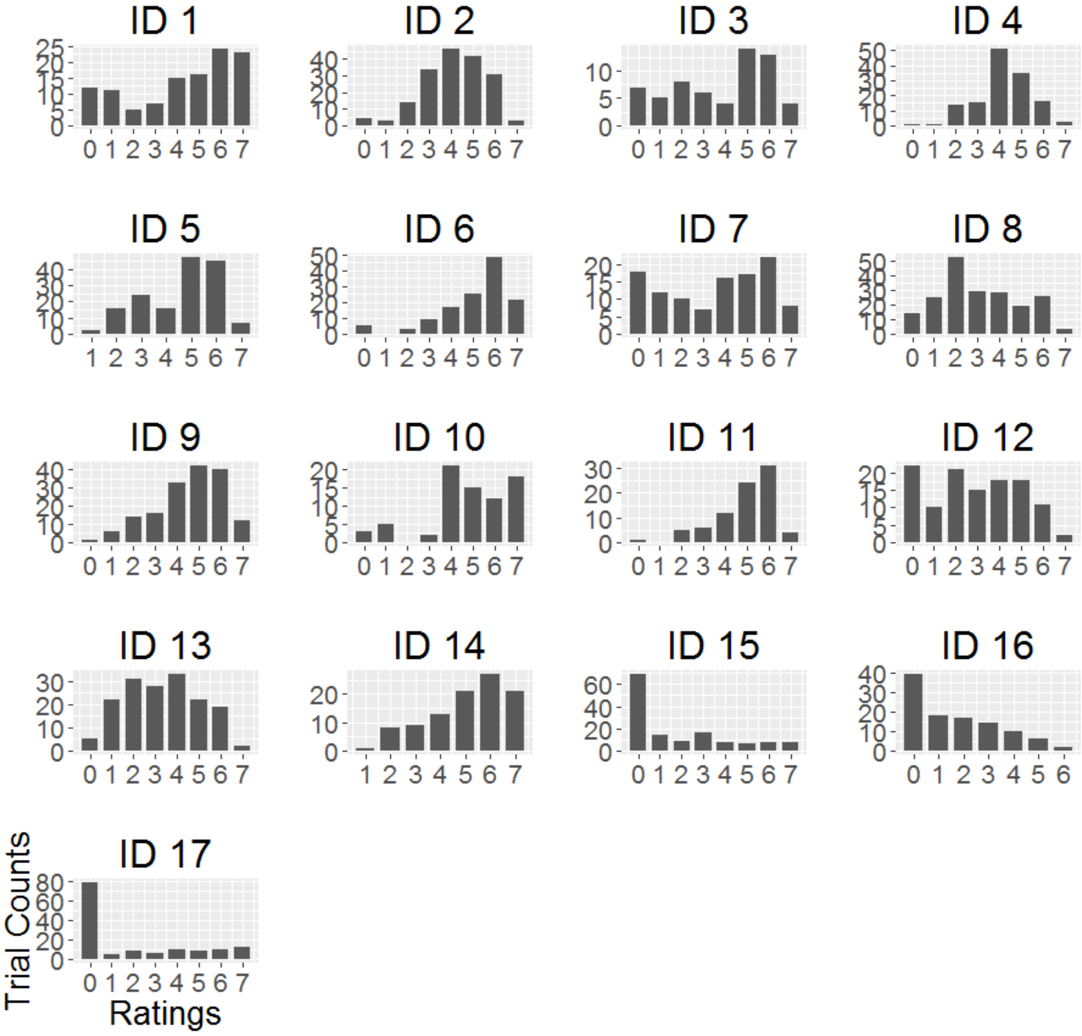

Figure A 9. Single subject readiness rating distribution.

Appendix B – Chapter 3 supplementary material

Model selection procedure and statistical details

As described in the Statistical Analysis section of the methods, we used linear mixed-effects models to test the effects of our explanatory variables on the probability of participants reporting awareness. To select the model that best explained our observed results, we followed the random effect selection procedure suggested in (Matuschek et al. 2017).

In all models, a random intercept was included to account for the variability in the dependent variable across participants. Further, we included those random effects that significantly improved the model fit. To determine the optimal random effects structure, we fit a baseline model which included all explanatory variables and all possible interactions as fixed terms. We then iteratively compared this baseline models against models with one additional random slope using a chi-squared test. If the inclusion of a random slope significantly improved the model fit, the random slope was included in the final model. This approach has been suggested as a better option than including random slopes for all fixed effects, as it decreases the probability of Type II errors while maintaining the same power against type I errors, and has previously been used in the literature (e.g. Steinemann et al. 2018). All models were fit using restricted maximum likelihood estimation (REML) with the *glmer* function in the homonymous R package (Bates, D, Mächler, M, Bolker, B.M., Walker 2015).

Tables *B1* and Table *B2* provide the detailed results of the random effect selection procedure for both main analyses and the final inference statistics reported in the main text.

<table>
<tr><th colspan="4">Supplementary Table 1: model 1 random effects selection</th></tr>
<tr><th rowspan="2">Test individual random effects</th><th colspan="3">Baseline model:
P(yes) ~ 1 + RP + Action + RP:Action + (1+|sub)</th></tr>
<tr><th>X^2</th><th>DF</th><th>p-value</th></tr>
<tr><td>yes ~ 1 + RP + Action + RP:Action + (1+RP|sub)</td><td>11.79</td><td>2</td><td>0.0027 **</td></tr>
<tr><td>yes ~ 1 + RP + Action + RP:Action + (1+Action|sub)</td><td>17.18</td><td>2</td><td><0.001***</td></tr>
</table>

Table B 1. **Model 1 selection steps and statistical results of model comparison.** We determined the optimal random effects structure with REML estimation. Random slopes for both RP and Action significantly improved the fit of the baseline model and were therefore included in the model.

<table>
<tr><th colspan="4">Supplementary Table 2: model 2 random effects selection</th></tr>
<tr><th rowspan="2">Test individual random effects</th><th colspan="3">Baseline model:
P(yes) ~ 1 + RP + RT + RP:RT + (1+|sub)</th></tr>
<tr><th>X^2</th><th>DF</th><th>p-value</th></tr>
<tr><td>yes ~ 1 + RP + RT + RP:RT + (1+RP|sub)</td><td>0.2367</td><td>2</td><td>0.888</td></tr>
<tr><td>yes ~ 1 + RP + RT + RP:RT + (1+RT|sub)</td><td>6.9814</td><td>2</td><td>0.030*</td></tr>
</table>

Table B 2. Model 2 random effect selection steps and statistical results of model comparison. We determined the optimal random effects structure with REML estimation. Only the RT random slope significantly improved the fit of the baseline model and was therefore the only random effect included in the model.

Appendix C - Chapter 5 supplementary material

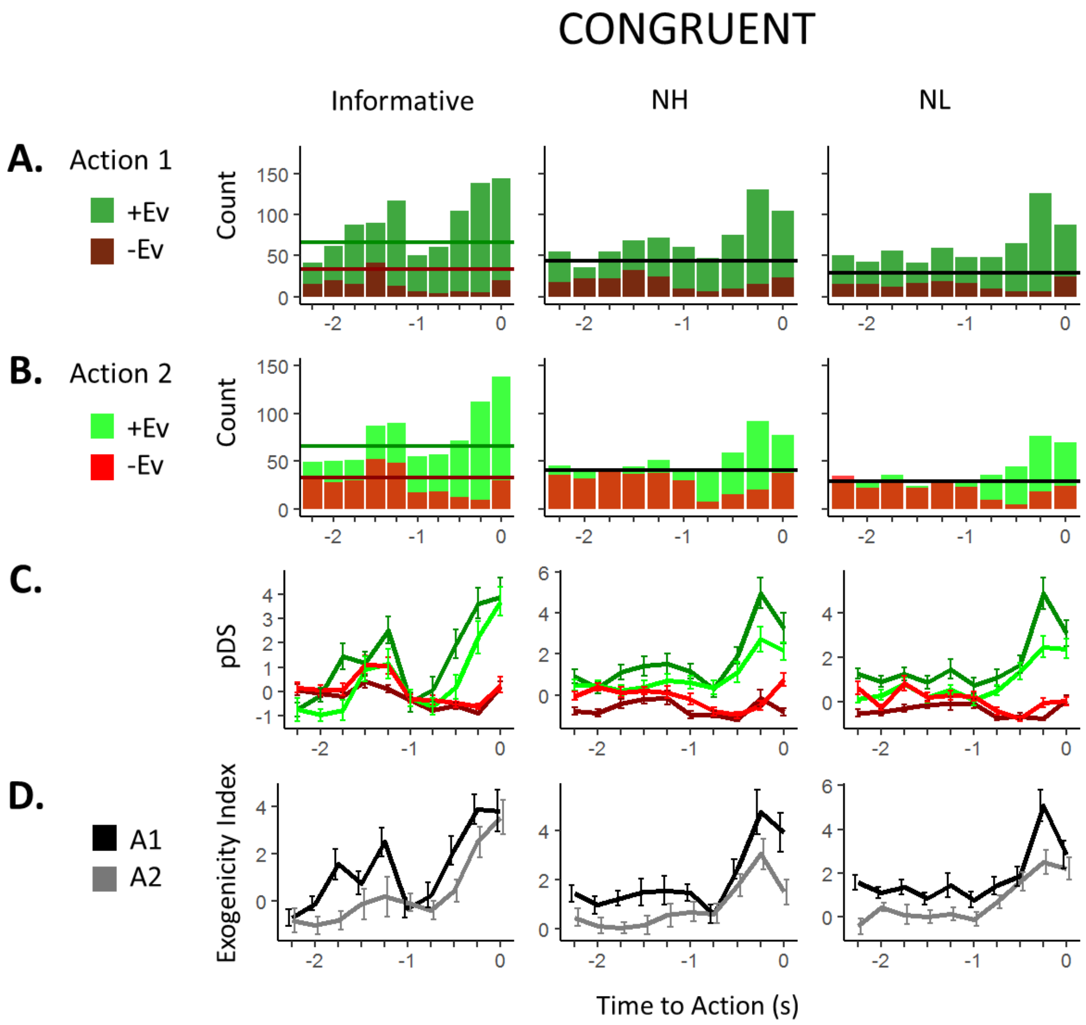

Figure C 1. Recency effect in congruent trials. **A)** Pooled distribution of +Ev (*green*) and –Ev (*red*) locked to the time of Action 1 (0 in the x axis). Horizontal lines indicate the expected nubmer of +Ev (*green* or *black*) and –Ev (*red* or *black*) in each condition, based on the probabilities and number of trials. **B)** Pooled distribution of +Ev (*green*) and –Ev (*red*) locked to the time of Action 2 (0 in the x axis). Horizontal lines indicate the expected nubmer of +Ev (*green* or *black*) and –Ev (*red* or *black*) in each condition, based on the probabilities and number of trials. **C)** Grand-averaged (± SEM) partial Deviation Score (pDS), for +Ev and –Ev separately, for Action 1 (*dark*) and Action 2 (*light*). The pDS is the difference between the observed and the expected number of +Ev and –Ev letters preceding Action 2 and the expected number for each letter position. **D)** Grand-averaged (± SEM) Exogenicity Index, before Action 1 (A1, *black*) and Action 2 (A2, *grey*). The Exogenicity Index is the difference between the +Ev pDS and the –Ev pDS.

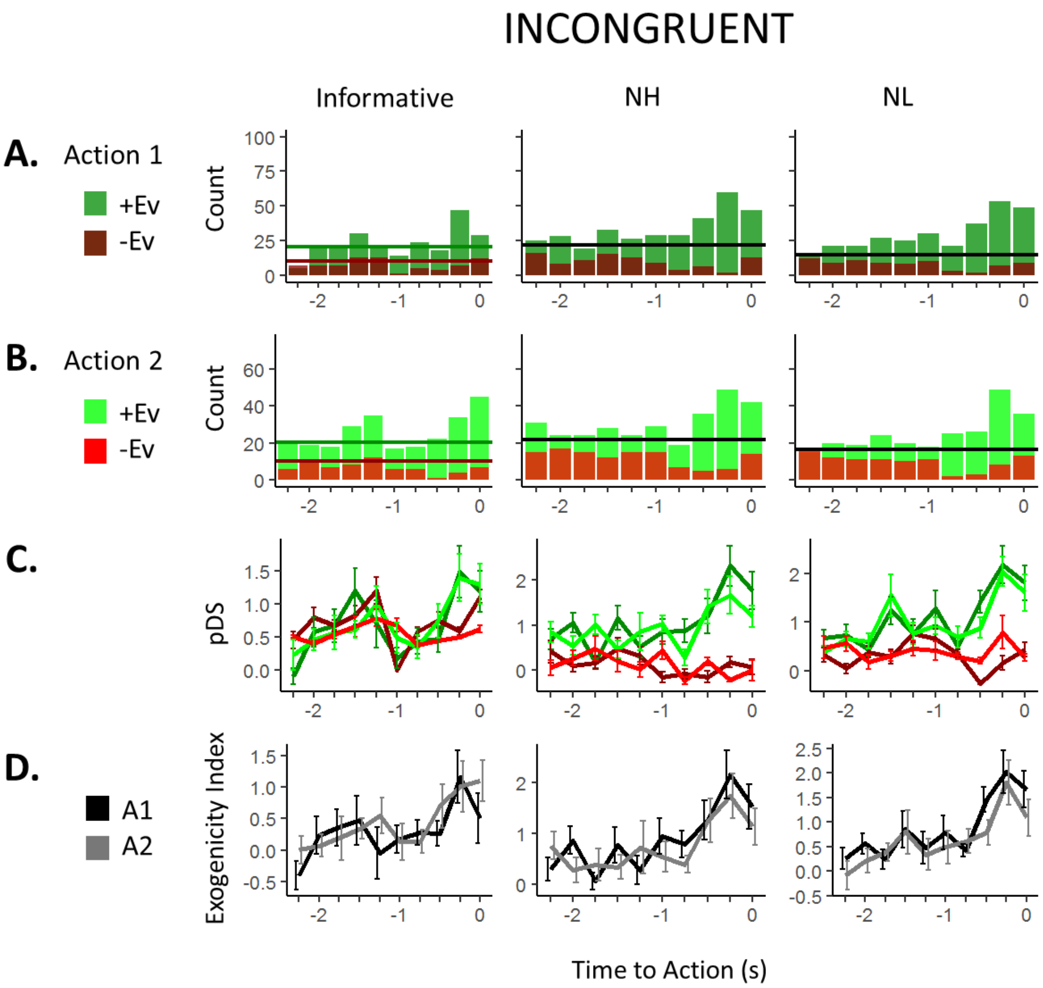

Figure C 2. Recency effect in incongruent trials. **A)** Pooled distribution of +Ev (*green*) and –Ev (*red*) locked to the time of Action 1 (0 in the x axis). Horizontal lines indicate the expected nubmer of +Ev (*green* or *black*) and –Ev (*red* or *black*) in each condition, based on the probabilities and number of trials. **B)** Pooled distribution of +Ev (*green*) and –Ev (*red*) locked to the time of Action 2 (0 in the x axis). Horizontal lines indicate the expected nubmer of +Ev (*green* or *black*) and –Ev (*red* or *black*) in each condition, based on the probabilities and number of trials. **C)** Grand-averaged (± SEM) partial Deviation Score (pDS), for +Ev and –Ev separately, for Action 1 (*dark*) and Action 2 (*light*). The pDS is the difference between the observed and the expected number of +Ev and –Ev letters preceding Action 2 and the expected number for each letter position. **D)** Grand-averaged (± SEM) Exogenicity Index, before Action 1 (A1, *black*) and Action 2 (A2, *grey*). The Exogenicity Index is the difference between the +Ev pDS and the –Ev pDS.

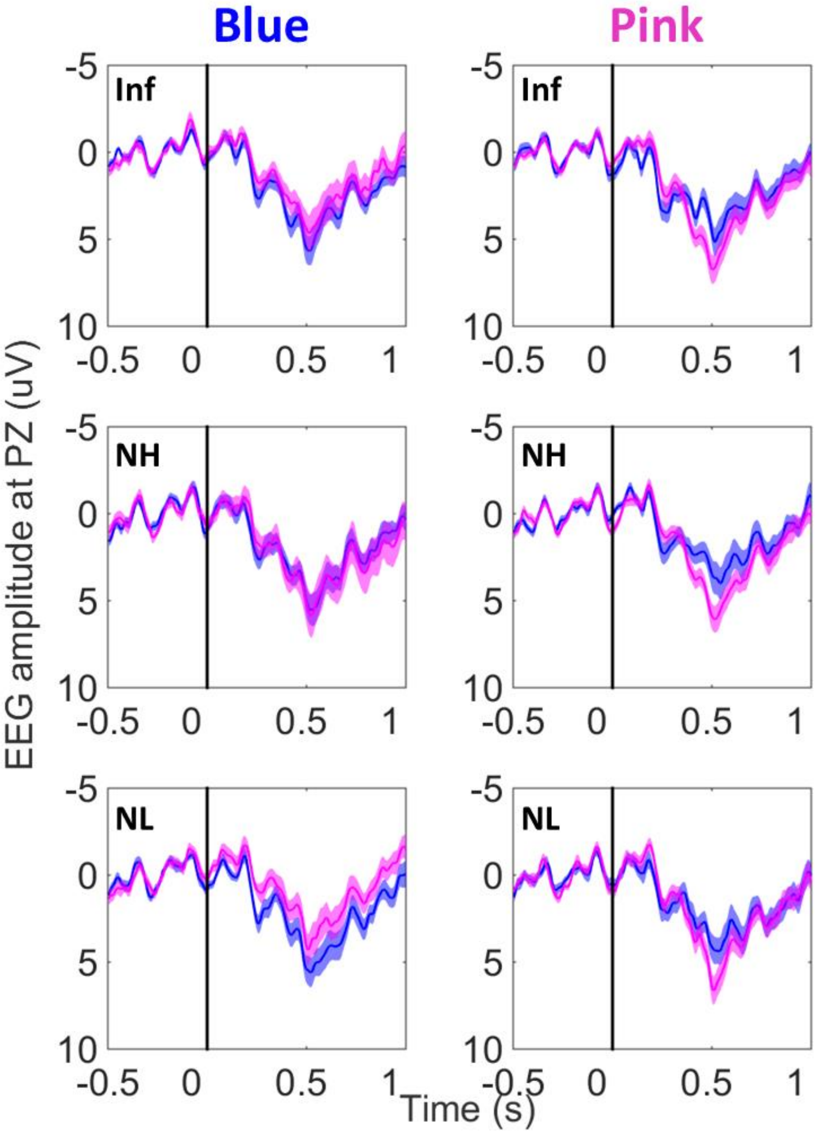

Figure C 3. Colour-specific p3 averages in congruent trials. Grand-averaged p3 amplitudes in congruent trials up to the time of Action 2 (collapsed for Action 1 and Action 2), for 'Blue' and 'Pink' decisions separately.