



A University of Sussex PhD thesis

Available online via Sussex Research Online:

<http://sro.sussex.ac.uk/>

This thesis is protected by copyright which belongs to the author.

This thesis cannot be reproduced or quoted extensively from without first obtaining permission in writing from the Author

The content must not be changed in any way or sold commercially in any format or medium without the formal permission of the Author

When referring to this work, full bibliographic details including the author, title, awarding institution and date of the thesis must be given

Please visit Sussex Research Online for more information and further details

**The role of body, action, and consciousness in adaptive
interactions with the world: An investigation of
instrumental learning.**

Lina I. Skora

Thesis submitted for the degree of Doctor of Philosophy
UNIVERSITY OF SUSSEX

July 2020

SUMMARY

Recent years have seen an increasing theoretical and empirical appreciation of bodily influences on adaptive processes. The body and action are proposed to play a central part in our interactions with the world through motivating behaviours, colouring perception with emotion, and shaping conscious experience. Instrumental learning is a fundamental substrate of adaptive behaviour, and can be used as a vehicle to understand the relationships between those processes. This work investigates the impact of bodily information and the need for conscious access in instrumental learning, as well as whether instrumentally learned, active associations shape conscious experience.

Chapters 2 and 3 ask whether cardiac information affects simple forms of adaptive behaviour, such as unconscious instrumental conditioning. The results show evidence for absence of unconscious learning, in contrast to previous reports, and absence of any learning-related cardiac activity without stimulus awareness. Together, those chapters show that instrumental conditioning may require conscious awareness. Chapters 4 and 5 further investigate the feasibility of unconscious instrumental learning. Chapter 4 is a Stage 1 Registered Replication of a prominent paradigm demonstrating unconscious instrumental learning, leveraging statistical and methodological advances. Chapter 5 constitutes a conceptual replication of the same paradigm in two modes of conditioning – trace and delay – demonstrating absence of successful instrumental conditioning without conscious awareness. Chapter 6 shifts focus from the body on the inside to the body on the outside, using instrumental learning to examine the effect of action on our conscious experience. The results demonstrate that access to consciousness is facilitated by our instrumental actions in the world.

Overall, this body of work extends the current understanding of instrumental learning as a fundamental component of adaptive behaviour, showing that conscious access is required to drive adaptive interactions with the world, and to further shape our conscious perception in line with action.

STATEMENT

This thesis is written in a 'paper-style' format, where each experimental chapter is self-contained, written in a style that is appropriate for publication in peer-reviewed journals in the field. They are tied together by the theoretical overview and general discussion. All empirical work within this thesis is original research that I carried out during the period of registration on the doctoral research programme at the University of Sussex. I am the lead author of all chapters, and held responsibility for conceptualisation, experimental design, data collection, data analysis, and writing. The work was supervised primarily by Dr Ryan Scott, as well as Prof Anil Seth and Dr Daniel Campbell-Meiklejohn on a project-specific basis. As supervisors, Dr Scott, Prof Seth, and Dr Campbell-Meiklejohn provided oversight of the experimental approach, and contributed to design of the studies, review of the written work, and editing of the manuscripts for publication.

Chapter 4 within this thesis was produced in collaboration with Prof Zoltan Dienes, who provided methodological guidance, and Dr James Livermore, who provided analytical support. Data collection for chapters 2, 3, and 6 was supported on a voluntary basis by undergraduate and postgraduate students within the School of Psychology, who used the data independently for their final year projects.

Research was funded by the Sackler Centre for Consciousness Science, supported by the Sir Mortimer and Theresa Sackler Foundation, and the School of Psychology at the University of Sussex.

I hereby declare that this thesis has not been, and will not be, submitted in whole or in part to another university for the award of any other degree.

Lina Skora

.....

Lina I. Skora
31.07.2020

ACKNOWLEDGEMENTS

First and foremost, I would like to thank my primary supervisor, **Dr Ryan Scott** – the most enthusiastic and supportive supervisor imaginable, who could always dispel my moments of panic or lack of confidence. A huge thanks to my other supervisors – **Prof Anil Seth**, for providing support, a stimulating environment, and fascinating theories to play with; and **Dr Dan Campbell-Meiklejohn**, for challenging supervision meetings and invaluable methodological advice. Thanks also to **Prof Zoltan Dienes** for sharing the knowledge of all things Bayesian.

Thanks to **my parents**, the original Doctors Skora, who have been an inspiration since I can remember. A great hug and thanks to my wonderful sister **Karolina**, for moral support and showing interest in even the most mundane aspects of my research (or politely feigning it).

A huge thanks to the wonderful people without whom this thesis would be nowhere near completion, and the process nowhere near as fun:

James Livermore, for always being there, as well as the countless conversations over wine, pep-talks, and even the loudest arguments over the nature of consciousness. I don't think you'll ever convince me about the Quantum Integrated MicroHomunculus Theory...

Bence Palfi, the best office-mate and companion in all the highs and lows of a PhD. Thanks for always being there, all the discussions and debates, and all the fun times!

George Deane, for all the theoretical sanity checks and fascinating conversations – before the PhD, during the PhD, and hopefully after too.

Lab mates **Jo Cutler**, **Jolyon Miles-Wilson** and **James Livermore**, as well as the **Sackler Centre colleagues**, especially **Keisuke Suzuki** and **Alberto Mariola**, for mental and academic support (and all the mutual rants).

Albertyna Paciorek, for our exciting collaboration and introducing me to a whole new line of research.

And all the friends who made this a stimulating and exciting time, and – knowingly or not – helped to keep me sane: **Antoine, Manuel, Vanessa, Marco, Guillaume, Alexa, Amber, Marusa, Vangelis, Kiki, Vlad, Jukka, Francesco, Riccardo**.

Lina

PUBLICATIONS

At the time of submission, the following chapters from this thesis have been accepted for publication, or are under review, in peer-reviewed journals:

Chapter 4 is a Stage 1 Registered Report, and has received in-principle acceptance at *Cortex*. Due to the University closure caused by the Covid-19 pandemic, data collection could not be completed during the period of registration on the doctoral research programme. This paper is available as:

Skora, L. I., Livermore, J. J. A., Dienes, Z., Seth, A. K., & Scott, R. B. (2020). Feasibility of unconscious instrumental conditioning: A registered replication. *Cortex* (in-principle accepted). <https://doi.org/10.31234/osf.io/p9dgn>

Chapter 5 has been submitted to *Cognition* as:

Skora, L. I., & Scott, R. B. (*under review*). Instrumental conditioning requires conscious awareness.

Chapter 6 has been submitted to *Neuroscience of Consciousness* as:

Skora, L. I., Seth, A. K., & Scott, R. B. (*submitted*). Sensorimotor predictions shape reported conscious visual experience in a breaking Continuous Flash Suppression task.

The following publication has been prepared and published during the duration of the doctoral studies, but is not part of this thesis:

Paciorek, A.*, & **Skora, L.*** (2020). Vagus nerve stimulation as a gateway to interoception. *Frontiers in Psychology*, 1569, 1-8.
<https://doi.org/10.3389/fpsyg.2020.01659>

*joint-first authorship

ABBREVIATIONS USED IN THIS THESIS

ACC – Anterior Cingulate Cortex
AIC – Anterior Insular Cortex
ANS – Autonomic Nervous System
b-CFS – breaking Continuous Flash Suppression
B – Bayes factor
CAN – Central Autonomic Network
CFS – Continuous Flash Suppression
ECG - Electrocardiography
EEG – Electroencephalography
ERN – Error-related Negativity
FEP – Free Energy Principle
fMRI – functional Magnetic Resonance Imaging
GVA – General Vagal Afferents
GLMM – Generalised Linear Mixed Model
HEP – Heartbeat Evoked Potential
HR – Heart Rate
HRL – Homeostatic Reinforcement Learning
IAcc – Interoceptive Accuracy
IGT – Iowa Gambling Task
Ne – Error-related Negativity
NTS – Nucleus of the Solitary Tract
OFC – Orbitofrontal Prefrontal Cortex
PE – Prediction error
Pe – Error-related positivity
PDT – Perceptual Discrimination Task
RL – Reinforcement Learning
SCR – Skin Conductance Response
SMH – Somatic Marker Hypothesis
S-R – Stimulus-response
VMPFC – Ventromedial Prefrontal Cortex

CONTENTS

STATEMENT	III
ACKNOWLEDGEMENTS	IV
PUBLICATIONS	V
ABBREVIATIONS USED IN THIS THESIS	VI
CONTENTS	VII
1. THEORETICAL OVERVIEW	1
1.2. EMBODIED APPROACHES TO THE STUDY OF ADAPTIVE BEHAVIOUR: RELEVANT THEORETICAL FRAMEWORKS	3
1.2.1. EMBODIED THEORIES OF EMOTION, MOTIVATION, AND DECISION-MAKING	3
1.2.1.1. The James-Lange theory of emotion	3
1.2.1.2. The Somatic Marker Hypothesis	4
1.2.1.3. Homeostatic control and allostasis	6
1.2.2. PREDICTIVE PROCESSING	9
1.2.2.1. Active inference	10
1.2.2.2. Interoceptive inference	11
1.3. THE BODY ON THE INSIDE	14
1.3.1. INTEROCEPTION	14
1.3.2. ANATOMICAL AND FUNCTIONAL ORGANISATION OF THE INTEROCEPTIVE SYSTEM	15
1.3.3. INTEROCEPTIVE EFFECTS ON PERCEPTION, COGNITION, AND AFFECT	16
1.4. THE BODY ON THE OUTSIDE	18
1.4.1. ACTION-ORIENTED APPROACHES	18
1.4.2. THE EFFECT OF ACTION ON PERCEPTION AND CONSCIOUS EXPERIENCE	20
1.5. INTERIM SUMMARY: THE EMBODIED MIND	21
1.6. THE ROLE OF THE BODY IN INSTRUMENTAL LEARNING	22
1.6.1. INSTRUMENTAL LEARNING	22
1.6.2. AUTONOMIC AND INTEROCEPTIVE PROCESSES IN LEARNING	24
1.7. UNCONSCIOUS LEARNING	26
1.7.1. EVIDENCE FOR UNCONSCIOUS LEARNING	26
1.7.2. THEORETICAL CONSIDERATIONS IN UNCONSCIOUS LEARNING	27
1.7.3. METHODOLOGICAL AND STATISTICAL ISSUES IN UNCONSCIOUS LEARNING RESEARCH	28
1.8. OVERVIEW OF THIS RESEARCH	30
1.8.1. CONCLUSIONS AND AIMS	30
1.8.2. OVERVIEW OF THE EMPIRICAL WORK	30

2. THE ROLE OF CARDIAC PRECISION IN UNCONSCIOUS INSTRUMENTAL LEARNING	32
2.1. INTRODUCTION	33
2.2. EXPERIMENT 1	35
2.2.1. METHOD	35
2.2.1.2. Participants	35
2.2.1.3. Stimuli and Materials	36
2.2.1.3. Procedure	36
2.2.2. DATA PRE-PROCESSING	38
2.2.2.1. Exclusion Criteria	38
2.2.3. RESULTS	39
2.2.3.1. Bayes Factor	39
2.2.3.2. Evidence of learning: Performance	39
2.2.3.3. Evidence of learning: Reaction times	41
2.2.3.4. Performance and interoceptive accuracy	42
2.2.4. CONCLUSIONS TO EXPERIMENT 1	43
2.3. EXPERIMENT 2	45
2.3.1. METHOD	45
2.3.1.1. Participants	45
2.3.1.2. Stimuli and Materials	45
2.3.1.3. Procedure	46
2.3.2. DATA PRE-PROCESSING	46
2.3.2.1. Exclusion Criteria	46
2.3.2.2. Interoceptive Accuracy	47
2.3.3. RESULTS	47
2.3.3.1. Evidence of learning: Performance	47
2.3.3.2. Evidence of learning: Reaction times	49
2.3.3.3. Learning and interoceptive accuracy	49
2.3.4. CONCLUSIONS TO EXPERIMENT 2	50
2.4. DISCUSSION	51
3. CARDIAC CONCOMITANTS OF UNCONSCIOUS INSTRUMENTAL LEARNING AND THE ROLE OF CARDIAC PRECISION	54
3.1. INTRODUCTION	55
3.2. METHOD	58
3.2.1. PARTICIPANTS	58
3.2.2. STIMULI AND MATERIALS	58
3.2.3. ELECTROCARDIOGRAPHY	59
3.2.4. PROCEDURE	59
3.3. DATA PRE-PROCESSING	61
3.3.1. EXCLUSION CRITERIA	61
3.3.2. INTEROCEPTIVE ACCURACY	62
3.3.3. ELECTROCARDIOGRAPHY	62
3.4. RESULTS	62
3.4.1. BAYES FACTORS	62
3.4.2. EVIDENCE OF LEARNING: PERFORMANCE	63

3.4.3. EVIDENCE OF LEARNING: REACTION TIMES	64
3.4.4. CARDIAC RESPONSES TO PERFORMANCE FEEDBACK	65
3.4.4.1. Cardiac response to rewarding vs punishing performance feedback	66
3.4.4.2. Cardiac response to rewarding and punishing cues	68
3.5. DISCUSSION	70
3.6. SUPPLEMENTARY MATERIAL	75
 4. FEASIBILITY OF UNCONSCIOUS INSTRUMENTAL CONDITIONING: A REGISTERED REPLICATION REPORT	 78
4.1. INTRODUCTION	79
4.2. PILOT: RELATIONSHIP BETWEEN LEVEL OF AWARENESS AND LEARNING ABOVE THE OBJECTIVE THRESHOLD	81
4.2.1. METHOD	82
4.2.1.1. Participants	82
4.2.1.2. Stimuli and materials	82
4.2.1.3. Procedure	83
4.2.2. ANALYSIS AND RESULTS	84
4.2.2.1. Data pre-processing	85
4.2.2.2. Awareness check	85
4.2.2.3. Main conditioning task	85
4.2.3. PILOT: DISCUSSION	86
4.3. EXPERIMENT 1: DIRECT REPLICATION	87
4.3.1. METHOD	87
4.3.1.1. Participants	87
4.3.1.2. Stimuli and materials	87
4.3.1.3. PROCEDURE	88
4.3.2. PLANNED ANALYSIS	89
4.3.2.1. Data pre-processing	89
4.3.2.2. First crucial test: awareness check	89
4.3.2.3. Second crucial test: main conditioning task	90
4.4. EXPERIMENT 2	90
4.4.1. METHOD	92
4.4.1.1. Participants	92
4.4.1.2. Stimuli and Materials	92
4.4.1.3. PROCEDURE	92
4.4.2. PLANNED ANALYSIS	93
4.4.2.1. Data pre-processing	93
4.4.2.2. Crucial test: main conditioning task	93
4.5. SUPPLEMENTARY MATERIAL	95
 5. INSTRUMENTAL CONDITIONING REQUIRES CONSCIOUS AWARENESS	 100
5.1. INTRODUCTION	101
5.2. EXPERIMENT 1	104
5.2.1. METHOD	105
5.2.1.1. Participants	105

5.2.1.2. Stimuli and Materials	105
5.2.1.3. Procedure	105
5.2.2. ANALYSES AND RESULTS	108
5.2.2.1. Exclusion criteria	108
5.2.2.2. Evidence of learning: Performance in the conditioning task	108
5.2.2.3. Exploratory analysis: Reaction times	109
5.2.3. CONCLUSIONS OF EXPERIMENT 1	109
5.3. EXPERIMENT 2	110
5.3.1. METHOD	110
5.3.1.1. Participants	110
5.3.1.2. Stimuli and Materials	111
5.3.1.3. Procedure	111
5.3.2. RESULTS	114
5.3.2.1. Exclusion criteria	114
5.3.2.2. Evidence of learning: Performance in the conditioning task	115
5.3.3. CONCLUSIONS OF EXPERIMENT 2	116
5.4. GENERAL DISCUSSION	117
6. <u>SENSORIMOTOR PREDICTIONS SHAPE REPORTED CONSCIOUS VISUAL EXPERIENCE</u>	121
6.1. INTRODUCTION	122
6.2. EXPERIMENT 1	124
6.2.1. METHOD	124
6.2.1.1. Participants	124
6.2.1.2. Stimuli and Materials	125
6.2.1.3. Procedure	125
6.2.2. DATA ANALYSIS AND RESULTS	129
6.2.2.1. Bayes Factors	129
6.2.2.2. Conditioning task	129
6.2.2.3. Breakthrough task	130
6.2.3. CONCLUSIONS OF EXPERIMENT 1	132
6.3. EXPERIMENT 2	133
6.3.1. METHOD	133
6.3.1.1. Participants	133
6.3.1.2. Stimuli and Materials	133
6.3.1.3. Procedure	134
6.3.2. RESULTS	134
6.3.2.1. Conditioning task	134
6.3.2.2. Breakthrough task	135
6.3.3. CONCLUSIONS OF EXPERIMENT 2	138
6.4. GENERAL DISCUSSION	139
7. <u>GENERAL DISCUSSION</u>	148
7.1. SUMMARY OF FINDINGS – REVISITING THE MAIN AIMS	149
7.2. THE ROLE OF THE BODY IN INSTRUMENTAL LEARNING	151
7.3. THE ROLE OF CONSCIOUSNESS IN INSTRUMENTAL LEARNING	153

7.4. THE ROLE OF INSTRUMENTAL ACTION IN SHAPING CONSCIOUS PERCEPTION	156
7.5. CONCLUSIONS	158
REFERENCE LIST	159
APPENDICES	190
APPENDIX 1: STIMULI USED IN CHAPTERS 2, 3, 5 (EXP.1)	191
APPENDIX 2: STIMULI USED IN CHAPTER 4 (PILOT)	192
APPENDIX 3: STIMULI USED IN CHAPTER 5 (EXP. 2)	193
APPENDIX 4: STIMULI USED IN CHAPTER 6	193

1.

THEORETICAL OVERVIEW

1.1. Introduction

In recent years, embodied approaches to the study of the mind have been gaining traction. While the chief role of the brain-body communication axis is to ensure survival through homeostasis, there is an increasing understanding that bodily signals drive motivated, adaptive behaviour, as well as a host of other, higher-order processes. Bodily information – that is, neural information about the ongoing state of the internal organs (interoception), as well as its position in space, and movements – have been shown to influence perceptual, cognitive and affective processes, as well as consciousness, sense of self, and even psychiatric conditions. Stronger interpretations of this view go further, proposing that the body constitutes the ‘first prior’ (Allen & Tsakiris, 2018), where our perception is coloured by, and our emotion and conscious selfhood arise directly from the brain’s interpretation of the body’s internal milieu (Seth, 2013).

Instrumental learning is a fundamental substrate of adaptive behaviour, both phylogenetically and ontogenetically, allowing agents to learn to approach rewarding, positive stimuli in their environment, and avoid bad or harmful ones. Previous work has identified bodily markers of learning, arguing that the autonomic nervous system provides an internal feedback monitoring system to support the learning process (Ullsperger, Danielmeier, & Jocham, 2014). The extent to which instrumental learning is affected by bodily information can thus shed light on the importance of the bodily signals in such adaptive processes, especially in primitive scenarios, where learning might occur without the influence of higher-order cognition. Nonetheless, the extent to which instrumental learning is possible in simple settings – such as without conscious awareness of the stimuli – is itself an unsettled case. Aside from the bodily influences, learning may require conscious awareness to be able to drive motivated instrumental responses.

This theoretical overview provides a review of the relevant embodied approaches to the study of adaptive behaviour, before reviewing the evidence for interoceptive (conceptualised as the body ‘on the inside’), and active and proprioceptive (the body ‘on the outside’) effects on higher-order processes. It then focusses on instrumental conditioning as a fundamental substrate of adaptive behaviour, and outlines its autonomic and interoceptive contributions. Finally, it evaluates the theoretical and empirical evidence for unconscious instrumental learning, touching upon methodological considerations in studying unconscious learning. The overview culminates with an outline of the aims of the present thesis, and a brief summary of the empirical work conducted.

1.2. Embodied approaches to the study of adaptive behaviour: Relevant theoretical frameworks

1.2.1. Embodied theories of emotion, motivation, and decision-making

1.2.1.1. The James-Lange theory of emotion

One of the earliest proposals that our perception of the world, and the resultant emotions and motivations might rely on the body is the James-Lange theory, put forward independently, in short succession, by William James (1884, 1891) and Carl Lange (1885)¹. Contrary to the general consensus available at the time, the theory asserted that consciously felt emotional states are the result of physiological changes associated with an event, rather than the cause of those changes. Under this view, perceiving an event in the environment immediately causes a host of relevant physiological changes, and the perception of those changes constitutes the feeling of emotion. Using James's famous example as an illustration, it is not the emotion of fear that causes us to flee from a bear encountered in the wild – instead, the emotion of fear manifests because our body is automatically engaged in the flight response. As such, the perception of physiological changes corresponding to the flight response constitutes the feeling of fear. Under the James-Lange theory, emotion is inherently embodied – in fact, James goes as far as saying that “a purely disembodied human emotion is a nonentity (...), emotion dissociated from all bodily feeling is inconceivable” (1884, p. 194). Nonetheless, the original theory is painted in broad strokes. James admitted that this account pertains only to coarse emotions, those coupled with strong physiological changes, such as anger, fear, love, hate, joy, grief, shame, or pride (James, 2001). In contrast, subtler emotions (such as “moral or intellectual” feelings) may not encompass bodily responses.

One of the most prominent criticisms of the theory was that there is no one-to-one mapping between physiological states and specific emotions (Cannon, 1915), even in the coarser subset. Instead, similar visceral changes may occur in vastly different circumstances – for instance, increased heart rate may accompany fear or excitement. This was illustrated by Schachter and Singer (1962) in a two-factor theory of emotion, proposing that emotion is constructed from physiological changes and their cognitive appraisal. This opened up a more comprehensive, constructionist view of emotional experience. While it may be debatable

¹ It is noteworthy that Lange's theory focused more specifically on vasomotor changes, rather than general physiological activation, as James's perspective did.

whether visceral changes are the direct cause of emotions, contemporary consensus is that emotional states encompass visceral, as well as behavioural and cognitive aspects (Cameron, 2002; Critchley & Harrison, 2013; Damasio, 1999; Gendron & Barrett, 2009). In this perspective, emotion emerges from the brain's interpretation of the body's internal state with relation to the external situation (Barrett & Bar, 2009; Gendron & Barrett, 2009), where this interpretation can be instinctive (akin to James, 1884), or involve other categorisation or attribution processes (e.g. stimulus properties, motives, emotion regulation, and more; Russell, 2003). This view has been recently incorporated within and expanded under the predictive processing approach (see section 1.2.3. Predictive processing).

1.2.1.2. The Somatic Marker Hypothesis

The Somatic Marker Hypothesis (SMH) proposes that decision-making is influenced by somatic (bodily) markers – physiological activity arising from the internal milieu of the body, including the viscera, as well as skeletal and smooth muscles (Damasio, 1994, 2004). Somatic markers are proposed to rise when appraising different behavioural options, especially under high uncertainty, to bias decisions through the proxy of emotional states. Similarly to the James-Lange theory, Damasio considers emotional states as the brain's representations of homeostatic changes in a given situation (Damasio, 1994, 1999, 2004). Under SMH, somatic markers are proposed to guide decision making by providing an indication of the value of the encountered stimulus through an associated emotional state – in other words, its value for the organism (Damasio, 1994, 2004). They can reflect the actual bodily state (the 'body' loop), as well as the brain's representations of whatever is expected to take place in the body in a given circumstance (the 'as-if' loop).

Most support for the SMH comes from comparing performance of patients with damage to the ventromedial prefrontal cortex (VMPFC) and orbitofrontal cortex (OFC) to healthy controls on the Iowa Gambling Task (IGT; Bechara, Damasio, Damasio, & Anderson, 1994; Bechara, Damasio, Tranel, & Damasio, 1997; Bechara, Tranel, & Damasio, 2000; Bechara, Tranel, Damasio, & Damasio, 1996). In the IGT, participants draw one card at a time from a choice of four possible decks. Each draw can result in a financial loss or gain, with fixed probabilities. Unbeknownst to the participants, two of the four decks carry high rewards if selected, but also high punishment (e.g. \$100). The remaining two carry lower rewards and lower punishment (e.g. \$50). Although the high-reward cards are initially appealing, they result in a large net loss if chosen frequently. In order to maximise long-term gain, participants should learn to draw from the lower-reward decks more frequently. On the IGT, healthy participants began to choose advantageously (i.e. from the lower-reward decks) more often, before reporting having explicit

awareness of the rules of the game. When pondering the more risky choice, those participants exhibited an elevated skin conductance response (SCR) – a response considered a somatic marker. In contrast, the lesion patients continued to choose disadvantageously (often even after reporting explicit knowledge of the rules), and never developed the anticipatory SCR.

This difference led to the conclusion that damage to the VMPFC interferes with normal processing of somatic signals, while other cognitive functions remain unaffected. Under SMH, VMPFC produces the somatic markers by integrating the information about the stimulus or event with information about the state of the body. When a stimulus is encountered, the emotional state associated with that stimulus on previous exposures is reproduced, either fully (the ‘body’ loop) or partially (‘as-if’ loop). This reproduction is the somatic marker, which becomes integrated with the current event to provide information about the prospective value of that stimulus for the organism – for example imminent loss, as in the case of drawing from the risky deck (Poppa & Bechara, 2018). As such, according to SMH, decision-making at the lowest-level is guided by neural substrates regulating emotional states and operating in service of homeostasis (Bechara, 2004; Reimann & Bechara, 2010).

The SMH has received considerable scrutiny and criticism over the years, addressing which comprehensively is beyond the scope of this review (for a thorough treatment, see Dunn, Dalgleish, & Lawrence, 2006). Most notably, it has been suggested that the notion of somatic markers guiding decision making may have been an unwarranted conclusion due to the methodological faults of the task. For instance, participants completing the IGT may have more knowledge about the game than they explicitly report, and the broad, open-ended questions used in the original paradigm to assess that knowledge may not be sufficiently sensitive to probe awareness, invoking unconscious influences on decision-making where the knowledge is actually conscious (Maia & McClelland, 2004, 2005).

Nonetheless, there has been a considerable amount of support for the physiological and neurobiological basis of the somatic markers. The appearance of larger anticipatory SCR preceding disadvantageous choices than that preceding advantageous choices has been replicated many times, in patients and healthy participants alike, providing support for the notion of somatic markers differentiating between advantageous and disadvantageous options (e.g. Bechara, Damasio, Damasio, & Lee, 1999; Campbell, Stout, & Finn, 2004; Carter & Smith Pasqualini, 2004; Crone, Somsen, Beek, & Van Der Molena, 2004; Oya et al., 2005; Suzuki, Hirota, Takasawa, & Shigemasa, 2003). Similarly, the VMPFC and OFC (notably, as parts of larger networks) have been consistently shown to be implicated in processing of rewards and punishments, as well as their expectation, and emotion-related learning (Critchley & Rolls, 1996; Elliott, Friston, & Dolan, 2000; Gottfried, O’Doherty, & Dolan, 2003; Kim, Shimojo, & O’Doherty,

2011; O'Doherty, 2004; O'Doherty, Deichmann, Critchley, & Dolan, 2002; O'Doherty, Kringelbach, Rolls, Hornak, & Andrews, 2001; Rolls, 2000). The VMPFC was also found to reflect fluctuations in SCR (Critchley, Elliott, Mathias, & Dolan, 2000). While the exact mechanism of action behind somatic markers remains a subject of ongoing work (Poppa & Bechara, 2018), it is now widely recognised that the physiological state of the body does affect not only motivation and decision making, but a range of other processes. As a result, the SMH is often incorporated into broader frameworks, such as interoceptive inference (see section 1.2.3. Predictive processing).

1.2.1.3. Homeostatic control and allostasis

The central assumption of the SMH is that physiological states of the body inform motivated behaviour and decision-making through the proxy of feelings and emotional states. This approach can be evaluated under the wider frameworks of homeostatic control and allostasis, where motivated behaviour serves to satisfy the homeostatic imperative – to maintain continuous survival through minimisation of negative internal states. On this framework, instrumental decisions such as approaching rewards and avoiding harm are informed by affective states reflecting the current physiological state of the organism.

Bodily afferents relay to the brain the information about the state of the organism, such as the need for nutrients or hydration, temperature, blood oxygen levels, inflammation, etc. To ensure continuous survival, those variables must be monitored and maintained within an appropriate viable range (called a set-point) despite fluctuations in the organism's external and internal environments – a concept referred to as homeostasis ("stability through constancy"; Cannon, 1929). Homeostatic regulation can be achieved through peripheral autoregulatory processes, such as buffering systems of the blood maintaining plasma pH (Dworkin, 1993), or central reflex mechanisms, where disturbances to the internal environment are communicated through visceral feedback to trigger a compensatory response to return the deviation to its viable range (for example, a reduction in heart rate following increases in blood pressure; Berntson & Cacioppo, 2007; Berntson, Cacioppo, & Bosch, 2016).

However, it has been argued that although achieving homeostasis through error-correction mechanisms is necessary, it is also inherently inefficient (Sterling, 2004, 2012; Sterling & Eyer, 1988). Rather than correcting a variable at a time to a stable set-point², an efficient

² Some commentators note that this critique stems from a misinterpretation of Cannon's (1929) description of homeostasis as "stability through constancy". In fact, Cannon never advocated an invariable set-point which would lock down a variable irrespective of other environmental or internal demands.

regulatory strategy for dynamic organisms is to constantly anticipate physiological demands and allocate internal resources to meet them. This concept is referred to as allostasis ("stability through change"; Sterling & Eyer, 1988). While it remains a debate whether allostasis extends, encompasses, or replaces homeostasis (Corcoran & Hohwy, 2018; Power, 2004), it is commonly agreed that regulating the internal environment is a complex process that requires higher-order top-down anticipatory control, related to both internal and external conditions, in addition to basic homeostatic reflexes operating through error-detection mechanisms (Berntson et al., 2016; Ramsay & Woods, 2014; Schulkin, 2003b, 2003a; Sterling & Eyer, 1988). As such, the set-point bounds are dynamically adjusted depending on anticipated physiological demands. For example, blood pressure is expected to vary throughout the day in line with multiple factors, such as the circadian rhythm or type and level of activity. Effectively, allostasis can be considered as a mechanism to achieve homeostasis (Day, 2005; McEwen, 1998), allowing for anticipation of upcoming circumstances.

So how does the homeostatic imperative give rise to adaptive, motivated behaviour? Both homeostatic regulation and allostasis operate by reorienting the organism's physiological resources to deal with changes and challenges – in a more reactive or predictive fashion, respectively. When homeostatic regulation cannot maintain physiological variables within appropriate bounds through internal, reactive regulation, an organism can adjust overt behaviour with respect to the external world by means of 'behavioural' allostasis (Corcoran & Hohwy, 2018; Gu & Fitzgerald, 2014) – for example, engage in motivated behaviour to actively seek food, water, or shelter from the heat. As such, allostatic mechanisms directly drive adaptive behaviour, prompting organisms to act to fulfil appetitive or self-protective (or even social) motivational drives in order to maintain homeostasis in the longer-term (Schulkin, 2003b, 2003a, 2011).

This motivation has been proposed to be engendered through affective states – subjective feelings such as hunger, thirst, dyspnoea, or sickness, earlier also conceptualised as 'drives' (Hull, 1943, 1952; Hulme, Morville, & Gutkin, 2019), which reflect deviation from viable states and indicate the need to adjust behaviour in order to return the body to homeostatic bounds (Craig, 2009; Critchley, 2005; Critchley & Garfinkel, 2017; Critchley & Harrison, 2013; Quadt, Critchley, & Garfinkel, 2018b). In this light, visceral states communicating homeostatic errors emerge as a mechanism driving subjective emotional states, which in turn drive overt motivated behaviour to reduce the negative state by biasing decision-making to achieve advantageous, adaptive outcomes for the organism via allostasis (e.g. approaching food to satisfy hunger, or avoiding a sickness-provoking plant). Motivated behaviour is thus firmly

embodied, allowing to fulfil the imperative to maintain homeostasis (Lowe, Morse, & Ziemke, 2017).

While this account may initially seem restricted to motivated behaviour under current or anticipated homeostatic imbalance, the mechanism also extends to explain longer-term adaptive behaviour through instrumental or reinforcement learning. It is axiomatic that the purpose of adaptive behaviour is to remain alive through achieving physiologically relevant rewards, such as food, water, shelter, or opportunities to mate. Agents can learn the homeostatic consequences of their actions (i.e. disturbance or restoration), and orient their behaviour on future occasions to avoidance or approach, respectively, corresponding to instrumental learning (Balleine, 2011; Dickinson & Balleine, 1994; Dworkin, 1981; Keramati & Gutkin, 2014; see section 1.6.1. Instrumental learning, for a more detailed analysis). Homeostatic reinforcement learning (HRL) accounts show that when reward value is computed as a function of internal physiological state, agents seek to maximise rewards in order to prevent potential homeostatic disturbances, rather than merely react to them, as homeostatic and allostatic accounts would predict (Hulme et al., 2019; Keramati & Gutkin, 2011, 2014; Korn & Bach, 2015).

The process of driving adaptive behaviour through affective states is proposed to be accomplished by the interoceptive system – a sensory system for communicating internal physiological information to the brain for coordinated regulation, sensing and representing it in the brain as feeling states (Cameron, 2002; Sherrington, 1948; Tsakiris & Critchley, 2016; the interoceptive system will be described in detail in section 1.3.2. Interoception). The key player of the interoceptive network is the insular cortex, implicated in encoding and representing visceral information, emotion, as well as decision-making (Craig, 2009; Critchley, 2005; Singer, Critchley, & Preuschoff, 2009). The insula's role in representing information about the internal state of the body enables to contextualise the current situation with reference to the agent's physiological state, provoking affective states (Craig, 2013; Gu & Fitzgerald, 2014; Gu, Hof, Friston, & Fan, 2013; Rosen & Schulkin, 2004). It then feeds this information forward to other brain systems involved with valuation and decision-making, including the OFC and VMPFC (Rangel, Camerer, & Montague, 2008; Rushworth, Kolling, Sallet, & Mars, 2012; Rushworth, Noonan, Boorman, Walton, & Behrens, 2011). Through this mechanism, representations of the internal state endow stimuli and choices with motivational value, informing and driving instrumental decision-making.

1.2.2. Predictive processing

Over the past two decades, predictive processing has been emerging as one of the leading approaches to understanding brain function. Rooted in von Helmholtz's idea of unconscious inference (von Helmholtz, 1867), predictive processing (also conceptualised as the 'Bayesian brain' hypothesis) sees the brain as a prediction machine, constantly predicting the causes of sensory signals. The framework has provided rich theoretical and empirical accounts of perception and action, as well as cognition, emotion, and the sense of self. With relevance to this thesis, predictive processing and its corollaries addressed below (active inference and interoceptive inference) offer a comprehensive, biologically plausible account of perception and action as tightly intertwined and inherently embodied under the homeostatic imperative.

Under predictive processing, information about the world arriving through the senses is continuously interpreted in relation to prior expectations or beliefs about its causes. Because the brain has no direct access to the causes of the incoming information, they are referred to as 'hidden'. The brain infers the hidden causes through a probabilistic process, proposed to operate under the principles of Bayesian inference, where the incoming information (likelihood) is combined with the expectations or beliefs about its causes (prior prediction) to arrive at a best guess (posterior prediction) of the causes of that information (Clark, 2013; Hohwy, 2013, 2020; see Figure 1).

Biologically, predictive processing is often formalised through a predictive coding scheme (Rao & Ballard, 1999). The scheme postulates that the predictions are instantiated by probabilistic generative models encoded in cortical hierarchies. The inference process is performed in a descending fashion, where higher (deeper) levels in the hierarchy generate predictions about representations at lower levels. This activity encodes probability distributions of the states of the world causing specific sensory inputs. Incoming sensory signals are compared against such descending predictions to give rise to prediction errors (PEs) at each hierarchical level of processing, indicating a mismatch between the predicted input and the actual input. PEs are then passed up the hierarchy, allowing to update the higher-level predictions about the hidden causes. This adjustment of the higher-level prediction cancels out the PE at the lower level, and the process continues in cycles throughout the hierarchy (Kanai et al., 2015; Rao & Ballard, 1999). In this way, the hierarchies attempt to continuously predict the most likely causes of incoming sensory signal and "explain it away", so that only minimal PE is left to be propagated back up the hierarchy – corresponding to perceptual inference (Clark, 2013; Hohwy, 2013). Perception thus emerges as a product of that inferential process, where incoming sensory signals are quashed by top-down predictions. The upshot of this process is that perception is

strongly determined by prior predictions used by the brain to attempt to explain sensory information.

Perceptual inference is optimised by precision-weighting of the PEs. Precision refers to the reliability or relevance assigned to the incoming sensory signal. In statistical terms, precision reflects the inverse variance of the probability distribution (highly precise distributions have little variance). Precision-weighting allows to regulate the interaction between the top-down and bottom-up signals through adjusting gain on PE-conveying neuronal populations, such that highly precision-weighted PEs have a greater influence in updating the predictions (see Figure 1). Precision-weighting itself is thought to be associated with attention (Feldman & Friston, 2010; Friston, 2009; Hohwy, 2012), where attention tunes the influence of incoming information by assigning higher or lower precision to the corresponding sensory channel. In practice, directing attention to the visual channel would increase the gain on the incoming visual information, increasing its precision.




Figure 1. An illustration of Bayesian inference (left plot) and the influence of higher sensory precision (middle plot) and lower prior precision (right plot). The graphs represent probability distributions associated with sensory evidence (likelihood), prior beliefs, and posterior beliefs. The widths of the distributions represent their variance. Under Bayesian inference, sensory evidence combines with prior beliefs to form a posterior belief (left). When the sensory evidence is assigned high precision (it has high inverse variance), posterior belief will be dominated by the sensory input, proportionately to its precision (middle). The same can happen if prior beliefs are assigned low precision (right). Conversely, if sensory evidence is assigned low precision, posterior belief will be dominated by prior belief. The schematic is adapted from Adams, Stephan, Brown, Frith, & Friston, 2013.

1.2.2.1. Active inference

The hierarchical predictive approach to perception outlined above can be generalised to encompass action, considerably extending the framework beyond its account of perception as a passive process of inferring the most likely causes of sensory stimulation. On this extended account, PEs can be minimised either by updating predictions (in the process of perceptual inference), or by executing an action to change the sensory inputs so that they match the

predictions (active inference). This implies that perception is intimately tied with action, where action can select and shape new sensory inputs (Clark, 2013).

According to active inference, agents actively execute actions in order to realise prior expectations and minimise proprioceptive PEs. The objective of action is to fulfil top-down expectations by either actively seeking new information so as to refine the predictive models of the hidden causes (in an epistemic fashion, e.g. through saccadic eye movements), or generating the expected sensory consequences to match the predicted ones (in an instrumental fashion, e.g. moving a limb to the predicted position) (Friston, 2009). In this way, motor control amounts to the top-down prediction of sensory consequences, which then engages action (through motor reflexes) so as to change the bottom-up proprioceptive or sensory inputs, minimising the PE (Adams, Shipp, & Friston, 2013; Clark, 2015a; Friston, 2009, 2011; Friston et al., 2010). On this account, perception and action are implemented by the same computational architecture, and are tightly intertwined – top-down perceptual predictions inform actions which then elicit the predicted sensory stimulation (Clark, 2015b; Friston, Adams, Perrinet, & Breakspear, 2012; see Figure 2). This has been suggested to be necessary for predictive control (Lauwereyns, 2012; Seth, 2015; Seth et al., 2016) and homeostatic integrity of an organism (Friston, 2010; Maturana & Varela, 1980; Seth & Tsakiris, 2018). Most importantly, in this way, action is not simply a response to stimuli or events in the world, but a means of selecting the next input through a rolling cycle of predicting the next state and bringing it about.

1.2.2.2. Interoceptive inference

In a manner parallel to active inference, the framework has also been applied to interoception, the sense of internal physiological state of the body. Under interoceptive inference, the brain performs the same active process of PE minimisation internally, with respect to interoceptive PEs, as it does externally with proprioceptive PEs. Interoceptive afferent signals communicate bottom-up PEs, which are met by cascading top-down predictions about the internal state of the body, where the lower levels reflect homeostatic variables, and higher levels integrate the interoceptive, proprioceptive, and exteroceptive cues (Seth, 2013; Seth & Friston, 2016; Seth, Suzuki, & Critchley, 2012). Interoceptive PEs can be suppressed either by updating the predictions, or by engaging autonomic control mechanisms to adjust the internal parameters in line with prediction (similarly to engaging motor reflexes under active inference; see Figure 2).

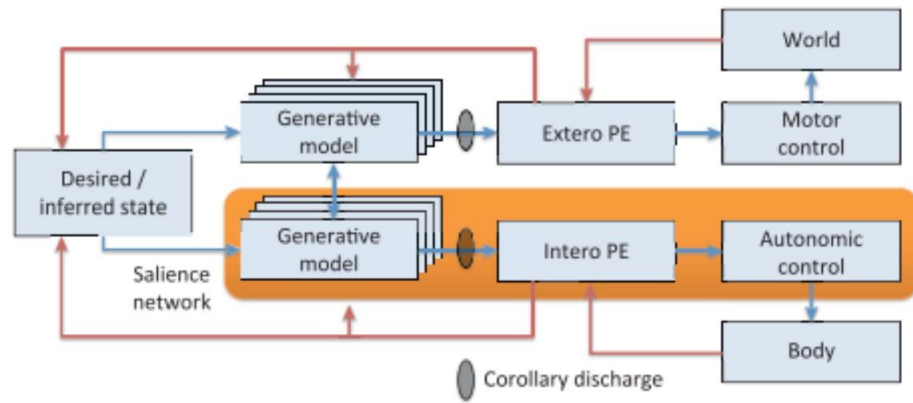


Figure 2. An illustration of active inference (upper section) and interoceptive inference (lower section, highlighted in orange). Blue arrows signify top-down, and red arrows signify bottom-up predictions. Illustration obtained from Seth, 2013.

Interoceptive inference has been proposed to be crucial mechanism for adaptive control and predictive regulation of the internal state, underpinning homeostasis and allostasis. Ascending interoceptive afferents communicate the current internal state of the body, which is compared against the descending predictions providing a homeostatic set-point (i.e. the appropriate range of homeostatic values). Effector systems of the autonomic nervous system can then be engaged to minimise the PE, corresponding to homeostatic control (for example, by mobilising glucose reserves in case of hypoglycaemia). Alternatively, the predictions can be fulfilled through engaging in allostasis (in this example, preparing a meal; importantly, precision of the low-level interoceptive signal must be attenuated so that action execution ensues, rather than updating the predictions of not eating) (Seth & Friston, 2016). Interoceptive, proprioceptive and exteroceptive predictions thus act in concert to ensure homeostasis – interoceptive PEs inform the current motivational state, and proprioceptive and exteroceptive predictions specify allostatic goals (Pezzulo, Barca, & Friston, 2015). As described in an earlier section (1.2.1.3. Homeostatic control and allostasis), a predictive model of the internal state can thus drive value-based choices by providing a reference to the internal state of the body in a predictive fashion – the expected internal state should a given behaviour be performed (Gu & Fitzgerald, 2014).

This homeostatic imperative under the predictive processing account is formalised by the Free Energy Principle (Friston, 2009, 2010; Friston et al., 2010). According to FEP, organisms must minimise their internal entropy in order to continue living – that is, they must inhabit only a range of states that ensure survival, corresponding to homeostatic bounds. Venturing out of those bounds, for instance by allowing blood pressure or the body temperature to rise above the viable range, or depriving the organism of nutrients, would inevitably cause death. According to FEP, living organisms resist the tendency towards entropy, imposed by the second law of

thermodynamics, by minimising surprisal (information-theoretic surprise) or – simplifying – variational free-energy. Under certain mathematical assumptions (which are beyond the scope of this thesis), free-energy corresponds to long-term average of PE. As such, PE minimisation (exteroceptive, proprioceptive, and interoceptive) is a mechanism geared at resisting the tendency towards entropy, achieving the fundamental objective of staying alive.

According to this view, the body has been proposed to constitute the “first prior” (Allen & Tsakiris, 2018). Continued existence is inherently dependent on regulating the interactions with the world so that they fulfil the homeostatic imperative – that is, finding oneself in states that are conducive to staying alive, rather than bad, “surprising” states. Organisms manage this feat by making use of internal models weaving in interoceptive and exteroceptive cues to regulate the interaction with the world under the homeostatic imperative (a view that stems from cybernetics; Seth, 2015). Homeostatic priors are afforded a priority status within the predictive hierarchy (through a high precision-weighting) to determine inference and behaviour based on the subjective needs of the organism. Hence, perceptual and value-based salience are relative to the impact a stimulus or decision might have on the organism’s condition (Friston et al., 2015; Pezzulo et al., 2015). Under this perspective, the internal model of the body (or the “first prior”) provides a reference frame to facilitate interactions with the world through a subjective sense of being a ‘self’ (Allen & Tsakiris, 2018; Apps & Tsakiris, 2014; Craig, 2009; Seth, 2015; Seth & Tsakiris, 2018).

Interoceptive inference was also proposed to be the mechanism giving rise to subjective feelings of emotion. On this view, emotion arises from the inference of the hidden causes of bodily signals, through the integration of interoceptive, exteroceptive and proprioceptive predictions reflecting is the organism’s current context (Barrett, 2016; Barrett & Simmons, 2015; Clark, 2016; Pezzulo, 2014; Seth, 2013; Seth & Critchley, 2013), thus providing a mechanistic account of two-factor theories like the Schachter-Singer (1962) account (visited in section 1.2.2. James-Lange theory of emotion). Emotional content is shaped by the hierarchy of generative models predicting the bodily responses and possible actions in the current environmental context. Thus, the process of integrating interoceptive with exteroceptive information inherently affects how we perceive and interpret the external world – through the lens of internal states. A relatable example is provided by Pezzulo (2013): because of the high precision afforded to interoceptive signals, if the body is in a state of arousal, such as after watching a horror movie, the sound of a squeaking window at night (exteroceptive information) might be interpreted as a burglar or a bogeyman – rather than, more plausibly, wind.

Overall, predictive models of interoceptive signals play a fundamental role in regulation of the internal state, maintaining the body in viable states through homeostatic control and

allostasis. Consequently, they also contribute to shaping emotion, perception, and motivated behaviour. Both active and interoceptive inference thus highlight the complex interaction between the brain, body and the world, illuminating the general structure of our interactions with the world as firmly embodied under a single, dynamic predictive regime (Clark, 2016).³

1.3. The body on the inside

1.3.1. Interoception

Broadly, interoception refers to the sensing, and interpreting, and integrating the neural signals communicating the body's moment-to-moment internal condition (Cameron, 2002; Khalsa et al., 2017; Sherrington, 1948). A comprehensive definition, highlighting the levels of complexity of the interoceptive system, is provided by Quadts et al., 2018 (p.112):

“Interoception is an umbrella term encompassing (1) afferent (body-to-brain) signalling through distinct neural and humoral (immune and endocrine) channels; (2) neural encoding, representation, and integration of this information concerning internal bodily state; (3) the influence of such information on other perceptions, cognitions, and behaviors; (4) and the psychological expression of these representations as consciously accessible physical sensations and feelings.”

The first, lowest level refers to the afferent signal carrying continuous information about the internal state of the body (including the cardiovascular, gastric, respiratory, or inflammatory information), its neural pathways, as well as neural readouts, such as the heartbeat-evoked potential (HEP; Montoya, Schandry, & Müller, 1993) or oscillations related to gastric activity (Rebollo, Devauchelle, Béranger, & Tallon-Baudry, 2018). This information is then integrated in the brain to represent the homeostatic condition of the organism, and distributed to other regions, thus informing homeostatic control, allostasis (addressed in previous sections), as well as other perceptual, affective and cognitive processes, including emotional states and motivated behaviour (Berntson, Gianaros, & Tsakiris, 2018; Critchley & Garfinkel, 2017). At the highest level is the psychological expression of the interoceptive information, referring to conscious awareness of the interoceptive signals from different channels. This psychological manifestation is itself conceptualised across three levels: interoceptive accuracy (referring to objective accuracy of perceiving a bodily signal), sensibility (referring to subjective beliefs about the ability

³ While it is beyond the scope of this theoretical overview to evaluate the varied approaches to predictive processing in detail, it is important to note that while active and interoceptive inference are fundamentally embodied by nature (as both take root in the Free Energy Principle), some predictive accounts do not take that stance and instead focus on purely internal, modular representations (see Allen & Friston, 2018; Clark, 2015b; Dolega, 2017).

to perceive one's own bodily signals), and awareness (referring to the correspondence between accuracy and confidence, i.e. a metacognitive measure of interoceptive performance; Garfinkel, Seth, Barrett, Suzuki, & Critchley, 2015; Quadts, Critchley, & Garfinkel, 2018a, but note that they were shown to reflect independent processes; Forkmann et al., 2016).

1.3.2. Anatomical and functional organisation of the interoceptive system

Interoceptive information encompasses distinct types of signal from the body's viscera (internal organs, tissues, and cells), including visceromotor, nociceptive, thermal, metabolic, hormonal, and immunological information (Berntson et al., 2018; Craig, 2014; Critchley & Harrison, 2013). Information from the viscera (organs, tissues and cells located in the thoracic and abdominal cavities) constitutes the General Visceral Afferents (GVA), which convey information to the brain through two main pathways. Motivational GVA, such as hunger, satiety, thirst, nausea, or respiratory sensations, are relayed through cranial nerves, most notably the vagus nerve, the longest and most widely distributed cranial nerve. GVA carrying nociceptive, thermal and chemosensory information (largely communicating tissue damage) travel through spinal afferent pathway, projecting to the dorsal horn of the spinal cord, and to the spinal lamina 1 (Craig, 2014; Saper, 2002). Both pathways terminate in the nucleus of the solitary tract (NTS).

The NTS is the first interoceptive relay port, sensing and regulating the physiological state of the body. At the early stages, visceral information also inputs to the periaqueductal gray matter, parabrachial nucleus, hypothalamus, and the ventromedial posterior (viscerosensory) thalamus (Critchley & Harrison, 2013). Through the thalamic relays, visceral information is then propagated to a broader network of regions implicated in perception, cognition, adaptive behaviour and decision-making, including the amygdala, the insular cortex, the cingulate cortex, and frontal / orbitomedial regions (Blessing, 1997; Cechetto & Shoemaker, 2009; Critchley & Harrison, 2013), constituting a wider interoceptive network for representing, regulating, and feeding visceral information forward (Kleckner et al., 2017). Through the wide distribution, the system overlaps with the Central Autonomic Network (CAN; Benarroch, 1993), the emotional motor system, and the classical limbic system (Strigo & Craig, 2016).

The primary cortical site for visceral processing is the insular cortex. Primary viscerotopic representations are localised in the dorsal-posterior regions of the insula, and the re-represented and integrated in the mid- and anterior part (anterior insular cortex, AIC), suggesting a posterior-to-mid-to-anterior organisation of interoceptive representations in the human insula (Craig, 2009). The AIC specifically is involved with integration of the bodily signals into cortical representations of bodily responses and changes in the internal states (Critchley, 2005; King, Menon, Hachinski, & Cechetto, 1999; Singer et al., 2009), as well as their conscious

representation (Craig, 2009; Gu et al., 2013; Wang et al., 2019). The AIC was also proposed to serve the function of comparing top-down interoceptive predictions against bottom-up interoceptive PEs under the predictive processing framework (Barrett & Simmons, 2015; Gu et al., 2013; Seth, 2013), a notion supported by functional magnetic resonance (fMRI) studies showing AIC activation in anticipation of painful or affective tactile stimuli (Lovero, Simmons, Aron, & Paulus, 2009; Ploghaus et al., 1999). The insula is also highly interconnected with areas supporting cognition, emotion, and motivation, including the prefrontal regions, the cingulate cortex, and the amygdala (Oppenheimer & Cechetto, 2016).

With special relevance to this thesis, the insula is also implicated in dynamic cardiac regulation, gauging interoceptive and exteroceptive cues to adjust cardiac outputs to ensure appropriate adaptive responses (Hassanpour et al., 2018; Oppenheimer & Cechetto, 2016; this aspect will be revisited in section 1.6.2. Autonomic and interoceptive processes in instrumental learning). Insular activity was shown to be enhanced in response to asynchronous, relative to synchronous, cardiac feedback (i.e. tones presented out of or in synchrony with participants' heart rate), suggesting that the region detects and computes physiological mismatches in cardiac signal (Gray, Harrison, Wiens, & Critchley, 2007; see also Banellis & Cruse, 2020; Pfeiffer & De Lucia, 2017; van Elk, Lenggenhager, Heydrich, & Blanke, 2014). The HEP (an electroencephalographic (EEG) readout of cardiac activity) was also localised to the insula (as well as the amygdala and the sensorimotor cortex; García-Cordero et al., 2017; Kern, Aertsen, Schulze-Bonhage, & Ball, 2013), and proposed to reflect precision of the interoceptive signal (Petzschnner et al., 2019). Together, those findings suggest that interoceptive representation of cardiac activity and their potential reliability and usefulness for further processes are reflected in the insula.

1.3.3. Interoceptive effects on perception, cognition, and affect

The last twenty years have seen a wealth of research addressing interoceptive effects on a host of other processes, over and above behaviours operating directly in service of homeostasis and allostasis. While a detailed review of the rich and differentiated body of work is outside of the scope of this thesis, it will be beneficial to briefly summarise the findings pertaining to this topic, and the challenges around it (note that the literature relevant to learning will be reviewed in more detail in section 1.6.2. Autonomic and interoceptive processes in instrumental learning).

A vast proportion of research on interoceptive effects on other processes has focused on cardiac activity, largely due to the relative ease of measurement. For example, individual baroreflex sensitivity was shown to affect cognitive performance and pain sensitivity (Duschek, Werner, & Reyes del Paso, 2013). Within the subfield, cardiac cycle is often exploited as a means

of evaluating processing during the physiological activity of the heart at systole (when baroreceptors are active and blood is ejected from the heart) compared to diastole (between the heartbeats, when baroreceptors are quiet). The phase of the cardiac cycle has been found to differentially affect emotional processes, with stimulus presentation at systole enhancing the processing of fearful (motivationally-relevant) stimuli (including detection, intensity ratings, and processing in the amygdala) in comparison to other emotions. Systole was also shown to enhance active visual search (Galvez-Pol, McConnell, & Kilner, 2018; Kunzendorf et al., 2019), visual selection accuracy (Pramme, Larra, Schachinger, & Frings, 2016), and response inhibition (Rae et al., 2018). In contrast, systolic presentation attenuates memory for words, pain perception, or startle responses (Garfinkel & Critchley, 2016). Individual differences in interoceptive accuracy, sensibility or awareness have been shown to modulate decision-making (Dunn et al., 2010), cardiac responses to emotional images (Pollatos, Herbert, Matthias, & Schandry, 2007), susceptibility to the rubber hand illusion (Tsakiris, Tajadura-Jiménez, & Costantini, 2011), or loss aversion (Sokol-Hessner, Hartley, Hamilton, & Phelps, 2015).

However, it is important to note that the cardiac cycle is an oscillatory signal, present at all times. While the results showing its selective impact on higher-order processes illustrate the strong interaction between the brain and the body, it is difficult to assess whether the effect is functional, or simply a consequence of fluctuations in neural noise. This highlights the need for more targeted, controllable means of manipulating interoceptive signals to allow for assessment of causality.

An attempt at manipulating how cardiac signal is used by the brain has been made by manipulating its perceived synchronicity with real-time cardiac feedback. An auditory stream synchronous to a participants' own heartbeat (e.g. a stream of tones such as beeps, played in time with heartbeat) can enhance fairness judgments (Lenggenhager, Azevedo, Mancini, & Aglioti, 2013) and error processing (Łukowska, Sznajder, & Wierzchoń, 2018), while an asynchronous stream disrupts it. Similar results were obtained with a visual synchronous stream affecting body ownership (Aspell et al., 2013; Suzuki, Garfinkel, Critchley, & Seth, 2013). This suggests that bodily (in this case, cardiac) information might functionally inform some higher-order processes, especially those which have an embodied component, such as the sense of one's own body or empathy. More formally, it has been proposed that synchronicity increases the precision of cardiac information reaching the brain (i.e. cardioceptive precision; Petzschner, 2017), although it is still a largely hypothetical assumption. Nonetheless, as mentioned in the previous section, the insula has been shown to be sensitive to mismatches in auditory cardiac feedback (Gray et al., 2007), with similar results in EEG readouts of cardiac processing (Banellis & Cruse, 2020; Pfeiffer & De Lucia, 2017; van Elk et al., 2014). This suggests that external cardiac

feedback, although an undeniably exteroceptive signal, may be a promising avenue for investigating the effects of cardiac precision. Interestingly, only one study thus far has applied this method to observe its role in performance (Łukowska et al., 2018), showing that asynchronous feedback disrupted performance-error related activity. This has implications for the study of the role of the body in instrumental learning, and will be addressed in greater detail in a later section (1.6.2. Autonomic and interoceptive processes in instrumental learning).

It is also noteworthy that the field has attracted some methodological criticism, mostly related to the validity and generalisability of the tasks used to assess individual cardiac ability. Those tasks typically involve counting one's own heartbeat at rest, without touching the heart or pulse spots (Schandry, 1981), referred to as the Heartbeat Counting Task, or discriminating whether a stream of discrete sounds (e.g. beeps) is in or out of synchrony with one's heartbeat (Heartbeat Discrimination Task; Katkin, Blascovich, & Goldband, 1981). Both tasks have attracted considerable critique in terms of validity and generalisability (Brenner & Ring, 2016). The Schandry task especially was proposed to be fundamentally biased and rely on processes unrelated to actual interoception, such as participants' prior knowledge of average heart rates or even counting time units (Desmedt, Luminet, & Corneille, 2018; Ring & Brenner, 2018; Ring, Brenner, Knapp, & Mailloux, 2015; Zamariola, Maurage, Luminet, & Corneille, 2018).

1.4. The body on the outside

1.4.1. Action-oriented approaches

The interaction of action, perception, and cognition has long been a topic of enquiry (Gibson, 1979; James, 1891). The active nature of perception and cognition was already highlighted over a hundred years ago, for example by Dewey (1896), who emphasised that what we experience with the senses is secondary to action, and determined by movement of the body, head, and eyes. Despite that, the long dominance of cognitivism firmly established an opposite view, where cognition is considered context-invariant, modular, and substrate-neutral (e.g. Fodor, 1983), and operating through mental representations. In this view, cognition determines action, and perception is largely independent of action.

However, this approach is gradually being abandoned in favour of an action-oriented perspective, where cognition and action are closely intertwined (Clark, 1998; O'Regan & Noë, 2001; Valera, Thompson, & Rosch, 1992). Under this perspective, cognitive systems are inseparable from embodiment, and any processing is a dynamic, context-sensitive process that is largely driven by action of the agent in the world. Similarly, it has been proposed that the conscious, first-person perspective is organised for supporting adaptive, flexible action, where

consciousness provides an egocentric frame of reference for agents' interactions with the world (Merker, 2013). Action-oriented frameworks have been yielding considerable theoretical and practical advances in philosophy, robotics, psychology, cognitive science, and the neurosciences, and thus constitute a rich field of research, vastly beyond the scope of this thesis (a state of the art is provided in Engel, Friston, & Kragic, 2015). The focus here will be placed specifically on action-oriented approaches to consciousness rooted in predictive processing.

Action-oriented predictive processing has been applied to provide a new perspective on perception and conscious experience (although note that predictive processing itself is not a theory of consciousness). Under predictive processing, conscious experience has been proposed to be shaped by the posterior predictions (the generative model) that best suppress PEs (see section 1.2.3. Predictive processing) – that is, the model that best 'explains away' the incoming sensory input (Hohwy, 2013; Seth et al., 2016).

However, agents do not passively perceive their environment – they actively engage with it. As such, perceptual experience should be influenced not only by passive predictions about the world, but more generally by predictions encompassing how the world changes based on the agent's actions – that is, on predictions about sensorimotor contingencies (Clark, 2015b; O'Regan & Noë, 2001; Seth, 2014). This view is strongly related to active inference (section 1.2.3.1. Active inference), where action is construed not as a response to perceptual inputs, but as a way of selecting the next input through PE minimisation. In order to successfully minimise PEs, the brain should not only predict the causes of sensory inputs, but employ those predictions to also predict how sensory input would change as a result of specific actions (Seth, 2014), a view termed predictive perception of sensorimotor contingencies.

According to this perspective, the predictions that best explain the incoming sensory input (i.e. best minimise the PEs) are not necessarily those which are the most veridical, representing the world in the truest sense, but those which are best for supporting adaptive interactions with the world (Clark, 2015b; Seth, 2014, 2015). As such, by executing actions in order to fulfil sensorimotor predictions, organisms are effectively constructing their sensory experience in line with those predictions (Clark, 2015a, 2015b; Seth, 2014). Consequently, predictive processing agents do not need to model their environment in the most comprehensive way – instead, they model it selectively, in an action-oriented fashion, enabling interactions with the environment in a smooth, streamlined manner (Baltieri & Buckley, 2017; Land, 2012; Seth, 2015; Seth & Tsakiris, 2018; Tschantz, Seth, & Buckley, 2019).

1.4.2. The effect of action on perception and conscious experience

There already is some evidence showing that action can modulate conscious perception. Voluntary action has been found to help disambiguate a bistable or otherwise ambiguous percept, if the action is congruent with an aspect of that percept, for example when the direction of movement executed by participants when viewing the stimulus corresponds to the direction of movement of the stimulus itself (Beets et al., 2010; Di Pace & Saracini, 2014; Dogge, Gayet, Custers, & Aarts, 2018; Maruya, Yang, & Blake, 2007; Mitsumatsu, 2009). Those findings suggest that a match between action kinematics and stimulus kinematics (such as motion directions) facilitate perception relative to mismatches. A comparable effect has been found for proprioception, where correspondence between the target and proprioceptive information can bias or sharpen the visual percept (Butz, Thomaschke, Linhardt, & Herbort, 2010; Salomon, Lim, Herbelin, Hesselmann, & Blanke, 2013; Yon, Gilbert, De Lange, & Press, 2018).

Those results are often interpreted as action and perception sharing common mental representations, where actions are controlled by internal representations of the desired goals coded as expected sensory outcomes (Hommel, Muesseler, Aschersleben, & Prinz, 2001; Prinz, 2003). In this view, action biases perception by activating the corresponding outcome representations. However, this view does not preclude an influence of attentional cueing, where the executed movement direction would cue perception of anything moving in the same direction, or simple congruency effects, where any match facilitates perception over a mismatch (e.g. an effect also found for olfactory and visual percepts; Zhou, Jiang, He, & Chen, 2010)

Action-oriented predictive processing approaches, described in the section above, considerably extend this view, proposing that it is the prediction of sensorimotor contingencies that brings about the conscious experience of the related stimulus. This approach provides testable hypotheses, although the direct effect of predictions about sensorimotor contingencies has not yet been investigated. A recent study on sensorimotor contingencies by Suzuki, Schwartzman, Augusto, and Seth (2019) is a step in that direction, showing that mere congruencies (e.g. veridical or reversed sensorimotor coupling of motion directions in participants' interactions with stimuli) had no effect on time of breakthrough from interocular suppression, but live temporal contingencies did (i.e. live versus replayed interactions). This result suggests that there exists a dynamic causal coupling between actions and their sensory consequences.

1.5. Interim summary: The embodied mind

The picture emerging from the above analysis is that of a strong influence of the body (both on the inside and the outside) on perception, emotion, cognition, and conscious experience. On the most basic level, the imperative for maintaining homeostasis requires living organisms to satisfy motivational drives (e.g. fulfilling hunger, running away from danger), which can only be achieved through goal-directed interaction with the environment. Motivational drives were proposed to produce subjective feeling states, reflecting the integrated model of the organism's internal state and the environmental context the agent is in, as well as predictions of homeostatic objectives which can be fulfilled immediately or in the future. Those feelings states have adaptive value – they inform decision-making by imbuing choices and situations with value, so as to achieve positive states (such as attaining reward, reducing the unpleasant drive) and avoid negative ones. Predictive-processing accounts (especially the FEP) take this view further, proposing that perception, action, and emotion are firmly embodied through the imperative to minimise surprising states and stay alive. Under this framework, conscious perception is shaped by predictions of how the world changes in relation to our actions.

A vast amount of research shows various ways interoceptive signals (predominantly the heart) influence perception, cognition, and affect. A separate line of research also shows the effects of action and proprioception on perception and conscious experience. However, direct causality or functional mechanisms are still hard to establish. Predictive processing approaches can furnish hypotheses that help address those questions, as will be evident in some of the empirical chapters of this thesis.

Nonetheless, despite the development of rich theoretical frameworks and the wealth of empirical research on bodily influences, the embodied imperative is not without controversy. It has been argued that for bodily information to directly cause emotion (as under the James-Lange theory) or directly inform decision-making (as under the SMH) would be inefficient as peripheral arousal is noisy and non-specific (Rolls, 2005). Related to SMH, rat models demonstrated that the specific physiological state does not have to be reproduced in order to motivate a corresponding behaviour, as is predicted by SMH (Balleine, Ball, & Dickinson, 1994; Balleine, Davies, & Dickinson, 1995). Instead, what is crucial to drive adaptive behaviour is the *representation* of the given motivational state and the values of outcomes for the organism. Indeed, lesions to the insula and amygdala (both parts of the general interoceptive network and implicated in the representation of bodily states and their relevance to the organism) considerably impaired instrumental performance in rats (see Balleine, 2011 for an overview). Finally, with emotions occurring fairly late in the processing of a stimulus or event, emotions

were also proposed to be better viewed as a consequence, rather than a substrate (Moors, 2009).

However, it is important to recognise that more recent theories of embodied emotion and decision-making (Allen & Tsakiris, 2018; Barrett, 2016; Gu & Fitzgerald, 2014; Seth, 2013) do not assert that bodily activation is the direct cause of emotion or decision-making – instead, it is the process of inference on the causes of ongoing bodily sensations (as noisy as they are) that gives rise to a subjective affective state (in that sense, emotion is a consequence). This process combines exteroceptive and interoceptive information to then inform emotion, and feed into valuation and motivating adaptive behaviour, as evidenced by the neuroimaging and neuroanatomy evidence reviewed above. As such, those frameworks are not incompatible with the position of the critics above (although of course they are not exempt from other debates, e.g. about the neural implementation of the predictive models). Nonetheless, the details of the exact psychophysiological mechanism are still unknown.

1.6. The role of the body in instrumental learning

1.6.1. Instrumental learning

Learning is a crucial component of adaptive behaviour and, consequently, survival. Learning allows animals to predict the occurrence of important stimuli or events in their environment, and to interact with the environment to obtain adaptive outcomes and avoid harmful ones. The former is typically investigated with Pavlovian (classical) conditioning procedures, where a given stimulus acquires the capacity to provoke anticipatory, reflexive responses through exposure together with a direct reinforcer, such as anticipatory salivating in response to seeing a stimulus previously paired with food (Pavlov, 1927; Rescorla, 1988). In Pavlovian conditioning, any observed changes in behaviour are a reflection of the learned association predictive of an outcome, and outcomes are delivered regardless of the animal's behaviour. While learned reflexive responses can be highly adaptive, they are not sufficient for survival in a dynamic environment requiring the agent to adapt by acquiring the ability to act on the environment itself (Balleine & Dickinson, 1998; Dayan & Balleine, 2002). This vital ability is referred to as instrumental learning.

Instrumental behaviour can be defined as actions whose “acquisition and maintenance depend on the consequences for the animal or, in other words, on the fact that the action is instrumental in causing some outcome” (Dickinson, 1994, p.47). Only instrumental behaviour provides the capacity to control the environment in the service of the individual's needs or wants, making it the fundamental aspect of adaptive behaviour. Instrumental learning is

typically investigated through free-operant or instrumental (both terms refer to the same process) conditioning (Skinner, 1932), for example where animals learn to approach stimuli predictive of positive outcomes and to avoid stimuli predictive of negative outcomes (see Staddon & Cerutti, 2003 for a detailed review of the free-operant procedures and reinforcement schedules).

Control of instrumental behaviour is a complex process, involving not only learning of the action-outcome associations, but also assigning motivational value to the outcomes to invigorate future goal-directed behaviour, including action initiation and choices between different actions (Balleine, Delgado, & Hikosaka, 2007; Balleine & O'Doherty, 2010; Balleine & Ostlund, 2007). Instrumental learning was proposed to be mediated by at least three different mechanisms (Balleine, 2011; Balleine & Dickinson, 1998), including stimulus-response (S-R) learning, evaluative conditioning, and incentive learning.

Early accounts cast instrumental learning purely in terms of learning S-R associations (e.g. Thorndike's law of effect, 1898, 1911). S-R accounts proposed that presentation of a positive reinforcement (reward) after an instrumental action is performed strengthens the association between the stimulus present when the action was performed and the response itself.⁴ However, S-R theory assumed that animals have no knowledge or representations of the consequences of their behaviour. As such, it could not account for how future responses are made based on the relevance of the anticipated consequences to the current motivational state (in this sense, the theory accounts for habitual behaviour better than motivated behaviour).

Later developments proposed that outcomes are also encoded during instrumental learning, and acquire value through evaluative conditioning and incentive learning (Balleine, 2011; Dickinson & Balleine, 1994; Dickinson & Dawson, 1987). According to this approach, pairings between a given stimulus and outcome (e.g. food providing a positive experience of reducing hunger) endow the stimulus with a higher value, corresponding to evaluative conditioning. In an instrumental scenario, a specific motivational state (e.g. hunger) causes an agent to assign a higher value to positive outcomes of their actions (e.g. obtaining food) through evaluative conditioning. This relationship is then represented abstractly as incentives or values assigned to specific outcomes (i.e. rewards), which are able to invigorate instrumental performance upon encountering the stimulus again, whether in that motivational state or not (Balleine, 2011; Berridge, 2000; Dickinson, 1994; Dickinson & Balleine, 1994). Through this

⁴ It is interesting to note that Thorndike's original formulation included physiological drives, which he called 'satisfaction' and 'annoyance', loosely corresponding to reward and punishment, which motivate future instrumental behaviour (1989). He later abandoned those concepts in favour of a more behaviourist approach.

mechanism, agents learn to act on stimuli associated with rewarding outcomes, and avoid those associated with punishing outcomes.

The process of incentive learning is closely related to sensory or affective representations of the bodily state. Recent neuroimaging studies (reviewed briefly in sections 1.2.1.3. Homeostatic control and allostasis, and 1.3.2. Anatomical and functional organisation of the interoceptive system) provide support for this account, showing extensive functional interrelationships between neural regions encoding the state of the body and its conscious representation with areas involved in reward, value representations, and decision-making.

1.6.2. Autonomic and interoceptive processes in learning

As described above, basic forms of learning such as classical and instrumental conditioning have been long understood as a strongly embodied, homeostasis-driven process, allowing agents to anticipate biologically relevant events or resources and to act on the world to obtain them. In parallel, bodily signals have also been proposed to reflect an internal performance monitoring system or an orienting response, supporting the learning process.

Monitoring performance for errors is a vital part of successful decision-making and learning. Efficient performance monitoring allows for rapid detection of errors in performance and for appropriate and flexible adjustment of future behaviour. Performance monitoring is an inherent component of instrumental learning, where actual outcomes are compared to expected outcomes, with mismatches (e.g. following from an error) informing the system of the necessity for behavioural adjustment (often investigated with reinforcement learning (RL) schemes; Dayan & Niv, 2008; Sutton & Barto, 1998). This is a complex process, involving not only monitoring external events for feedback such as action consequences, but also evaluating the motivational or adaptive relevance of the outcomes (i.e. whether they are rewarding or punishing).

In this light, it has been proposed that it is vital that performance is also monitored internally (Hajcak, McDonald, & Simons, 2003; Müller, Möller, Rodriguez-Fornells, & Münte, 2005; Sokolov, 1963; Ullsperger, Danielmeier, & Jocham, 2014). Performance monitoring, evaluation of outcomes and options, and behavioural adjustment are achieved by a large-scale network associated with response organisation and selection, which also serves to modulate psychophysiological resources in line with dynamic environmental demands (Hajcak, McDonald, et al., 2003; Thayer & Lane, 2000; Ullsperger et al., 2014; Ullsperger & Von Cramon, 2004). This network has been suggested to encompass the anterior executive regions and the CAN, including the amygdala, the insular cortex, the anterior cingulate cortex (ACC), hippocampus, and the NTS (Benarroch, 1993, 2004; Devinsky, Morrell, & Vogt, 1995; Thayer & Lane, 2000). It

is also widely distributed, including feedback and connections with the striatum (involved with reward processing), as well the OFC and VMPFC, the prefrontal regions involved with valuation and decision-making.

Correlates of error- or feedback-related processing have been observed in error-related negativity (ERN or Ne) and positivity (Pe) components in electroencephalography (EEG) (Falkenstein, Hoormann, Christ, & Hohnsbein, 2000; Falkenstein, Hoormann, & Hohnsbein, 1999; Nieuwenhuis, Ridderinkhof, Blom, Band, & Kok, 2001; Overbeek, Nieuwenhuis, & Ridderinkhof, 2005), which are thought to originate in the ACC. The ACC is the crucial component of the error-monitoring network, and is involved in detection of errors and response conflicts, signalling the need for behavioural adjustments (Holroyd & Coles, 2002; Ridderinkhof, Nieuwenhuis, & Bashore, 2003) and reflecting 'surprise' signals (mismatches between expected and obtained outcomes; Hauser et al., 2014).

The internal performance monitoring mechanism is also reflected in a range of autonomic responses to errors in performance. Errors (of commission as well as omission) provoke pupil dilation (Critchley, Tang, Glaser, Butterworth, & Dolan, 2005; Wessel, Danielmeier, & Ullsperger, 2011), increased skin conductance response (Crone, Somsen, Beek, & Van Der Molena, 2004b), as well as – with the greatest relevance to this thesis – a deceleration in heart rate (HR; Crone et al., 2003). Collectively, those responses were linked to the orienting response of the ANS – a reflexive response geared at orienting the agent to changes in the environment in order to adapt the organism to any perturbations (Buser & Rougeul, 1961; Pavlov, 1927; Sokolov, 1963).

HR deceleration has been observed to follow immediately after an error in performance (e.g. in a learning task), where the heart tends to decelerate more in response to negative feedback than to positive feedback (Crone, Bunge, de Klerk, & van der Molen, 2005a; Crone et al., 2003; Somsen, Van Der Molen, Jennings, & Van Beek, 2000; van der Veen, van der Molen, Crone, & Jennings, 2004a). This deceleration in response to errors has been proposed to reflect the processing of the error and the need for behavioural adjustment. Importantly, this response is only evident when feedback is valid – that is, when it carries reliable, usable information about performance (Crone et al., 2003; Groen, Wijers, Mulder, Minderaa, & Althaus, 2007; Mies, Van der Veen, Tulen, Hengeveld, & Van der Molen, 2011). Feedback is considered valid or informative when it can be related to the nature of the stimulus – that is, when the stimulus reliably predicts the outcome, and the feedback reflects the match or mismatch between the predicted and actual outcomes. In addition, HR deceleration has been found to diminish as learning progresses, instead becoming more pronounced in anticipation of negative feedback presentation than following it (Crone et al., 2004b; Groen et al., 2007), or even appearing as

early as the negative stimulus itself (Kastner, Kube, Villringer, & Neumann, 2017). This effect has been taken to imply that there is a shift from external, response-based performance monitoring to internal, prediction-based monitoring over the course of learning. Crucially, the strength of the phasic cardiac response to negative feedback has been found to correlate with the strength of the prediction error signal (derived from a RL algorithm; Kastner et al., 2017), which alerts the learner to the need to adjust behaviour.

Elsewhere, performance on decision-making tasks has also been found to be affected by individual interoceptive ability. Although direct evidence on instrumental learning is missing, it has been shown that people displaying better performance on interoceptive accuracy or awareness show improvements in risky decision-making tasks, such as the Iowa Gambling Task also used in SMH research (Dunn et al., 2010; Werner, Jung, Duschek, & Schandry, 2009; Werner et al., 2013), or in performance on a real-life trading floor (Kandasamy et al., 2016). Those results suggests that accurate perception of bodily information may mediate the relationship between bodily changes and cognitive-affective processing, such as decision-making.

1.7. Unconscious learning

1.7.1. Evidence for unconscious learning

Whether learning can proceed without conscious awareness of the target stimulus has long been a topic of intensive research and theoretical debate. Typically, unconscious learning is investigated with simple associative learning paradigms, where participants learn stimulus-stimulus or stimulus-outcome associations, or priming tasks. Most often, stimuli are rendered unconscious using masking or other suppression techniques. Depending on the paradigm, learning is usually indexed with presence or absence of the conditioned response, reaction times (i.e. where shorter reaction times to congruent vs incongruent pairings are indicative of learning), or the priming effect (where the prime influence processing of the subsequent stimulus, e.g. its perception).

Research into unconscious learning suggests that simple forms of associative learning and priming can take place without conscious awareness of the stimuli. This has been found for classical conditioning with the startle eye blink response (Clark & Squire, 1998), fear conditioning (Knight, Nguyen, & Bandettini, 2003; Raio, Carmel, Carrasco, & Phelps, 2012), emotional learning (Olsson & Phelps, 2004), visuospatial learning (Rosenthal, Kennard, & Soto, 2010; Seitz, Kim, & Watanabe, 2009a), associative learning between sensory modalities (Faivre, Mudrik, Schwartz, & Koch, 2014b; Scott, Samaha, Chrisley, & Dienes, 2018), and priming (Van Den Bussche, Van Den Noortgate, & Reynvoet, 2009).

However, it has been pointed out that the role of conscious awareness in learning might vary between different learning procedures and the processing demands they impose (Knight et al., 2003). Most notably, unaware learning (as well as priming) is typically observed when stimuli are presented at short spatiotemporal intervals, or concurrently (van Gaal, de Lange, & Cohen, 2012). For example, unaware classical conditioning was demonstrated in delay scenarios (when stimuli to be associated overlap temporally), but not in trace scenarios (where they are separated by an interval; Clark & Squire, 1998).

In contrast to simple forms of associative learning, instrumental learning can be classified as a considerably more complex process. Agents must not only learn the associations between different stimuli, or stimuli and their outcomes, but also deploy action selectively (e.g. approach or avoid), and adapt their behaviour in the long-term. This involves integrating information over a long temporal scale and distinct modalities, as involved in processing the visual input, extracting its predictive value, deploying a selective response, processing the reinforcement, and comparing the expected outcome with the actual outcome in order to update the representations of stimulus value. As such, it is a considerably more complex process than the aforementioned simpler forms of associative learning, which do not require selective decisions on whether to act or not, or a behavioural adaptation from trial to trial.

Nonetheless, there has been some evidence showing that instrumental learning can proceed unconsciously. It was shown that participant can learn to approach stimuli predictive of a monetary reward and to avoid stimuli predictive of punishment in an amended Go/NoGo task, where the stimuli were presented subliminally under forward-backward masking (Mastropasqua & Turatto, 2015; Pessiglione et al., 2008). This evidence suggests that goal-oriented behaviour requiring behavioural adaptation does not require conscious access. While this result has firmly established itself in the consciousness literature (despite scant support), there is also recent emerging evidence suggesting that more complex forms of learning, including instrumental conditioning (Reber, Samimizad, & Mormann, 2018), contingency learning (Travers, Frith, & Shea, 2018), as well as fear conditioning (Mertens & Engelhard, 2020) may not be possible unconsciously. With relevance specifically to instrumental conditioning, Reber and colleagues (2018) found that the ability to learn instrumental associations is directly related to stimulus visibility, with controlled subliminal stimulus presentation resulting in an inability to deploy behaviour instrumentally.

1.7.2. Theoretical considerations in unconscious learning

Past theoretical and empirical accounts of conscious versus unconscious processing suggest that the increased level of complexity required for selective behaviour may indeed require conscious

access to be established. While there is no clear agreed theory demarcating what conscious versus unconscious mental states may represent (Breitmeyer, 2015; Dupoux, Gardelle, & Kouider, 2008; Kouider & Dehaene, 2007), there have been attempts at outlining conditions under which unconscious processing can take place (Mudrik, Faivre, & Koch, 2014). A number of theoretical contributions propose that consciousness is related to long-lasting, long-range connections between distinct brain regions, supporting recurrent information integration across distinct cognitive modules (Dehaene & Naccache, 2001; Dehaene & Changeux, 2011; Dehaene, Charles, King, & Marti, 2014; Lamme, 2006; Mudrik, Faivre, & Koch, 2014). As such, low-level or short-range (spatial or temporal) information integration might be possible without conscious awareness. However, consciousness might be necessary at increased levels of complexity – including (but not limited to) semantic knowledge, complex visual processing, decision-making, and problem-solving (Baars, 2002; Treisman, 2003), all of which involve integration of information across longer spatiotemporal intervals or larger spatial distance. Indeed, neuroimaging evidence has found conscious processing to be characterised by global, long-range spread of activity, in contrast to more localised, shorter-range projections when processing is unconscious (Baars, Ramsøy, & Laureys, 2003; Dehaene et al., 2014, 2001; Dehaene & Naccache, 2001; Melloni et al., 2007).

As mentioned earlier, the tasks used to probe instrumental learning (Pessiglione et al., 2008; Mastropasqua & Turatto, 2015) involve a fairly complex process of integrating information over a large temporal scale and distinct modalities, necessary to process the visual input, extract its predictive value, deploy selective action in response to the predictive cue, and process the reinforcement. Their complexity is then significantly greater than in the aforementioned classical conditioning or associative learning scenarios, where there are fewer events, often in closer temporal proximity. Assuming that a subliminally presented cue is not capable of evoking large-scale activity to be integrated with subsequent processes, instrumental learning without conscious awareness should be prevented.

1.7.3. Methodological and statistical issues in unconscious learning research

A key challenge in any research into unconscious influences on behaviour lies in reliably asserting that processing is genuinely unconscious (Newell & Shanks, 2013; Rebuschat, 2013; Timmermans & Cleeremans, 2015). Although it is frequent practice in this line of research to infer unconscious processing when a behavioural measure (e.g. conditioning, priming, etc.) is above chance, while a separate measure of awareness is non-significantly different from chance

performance (e.g. a non-significant result in a discrimination task), this approach has drawn criticism (Dienes, 2015; Vadillo, Konstantinidis, & Shanks, 2016). A non-significant result alone cannot disambiguate between no evidence for an effect (i.e. insensitive data, e.g. due to the small sample size) and absence of an effect (i.e. support for the null hypothesis). As such, finding that performance on an awareness check does not significantly differ from chance is not enough to assert true absence of awareness – an assertion which must be fulfilled to enable any inferences about the effect of interest, such as presence of unconscious conditioning (Dienes, 2015; Shanks, 2017). In a meta-analysis, Vadillo, Konstantindis and Shanks (2016) argue that the seemingly chance-performance on this type of awareness test is more likely to reflect a false negative, especially in low-powered studies with a small sample. Those considerations cast doubt on the practices used in unconscious learning research (especially the older studies), and consequently, the validity of their results.

However, this failing can be rectified in several ways. One is to ensure that the methods to assess awareness are relevant, sensitive, and immediate (Berry & Dienes, 1993; Newell & Shanks, 2013; Shanks & St. John, 1994). The relevance criterion posits that the assessment of awareness should only target information relevant to the behaviour at hand; sensitivity posits that the assessment should be made under similar retrieval conditions (e.g. the same stimuli are used on the awareness test and on the performance test; Newell & Shanks, 2013). Optimal experimental conditions under those criteria could be achieved by, for instance, a closer similarity between the awareness test and the measure of interest. The immediacy criterion refers to the need for the awareness assessment to be made concurrently with or immediately after the behaviour of interest. Here, an optimal task could involve a judgment of awareness immediately after performance, on a trial-by-trial basis (in this sense, it would also account for relevance and sensitivity).

Another part of the solution to this problem is to apply appropriate statistical methods, such as the Bayes factor, which allows to determine whether a null result indicates true support for the null (e.g. absence of awareness) over the alternative hypothesis (presence of awareness), or whether the data are insensitive (Dienes, 2014, 2016; Sand & Nilsson, 2016). In conjunction, the above methods can be used to increase the sensitivity to detect the effect of interest in unconscious learning paradigms, reduce type I error, and draw substantiated statistical conclusions.

1.8. Overview of this research

1.8.1. Conclusions and aims

Instrumental learning is a fundamental substrate of adaptive behaviour, both phylogenetically and ontogenetically, allowing agents to learn to approach rewarding, positive stimuli in their environment, and avoid bad or harmful ones. Past research indicates that adaptive behaviour – including instrumental learning – is strongly dependent on the condition of the body. Under this perspective, the representations of the internal state of the body provide motivation for adaptive action in order to obtain rewards and avoid punishment (e.g. achieving homeostasis and allostasis, and avoiding dysregulation, respectively). Additionally, the autonomic nervous system has been proposed to continuously monitor performance for errors, a process reflected in electroencephalographic and autonomic readouts of errors, such as erroneously approaching a punishing stimulus. Shifting focus to the body on the outside, action and proprioception also appear to directly affect conscious experience, by providing an egocentric frame of reference for adaptive interactions with the world. Based on this literature, there are three main aims in the present thesis.

The first aim concerns the role of bodily information in instrumental learning. The extent to which instrumental learning is affected by the salience of bodily information (i.e. interoceptive precision) can shed light on the importance of bodily signals in adaptive processes, especially in primitive scenarios, where learning might occur without the influence of higher-order cognition.

Nonetheless, the extent to which instrumental learning is possible in simple scenarios – such as without conscious awareness of the stimuli – is itself an unsettled case. Aside from the bodily influences, learning may require conscious awareness to be able to drive motivated instrumental responses. Investigating the role of conscious awareness in instrumental learning is the second aim of the present thesis.

The third and final aim concerns the role of action in conscious experience. According to action-oriented predictive approaches to consciousness, it is the prediction of our interactions with the world – the sensorimotor contingencies – that directly shapes conscious experience. Yet, direct investigations of this hypothesis are still lacking.

1.8.2. Overview of the empirical work

The first two empirical chapters investigate the role of interoceptive precision in unconscious instrumental learning. Chapter 2 attempts to answer the question whether precision of the cardiac signal can affect simple forms of adaptive learning, such as unconscious instrumental

conditioning. Surprisingly, the results show evidence for absence of unconscious learning, in contrast to previous reports, making it difficult to interpret the impact of the interoceptive manipulation. Chapter 3 thus aimed to replicate the result, with an added measure of cardiac activity to investigate whether cardiac markers of learning might be present even in absence of stimulus awareness, again showing evidence of absence of unconscious learning, and of any learning-related cardiac activity. Together, those chapters show that instrumental conditioning might require conscious awareness.

Chapters 4 and 5 investigate the feasibility of unconscious instrumental conditioning specifically. Chapter 4 is a Stage 1 Registered Replication Report, a two-experiment direct replication of the prominent paradigm demonstrating unconscious instrumental learning, leveraging statistical and methodological advances in the field. Chapter 5 constitutes a conceptual replication of the same paradigm, proceeding in two modes of conditioning – trace and delay – demonstrating absence of successful instrumental conditioning without conscious awareness.

The final chapter, chapter 6, shifts focus from the body on the inside to the body on the outside, and examines the effect of action on our conscious experience. Instrumental learning was used to build novel and unique associations between stimuli and actions, which were then investigated in a breaking-Continuous Flash Suppression task. The results demonstrate that valid, congruent sensorimotor predictions directly facilitate access to conscious awareness, suggesting that our perceptual experience is affected by our actions in the world.

Overall, this body of work extends the current understanding of instrumental learning as a fundamental component of adaptive behaviour, showing that conscious access is required to drive active, adaptive interactions with the world, and to further shape our conscious perception in line with action.

2.

THE ROLE OF CARDIAC PRECISION IN UNCONSCIOUS INSTRUMENTAL LEARNING

Abstract

The ability to learn the associations between stimuli and outcomes, for example through instrumental conditioning, is essential for adaptive behaviour. Past research suggests that performance in such tasks might be affected by the autonomic nervous system, which aids internal feedback monitoring via an orienting response which follows motivationally salient events, such as mistakes in performance. Yet, it remains unclear whether such autonomic markers are driving learning, or are simply a correlate of salient events. Here, we manipulated the precision of cardiac information in an attempt to functionally affect the reliability of autonomic information during the unconscious instrumental learning process (approximating a primitive learning scenario without access to higher-order cognition). Precision was enhanced with synchronous real-time cardiac feedback, where tones were played in rhythm with participants' actual heartbeat. Conversely, precision was disrupted with asynchronous feedback. In the control condition, no cardiac feedback was provided. In two experiments, the manipulation did not reliably affect performance in the unconscious instrumental learning task. However, absence of learning in the control condition shows that instrumental conditioning may be more difficult to attain than previously thought, allowing for limited interpretation of the effectiveness of our manipulation. The chapter culminates with a discussion of this unexpected outcome and its implications.

2.1. Introduction

Learning to associate stimuli or actions with the possibility of reward or punishment is critical for survival, allowing humans and other animals to regulate their behaviour towards attaining positive outcomes (such as food, water, or pleasure) and avoiding harm. In the simplest form, classical (Pavlovian) conditioning, agents learn to react through direct experience with deterministic outcomes, positive or negative (Pavlov, 1928). However, the ability to extend those associations and learn to exert direct control over the environment, as in the case of instrumental (operant) conditioning, where actions resulting in positive consequences are reinforced, is imperative for survival in a dynamic world (Balleine & Dickinson, 1998b). Such forms of learning, combining classical and instrumental conditioning (Cameron, 2002; Dayan & Berridge, 2014), allow an individual to make choices dependent on value expectations from past experience, either as stimulus-response habits, or more flexible, adaptable actions (Daw, Niv, & Dayan, 2005; Dayan & Daw, 2008; Dezfouli & Balleine, 2012).

The fundamental, low-level nature of associative learning is highlighted by the findings that, under certain conditions, associations can be learned without conscious awareness of the contingencies, or even of the stimuli themselves. For example, simple conditioning (e.g. appetitive or aversive) was demonstrated when the predictive stimuli are rendered subliminal (Knight et al., 2003; Morris, Ohman, & Dolan, 1998; Olsson & Phelps, 2004; Seitz, Kim, & Watanabe, 2009b). Furthermore, perhaps more interestingly, associative learning was also shown to extend into the instrumental domain, and guide choices towards rewarding and away from punishing stimuli in absence of awareness of the stimuli themselves (Mastropasqua & Turatto, 2015; Pessiglione et al., 2008; but see Reber, Samimizad, & Mormann, 2018). This and other evidence suggests that motivational value can be represented and acted upon unconsciously (Bijleveld, Custers, & Aarts, 2012a; Pessiglione et al., 2007)

Effective learning of any association relies on continuous performance monitoring in order to detect errors and adjust behaviour accordingly. It has been suggested that this mechanism is partly internal, and is reflected in the autonomic nervous system (Critchley, 2005). The most prominent examples come from the error-related negativity (ERN) and positivity (Pe), electroencephalographic event-related potentials observed following errors in performance (Hajcak, McDonald, & Simons, 2004, 2003). A similar pattern was found in skin conductance response (SCR; Hajcak et al., 2003), and in cardiac signal, whereby the heart rate decelerates more following negative performance feedback (e.g. punishment) than following positive feedback (e.g. reward; Crone et al., 2003; Crone, Somsen, Beek, & Van Der Molena, 2004; Mies, Van der Veen, Tulen, Hengeveld, & Van der Molen, 2011; van der Veen, van der Molen, Crone, & Jennings, 2004). ERN, Pe, SCR, and heart rate deceleration alike have been proposed to

constitute an internal source of performance feedback, reflecting the mismatch between the representation of the response and the actual response, and have been reported even in absence of explicit awareness of the error (Nieuwenhuis et al., 2001). Indeed, such performance-dependent autonomic responses have been suggested to constitute an orienting response (Critchley, 2005; Sokolov, 1963), allowing the agent to adapt to environmental demands.

In a more general sense, these findings fall into the broader spectrum of research on autonomic influences on cognition. Research suggests that a host of both conscious and unconscious processes can be significantly affected by the state of the body, termed interoception, encompassing visceral, cardiac and cutaneous afferents (Critchley & Garfinkel, 2018). Famously, bodily signals were shown to affect decision-making by “marking” advantageous choices before explicit awareness (somatic marker hypothesis; Bechara, Damasio, Tranel, & Damasio, 1997; Carter & Smith Pasqualini, 2004). Elsewhere, interoceptive signals (especially cardiac afferents) have been proposed to affect a range of processes. For example, the cardiac cycle was shown to modulate access to visual or somatosensory awareness (Motyka et al., 2019; Salomon et al., 2016), visual search (Galvez-Pol et al., 2019), or processing of salient stimuli such as pain (Gray, Minati, Paoletti & Critchley, 2010) and fearful faces (Garfinkel et al., 2014). These effects are sometimes modulated by individual accuracy at discerning the interoceptive signals, termed interoceptive accuracy (IAcc). However, despite the modulatory effect, those studies concern a continuous oscillatory signal from the heart, and thus provide limited evidence regarding the functional, adaptive effect of cardiac signal on behaviour. As such, the question of the extent to which cardiac signal plays a functionally important part in behaviour remains open.

Combining those lines of research together, we hypothesise that cardiac signals functionally contribute to learning the stimulus-outcome associations, especially for stimuli carrying considerable motivational or affective value, such as reward and punishment. This idea is driven by the fact that learning appropriate responses to motivationally salient stimuli operates in service of homeostasis and allostasis, which, in turn, helps to maintain the body within viable states (Bijleveld, Custers, & Aarts, 2012b; Pezzulo et al., 2015). As such, low-level bodily markers – such as aforementioned autonomic internal feedback mechanisms – would support behaviour targeted at attaining reward and punishment, hence promoting survival. If this intuition is right, primitive forms of associative learning that can proceed without awareness and access to higher-order cognition should be similarly affected by cardiac interoceptive signals marking the motivational value of encountered stimuli.

Because direct, selective manipulation of cardiac signals is challenging, some research has employed external manipulations targeted at disrupting the integration of interoceptive

signals in the brain. For example, providing subjects with real-time external cardiac feedback, in or out of synchrony (i.e. hearing the heartbeat through the headphones, or seeing a stimulus flashing) has been shown to affect performance in a visual discrimination task (Łukowska et al., 2018), fairness judgments in the ultimatum game (Lenggenhager et al., 2013), or body ownership in the rubber-hand illusion (Suzuki et al., 2013). While external cardiac feedback is undeniably an exteroceptive signal, it has been proposed that asynchronous feedback can serve as a disruption to interoceptive information proper at the integration stage (Gray et al., 2007; van Elk et al., 2014), affecting the precision of the incoming signal, and its reliability (Petzschner et al., 2019; Seth, 2013; Seth & Friston, 2016). In contrast, synchronous feedback would amplify the cardiac signal, increasing precision of the information.

Here, such manipulation was employed in order to assess whether disrupting and amplifying task-related cardiac information can impact the formation of motivationally-relevant stimulus-response associations in a subliminal conditioning task⁵ (adapted from Pessiglione et al., 2008). A subliminal learning task was chosen in order to approximate a simple adaptive learning scenario without access to higher-order cognition. Assuming that the information coded in cardiac frequency provides a useful marker of internal performance monitoring (e.g. Crone et al., 2004), we hypothesise that disrupting that information by providing cardiac feedback asynchronously to actual heartbeat (thus decreasing the precision of the signal) should disrupt unconscious learning, as task-related interoceptive information (i.e. the deceleration following negative performance feedback) would lose reliability. In contrast, amplifying the cardiac feedback in synchrony with actual heartbeat should also amplify the task-related cardiac information, thus improving learning. Finally, individual interoceptive accuracy should affect the strength of the manipulation, such that individuals high in IAcc should be more susceptible.

2.2. Experiment 1

2.2.1. Method

2.2.1.2. Participants

97 participants (22 males) with a mean age of 22.5 years ($SD = 4.2$) were recruited for participation via the University of Sussex online recruitment system, and an internal mailing list. All participated in exchange for course credit or £6. All participants reported having normal or corrected-to-normal vision, and no current or history of cardiac or neurological illness. Ethical approval was granted by the Science and Technology Cross-School Research Ethics Committee

⁵ For the purpose of this paper, conditioning and learning will be used interchangeably to refer to the same process of learning stimulus-outcome associations.

at the University of Sussex, and the study was conducted in accordance with the Declaration of Helsinki.

2.2.1.3. Stimuli and Materials

The experiment was conducted using Matlab 2017b (MathWorks, 2017), running Psychophysics Toolbox (Brainard, 1997). All stimuli were presented on a Samsung 2233RZ LCD monitor (1680 by 1050 pixels) with a refresh rate of 120Hz, with the aim of ensuring fast and precise stimulus presentation in line with previous recommendations (Wang & Nikolić, 2011).

The target stimuli included twelve neutral symbols obtained from Agathodaimon font in the main task, and two circular shapes in the perceptual discrimination task used for threshold finding (see Appendix 1). All were 180x180 pixels in size, and presented in light grey (RGB: 217 217 217; HSV: 0° 0% 94.9%) on white background. The stimuli were forward and backward masked with black and white noise masks, also 180x180 pixels in size, with block size of 3x3 pixels. The forward and backward masks were different, but kept consistent between blocks and participants. Low contrast cues and the type of mask were deliberately chosen in order to increase the duration of presentation without conscious awareness, following Scott, Samaha, Chrisley and Dienes (2018).

Auditory cardiac feedback was produced using custom software developed for interoceptive feedback (Suzuki et al., 2014; previously checked for accuracy against an electrocardiogram). A Nonin soft finger USB pulse oximeter, placed on the ring finger of their non-dominant hand throughout the duration of the task. In the synchronous (SYNC) feedback condition, participants could hear beeps through the headphones that corresponded to their actual heartbeat, in real time, with each beep delivered at the point of R-peak, computed by averaging 5 previous interbeat intervals (Suzuki et al., 2014). In the asynchronous (ASYNC) feedback condition, the beeps were presented asynchronously, and reflected the detected heart rate with a 1500ms delay. In the CONTROL condition, no auditory feedback was delivered.

2.2.1.3. Procedure

Threshold setting.

Participants were seated with their chin on a chin rest placed at 45cm distance from the screen. Each session began with the threshold of visual awareness determined individually for each participant using a masked perceptual discrimination task. Each trial began with a fixation cross (500ms), followed by a mask (300ms), a target cue (either a symmetrical circular shape or an asymmetrical circular shape, starting at 600ms), and another mask (300ms). After each sequence, participants were asked to report whether the target cue was symmetrical or

asymmetrical by pressing corresponding arrows. Next, they were asked to assess whether they had any confidence that they saw the cue, or if they were guessing, also using corresponding arrows (following Scott et al., 2018). They were explicitly instructed to press 'some confidence' if they had any visual experience, even the faintest, of the cue, and 'total guess' only if they felt they did not see the cue and were responding randomly. Each time a correct response was made with confidence, the display duration of the target cue was reduced by 50ms on the following trial. When a duration of 100ms was reached, or the first guess response was made, the display duration returned to the previous level (+50ms), and subsequently reduced in 8.35ms steps on the following trials, corresponding to a single screen refresh duration for a 120Hz monitor. A reduction in exposure duration continued to be made after each non-guess response but not after guess responses. This process continued until participants indicated guessing on six consecutive trials, regardless of the accuracy of responses. The cue display duration in those trials was recorded as their individual unconscious threshold.

Main task.

The main task was adapted from the subliminal instrumental conditioning task used previously (Mastropasqua & Turatto, 2015; Pessiglione et al, 2008), in which participants make speeded go or no-go responses to the masked cues. Here, each trial consisted of a fixation cross (500ms), mask (300ms), target cue (display duration determined in the perceptual discrimination task), and mask (300ms), followed by a decision prompt in the form of a question mark, during which participants had 2 seconds to make a response (see *Fig.1.* for task chronology). Pressing the space bar (Go) indicated a decision to take the risk, at which point the participant could win 1 token (golden token displayed on the screen) or lose 1 token (a red cross over the golden token displayed), depending on the type of cue presented between the masks. Not pressing the space bar (NoGo) indicated a safe choice, which always resulted in a null outcome (greyed-out token). The task contained 6 blocks of 100 trials, with 50 rewarding and 50 punishing trials in randomised order. For each block, two target cues from the pool of 12 were randomly assigned to be rewarding or punishing, without replacement, ensuring that each block contained a novel pair. Participants were incentivised to maximise their earnings in the task through a prize draw, with the number of entries contingent on the amount of tokens won.

Due to the adaptation of the visual system to masking, after completing each block participants were asked to determine if they had any visual experience of the target cue on a binary (yes/no) scale. If they responded 'yes', the display duration was decreased by one step (8.35ms) in the following blocks. If they responded 'no', the duration was maintained.

The six blocks were randomly divided between three auditory cardiac feedback conditions (SYNC, ASYNC, CONTROL), resulting in two blocks per condition, in a randomised order.

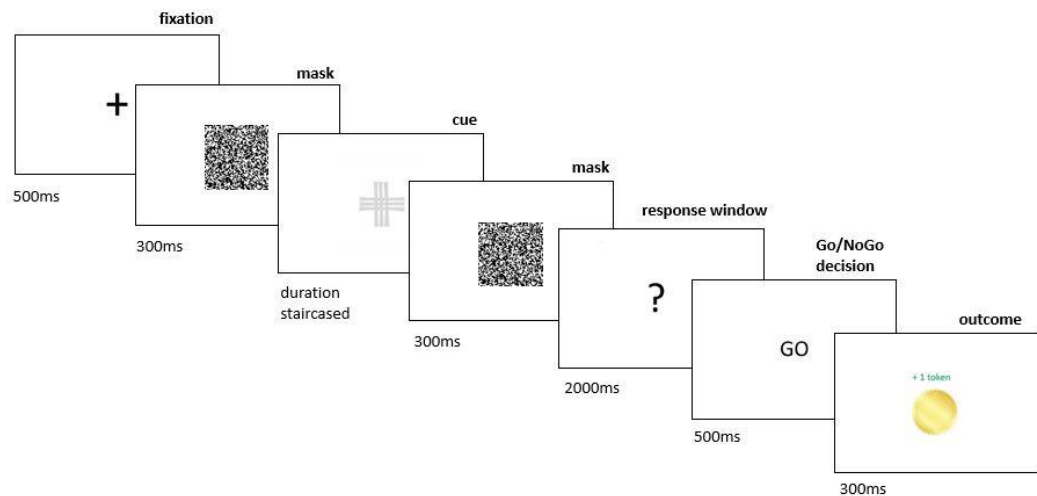


Figure 3. Subliminal instrumental conditioning task (main task). Chronological screenshots depict a single trial sequence, with durations in milliseconds. After cue presentation using forward-backward masking, participants had 2 seconds to make a Go response with the spacebar, or refrain from responding (NoGo). Following the response, the decision and feedback were displayed on the screen. In the example shown, a participant responded Go, which was the correct response for the cue presented, thus was rewarded with one gold token.

Interoceptive accuracy.

Finally, individual cardiac interoceptive accuracy was measured using the heartbeat discrimination task (Whitehead, Drescher, Heiman, & Blackwell, 1977). In the discrimination task, participants were asked to judge whether the auditory cardiac feedback, again delivered through the headphones, was synchronous or asynchronous with their own heartbeat (16 trials of 10 seconds each, 8 with synchronous, and 8 with asynchronous feedback). Asynchronous feedback was produced by manipulating the heart-rate to be either 70% or 130% of the frequency the recorded heart-rate, with the task employing 50% of both types (Suzuki et al., 2013). In both tasks, participants were explicitly instructed to sit still and avoid tracking their pulse physically. Following the interoceptive accuracy task, participants were thanked and debriefed.

2.2.2. Data pre-processing

2.2.2.1. Exclusion Criteria

A two-step process was applied to exclusion criteria using the block-by-block binary judgments of visual awareness of the cues. Firstly, participants were excluded from analysis entirely if they

reported awareness of the cues on over 50% (4-6) of the blocks. This resulted in 21 excluded subjects. The second step was performed on the remaining 76 participants. For participants who indicated awareness on fewer than 50% (1-3) of the blocks, only those blocks were removed from further analysis, resulting in partial exclusions for 29 participants (48 blocks (11%); SYNC: 11, ASYNC: 20, CONTROL: 17). 4 participants for whom the partial exclusions resulted in an entire condition missing were also removed, yielding usable data for 72 participants.

For the reaction time (RT) analysis, RTs under 100ms or greater than 2 standard deviations from individual means were excluded. Two additional participants who lost more than 25% of RT trials (150) to those criteria were excluded, yielding a sample of 70. Note that NoGo trials yielded no RTs, thus were not included in the analysis.

2.2.2.2. Interoceptive Accuracy

Interoceptive accuracy (IAcc) scores were computed for each participant by averaging the number of correct “synchronous”/“asynchronous” responses to the auditory cardiac feedback in the discrimination task.

2.2.3. Results

2.2.3.1. Bayes Factor

For hypothesis testing, Bayes Factors (Bs) will be reported alongside p -values for all comparisons made. Bs can help to disambiguate non-significant results as either indicating support for the null hypothesis (H_0 , positing no effect) or indicating insensitive data (i.e. the data are not in favour of either H_0 or H_1 , which uses an estimated raw effect size as the standard deviation of its distribution; Dienes, 2014). By convention, Bs smaller than 1/3 indicate evidence for H_0 . Bs larger than 3 indicate evidence for H_1 . Bs between those values indicate insensitive data.

2.2.3.2. Evidence of learning: Performance

In the conditioning task, participants were found to execute significantly more Go responses ($M = 61\%$) than NoGo responses ($M = 39\%$; $t(71) = 8.23$, $p < .001$), regardless of cue and condition, indicating a general response bias towards Go. In order to account for this, type I d' (a Signal Detection Theoretic measure of sensitivity to signal versus noise; Stanislaw & Todorov, 1999) was computed for each condition, treating Go responses to rewarding cues as Hits, and Go responses to punishing cues as False Alarms. The resulting measure of sensitivity can be taken as evidence of successful learning (i.e. discrimination between the cues) if it is significantly above 0. In a one-sample t -test against 0 (no sensitivity, thus no ability to discriminate between the

cues), the total d' was significantly above 0 ($M = 0.05$, $SE = 0.02$, $t(70) = 2.72$, $p = 0.009$, $B_{H(0,0.7)} = 1.90$), suggesting that, on average and irrespective of condition, participants were significantly above chance at discriminating the cue, but the result is insensitive under Bayesian criteria with respect to the level of learning previously obtained. For computing B, the predictions of H1 (learning is present) were modelled as a half-normal distribution centred on 0, with an SD equal to 0.7 (the expected effect size if learning is present; derived from Pessiglione et al., 2008).

The presence of instrumental learning by condition was assessed using a one-way repeated-measures analysis of variance (ANOVA), with learning (indexed by d') as the dependent variable, and cardiac feedback condition (SYNC, ASYNC, CONTROL) as the within-subject factor. All assumptions of a repeated-measures ANOVA were satisfied. The ANOVA revealed no main effect of condition on d' ($F(2,142) = 1.16$, $\eta_p^2 = 0.02$, $p = 0.317$). However, an interesting pattern was found when assessing the presence of learning in each condition separately in one-sample t-tests against 0. d' was found to be significantly above 0 in the SYNC cardiac feedback condition ($M = 0.09$, $SE = 0.03$, $t(70) = 2.69$, $p = 0.009$, $B_{H(0,0.7)} = 3.07$), but not in the ASYNC condition ($M = 0.02$, $SE = 0.03$, $t(70) = 0.77$, $p = 0.442$, $B_{H(0,0.7)} = 0.07$). In the CONTROL condition, which should constitute a learning baseline with no cardiac feedback given, there was no evidence of learning, but the result is marginally insensitive ($M = 0.06$, $SE = 0.04$, $t(70) = 1.66$, $p = 0.101$, $B_{H(0,0.7)} = 0.39$). See Figure 4 for a graphical representation.

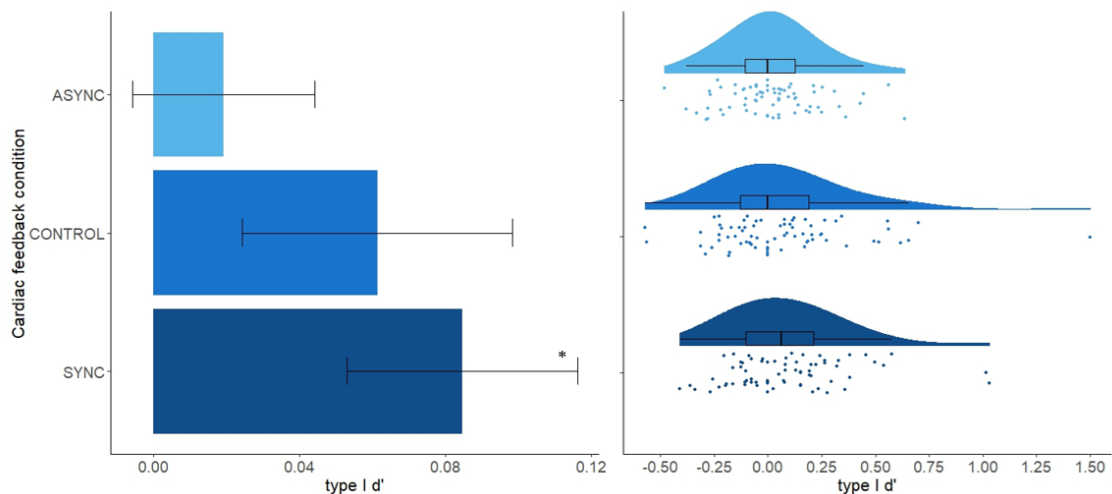


Figure 4: Type I d' values across the three cardiac feedback conditions in Experiment 1. Left panel: mean d' by condition (± 1 SEM). Star indicates significant difference from 0 at $p < .01$. Right panel: distributions of data with boxplots.

Subsequently, an exploratory analysis was conducted in order to better understand the temporal nature of learning. In a normal learning scenario, performance should begin around a chance level, and steadily improve with time and accumulation of experience. As such, each trial

block was split into 5 bins of 20 trials, and entered into a 3 (condition) by 5 (bins) repeated-measures ANOVA. The result revealed no interaction of condition and bin ($F(8,568) = .79$, $\eta_p^2 = 0.01$, $p = 0.590$), and no main effect of bin ($F(4,284) = 0.78$, $\eta_p^2 = 0.01$, $p = 0.537$) or condition ($F(2,142) = 1.26$, $\eta_p^2 = 0.02$, $p = 0.287$). Thus, there was no evidence that performance was affected by accumulation of experience over time. Nonetheless, visual inspection (Figure 5) reveals that performance in ASYNC condition was diverging from SYNC and CONTROL by the end of the learning session.




Figure 5: Type I d' (+/- 1 SEM) values across the three cardiac feedback conditions in Experiment 1, across 5 bins of 20 trials each.

2.2.3.3. Evidence of learning: Reaction times

For the RT analysis, an $RT_{\text{difference}}$ index was computed by subtracting RTs to rewarding cues from RTs to punishing cues. As such, positive values indicate that participants took a longer time to respond to punishing cues than to rewarding cues, in line with RT-oriented indicators of learning (e.g. Atas, Faivre, Timmermans, Cleeremans, & Kouider, 2014). Zero indicates that there was no difference between the two. Overall, irrespective of condition, $RT_{\text{difference}}$ was not significantly different from 0 in a one-sample t-test ($M = 3.63\text{ms}$, $SE = 4.60$, $t(69) = 0.79$, $p = 0.433$, $B_{H(0,34)} = 0.28$), suggesting that participants were not responding slower to punishing cues than to rewarding cues. B was computed with the predictions of H1 (learning is present) modelled as a half-normal distribution centred on 0, with an SD equal to 34ms (the expected effect size obtained from a past study which found an RT difference in the absence of performance effects in a similar task; Atas et al., 2014).

The index was then submitted as a dependent variable into a repeated measures ANOVA with condition (SYNC, ASYNC, CONTROL) as the within-subject factor. The ANOVA revealed no

main effect of condition ($F(2,138) = 0.42$, $\eta_p^2 = 0.01$, $p = 0.660$). In one-sample t -tests for each condition separately, the $RT_{\text{difference}}$ index was not significantly different from 0 in any condition (SYNC: $M = 5.43\text{ms}$, $SE = 8.02$, $t(69) = .68$, $p = 0.501$, $B_{H(0,34)} = 0.42$; ASYNC: $M = 4.63\text{ms}$, $SE = 8.73$, $t(69) = 0.53$, $p = 0.598$, $B_{H(0,34)} = 0.39$; CONTROL: $M = -4.37\text{ms}$, $SE = 7.91$, $t(69) = -0.55$, $p = 0.58$, $B_{H(0,34)} = 0.15$). See Figure 6 for a graphical representation.

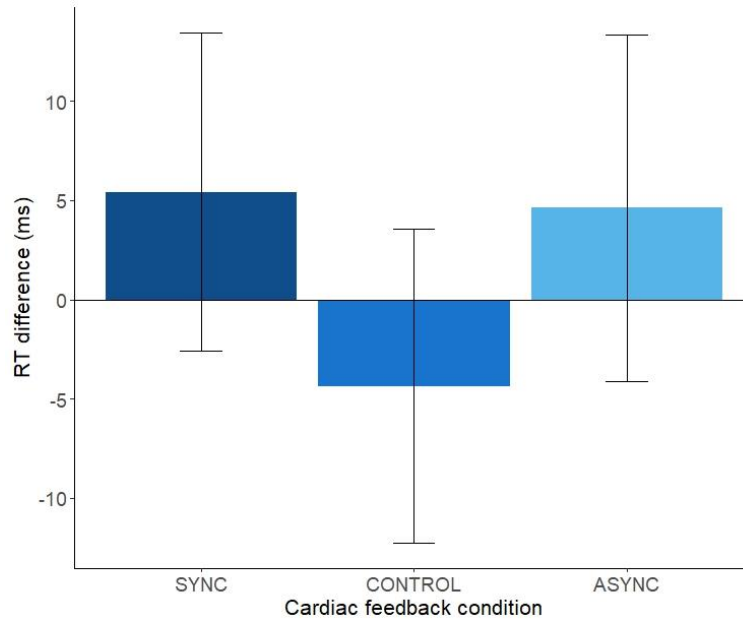


Figure 6: Mean $RT_{\text{difference}}$ index (± 1 SEM) across the three cardiac feedback conditions in Experiment 1. Positive values indicate that participants took a longer time to respond to punishing cues than to rewarding cues, in line with RT-oriented indicators of learning.

2.2.3.4. Performance and interoceptive accuracy

The average score on interoceptive accuracy in the cardiac discrimination task was 0.53 ($SD = 0.15$). In order to assess the effect of individual cardiac interoceptive ability, a second analysis was performed using a generalised linear mixed-effects model (GLMM). In this scenario, a GLMM allows to model the influence of IAcc by allowing each individual's slope to differ, in line with their IAcc score. All models were fit using lme4 package (Bates, Mächler, Bolker, & Walker, 2015) in R (R Core Team, 2018), by maximum likelihood with a binomial distribution and a logit link function. Treatment (dummy) coding was applied. Optimiser adjustments were applied to counter the non-convergence warnings.

The model included correct responses (i.e. Go to rewarding cues and NoGo to punishing cues) as a response variable (note that this is different from d' reported earlier – d' cannot be included in this model as it already is an aggregate measure). Full model specification included condition (SYNC, ASYNC, CONTROL), IAcc (continuous), and their interaction as fixed effects, and

subject-specific random slopes for IAcc as random effect. Random intercepts were suppressed, in line with recommendations to address convergence issues (Matuschek, Kliegl, Vasishth, Baayen, & Bates, 2017; Singmann & Kellen, 2019). See *Table 2* for estimated fixed effects regression coefficients. The analysis of deviance on the model, conducted using the Anova function in car package (Companion to Applied Regression; Fox & Weisberg, 2011), revealed no main effect of condition ($\chi^2 = 2.80$, $df = 2$, $p = 0.25$) or IAcc ($\chi^2 = 0.35$, $df = 1$, $p = .555$) and no significant interaction between condition and IAcc ($\chi^2 = 3.54$, $df = 2$, $p = 0.171$). In order to assess whether IAcc made a significant difference to model fit, this model was then compared to a simpler one, without IAcc as a fixed effect. Model comparison revealed that IAcc was not a significant addition to the model ($BIC_{IAcc} = 54594$, $BIC_{simple} = 54627$, $\chi^2 = 3.82$, $df = 3$, $p = 0.281$).

Table 1. Regression coefficients for the fixed effects from the generalised linear mixed model (GLMM) in Experiment 1. The intercept refers to the control condition. $N = 72$, number of observations = 39400. Note that parameters are given on the logit (not response) scale.

	<i>Estimate</i>	<i>Std. Error</i>	<i>z</i>	<i>p</i>
Intercept (CONTROL)	0.09	0.07	1.26	0.202
SYNC	-0.14	0.09	-1.51	0.131
ASYN	-0.10	0.09	-1.09	0.276
IAcc	-0.11	0.14	-0.81	0.417
SYNC: IAcc	0.31	0.17	1.88	0.060
ASYN: IAcc	0.17	0.17	0.98	0.326

2.2.4. Conclusions to Experiment 1

Experiment 1 sought to assess whether disrupting and amplifying task-related cardiac information with use of aurally delivered real-time cardiac feedback can influence learning in a simple, deterministic subliminal instrumental conditioning task. We hypothesised that synchronous cardiac feedback would improve learning with respect to the control condition (where no cardiac feedback was delivered), as task-related internal orienting or performance-feedback responses would become more precise. In contrast, asynchronous feedback was predicted to impair learning, due to the impaired precision of cardiac information.

Despite the presence of learning on average, as indexed by d' , a closer look at performance within each of the cardiac feedback conditions revealed that the effect is driven by presence of learning in the SYNC condition, but not in ASYN or CONTROL conditions. Interestingly, no clear evidence of learning was found in the CONTROL condition, which should have served as a baseline for evaluating learning in the remaining two conditions (note the insensitive data – although the Bayes Factor is closer to the H_0 bound). With this unclear result, as well as a small effect size (d' of 0.06 in comparison to 0.7 in Pessiglione et al. (2008) and 0.45

in Mastropasqua & Turatto (2015)), it is difficult to make inferences about the impact of the cardiac feedback manipulation.

As such, the most pertinent question is why participants in the present study failed to learn the stimulus-action contingency to the extent reported previously in nearly identical subliminal learning experiments (Pessiglione et al., 2008; Mastropasqua & Turatto, 2015). One candidate explanation is that the minor methodological differences in stimulus presentation affected the result. While the original Pessiglione et al (2008) design used white-on-black stimuli, presented for 33ms (corresponding to the objective threshold of conscious perception), the present experiment used lower-contrast grey-on-white stimuli, with a display duration determined individually for each participant (200ms on average). However, individual exposure durations and lower-contrast were introduced in order to allow for longer exposure durations without the cue being consciously perceived, which in turn should allow more time for the stimulus to be processed. As such, these changes should have, if anything, made it easier to obtain the effect, rather than obliterate it.

Another change introduced in the present experiment was the threshold of conscious perception. In contrast to the original Pessiglione design, which targeted the objective threshold (whereby responses on an awareness metric are objectively at chance), the present experiment attempted to target the subjective threshold (the level at which participants *believe* their performance is at chance, thus providing a closer index of subjective experience (Cheesman & Merikle, 1984, 1986; Dienes & Seth, 2010; Merikle & Daneman, 2000). It has been noted previously that subliminal effects found below the objective threshold are smaller and short-lived (Dienes, 2004, 2007; Draine & Greenwald, 1998; Lau & Passingham, 2006). Again, this change should have made it more likely for the effect to be detected.

Note, however, that in the threshold-setting task of the present experiment, participants were instructed to respond randomly and indicate guessing when they felt they had no confidence in seeing the cue presented. Typically, the subjective threshold is targeted by instructing to indicate guessing when participants feel they have no confidence in their judgment about the nature of the cue (here, whether it was symmetrical or not), as opposed to confidence in their awareness. This difference might have caused the thresholds to be determined as lower than the subjective level, bringing it closer to the objective level used in Pessiglione et al. (2008), albeit determined differently. Because subjective threshold implies a more plausible theory of consciousness than objective one (Seth & Dienes, 2010), this issue is rectified in Experiment 2.

Finally, as with any experimental investigation into unconscious processing, it is imperative to demonstrate that perception is truly outside of conscious awareness (Dienes,

2015; Newell & Shanks, 2013). While the present experiment attempted this by finding each participant's threshold individually and conservatively excluding an entire block if any cue awareness was reported, it does not preclude the possibility that participants might have seen the cues occasionally, yet failed to report it at the end of the blocks (an issue also pertinent to the original Pessiglione et al. (2008) design). As such, it is possible that some undetected conscious knowledge could have contributed to the learning process. Even if the visual experience was transient or degraded, the present study provided no means of assessing this on a trial-by-trial basis. This methodological flaw is addressed by Experiment 2.

2.3. Experiment 2

In order to address the methodological issues of Experiment 1, Experiment 2 replicated the paradigm and introduced three changes aimed at increasing the sensitivity of the task. Firstly, the threshold setting task instructions were amended to appropriately target the subjective threshold of conscious perception. Secondly, the main task included a trial-by-trial awareness check, allowing to exclude all trials with conscious knowledge of the type of cue presented. Finally, the number of trials in the cardiac discrimination task was increased from 16 to 40. All other details remained unchanged.

2.3.1. Method

2.3.1.1. Participants

46 participants (15 males, 1 unknown) with a mean age of 23.7 years ($SD = 9.84$) were recruited for participation via the University of Sussex online recruitment system, and an internal mailing list. Sample size was determined using the Bayesian Stopping Rule, with data collection continuing until a sensitive result was obtained in the main (learning) task (see *Analyses and Results* for detail). All participated in exchange for course credit or £6. All participants reported having normal or corrected-to-normal vision, and no current or history of cardiac or neurological illness. Ethical approval was granted by the Science and Technology Cross-School Research Ethics Committee at the University of Sussex, and the study was conducted in accordance with the Declaration of Helsinki.

2.3.1.2. Stimuli and Materials

Stimuli, materials, equipment and software used were identical to Experiment 1.

2.3.1.3. Procedure

Threshold setting.

Experiment 2 addressed the issue of the threshold likely being determined at a lower level than intended in Experiment 1. As in Experiment 1, participants were presented with a single cue (symmetrical or asymmetrical; starting at 600ms), forward-backward masked (300ms). Following the presentation, they were asked to report if the cue was symmetrical or not, and assess whether they had any confidence in their decision (as opposed to confidence that they saw the cue or not). They were explicitly instructed to press 'some confidence' if they had any confidence in what the cue was, even just a hunch, and 'total guess' only if they had no idea and were responding randomly. The staircasing procedure remained identical to Experiment 1, with the threshold being set when participants reported guessing on 6 consecutive trials, irrespective of decision accuracy.

Main task.

In order to be able to assess cue awareness with greater immediacy and sensitivity (Berry & Dienes, 1993; Shanks & St. John, 1994), Experiment 2 included a trial-by-trial awareness check. Following the Go or NoGo response on each trial, participants were asked to report whether the masked stimulus was symmetrical or asymmetrical by pressing corresponding arrows, followed by a binary assessment of confidence ('some confidence' or 'total guess'), in a manner identical to the threshold setting task. If they responded correctly with confidence three times in a row, cue display time was reduced by 8.35ms, corresponding to single screen refresh duration. All other task details remained identical to Experiment 1.

Interoceptive accuracy.

Identical to experiment 1. Following recommendations (Brener & Ring, 2016), the number of trials was increased from 16 to 40.

2.3.2. Data pre-processing

2.3.2.1. Exclusion Criteria

In order to ascertain that analyses are conducted only on trials that were truly unconscious, all individual trials where participants made a correct symmetry judgment with confidence were marked as aware and excluded (12.05% of all trials). Five participants who were aware on more than 25% (150) of all trials were excluded from further analysis. One participant was excluded due to failure to make any Go responses, yielding a final sample of 40 participants.

For the reaction time (RT) analysis, RTs under 100ms or greater than 2 standard deviations from individual means were excluded (8.76%). One participant who lost more than 25% of RT trials (150) to those criteria was excluded, yielding a final sample of 39 participants for the RT analysis. Note that NoGo trials yielded no RTs, thus were not included in the analysis.

2.3.2.2. Interoceptive Accuracy

Interoceptive accuracy (IAcc) scores from the heartbeat discrimination task were computed for each participant, by averaging the number of correct responses.

2.3.3. Results

2.3.3.1. Evidence of learning: Performance

Similarly to Experiment 1, participants made more Go ($M = 57\%$) than NoGo responses ($M = 43\%$; $t(39) = 2.07$, $p = 0.045$), indicating a general response bias towards Go. Type I d' was computed in order to account for this. Boxplot inspection and studentized residuals test revealed one outlier in the ASYNC condition, which was removed. At the group level, d' scores were entered into a one-sample t-test against 0, which indicates no ability to discriminate the stimuli. B was computed for the difference, with the predictions of H_1 modelled as a half-normal distribution, mean specified as 0, and d' of 0.7 as the SD of the mean (corresponding to the expected effect size from Pessiglione et al., 2008). In line with the Bayesian Stopping Rule, data collection continued until a sensitive result was found in support of either H_0 (absence of learning) or H_1 (presence of learning). The result indicates that the average (all conditions) d' was not significantly different from 0 ($M = 0.01$, $SE = 0.03$; $t(39) = 0.38$, $p = 0.706$, $B_{H(0,0.7)} = 0.06$). As such, we found evidence of absence of subliminal instrumental conditioning on average.

Presence of instrumental learning by condition was assessed using a one-way repeated measured ANOVA, with learning (indexed by d') as the dependent variable, and cardiac feedback condition (SYNC, ASYNC, CONTROL) as the within-subject factor. All assumptions of a repeated-measures ANOVA were satisfied. Again, ANOVA revealed no main effect of condition on d' $F(2,76) = 0.59$, $\eta_p^2 = 0.02$, $p = 0.566$). In one-sample t-tests against d' of 0 (no sensitivity) for each condition separately, we found no evidence of learning in any of the conditions (SYNC: $M = -0.04$, $SE = 0.042$, $t(39) = -0.90$, $p = 0.374$, $B_{H(0,0.7)} = 0.03$; ASYNC: $M = 0.01$, $SE = 0.05$, $t(38) = 0.23$, $p = 0.822$, $B_{H(0,0.7)} = 0.08$; CONTROL: $M = 0.02$, $SE = 0.05$, $t(39) = 0.37$, $p = 0.714$, $B_{H(0,0.7)} = 0.09$).




Figure 7: Type I d' values across the three cardiac feedback conditions in Experiment 2. Left panel: mean d' by condition (± 1 SEM). Right panel: distributions of data with boxplots.

A 3 (condition) by 5 (bins of 20 trials) repeated-measures ANOVA, with d' as the dependent variable, revealed a main effect of bin $F(4,156) = 3.40$, $\eta_p^2 = 0.08$, $p = 0.011$, but no interaction between condition and bin ($F(8,312) = 0.55$, $\eta_p^2 = 0.01$, $p = 0.819$), and no main effect of condition ($F(2,78) = 0.77$, $\eta_p^2 = 0.02$, $p = 0.467$). In pairwise comparisons (Sidak-adjusted) between the bins, only bin 3 ($M = -0.13$, $SE = 0.05$) and bin 5 ($M = 0.12$, $SE = 0.06$) differed significantly ($p = 0.028$), suggesting that overall performance improved from bin 3 (middle of task) to bin 5 (end of task; see Figure 8). Note that Bs were not computed for those comparisons due to the difficulty of obtaining a reasonable expected effect size.




Figure 8: Type I d' (± 1 SEM) values across the three cardiac feedback conditions in Experiment 2, across 5 bins of 20 trials each.

2.3.3.2. Evidence of learning: Reaction times

Identically to Experiment 1, an $RT_{\text{difference}}$ index was computed by subtracting RTs to rewarding cues from RTs to punishing cues (with positive values indicating longer response times to punishing than to rewarding cues). Overall, $RT_{\text{difference}}$ was not significantly different from 0 in a one-sample t-test ($M = 2.80\text{ms}$, $SE = 9.35$, $t(39) = 0.30$, $p = 0.767$, $B_{H(0,34)} = 0.34$), suggesting that participants were not responding slower to punishing cues than to rewarding cues, as would be expected had learning occurred (note the marginally insensitive Bayes Factor, indicating that by conventional thresholds of evidence the data cannot be taken to support the null hypothesis of no difference).

In a repeated-measures ANOVA with condition (SYNC, ASYNC, CONTROL) as the within-subject factor, we found no evidence for an effect of condition on $RT_{\text{difference}}$ ($F(2,74) = 0.87$, $\eta_p^2 = 0.02$, $p = 0.402$). In one-sample t-tests for each condition separately, the $RT_{\text{difference}}$ index was not significantly different from 0 in either (SYNC: $M = -19.11$, $SE = 11.18$, $t(37) = -1.71$, $p = 0.096$, $B_{H(0,34)} = 0.12$; ASYNC: $M = -10.96\text{ms}$, $SE = 22.44$, $t(37) = -0.49$, $p = 0.628$, $B_{H(0,34)} = 0.41$; CONTROL: $M = -11.03\text{ms}$, $SE = 17.02$, $t(37) = 0.65$, $p = 0.521$, $B_{H(0,34)} = 0.76$; but note insensitive data in ASYNC and CONTROL; see Figure 9).

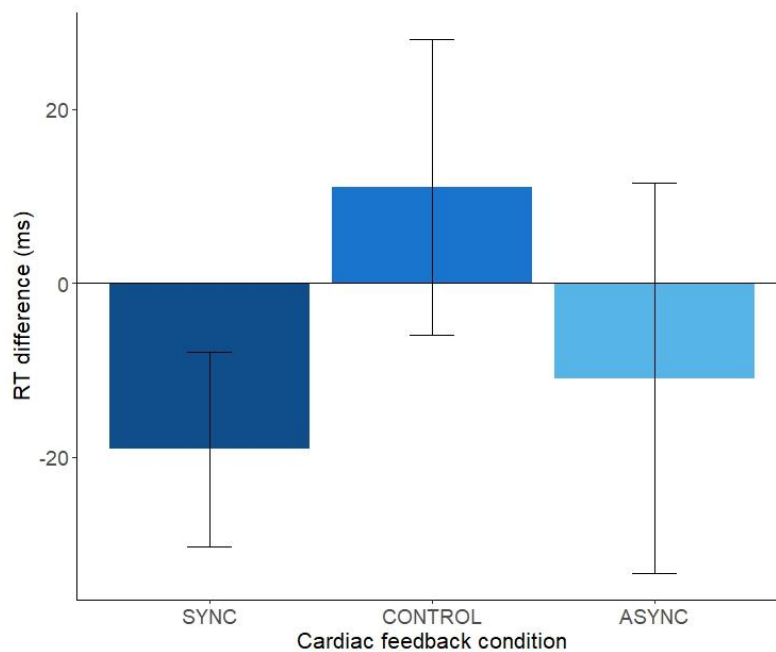


Figure 9: Mean $RT_{\text{difference}}$ index (± 1 SEM) across the three cardiac feedback conditions in Experiment 2.

2.3.3.3. Learning and interoceptive accuracy

The average score on interoceptive accuracy in the cardiac discrimination task was 0.51 ($SD = 0.12$). The effect of individual cardiac interoceptive accuracy was modelled with a GLMM

specified identically to Experiment 1. See *Table 3* for estimated fixed effects regression coefficients. The analysis of deviance on the model, conducted using the Anova function in car package (Fox & Weisberg, 2019), revealed no main effect of condition ($\chi^2 = 1.14$, $df = 2$, $p = 0.565$) or IAcc ($\chi^2 = 0.67$, $df = 1$, $p = 0.413$) and no significant interaction between condition and IAcc ($\chi^2 = 0.06$, $df = 2$, $p = 0.971$). In model comparison, a simpler model was again favoured, suggesting that IAcc was not a significant predictor ($BIC_{IAcc} = 30407$, $BIC_{simple} = 30378$, $\chi^2 = 0.73$, $df = 3$, $p = 0.867$).

Table 3. Regression coefficients for the fixed effects from the generalised linear mixed model (GLMM) in Experiment 2. The intercept refers to the control condition. N = 40, number of observations = 21925. Note that parameters are given on the logit (not response) scale.

	<i>Estimate</i>	<i>Std. Error</i>	<i>z</i>	<i>p</i>
Intercept (CONTROL)	-0.03	0.13	-0.31	0.759
SYNC	-0.06	0.15	-0.43	0.666
ASYNC	-0.01	0.15	-0.05	0.957
IAcc	0.13	0.25	0.53	0.599
SYNC: IAcc	0.07	0.28	0.24	0.812
ASYNC: IAcc	0.01	0.28	0.07	0.942

2.3.4. Conclusions to Experiment 2

Experiment 2 attempted to improve the methodology of Experiment 1 in order to assess whether disrupting or amplifying task-related cardiac information can influence learning in a simple, deterministic subliminal instrumental conditioning task. In the present study, the stimuli were presented reliably below the subjective threshold of conscious perception, and a trial-by-trial awareness check was used in order to eliminate all trials where participants displayed cue awareness and ensure only truly subliminal trials were analysed.

The results of Experiment 2 demonstrate that no subliminal conditioning was achieved under those conditions. This was the case on average, as well as when evaluating the three cardiac feedback conditions separately. In fact, after excluding all trials where participants were aware of the cue, the effect size found here was still smaller than in Experiment 1 (d' of 0.01 overall and 0.02 in the control condition, in comparison to 0.05 overall 0.06 in the control condition of Experiment 1). Similarly, we found no evidence of learning in the RT data (although note that the data is insensitive, preventing a strong conclusion from being drawn), and no evidence of an influence of individual interoceptive accuracy on learning. Again, lack of learning at baseline prevents any conclusions about the effectiveness of the interoceptive manipulation used.

2.4. Discussion

The ability to learn the associations between stimuli and outcomes is the cornerstone of adaptive behaviour. Previous research suggests that this process might be partially driven by an internal feedback monitoring mechanism in the autonomic nervous system, which allows for an orienting response following motivationally salient events, such as errors in performance (Crone et al., 2004b; Łukowska et al., 2018). Yet, it remains unclear whether such autonomic or bodily markers are functionally driving learning, or are simply an autonomic correlate of adaptively salient events. Here, we took a step towards answering that question by designing a manipulation targeted at disrupting and amplifying the precision of cardiac information. Participants listened to the real-time rhythm of their own heartbeat delivered asynchronously and synchronously while performing a subliminal conditioning task, chosen in order to approximate a simple adaptive learning scenario without the need for higher-order cognition. We hypothesised that if the information coded in cardiac frequency (i.e. cardiac deceleration following negative performance feedback; Crone et al., 2004) is functionally useful for instrumental learning, disrupting its precision should disrupt performance, while amplifying its precision should improve it.

Unfortunately, our ability to draw conclusions regarding the impact (or lack thereof) of the cardiac feedback manipulation is limited due to the lack of evidence of learning in the first place. Our manipulation relied on the assumption that instrumental conditioning is feasible without explicit awareness of the stimuli (Atas et al., 2014; Mastropasqua & Turatto, 2015; Pessiglione et al., 2008). In Experiment 1, we found strong evidence of learning in the synchronous cardiac feedback condition, but the methodological limitations cast doubt on the reliability of this result. Experiment 2 rectified those by correcting the individual threshold-setting task, as well as introducing trial-by-trial awareness checks, allowing to exclude all trials where participants might have been aware of the cues. The results of Experiment 2 demonstrate that when participants had no conscious awareness of the rewarding or punishing nature of the stimuli, they were unable to discriminate between them and adjust their behaviour accordingly (choose to Go or NoGo, respectively). Bayes Factors were used to supplement null hypothesis significance tests, allowing us to assert that the null result obtained indicates lack of learning.

In this light, we must echo the question posed in conclusions of Experiment 1, and ask why participants failed to learn the stimulus-outcome contingency, contrary to previous research using the same task (Mastropasqua & Turatto, 2015; Pessiglione et al., 2008). Arguably, the methods employed in Experiment 2 to ensure processing is genuinely below the threshold of conscious perception are more rigorous than those used previously. Here, we excluded all

trials where there was evidence of cue awareness, and all subjects who demonstrated awareness on over 25% of trials, thus analysing performance only in genuinely subliminal cases. In contrast, previous studies assessed absence of awareness by comparing performance between the conditioning task itself, and a separate discrimination task. Although it is frequent practice in implicit learning research to infer unconscious processing when a behavioural measure (e.g. conditioning) is above chance, while an explicit measure of awareness yields chance performance (e.g. a null result in a discrimination task), this approach has been heavily criticised. Firstly, a non-significant result alone cannot disambiguate between no evidence of an effect (i.e. insensitive data, e.g. due to the small sample size) and lack of an effect (i.e. support for the null hypothesis). As such, a non-significant result in a discrimination task serving as an awareness check does not indicate that awareness was indeed absent, casting doubt on the practice used in the original study (Pessiglione et al., 2008). In a meta-analysis, Vadillo, Konstantindis and Shanks (2016) argue that seemingly chance-performance on this type of awareness tests is more likely to reflect a false negative, especially in low-powered studies with a small sample. This issue is addressed more fully in Chapter 3 of this thesis.

It is also possible that the conflicting findings are due to the differences in task and stimulus presentation parameters. The high-contrast, white-on-black cues used in the original task (Pessiglione et al., 2008) could increase the likelihood of becoming visible on occasion, especially given that individual differences in perceptual thresholds were unaccounted for. Similarly to the issue in Experiment 1, lack of trial-by-trial awareness checks in the previous studies risks the possibility that participants had some awareness of the cues or their parts (e.g. a distinctive edge), even if brief or degraded, and failed to report it when prompted at the end of the task. To avoid those pitfalls, we used low-contrast, light grey-on-white stimuli, in order to maximise the exposure duration under masking, and maximise the likelihood of being processed while still being consciously undetected. We also used a staircased perceptual discrimination task in order to determine the threshold for each individual (following the robust procedure of Scott et al., 2018), and presented the stimuli just below the determined subjective threshold.

It might also be the case that instrumental conditioning requires some degree of awareness. While evidence exists that other forms of simple associative learning can take place unconsciously, even under conservative awareness criteria (Dupoux, Gardelle, & Kouider, 2008; Scott et al., 2018) it is unclear whether the same is possible for more complex processes, such as numerical or semantic processing (Hasselman et al., 2015), or learning across larger spatiotemporal intervals (Clark & Squire, 1998; Mudrik et al., 2014), but note that this also appears to be subject to task parameter differences such as similarity between stimuli; Gaillard et al., 2007; Reber, Luechinger, Boesiger, & Henke, 2012). The present task might be subject to

the latter problem. In order to effectively learn when to act and when to refrain from acting, participants must learn from three temporally distinct events: the stimulus itself, their action, and the consequences. It is plausible that the subliminal cue cannot evoke a sufficiently large-scale activity to be integrated with subsequent information across the length of the trial (up to 4 seconds in total). In contrast, in the aforementioned simpler forms of associative learning, the information to be associated typically follows immediately. As such, it is possible that some degree of consciousness might be necessary for behaviour requiring goal-directed, selective action (Dehaene & Naccache, 2001; Kouider & Dupoux, 2001). Support for this idea comes from Reber et al. (2018), who, employing yet another means to degrade cue awareness, also showed that instrumental conditioning failed to take place below the threshold. Chapters 3 and 4 of this thesis address this issue more extensively, and explore the conditions under which instrumental conditioning might be possible.

Finally, it is conceivable that our methodology, although more rigorous, made it more difficult for learning to happen. For one, with the trial-by-trial awareness check, we introduced a further variable time lag (about 2 seconds) between the trials. This might have further disrupted the already precarious process of temporal integration discussed above. However, given that this disruption occurs between trials, it could be argued that it would only affect the long-term consolidation of the associations learned within trials. If so, a sufficient number of trials should eliminate the issue. In the original study, evidence of learning plateaued after around 60 trials, uninterrupted by an awareness check. As such, it is possible that the 100 trials per block were not sufficient to produce conditioning in our case. Yet, even in the final bin (80-100 trials), the index of learning in the present experiment was considerably smaller than that reported at plateau in the previous experiments.

To summarise, in two experiments we found no evidence of successful instrumental learning at baseline. We propose that forms of learning involving selective action might be more challenging to evoke under sub-threshold conditions than previously thought. Future experiments should methodically explore the conditions under which unconscious instrumental learning succeeds or fails (see Chapters 3 and 4 of this thesis for a more thorough treatment of this issue). Unfortunately, the present result limited our ability to assess the effectiveness of the manipulation of interest, which attempted to affect learning by amplifying or disrupting cardiac precision. Taking this forward, Chapter 2 explores whether there is any indication of learning-related cardiac activity in the same task.

3.

CARDIAC CONCOMITANTS OF UNCONSCIOUS INSTRUMENTAL LEARNING AND THE ROLE OF CARDIAC PRECISION

Abstract

Performance monitoring is a vital aspect of successful learning and decision-making. Past research shows that performance monitoring is reflected in autonomic nervous system activity, including a deceleration in heart rate following errors in performance. Such autonomic responses were proposed to constitute an orienting response that indicates a need for behavioural adjustments. However, it is unclear whether autonomic signals directly drive performance, or are simply a correlate of salient events such as errors. Similarly, it is unclear to what extent those signals are present without conscious awareness of the task stimuli. In this experiment, we used unconscious instrumental learning as an approximation of a primitive adaptive learning scenario without access to higher-order cognition to investigate the functional role of cardiac information in the learning process. We manipulated the precision of cardiac information as participants performed the learning task, and collected electrocardiography to measure their cardiac activity throughout the learning process. The results demonstrate absence of instrumental learning without stimulus awareness, precluding a meaningful interpretation of the effectiveness of the cardiac manipulation. We also found absence of learning-related cardiac deceleration, suggesting that the autonomic performance monitoring mechanism might not be engaged when stimuli are perceived unconsciously. Those results suggest that conscious awareness may be necessary for successful decision-making.

3.1. Introduction

Performance monitoring is a critical aspect of successful decision-making. Efficient monitoring involves the ability to swiftly detect errors in performance in order to adjust future behaviour accordingly. This can be achieved by monitoring external events for feedback (e.g. consequences of action, such as reward or punishment). In the instrumental conditioning or learning process⁶, actual outcome is compared to expected outcome, with mismatches (e.g. following from an error) informing the learner of the necessity for behavioural adaptation. However, it has been proposed that performance is additionally monitored internally by the autonomic nervous system (ANS) (Hajcak, McDonald, et al., 2003; Müller et al., 2005; Sokolov, 1963). Correlates of error- and feedback-related activity have been observed in error-related negativity (ERN or Ne) and positivity (Pe) components in electroencephalography (EEG) (Falkenstein et al., 2000; Nieuwenhuis et al., 2001; Overbeek et al., 2005), skin conductance response (Crone et al., 2004a), and in evoked cardiac deceleration (Crone et al., 2003; Somsen et al., 2000). Together, they have been suggested to constitute an internal feedback monitoring system which reflects mismatches between an internal representation of the response or its predicted outcome, and the actual outcome.

Cardiac deceleration in response to feedback has been extensively studied in the context of learning. The heart tends to decelerate more in response to negative feedback (initiated in anticipation of feedback and continuing until after its presentation), reflecting the processing of error and the need for behavioural adjustment, than in response to positive feedback, which elicits faster recovery (Crone, Bunge, de Klerk, & van der Molen, 2005b; Crone et al., 2003; Somsen et al., 2000; van der Veen et al., 2004a). Notably, this response is only evident when feedback is valid – that is, when it carries reliable, usable information about performance (Crone et al., 2003; Groen et al., 2007; Mies et al., 2011). Feedback is considered valid or informative when it can be related to the nature of the stimulus – that is, when the stimulus reliably predicts the outcome, and the feedback reflects the match or mismatch between the predicted and actual outcomes. In addition, this deceleration response has been found to diminish as learning progresses, instead becoming more pronounced in anticipation of negative feedback presentation than following it (Crone et al., 2004a, 2003; Groen et al., 2007), or even appearing as early as the negative stimulus itself (Kastner et al., 2017). This effect has been taken to imply that there is a shift from external, response-based performance monitoring to internal, prediction-based monitoring over the course of learning. Crucially, the strength of

⁶ For the purpose of this paper, conditioning and learning will be used interchangeably to refer to the same process of learning stimulus-outcome associations.

the phasic cardiac response to negative feedback has been found to correlate with the strength of the prediction error signal (derived from a reinforcement learning algorithm; Kastner et al., 2017), which alerts the learner to the need to adjust behaviour.

While feedback-related cardiac deceleration appears reliably in standard learning paradigms (e.g. probabilistic learning), the extent to which internal feedback monitoring mechanisms operate in unconscious learning remains unknown. It has been previously demonstrated that instrumental learning can occur without conscious awareness of the outcome-predictive cues themselves, which also precludes conscious associations between the cues and their contingent outcomes (Mastropasqua & Turatto, 2015; Pessiglione et al., 2008). In the original Pessiglione et al (2008) task, subjects learned to approach rewarding stimuli and avoid punishing ones without ever seeing them consciously. At this point, it is noteworthy that the capacity for complex forms of unconscious learning, such as those involving adjustment of instrumental responses in the absence of cue or contingency awareness, has recently been put into question (Mertens & Engelhard, 2020; Reber, Samimizad, & Mormann, 2018; Travers, Frith, & Shea, 2018; see also Chapters 2, 4, and 5 of this thesis). Nonetheless, error-related cardiac deceleration has been demonstrated to occur in an unconscious stimulus discrimination task, where participants judged the orientation of a masked Gabor patch, albeit to a lesser extent than visible stimuli (Łukowska et al., 2018). As such, it appears that internal feedback monitoring is indeed active in the absence of reportable awareness of the nature of the cue.

If this assumption is true, it is plausible that there could be cardiac behaviour indicative of differentiating between the stimuli, before any behavioural manifestation of learning (i.e. instrumental adjustment of behaviour). Although in the aforementioned conditioning task feedback is always consciously perceived, error-related cardiac deceleration would only be expected if it is informative – that is, if the mechanism can associate the nature of the unconscious cue with the feedback presented. As such, differentiating whether the stimulus is rewarding or punishing is essential for feedback to provide any information (i.e. I made an error approaching a punishing cue). To explore this, the first objective of the present experiment is to replicate the Pessiglione et al (2008) unconscious conditioning paradigm, including a continuous measure of cardiac activity via an electrocardiogram (ECG) to assess presence of cardiac deceleration both in response to feedback, and to the predictive cue itself.

However, the evidence described earlier is insufficient to determine whether cardiac activity actually drives learning, or is simply a correlate or a marker of the learning process (although its correlation with the magnitude of prediction error in a reinforcement learning model is a promising start; Kastner et al., 2017). One of the ways to assess that would be to directly manipulate the precision or reliability of the cardiac signal, affecting the quality or

usefulness of the learning-related cardiac information in the brain. Because direct, selective manipulation of the cardiac signal is challenging, some research has employed external manipulations targeted at disrupting the integration of interoceptive (internal) signals in the brain. One such approach is auditory cardiac feedback – providing subjects with an auditory stream either in direct synchrony, in real-time with their heartbeat, or out of synchrony. While external cardiac feedback is undeniably an exteroceptive signal, it has been shown that the brain processes synchronous cardiac feedback in a manner similar to other self-generated sounds (van Elk et al., 2014), and that it reacts to unexpected omissions from a cardiac-synchronous auditory stream, but not from an asynchronous one (Pfeiffer & De Lucia, 2017). It has been proposed that asynchronous feedback can serve as a disruption to interoceptive information proper at the integration stage, affecting the precision of the incoming signal, and its reliability (Gray et al., 2007; Petzschner, 2017; van Elk et al., 2014). In contrast, synchronous feedback would increase the precision of cardiac information. The construct of precision is formalised in the interoceptive inference framework (Owens, Allen, Ondobaka, & Friston, 2018; Owens, Friston, et al., 2018; Petzschner et al., 2019; Seth & Friston, 2016; Seth, 2013), which proposes that predictions (also called beliefs) about internal states of the body are continuously updated based on the afferent neural signal from the body. Put simply, the brain makes continuous inferences about the state of the body based on prior expectations and the incoming signal, both of which are weighted according to their expected relative precision in a given context. A mismatch between the prior prediction and the actual signal, called a prediction error, would update the beliefs. Importantly, belief update is increased if the prediction errors are precise, and reduced if the predictions are precise (Petzschner et al., 2019).

In the context of learning, if cardiac signal functionally informs the learning process, then amplified, high-precision signal should improve the ability to learn, and disrupted signal should impair it (see also Chapter 2 of this thesis). For instance, performance error-related cardiac deceleration that is deemed very imprecise should have less impact as a learning-related marker than signal that is precise. In this view, cardiac signals would impact instrumental behaviour insofar as they modulate the inference in line with their precision. Thus, the second objective of the present experiment is to investigate the effect of cardiac precision on instrumental learning.

To summarise, we present an unconscious instrumental conditioning task modelled on Pessiglione et al. (2008), with an added manipulation of cardiac precision via auditory cardiac feedback and a continuous ECG measure. We focus on 3 questions: 1) Is there evidence for instrumental learning and is it affected by the precision of cardiac signal, manipulated with auditory cardiac feedback? If cardiac information is functionally important for learning,

performance should be affected by amplifying or disrupting cardiac precision, relative to baseline. 2) Is there evidence of performance feedback-related cardiac deceleration in the absence of cue awareness? The presence of this pattern would indicate that the cue-feedback association is still being processed without explicit awareness of the contingency. 3) Is there evidence of feedback-related deceleration decreasing, and cue-related deceleration increasing, in line with previously reported markers of learning? Again, the presence of this pattern might suggest that the heart is incorporating the informative value of the cues and feedback.

3.2. Method

3.2.1. Participants

40 participants (13 males, 2 unknown) with a mean age of 25 years ($SD = 3.27$, range = 21-33, age for 7 participants unknown) were recruited for participation via the University of Sussex online recruitment system. All participated in exchange for course credit or £9.00. All participants reported having normal or corrected-to-normal vision, and no current or history of cardiac or neurological illness. Data for one participant was removed due to software malfunction. Ethical approval was granted by the School of Psychology ethics committee at the University of Sussex, and the study was conducted in accordance with the Declaration of Helsinki.

3.2.2. Stimuli and Materials

The experiment was conducted using Matlab 2018b (MathWorks, 2018), running Psychophysics Toolbox (Brainard, 1997). All stimuli were presented on a Samsung 2233RZ LCD monitor (1680 by 1050 pixels) with a refresh rate of 120Hz, with the aim of ensuring fast and precise stimulus presentation in line with previous recommendations (Wang & Nikolić, 2011). The target stimuli included twelve neutral symbols obtained from Agathodaimon font in the main task, and two circular shapes in the perceptual discrimination task used for threshold finding (see Appendix 1). All were 180x180 pixels in size, and presented in light grey (RGB: 217 217 217) on white background. The stimuli were forward and backward masked with black and white noise masks, also 180x180 pixels in size, with block size of 3x3 pixels. The forward and backward masks were generated afresh on each trial by randomly scrambling the noise image. Low contrast cues and the type of mask were deliberately chosen in order to increase the duration of presentation without conscious awareness, following Scott, Samaha, Chrisley and Dienes (2018). Responses were collected with a standard keyboard.

Auditory cardiac feedback was produced using custom software developed for interoceptive feedback (Suzuki, Garfinkel, Critchley, & Seth, 2013; previously checked for accuracy against an electrocardiogram). The program detected the R-peaks of the cardiac QRS complex (adjusted by an average 400ms transit time between the finger and the heart) with a Nonin soft finger USB pulse oximeter, placed on the ring finger of their non-dominant hand throughout the duration of the task. In the synchronous (SYNC) feedback condition, participants could hear beeps through the headphones that corresponded to their actual heartbeat, in real time, with each beep delivered at the point of R-peak, computed by averaging 5 previous interbeat intervals (Suzuki et al., 2013). In the asynchronous (ASYNC) feedback condition, the beeps were presented asynchronously, and reflected the detected heart rate with a 1500ms delay. In the CONTROL condition, no auditory feedback was delivered.

3.2.3. Electrocardiography

Electrocardiogram (ECG) was recorded for the duration of the main task using Biopac MP36, running Biopac Student Lab 3.7.7 (Biopac Systems, 2012), with a sampling rate of 500Hz. The data was acquired using three disposable Ag/AgCl ECG electrodes, two placed below the left and right collarbones, and one on the left back, below the ribs.

3.2.4. Procedure

Threshold setting.

Participants were seated with their chin on a chin rest placed at 50cm distance from the screen. Each session began with the determination of the threshold of visual awareness individually for each participant, using a masked perceptual discrimination task. Each trial began with a fixation cross (500ms), followed by a mask (300ms), a target cue (either a symmetrical circular shape or an asymmetrical circular shape, starting at 600ms), and another mask (300ms). After each sequence, participants were asked to determine whether the target cue was symmetrical or asymmetrical by pressing corresponding arrows. Next, they were asked to assess their level of confidence in that judgment, also using corresponding arrows (following Scott et al., 2018). They were explicitly instructed to press 'some confidence' if they had confidence in their judgment, even a hunch, and 'total guess' only if they had no idea what the cue was, and were responding randomly. Each time a correct response was made with confidence, the display duration of the target cue was reduced by 50ms on the following trial. When a duration of 100ms was reached, or the first guess response was made, the display duration returned to the previous level (+50ms), and subsequently reduced in 8.35ms steps on the following trials, corresponding

to a single screen refresh duration on a 120Hz monitor. A reduction in exposure duration continued to be made after each non-guess response but not after guess responses. This process continued until they indicated guessing on 5 consecutive trials, regardless of the accuracy of responses. The cue display duration in those trials was recorded as their individual unconscious threshold.

Main task.

The main task was adapted from the unconscious instrumental conditioning task used previously (Pessiglione et al, 2008), in which participants make speeded go or no-go responses to the masked cues. Here, each trial consisted of a fixation cross (here, 1500ms in order to record HR at baseline), forward mask (300ms), target cue (display duration determined in the perceptual discrimination task), backward mask (1000ms), and blank screen jitter (500-700ms), followed by a decision prompt in the form of a question mark, during which participants had 2 seconds to make a response (see *Fig.1.* for task chronology). Pressing the space bar (Go) indicated a decision to take the risk, at which point the participant could win 1 token (golden token displayed on the screen) or lose 1 token (a red cross over the golden token displayed), depending on the type of cue presented between the masks. Not pressing the space bar (NoGo) indicated a safe choice, which always resulted in a null outcome (greyed-out token displayed). Feedback was displayed for 2000ms immediately following the decision (or after the 2s of decision prompt elapsed, in case of NoGo). The task contained 3 blocks of 120 trials, consisting of 60 rewarding and 60 punishing trials in randomised order. For each block, two target cues from the pool of 6 were randomly assigned to be rewarding or punishing, without replacement, in order to ensure each block contained a novel pair. Participants were incentivised to maximise their earnings in the task through a prize draw, with the number of entries contingent on the amount of tokens won.

Following the Go or NoGo response on each trial, participants were asked to determine whether the masked stimulus was symmetrical or asymmetrical by pressing corresponding arrows, followed by a binary assessment of confidence ('some confidence' or 'total guess'), in a manner identical to the threshold setting task. If they responded correctly with confidence three times in a row, cue display time was reduced by 8.35ms, corresponding to single screen refresh duration.

The three blocks were randomly assigned to the auditory cardiac feedback condition (SYNC, ASYNC, CONTROL), ran in a randomised order for each participant.

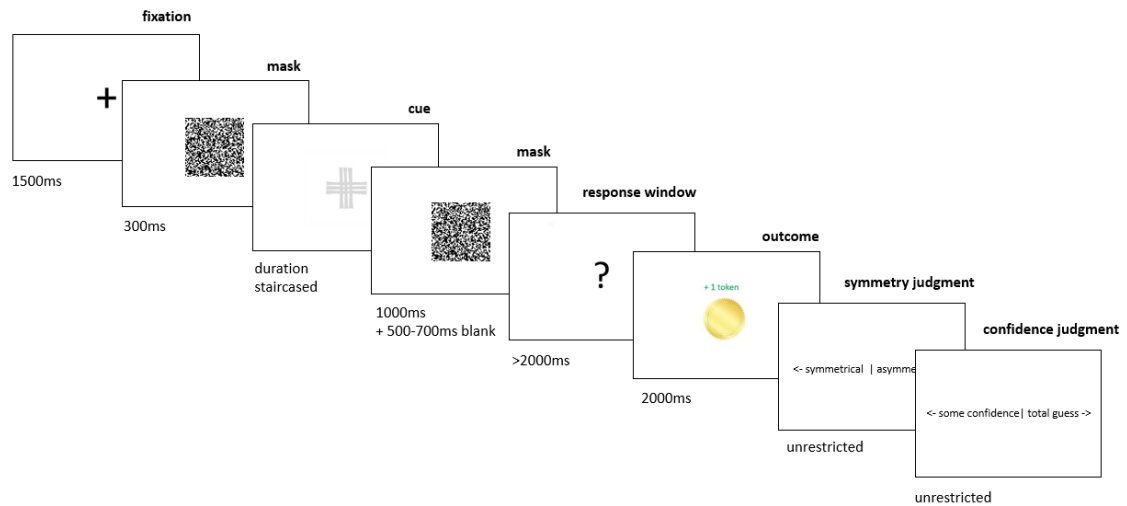


Figure 10: Unconscious instrumental conditioning task (main task). Chronological screenshots depict a single trial sequence, with durations in milliseconds. After cue presentation using forward-backward masking, participants had 2 seconds to make a Go response with the spacebar, or refrain from responding (NoGo). Following the response, feedback was immediately displayed on the screen. In the example shown, a participant responded Go, which was a correct response for the cue presented, thus was rewarded with one gold token.

Interceptive accuracy.

Individual cardiac interoceptive accuracy was measured using the heartbeat discrimination task (Whitehead, Drescher, Heiman, & Blackwell, 1977). In the task, participants were asked to judge whether the auditory cardiac feedback, again delivered through the headphones, was synchronous or asynchronous with their own heartbeat (40 trials of 10 seconds each, 20 with synchronous, and 20 with asynchronous feedback). Asynchronous feedback was produced by manipulating the heart-rate to be either 70% or 130% of the frequency the recorded heart-rate, with the task employing 50% of both types (Suzuki et al., 2013). In both tasks, participants were explicitly instructed to sit still and avoid tracking their pulse physically. Following the interoceptive accuracy task, participants were thanked and debriefed.

3.3. Data pre-processing

3.3.1. Exclusion Criteria

In order to ascertain that analyses are conducted only on trials that were truly unconscious, all individual trials where participants made a correct symmetry judgment with confidence were marked as aware and excluded (16.9% of all trials). The number of aware trials was not found to differ between conditions in a repeated measures ANOVA (SYNC: $M = 65.7$, $SD = 22.9$; ASYNC: $M = 70.7$, $SD = 25.6$; CONTROL: $M = 66.6$, $SD = 26.1$; $F(2,76) = 1.31$, $p = 0.277$). Nine participants

who were aware on more than 25% (90) of all trials were excluded from further analysis. Three participants who made only Go or only NoGo responses in any one or more blocks (e.g. due to a failure to understand the task), were also excluded, yielding a final sample of 27. Outlier analysis using boxplot inspection revealed no other outliers.

For the reaction time (RT) analysis, RTs under 100ms or greater than 2 standard deviations from individual means were excluded (0.9%). Note that NoGo trials yielded no RTs, thus were not included in the analysis.

3.3.2. Interoceptive Accuracy

Interoceptive accuracy (IAcc) scores from the heartbeat discrimination task were computed for each participant, by averaging the number of correct responses. They will not be analysed in the context of this paper.

3.3.3. Electrocardiography

Initial pre-processing of the ECG data was done in Biopac Student Lab 3.7.7. The ECG data was filtered offline with a high-pass filter (1Hz), and R-peaks of the QRS complex were detected for the length of the task. Heart rate (HR) in beats per minute (BPM) was subsequently computed from the R-R intervals. The data, complete with task event markers, was then exported into Matlab, where average BPM was computed for each subject at each event of interest (B0 at baseline, cue presentation, and feedback presentation), as well as one beat before (B-1), and one (B1) and two beats (B2) following each event. Events where the average BPM deviated from the subject's mean by more than 3 *SDs* were excluded from analysis (as reduced or inflated values could reflect electrode displacement, excessive motion, or other artefacts). A total of 3148 events (1%) were removed (note that this is from a total of 4 beats computed around 9 events in the task). Two participants were excluded due to large unusable segments of data. Data was complete for all participants included in the behavioural analysis.

3.4. Results

3.4.1. Bayes Factors

For hypothesis testing, Bayes Factors (Bs) will be reported alongside *p*-values for all comparisons made. Bs can help to disambiguate non-significant results as either indicating support for the null hypothesis (H_0 , positing no effect) or indicating insensitive data (i.e. the data are not in favour of either H_0 or H_1 , which uses an estimated raw effect size as the standard deviation of

its distribution; Dienes, 2014). By convention, B_s smaller than 1/3 indicate evidence for H_0 . B_s larger than 3 indicate evidence for H_1 . B_s between those values indicate insensitive data.

3.4.2. Evidence of learning: Performance

Type I d' (a Signal Detection Theoretic measure of sensitivity to signal versus noise; Stanislaw & Todorov, 1999) was computed for each condition, treating Go responses to rewarding cues as Hits, and Go responses to punishing cues as False Alarms. The resulting measure of sensitivity can be taken as evidence of successful learning (i.e. discrimination between the cues) if it is significantly above 0. In a one-sample t-test against 0 (no sensitivity, thus no ability to discriminate between the cues), the total d' was not significantly different from 0 ($M = -0.04$, $SE = 0.04$, $t(26) = -0.90$, $p = 0.377$, $B_{H(0,0.7)} = 0.03$), suggesting that, on average and irrespective of condition, participants were not able to learn the cue-outcome association. For computing B , the predictions of H_1 (learning is present) were modelled as a half-normal distribution centred on 0, with an SD equal to 0.7 (the expected effect size, previously found if unconscious learning is present; Pessiglione et al., 2008).

Looking at each condition separately, no evidence of learning was found in either SYNC ($M = -0.21$, $SE = 0.09$, $t(26) = -2.29$, $p = 0.031$, $B_{H(0,0.7)} = 0.04$), ASYNC ($M = 0.05$, $SE = 0.07$, $t(26) = 0.69$, $p = 0.498$, $B_{H(0,0.7)} = 0.19$) or CONTROL ($M = 0.07$, $SE = 0.08$, $t(26) = 0.97$, $p = 0.339$, $B_{H(0,0.7)} = 0.28$). A repeated-measures ANOVA, with learning (d') as the dependent variable and cardiac feedback condition (SYNC, ASYNC, CONTROL) as the within-subject factor, revealed no main effect of condition ($F(2,58) = 1.88$, $p = 0.16$, $\eta_p^2 = 0.16$). See Figure 11 for a graphical representation.

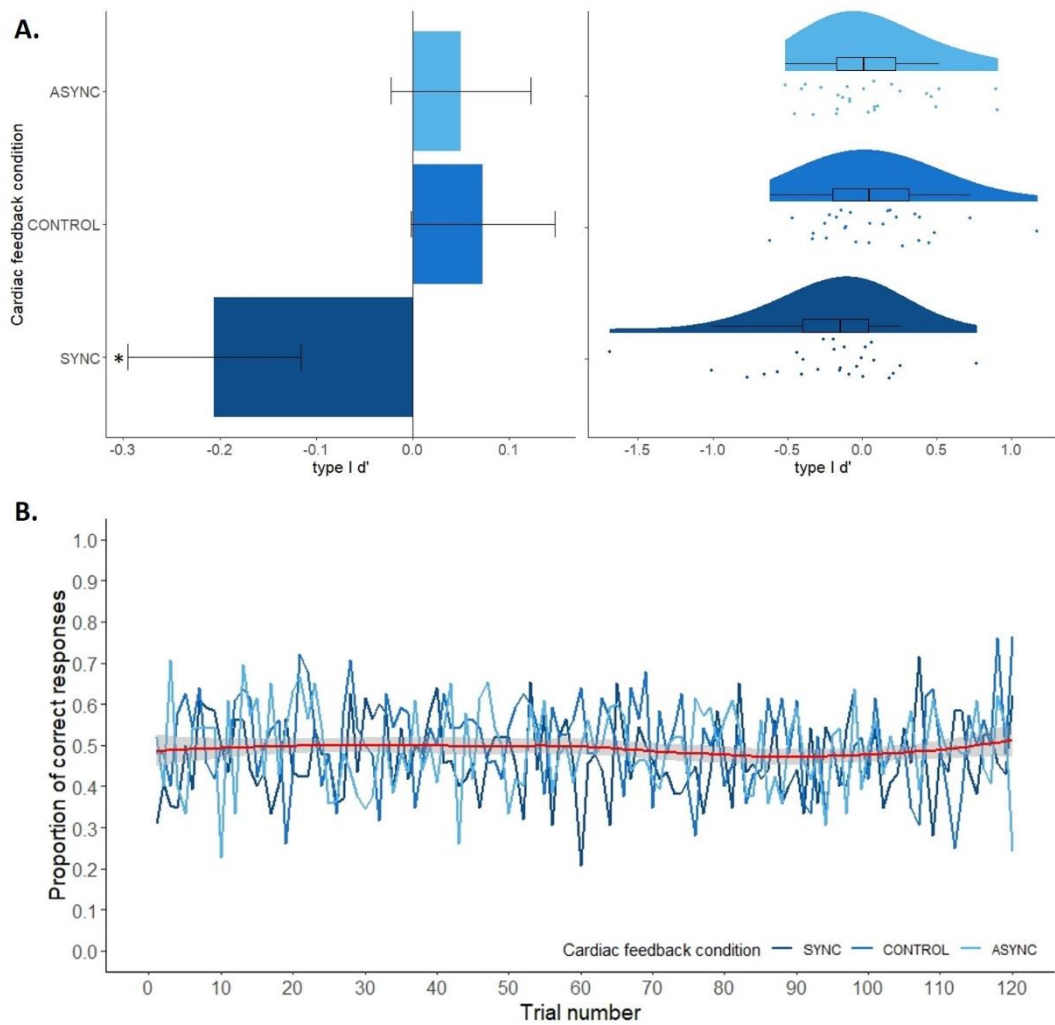


Figure 11: A: Type I d' values across the three cardiac feedback conditions. Left panel: mean d' by condition (± 1 SEM). Star indicates significant difference from 0 at $p < .05$. Right panel: distributions of data with boxplots. **B:** Average proportion of correct responses at each trial, in each cardiac feedback condition, with superimposed smoothed regression line in red, visualising erratic rates of correct responding averaging around the chance level (0.5).

3.4.3. Evidence of learning: Reaction times

For the RT analysis, an $RT_{\text{difference}}$ index was computed by subtracting RTs to rewarding cues from RTs to punishing cues. As such, positive values indicate that participants took a longer time to respond to punishing cues than to rewarding cues, in line with RT-oriented indicators of learning (e.g. Atas, Faivre, Timmermans, Cleeremans, & Kouider, 2014). Zero indicates that there was no difference between the two. Overall, irrespective of condition, $RT_{\text{difference}}$ was not significantly different from 0 in a one-sample t-test ($M = -1.26\text{ms}$, $SE = 11.71$, $t(26) = -0.11$, $p = 0.915$, $B_H(0,34) = 0.29$), suggesting that participants were not responding slower to punishing cues than to rewarding cues, against the RT-oriented indicators of learning. B was computed with the

predictions of H1 (learning is present) modelled as a half-normal distribution centred on 0, with an SD equal to 34ms (the expected effect size obtained from a past study which found an RT difference in the absence of performance effects in a similar task; Atas et al., 2014).

The index was then submitted as a dependent variable into a repeated measures ANOVA with condition (SYNC, ASYNC, CONTROL) as the within-subject factor. The ANOVA revealed no main effect of condition ($F(2,52) = 0.12$, $\eta_p^2 = 0.01$, $p = 0.892$). In one-sample t-tests for each condition separately, the $RT_{\text{difference}}$ index was not significantly different from 0 in either SYNC ($M = 9.61\text{ms}$, $SE = 21.40$, $t(26) = 0.45$, $p = 0.657$, $B_{H(0,34)} = 0.74$), ASYNC ($M = 7.83\text{ms}$, $SE = 19.24$, $t(26) = 0.41$, $p = 0.687$, $B_{H(0,34)} = 0.67$) or CONTROL ($M = -1.90\text{ms}$, $SE = 19.80$, $t(26) = -0.10$, $p = 0.924$, $B_{H(0,34)} = 0.46$). Note that all Bs indicate that the data are insensitive (i.e. a strong conclusion cannot be drawn). See Figure 12 for a graphical representation.

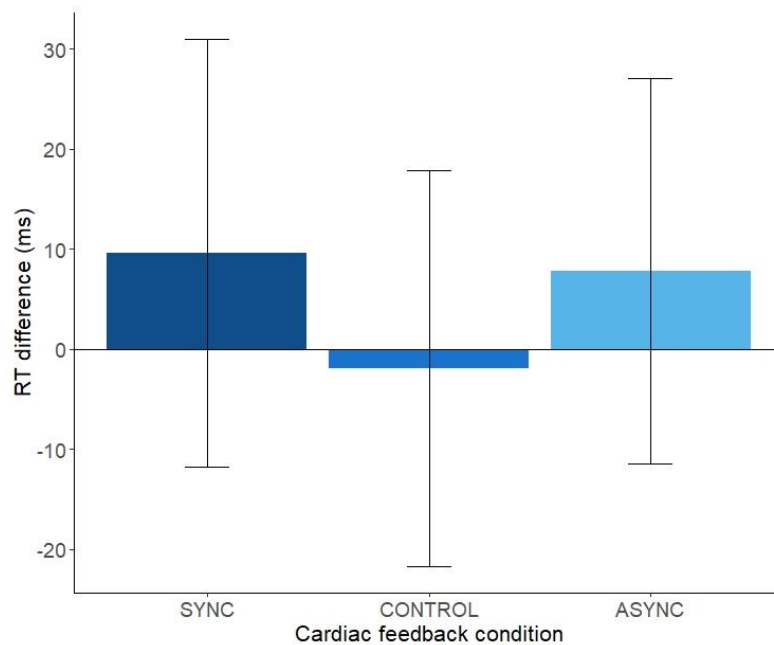


Figure 12: Mean $RT_{\text{difference}}$ index (± 1 SEM) across the three cardiac feedback conditions. Positive values indicate that participants took a longer time to respond to punishing cues than to rewarding cues, in line with RT-oriented indicators of learning.

3.4.4. Cardiac responses to performance feedback

The analysis of cardiac behaviour proceeded in two stages. First, following past research, we assessed whether the heart decelerates more following punishing performance feedback (i.e. in case of a Go response to a punishing cue) than following rewarding performance feedback (Go to a positive cue) in the unconscious instrumental conditioning task. Secondly, following the observations that over the course of learning the deceleration response starts to appear following the predictive cue itself, we assess whether the heart decelerates more following the

punishing cue than following the rewarding cue. As a side analysis, we also checked whether the cardiac response differed between the cardiac feedback conditions (see section 3.6 Supplementary Material).

Prior to analysis, we checked whether there were any fluctuations in HR across conditions. HR in BPM at baseline (1500ms fixation period), computed by averaging BPM0 at the time of fixation cross presentation in each trial), did not differ between the cardiac feedback conditions (SYNC: $M = 77.6$, $SD = 11.38$; ASYNC: $M = 77.41$, $SD = 11.02$; CONTROL: $M = 77.41$, $SD = 10.95$; $F(2,52) = 0.08$, $p = 0.922$, $\eta_p^2 = 0.003$). Due to the short interval between cue and performance feedback presentation after action execution (1500-1700ms), we also checked whether HR returned to baseline following any potential cue-evoked changes in HR. In a paired t-test, HR in BPM at baseline (computed as above) did not differ significantly from HR as computed at 1 beat before feedback presentation (B0 at fixation: $M = 75.58$, $SE = 1.89$; B-1 at feedback: $M = 75.50$, $SE = 1.91$; $M_{difference} = 0.71$, $SEM = 0.24$, $t(35) = 0.30$, $p = 0.766$).

3.4.4.1. Cardiac response to rewarding vs punishing performance feedback

In order to assess whether the heart decelerates more following punishing than rewarding feedback, HR (in BPM) across all conditions was submitted as a response variable into a linear mixed-effects model, fit using the lme4 package (Bates et al., 2015) in R (R Core Team, 2018). Linear mixed-effects models have an advantage over regular repeated-measures ANOVA in that they are more robust to imbalances in the data (e.g. randomly missing values), and allow to incorporate each participant's individual baseline (random intercepts) and responsiveness to the manipulation (random slopes). The model included feedback valence (rewarding/punishing), beat (B-1, B0, B1, B2) and their interaction as fixed effects (predictor variables). The random effects structure included subject-specific random intercepts and random slopes for feedback valence. Note that this random effects formulation was used following the parsimonious approach given a singular fit under maximal specification (i.e. subject-specific random intercepts and random slopes for the interaction of feedback valence and beat) (Matuschek et al., 2017). Treatment (dummy) coding was applied. The model was fit using maximum likelihood estimation. Note that only Go trials were used, as NoGo trials yielded no performance feedback and were not informative for task performance (they were construed as a safe choice or a pass).

See Table 4 for regression coefficients. Analysis of deviance on the model, conducted using the car package (Companion to Applied Regression; Fox & Weisberg, 2019) revealed a significant main effect of beat ($\chi^2 = 132.26$, $df = 3$, $p < 0.001$). There was also a main effect of feedback valence ($\chi^2 = 6.89$, $df = 1$, $p = 0.009$), but no interaction between feedback valence and beat ($\chi^2 = 4.38$, $df = 3$, $p = 0.223$). As such, contrary to our predictions, we found no support for

the idea that the heart responds with more deceleration to punishing than to rewarding feedback. Instead, cardiac deceleration was evident for both kinds of feedback (although rewarding feedback does appear to elicit a marginally faster acceleratory recovery; see *Fig. 13*). Regardless of feedback valence, cardiac deceleration initiated before feedback presentation and continued until one beat after feedback presentation, before recovering, in line with past evidence.

To test whether the effect differs between the first and second half of the task, the full block length was split into 2 bins of 60 trials, and bin number was added to the model as a fixed effect. As before, there was a main effect of beat ($\chi^2 = 132.64$, $df = 3$, $p < 0.001$). There was also a main effect of bin ($\chi^2 = 62.25$, $df = 1$, $p < 0.001$). All other effects and interactions remained non-significant, as in the previous model.

See Supplementary Material for the differences in HR response between the three cardiac feedback conditions.

Table 4. Regression coefficients for the fixed effects from the linear mixed model (performance feedback events). The intercept refers to B-1 for rewarding feedback. N = 27, number of observations = 22938.

	<i>Estimate (BPM)</i>	<i>Std. Error</i>	<i>df</i>	<i>t-value</i>	<i>p</i>
Intercept (B-1:REW)	77.62	2.09	27.11	37.20	< 0.001 ***
PUN	-0.17	0.19	91.48	-0.93	0.355
B0	-0.69	0.15	22884.41	-4.68	< 0.001 ***
B1	-0.56	0.15	22884.48	-3.82	< 0.001 ***
B2	0.41	0.15	22884.44	2.80	0.005 **
PUN:B0	-0.07	0.21	22884.40	-0.36	0.720
PUN:B1	-0.30	0.21	22884.42	-1.45	0.146
PUN:B2	-0.37	0.21	22884.40	-1.78	0.074




Figure 13. Change in cardiac activity (± 1 SEM) in response to performance feedback (rewarding, punishing). B-1 refers to BPM as measured 1 beat before feedback presentation. Stars indicate significant difference between the beat means, averaged across both feedback types due to no evidence for a main effect of feedback valence (**: <0.01 , ***: <0.001).

3.4.4.2. Cardiac response to rewarding and punishing cues

Subsequently, we repeated the analysis focusing on the cue as the event of interest, as opposed to performance feedback, following Kastner et al. (2017). The original model was rerun again, with cue valence (rewarding/punishing), beat (B-1, B0, B1, B2) and their interaction as fixed effects (predictor variables). All remaining model parameters remained the same.

See Table 5 for regression coefficients. Analysis of deviance on the model revealed a significant main effect of beat ($\chi^2 = 26.12$, $df = 3$, $p < 0.001$; see Figure 14). There was no main effect of cue valence ($\chi^2 = 0.32$, $df = 1$, $p = 0.573$) and no interaction between cue valence and beat ($\chi^2 = 0.21$, $df = 3$, $p = 0.976$). As such, we found that the heart did respond with a deceleration following cue presentation, but no evidence that this response is differentiated to rewarding vs punishing cues throughout the duration of the task.

However, the deceleration response to cues found by Kastner et al. (2017) was more pronounced later on in the learning process. As such, the full block length was split into 2 bins of 60 trials, and bin number was added to the model as a fixed effect. The expanded model revealed that the main interaction of interest (cue valence*beat*bin number) was not

significant ($\chi^2 = 0.19$, $df = 3$, $p = 0.980$), neither was the beat*bin number interaction ($\chi^2 = 0.36$, $df = 3$, $p = 0.948$) nor the cue valence*beat interaction ($\chi^2 = 0.21$, $df = 3$, $p = 0.975$). The cue valence*bin number interaction was significant ($\chi^2 = 10.82$, $df = 1$, $p = 0.001$). As earlier, there was also a main effect of beat itself ($\chi^2 = 26.18$ $df = 3$, $p < 0.001$) and of bin number ($\chi^2 = 82.34$, $df = 1$, $p < 0.001$). There was no main effect of cue valence ($\chi^2 = 0.31$ $df = 1$, $p < 0.578$).

Table 5. Regression coefficients for the fixed effects from the linear mixed model (cue events). The intercept refers to B-1 for rewarding cues across both bins. N = 27, number of observations = 38497.

	<i>Estimate</i>	<i>Std. Error</i>	<i>df</i>	<i>t-value</i>	<i>p</i>
Intercept (B-1:REW)	78.23	2.05	27.06	38.08	< 0.001 ***
PUN	-0.08	0.16	69.75	-0.47	0.643
B0	0.02	0.12	38443.07	-0.18	0.855
B1	-0.27	0.12	38443.13	-2.35	0.019 *
B2	-0.33	0.12	38443.09	-2.88	0.004 **
PUN:B0	-0.04	0.16	38443.06	-0.22	0.824
PUN:B1	0.04	0.16	38443.08	0.22	0.823
PUN:B2	0.01	0.16	38443.10	0.09	0.930




Figure 14: Change in cardiac activity (± 1 SEM) in response to cue valence across both bins in all conditions. B-1 refers to BPM as measured 1 beat before cue presentation. Stars indicate significant difference between the beat means, averaged across both cue types due to no evidence for a main effect of cue valence (*: < 0.05, **: < 0.01, ***: < 0.001).

3.5. Discussion

Successful learning and decision-making rely on efficient monitoring of performance feedback. Past research showed that performance feedback is both external (e.g. rewarding or punishing outcomes), as well as internal, manifested by feedback-related autonomic markers. Out of those, error-related cardiac deceleration has been studied in the context of learning, with evidence showing that the heart responds to negative (error-related) feedback with a more pronounced and longer deceleration than to positive feedback (Crone et al., 2003, 2004; Groen et al., 2007; Kastner et al., 2017). However, there is limited evidence regarding whether cardiac behaviour drives learning, or is simply a correlate of the learning process. Here, we employed an unconscious instrumental learning task with an added manipulation of cardiac precision through aurally delivered cardiac feedback, targeted at amplifying and disrupting the reliability of the performance-related cardiac information, as well as a continuous ECG. We sought to assess 1) whether there is evidence for unconscious instrumental conditioning, and whether it is affected by the precision of cardiac signal, manipulated with auditory cardiac feedback; 2) whether there is evidence for performance feedback-related cardiac deceleration in the absence of cue awareness; and 3) whether there is evidence of feedback-related deceleration diminishing, and cue-related deceleration increasing, in line with previously reported markers of learning. We address those questions in turn.

The behavioural results demonstrate that, on average, participants were unable to learn the association between the cues predictive of reward or punishment and their corresponding outcomes, without having conscious awareness of the cues. Trial-by-trial measure of cue awareness ensured that only truly unconscious trials were analysed, rectifying a flaw in the previous experiments (Mastropasqua & Turatto, 2015; Pessiglione et al., 2008). Evidence of learning was absent in the control condition, as well as in both the synchronous and asynchronous cardiac feedback conditions. Interestingly, synchronous cardiac feedback resulted in worse performance than the control condition. This runs contrary to our predictions, which stated that amplifying the precision of the cardiac channel by playing synchronous cardiac feedback should effectively increase the reliability of cardiac information, which should improve learning. However, due to the lack of evidence for learning at baseline, it is impossible to draw strong conclusions regarding the impact (or lack thereof) of the cardiac feedback manipulation.

This result speaks to recent proposals that instrumental learning might in fact require some degree of awareness (e.g. Reber et al., 2018). While there is evidence that some forms of simple associative learning can be achieved without stimulus awareness (Dupoux et al., 2008;

Scott et al., 2018), it is unclear whether the same is possible for more complex associations or learning across larger spatiotemporal intervals (Clark & Squire, 1998; Faivre, Mudrik, Schwartz, & Koch, 2014; Hasselman et al., 2015). In the present scenario, in order to learn effectively when to deploy an action and when to refrain from it, participants must integrate the information about the nature of the stimulus and the outcome it presents when action is executed, in a fairly long temporal interval. Yet, unconscious processing of any kind is argued to be short-lived and distributed only locally, preventing long-range associations (Baars et al., 2003; Dehaene et al., 2001; Melloni et al., 2007). As such, it might be the case that an unconscious cue is not processed extensively enough, or held in the working memory for long enough, to be integrated with subsequent information across the length of the trial. What follows is that a shorter temporal interval (e.g. delay conditioning; Clark & Squire, 1998) might enable instrumental learning (but see Chapter 5 of this thesis). Another possibility is that some degree of consciousness is inherently necessary for goal-oriented, selective action (Dehaene & Naccache, 2001; Kouider & Dupoux, 2001), but this question is beyond the scope of the present discussion.

However, it is noteworthy that the current paradigm might have disrupted the already fragile process of temporal integration. Large gaps between the events in the task had to be introduced in order to measure cardiac activity – critically, the backward mask was 1 second long, followed by another 500-700ms blank, separating the predictive cue from its outcome (in between which the action is executed, immediately followed by feedback). Those gaps were essential to allow the recording of cardiac activity following the event of interest uninterrupted by other processes. Yet, as mentioned before, the short-lived or short-range nature of unconscious processing might prevent learning across such large temporal scales. Consequently, extending the gap between the stimulus and outcome would inherently reduce the chance of successful learning. While this could be a serious limitation, the gap between stimulus and outcome is in fact comparable to the Pessiglione et al. (2008) design, where the response window was open for 2 seconds irrespective of whether the response was made or not. Regardless, past research already pointed to the infeasibility of unconscious learning under shorter inter-stimulus intervals (Reber, Samimizad, & Mormann, 2018; Chapters 2, 4 and 5 of this thesis).

The physiological results suggest that the heart failed to differentiate between rewarding and punishing performance feedback, contrary to past evidence showing that the heart decelerates more in response to punishing feedback (indicative of an error) than to rewarding one (Crone et al., 2005b, 2003; van der Veen et al., 2004a). Here, cardiac deceleration was evident for both types of feedback (albeit visual inspection suggests that recovery is slightly faster following rewarding one), starting a beat before feedback is presented and continuing

until one beat after, in line with past research (Groen et al., 2007). Failure to observe a differentiated response to rewarding and punishing feedback suggests that it was not informative – it did not reflect a mismatch between the expected outcome, predicted by the cue, and the actual outcome. This, in turn, suggests that the unconscious cue was not processed to the extent allowing for integration with feedback. That is, cue identity was not predictive of the expected outcome, so any feedback signifying actual outcome was rendered meaningless and uninformative, as the performance monitoring mechanism was unable relate the two.

One of the potential reasons behind this could be the same as the reason for the absence of learning – if the unconsciously presented cue is processed only locally, information about its identity cannot be integrated with subsequent processes, including those determining an instrumental response and those governing the performance monitoring mechanisms. If so, it would be unsurprising that the lack of ability for unconscious instrumental learning is paralleled by the lack of corresponding markers of learning. However, it is interesting that the previously reported unconscious discrimination task (Łukowska et al., 2018) found the expected error-related cardiac activity in the absence of cue-awareness (note that in their task no performance feedback was provided – HR deceleration was found following an error in type I decision). This might suggest that the nature of the cue was processed sufficiently to be fed into the internal performance monitoring system. If so, the crucial difference might indeed be timing, and our introduction of a lengthy (around 2 seconds) gap between the stimulus and feedback might have prevented their integration. If so, future research could rectify this flaw by requesting a response immediately following cue presentation (e.g. during the backward mask). However, this would limit the ability to record an uninterrupted ECG signal following cue presentation alone.

Regarding the cardiac deceleration following the predictive stimuli themselves, we failed to observe any differentiation in response to rewarding or punishing cues. However, the response found previously was learning-contingent – as learning progressed, subjects came to exhibit the deceleration response earlier, suggesting the predictive nature of the cue was learned, resulting in an increased reliance on the internal, predictive mechanism (Groen et al., 2007; Kastner et al., 2017). Here, we observed an absence of learning of the cue-outcome association, which means that the participants failed to learn which cue is predictive of which outcome. It is not surprising that lack of this knowledge was also reflected in the cardiac data, and the heart did not come to differentiate the cues by their rewarding or punishing nature.

It is also important to recognise that the current cardiac feedback manipulation technique might not be robust enough for the purpose of manipulating interoceptive precision. Presenting an auditory stream corresponding to participants' own heartbeats is undoubtedly an

exteroceptive signal, and there is still limited evidence regarding its processing in the brain. Recent EEG evidence demonstrates that the brain does process external synchronous cardiac feedback in a manner similar to other self-generated sounds, suggesting that it can be differentiated from cardiac-asynchronous sounds (van Elk et al., 2014). Elsewhere, unexpected omissions from a synchronous, but not asynchronous, auditory stream were found to elicit a surprise response (Pfeiffer & De Lucia, 2017). Still, the proposal that such manipulation is in fact affecting precision of the cardiac channel, as formalised in the interoceptive inference framework (Petzschnner, 2017; Petzschnner et al., 2019), remains theoretical. As such, different manipulations of cardiac precision should be explored and tested to fully understand its role in other processes. One such technique could be transcutaneous Vagus Nerve stimulation (tVNS), shown to directly amplify the afferent cardiac signal, and recently found to improve interoceptive accuracy (Villani, Tsakiris, & Azevedo, 2019).

As the cardiac precision manipulation is targeted at amplifying or disrupting the precision of the cardiac channel as represented in the brain, we assumed that it would affect the learning process, but it should not noticeably affect cardiac behaviour itself. Interestingly, previous research did report that error-related deceleration was abolished when participants listened to an auditory stream that was asynchronous to their own heartbeat (Łukowska et al., 2018). The authors proposed that participants' heart rate became aligned with what they perceived to be their own heartbeat, but was in fact a pre-recorded resting heart rate of another person, thus disrupting any error-related markers (though it is noteworthy that they did not conduct analyses to confirm this). This could indeed occur through the mechanism of prediction-error minimisation, as formalised in the active inference approach to interoceptive inference (Friston, 2010; Friston & Kiebel, 2009; Pezzulo, Rigoli, & Friston, 2015). Under active inference, the brain strives to continuously minimise prediction errors (mismatches between afferent sensory, or bodily, signals and the prior prediction) by adjusting internal parameters (e.g. homeostatic or other reflexes, including cardiac behaviour). A deviation in their actual state from prediction would provoke downstream changes to those parameters. As such, it is plausible that when the cardiac channel has low precision, prior predictions could provoke changes to the actual heart rate in line with the prediction. In the case of cardiac feedback, this could mean aligning the heartbeat with the perceived one, or engaging a prediction of what the heart "should" be doing. However, the precise, biologically plausible mechanism behind active inference in the interoceptive domain remains a heavily studied and debated topic (Adams, Shipp, et al., 2013; Barrett, Quigley, & Hamilton, 2016; Gentsch, Sel, Marshall, & Schütz-Bosbach, 2019; Owens et al., 2018). Hence, we restrained ourselves from making strong predictions about potential effects of our cardiac precision manipulation on actual cardiac behaviour, especially

given that its application in an unconscious learning paradigm is itself a novel and uncharted territory. Interested readers are invited to refer to the Supplementary Material of this chapter, where we present an exploratory analysis of performance feedback-related cardiac deceleration between the three cardiac feedback conditions.

Finally, previous research noted that HR measurements can be confounded by respiration rate, as respiration can affect the length of interbeat intervals (Berntson, Cacioppo, & Quigley, 1993). Whether to correct for respiration in HR measurements is an ongoing debate (Laborde, Mosley, & Thayer, 2017; Quintana, Heathers, Kemp, Tarvainen, & Billman, 2014). In the present experiment, we did not record respiration rates, so any potential influences of breath on cardiac deceleration, whether sharing a common basis or a confounding variable (Thayer, Loerbroks, & Sternberg, 2011) cannot be assessed.

To summarise, the present study investigated unconscious instrumental conditioning, its cardiac concomitants, and the influence of cardiac precision on learning. We found that participants were unable to learn the associations between predictive cues and their outcomes when the cue was presented unconsciously. We also found no evidence that cardiac responses could differentiate between rewarding and punishing performance feedback without cue awareness. Finally, we found no evidence that cardiac precision affected learning. Overall, we conclude that unconscious conditioning might be harder to achieve when properly controlled. Lack of learning at baseline prevents us from drawing stronger conclusions about the potential influence of cardiac precision on the learning process. Due to numerous timing constraints, which could heavily influence whether unconscious learning happens or not, it might not be possible to study the question of its cardiac concomitants in a sufficiently robust manner in the first place.

3.6. Supplementary material

The effect of cardiac precision manipulation through auditory feedback on cardiac response to punishing vs rewarding performance feedback.

In order to assess whether the cardiac response to performance feedback is affected by cardiac precision (manipulated through auditory cardiac feedback in three conditions; SYNC, ASYNC, CONTROL), the conditions were added as a fixed effect to the linear mixed effects model used in the main analysis. The expanded model failed to converge under the original random effects specification (subject-specific intercepts and slopes for feedback valence), so the random slopes were suppressed entirely, leaving subject-specific random intercepts.

See *Table S3.1* for a full table of regression coefficients. Analysis of deviance on the model revealed that the interaction of primary interest (condition*beat*feedback valence) was not significant ($\chi^2 = 1.25$, $df = 6$, $p = 0.974$), suggesting that cardiac precision manipulation might not affect the cardiac response to rewarding/punishing feedback itself. The model revealed a significant main effect of beat ($\chi^2 = 132.09$, $df = 3$, $p < 0.001$), confirming that the heart responded with deceleration to feedback presentation regardless of condition and feedback valence, main effect of feedback valence ($\chi^2 = 24.54$, $df = 1$, $p < 0.001$), and condition ($\chi^2 = 48.00$, $df = 2$, $p < 0.001$), as well as a significant interaction between condition and feedback valence ($\chi^2 = 7.28$, $df = 2$, $p = 0.026$). This pattern is intriguing, suggesting that it was the ASYNC condition that produced the largest differentiation between rewarding and punishing feedback across all beats (see *Figure S3.1.*) The remaining two-way interactions (condition*beat and feedback valence*beat) were not significant.

Table S3.1: Regression coefficients for the fixed effects from the linear mixed model (all conditions, performance feedback events). The intercept refers to B-1 for rewarding feedback in the control condition. N = 27, number of observations = 22938.

	<i>Estimate</i>	<i>Std. Error</i>	<i>df</i>	<i>t-value</i>	<i>p</i>
Intercept					
(CONTROL:REW:B-1)	77.43	2.07	27.38	37.41	< 0.001 ***
SYNC	0.47	0.25	22911.03	1.86	0.062
ASYNC	0.07	0.25	22911.04	0.27	0.790
PUN	-0.22	0.25	22911.05	-0.86	0.390
B0	-0.65	0.25	22911.00	-2.63	0.009 **
B1	-0.47	0.25	22911.00	-1.90	0.058
B2	0.55	0.25	22911.00	2.21	0.027 *
SYNC:PUN	0.14	0.36	22911.04	0.39	0.695
ASYNC:PUN	-0.04	0.36	22911.20	-0.10	0.922
SYNC:B0	-0.02	0.36	22911.00	-0.06	0.955
ASYNC:B0	-0.09	0.36	22911.00	-0.26	0.798
SYNC:B1	-0.09	0.36	22911.00	-0.26	0.796
ASYNC:B1	-0.19	0.36	22911.00	-0.52	0.602
SYNC:B2	-0.36	0.36	22911.00	-1.02	0.309
ASYNC:B2	-0.05	0.36	22911.00	-0.15	0.882
PUN:B0	-0.01	0.36	22911.00	-0.03	0.973
PUN:B1	-0.21	0.36	22911.00	-0.60	0.551
PUN:B2	-0.24	0.36	22911.00	-0.66	0.507
SYNC:PUN:B0	-0.09	0.51	22911.00	0.18	0.861
ASYNC:PUN:B0	-0.28	0.51	22911.00	-0.56	0.579
SYNC:PUN:B1	-0.05	0.51	22911.00	0.11	0.917
ASYNC:PUN:B1	-0.33	0.51	22911.00	-0.63	0.525
SYNC:PUN:B2	-0.05	0.51	22911.00	0.10	0.928
ASYNC:PUN:B2	-0.45	0.51	22911.00	-0.87	0.383

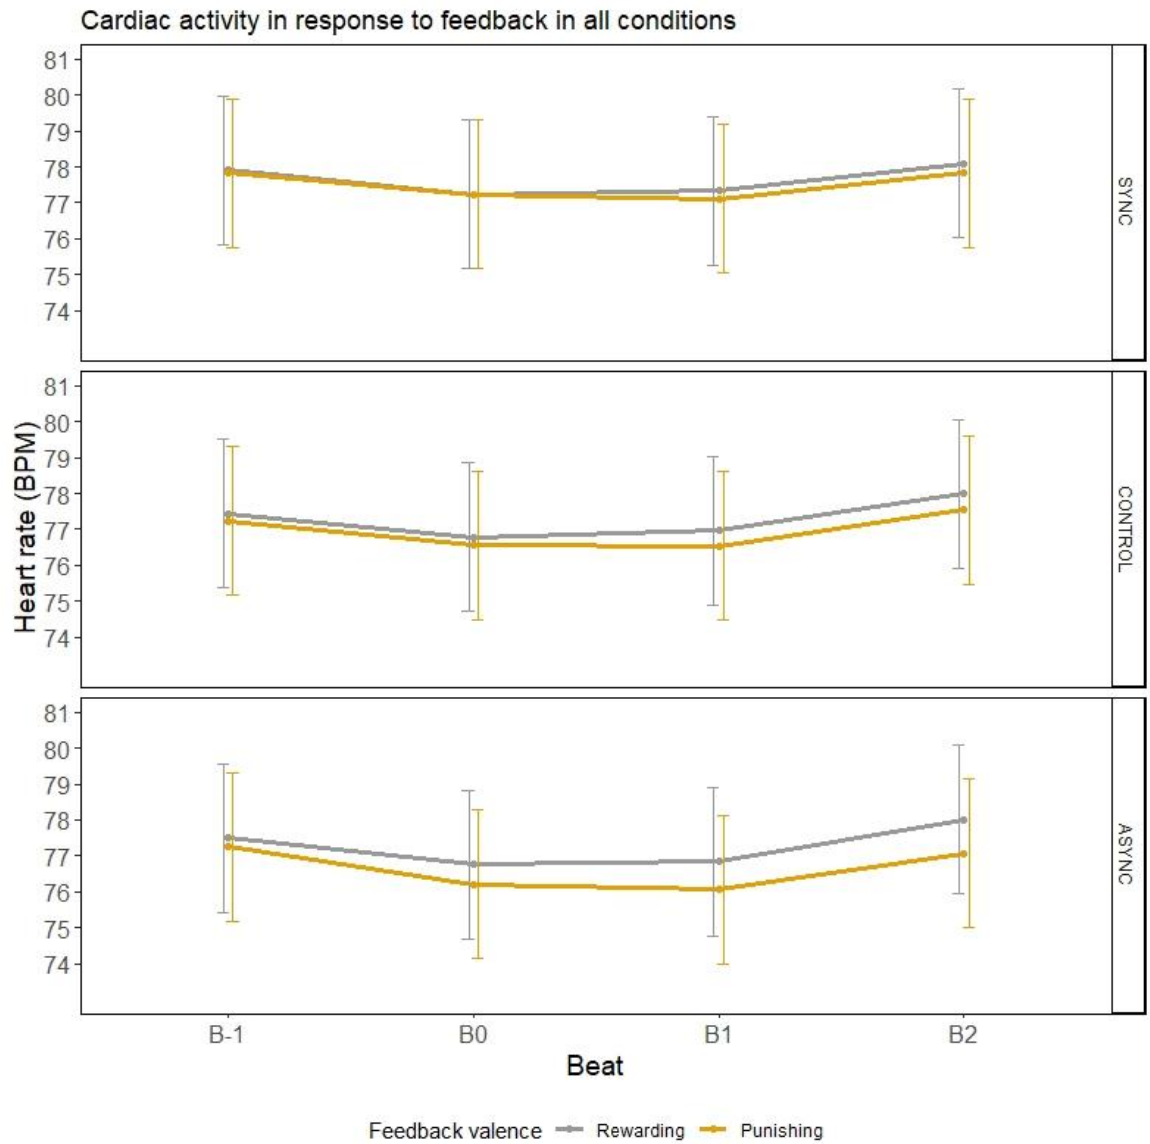


Figure S3.1: Change in cardiac activity (± 1 SEM) in response to performance feedback in all cardiac feedback conditions. B-1 refers to BPM as measured 1 beat before feedback presentation.

4.

FEASIBILITY OF UNCONSCIOUS INSTRUMENTAL CONDITIONING: A REGISTERED REPLICATION REPORT

Abstract

The extent to which high-level, complex functions can proceed unconsciously has been a topic of considerable debate. While unconscious processing has been demonstrated for a range of low-level processes, from feature integration to simple forms of conditioning and learning, theoretical contributions suggest that increasing complexity requires conscious access. Here, we focus our attention on instrumental conditioning, which has been previously shown to proceed without stimulus awareness. Yet, instrumental conditioning also involves integrating information over a large temporal scale and distinct modalities in order to deploy selective action, constituting a process of substantial complexity. With this in mind, we revisit the question of feasibility of instrumental conditioning in the unconscious domain. Firstly, we address the theoretical and practical considerations relevant to unconscious learning in general. Secondly, we aim to replicate the first study to show instrumental conditioning in the absence of stimulus awareness (Pessiglione et al., 2008), following the original design and supplementing the original crucial analyses with a Bayesian approach (Experiment 1). Should the replication be successful, we will attempt to replicate the effect yet again with improved methods to address the methodological issues related to sensitivity and immediacy (Experiment 2).

4.1. Introduction

Ever since the earliest demonstration of subliminal perception (Peirce & Jastrow, 1886), the extent to which information can be processed in the brain without conscious awareness has been a widely studied question. Unconscious processing has been demonstrated for many low-level processes such as feature detection and integration (e.g. integrating colour, shape and texture of an object into one coherent percept; Blake & Fox, 1974), as well as simple forms of learning, for instance emotional (Olsson & Phelps, 2004), visuospatial (Rosenthal et al., 2010) or associative (Scott et al., 2018). However, the extent to which unconscious processing is possible for higher-level, more complex functions remains a topic of debate (Axelrod & Rees, 2014; Mudrik et al., 2014; Sterzer et al., 2014). One such example is learning the contingencies between stimuli and outcomes, especially in instrumental scenarios, where the agent must learn from multiple temporally separated events: the stimulus itself, their action, and its consequence. This kind of learning has apparently been shown to be feasible in the absence of stimulus awareness (Mastropasqua & Turatto, 2015; Pessiglione et al., 2008). However, following recent evidence to the contrary (Reber et al., 2018) and discussions about the minimal conditions for unconscious processing (e.g. Mudrik et al., 2014), as well as developments in methods used to assess conscious awareness (Dienes, 2015; Rothkirch & Hesselmann, 2017; Shanks, 2017), we revisit the finding that instrumental learning can occur unconsciously. Here, we attempt to replicate the original result of Pessiglione and colleagues (2008), leveraging the developments in the field of unconscious processing to apply a more robust statistical approach (Experiment 1), and a more rigorous methodology (Experiment 2).

While there is no clear agreed theory demarcating what conscious versus unconscious mental states may represent (Breitmeyer, 2015; Dupoux et al., 2008; Kouider & Dehaene, 2007), there have been attempts at outlining conditions under which unconscious processing can take place (Mudrik et al., 2014). A number of theoretical contributions consider consciousness a necessary component for higher-level processing, including (but not limited to) semantic knowledge, complex visual processing, as well as problem solving and decision-making (Baars, 2002; Treisman, 2003). In those views, consciousness plays a role in enabling information to be integrated across distinct brain regions through long-range feed-back and feed-forward connections (Baars, 2002; Dehaene & Changeux, 2011; Dehaene, Sergent, & Changeux, 2003). In contrast, unconscious processing appears to be confined to separate areas, and does not result in a global spread of activity (Baars et al., 2003; Dehaene et al., 2001; Melloni et al., 2007).

In support of this view, previous research in the fields of associative learning and priming suggests that low-level or short-range spatiotemporal (Lin & He, 2009; Van Den Bussche et al.,

2009) and multisensory (Faivre et al., 2014b; Scott et al., 2018) information integration can proceed without conscious awareness of the stimuli (typically achieved with subliminal presentation methods such as masking or continuous flash suppression). Conversely, higher-level or longer-range spatial and temporal processing (e.g. for tasks requiring longer-term information maintenance or selective, flexible decision-making) should require conscious access (Dehaene & Naccache, 2001; Dehaene et al., 2003; Kouider & Dupoux, 2001). Previously reported instances of unconscious learning are in line with those assumptions. For example, classical conditioning can be achieved without awareness in delay scenarios (where stimuli to be integrated overlap temporally), but not in trace scenarios (where stimuli are temporally separated (Clark & Squire, 1998)), with similar results in other associative learning tasks (Knight et al., 2003; Raio et al., 2012; Seitz et al., 2009a).

As such, the idea that instrumental learning can proceed with unconsciously perceived stimuli is an interesting case. In the first experiment to demonstrate it (Pessiglione et al., 2008), participants learned to adjust their behaviour, through Go/NoGo decisions, in line with subliminally presented rewarding and punishing cues, learning to approach the rewarding and avoid the punishing stimulus without ever consciously perceiving them (constituting an example of trace instrumental learning). In order to successfully learn when to act and when to refrain from acting, participants had to learn from two temporally separated events across the length of the trial (up to 4 seconds): the subliminal stimulus itself, and its consequences, presented supraliminally as monetary reinforcement. Such a form of learning thus involves a fairly complex process of integrating information over a large temporal scale and distinct modalities, necessary to process the visual input, deploy selective action in response to the predictive cue, and process the reinforcement. The task is then considerably more complex than the aforementioned classical conditioning or associative learning scenarios, where there are fewer events, often in closer temporal proximity. Assuming that subliminally presented cue is not capable of evoking large-scale activity to be integrated with subsequent processes, the case of unconscious instrumental conditioning might appear at odds with the theory and previous experimental evidence covered above. Yet, instrumental conditioning is also one of the earliest and most fundamental forms of adaptive behaviour, both phylogenetically and ontogenetically. As such, the extent to which it requires conscious access is a question of considerable theoretical value.

A key challenge in any research into unconscious influences on behaviour lies in reliably asserting that processing is genuinely unconscious (Newell & Shanks, 2013; Rebuschat, 2013; Timmermans & Cleeremans, 2015). Although it is frequent practice in this line of research to infer unconscious processing when a behavioural measure (e.g. conditioning, priming, etc.) is above chance, while a separate measure of awareness is non-significantly different from chance

performance (e.g. a non-significant result in a discrimination task), this approach has been heavily criticised (Dienes, 2015; Vadillo, Konstantinidis, & Shanks, 2016). A non-significant result alone cannot disambiguate between no evidence for an effect (i.e. insensitive data, e.g. due to the small sample size) and absence of an effect (i.e. support for the null hypothesis). As such, finding that performance on an awareness check does not significantly differ from chance is not enough to assert true absence of awareness – an assertion which must be fulfilled to enable any inferences about the effect of interest, such as presence of unconscious conditioning in the original Pessiglione and colleagues study (Dienes, 2015; Shanks, 2017). This fallacy can be rectified in two ways: 1) ensuring that the methods are relevant, and sufficiently sensitive (Berry & Dienes, 1993, p.38; Shanks & St. John, 1994), and 2) with use of statistical methods, most prominently the Bayes factor, which enables to determine whether a null result indicates support for the null (e.g. awareness absent) over the alternative hypothesis (e.g. awareness present), or whether the data are insensitive (Dienes, 2014, 2016; Sand & Nilsson, 2016).

With these considerations in mind, we revisit the suggestion that instrumental learning can proceed without stimulus awareness. Experiment 1 will attempt to replicate the effect found by Pessiglione and colleagues (2008), following the original design and supplementing the original analyses with a Bayesian approach geared to determine a genuine absence of awareness, at least as measured by their test of awareness (whether this measure is a justified measure of awareness is an issue we will return to). Should the replication be successful, Experiment 2 will attempt to replicate the effect once again, this time with improved methods, to address the methodological issues related to the criteria of sensitivity and relevance in the original study.

In order to test whether stimuli that produce a certain level of learning are subliminal, one needs to know how much conscious perception would be needed to produce that level of learning (Dienes, 2015). Thus, a pilot study was conducted in which stimuli were presented moderately above the objective threshold in order to determine a relationship between the level of awareness (given the test of awareness used by Pessiglione et al, 2008) and learning. Thus, first we ran a pilot study to norm the relationship between learning and required awareness levels when the learning is based on conscious perception.

4.2. Pilot: Relationship between level of awareness and learning above the objective threshold

The pilot study aimed to assess both perceptual discrimination accuracy when awareness is present in a same/different discrimination task (as ensured with supraliminal stimulus

presentation), and the corresponding level of learning subsequently achieved in a Go/NoGo task with the same stimulus exposure duration. This will be assessed employing a methodology identical to that of the replication study, Experiment 1. The observed relationship between awareness and learning will be used to identify the rough appropriate effect size for Bayes Factor calculation in the corresponding task conducted without awareness. The pilot was pre-registered at <https://osf.io/rwnt7>.

4.2.1. Method

4.2.1.1. Participants

26 participants (3 males, $M_{AGE} = 22$, $SD_{AGE} = 4.3$) were recruited. Sample size was determined with G*power (Faul, Erdfelder, Lang, & Buchner, 2007), using a Cohen's d of 0.7 (a large effect size is justified given the supraliminal nature of the stimuli and the simplicity of the task), with 95% power. One participant was excluded after reporting to have misunderstood the learning task during debrief, yielding a final sample of 25 participants.

4.2.1.2. Stimuli and materials

The stimuli included 9 randomly selected characters from the Agathodaimon font presented in a white typeface on a black background (see Appendix 2), with a size of 70 x 70 pixels. For each participant, 3 were randomly assigned to the first perceptual discrimination task (PDT1; threshold-setting), 3 to PDT2 (awareness check), and 3 to the main leaning task (1 to be associated with rewarding, 1 with punishing, and 1 with neutral outcome). Two black-and-white visual noise masks of the same size as the stimuli were generated by scrambling one character image into 8.75 by 8.75 pixels squares. The same two masks were used for all participants in the same fashion (one preceding and one following the target stimulus). The outcome images were a circled £1 coin image for reward, a crossed-out £1 image for punishment, and a greyed-out coin for neutral.

The task was programmed using Matlab 2018b (MathWorks, 2018), running Psychophysics Toolbox (Brainard, 1997), and presented on a Samsung 2233RZ LCD monitor with a 120Hz refresh rate (following recommendations for precise visual presentation; Wang & Nikolić, 2011). Responses were collected with a standard keyboard.

4.2.1.3. Procedure

Perceptual discrimination task 1: Threshold finding.

This task aimed to establish a cue display duration that permitted conscious discrimination at above-chance levels without reaching ceiling. Participants were seated at a 50cm distance from the screen (ensured with a chinrest). Each trial began with a fixation cross (500ms), followed by presentation of two cues (display duration starting at 600ms), both forward-backward masked (67ms), separated by a 3s interval indicated by a fixation cross, following the method of Pessiglione et al (2008). Following the displays, participants were asked to indicate whether the cues presented were the same or different, and judge their confidence in that decision (on a binary scale between “some confidence” and “total guess”). Both responses were made using the arrow keys. Cue display duration started at 600ms, and dropped by 50ms with every correct and confident discrimination. Once participants reached 100ms or indicated guess for the first time, the display duration was increased by one increment (+50ms), and proceeded to decrease by smaller increments (8ms, corresponding to a single screen refresh duration on a 120Hz monitor). Once participants responded guess 6 consecutive times (irrespective of symmetry judgment accuracy), the corresponding display duration was taken to be their threshold of conscious perception.

The display duration was then set to be 16ms greater than the identified threshold and participants required to continue to make the same symmetry and confidence judgments for a minimum of one block of 10 further trials. If objective discrimination accuracy for those 10 trials was between 70 and 90% (above chance, indicating that participants can reliably discriminate the cues, but are not at ceiling), the task terminated and the duration was recorded as the display duration to be used in the main task. Note that confidence was discounted in this measure, and only objective accuracy was taken into account. If discrimination accuracy for these 10 trials was greater than 90% then the display duration was reduced by 8ms and the process repeated for another 10 trials until discrimination accuracy fell into the desired range (70-90%). Similarly, if discrimination accuracy for the 10 trials was below 70% the display duration was increased by 8ms and a further block of 10 trials completed until such time as the desired discrimination accuracy was achieved.

Main conditioning task.

In keeping with the original protocol (Pessiglione et al., 2008), participants were asked to choose between making a response by pressing a spacebar (Go), or refraining from a response (NoGo), to masked cues. In each block, one cue was paired with reward, one with punishment,

and one with the neutral outcome. Hence, participants could choose either to take a “risky” action (where they might win £1, lose £1, or have a neutral outcome depending on the proceeding cue) or to refrain from acting and thus ensure a neutral outcome.

Each trial began with a fixation cross (500ms), followed by a forward mask (67ms), one of the target cues (determined supra-threshold display duration), and backward mask (67ms). Subsequently, a question mark appeared on the screen, indicating that the response could be made. Regardless of the response (Go or NoGo), the response window remained open for 3000ms, after which the choice made (Go! or No!) was displayed (500ms), followed by the outcome (reward, punishment, or neutral; 2000ms). There was one block of 90 trials, with 30 rewarding, punishing and neutral trials each, in a randomised order.

Perceptual discrimination task 2: Awareness check.

The second and final discrimination task was used to assess the objective level of cue awareness, as indexed by same/different discrimination accuracy. No further adjustments to display duration were made, which remained at the level determined in the perceptual discrimination task. There was one block of 100 trials, with 50 same and 50 different trials in a randomised order.

4.2.2. Analysis and results

Bayes factors (B) were used to assess the strength of evidence for the alternative hypothesis, H_1 , over the null, H_0 (Wagenmakers et al., 2017). All Bayes factors, B, reported here represent the evidence for H_1 relative to H_0 ; to find the evidence for H_0 relative to H_1 , take $1/B$. Here, $B_{H(0, x)}$ refers to a Bayes factor in which the predictions of H_1 were modelled as a half-normal distribution with an SD of x (see Dienes & Mclatchie, 2017); the half-normal can be used when a theory makes a directional prediction where x scales the size of effect that could be expected. With the assumptions we used for modelling H_1 , as it happened, where an effect yielded a p value less than .02, the Bayes factor was above 6, though there is no guarantee of such a correspondence between B and p values (Lindley, 1957). To indicate the robustness of Bayesian conclusions, for each B, a robustness region will be reported, giving the range of scales that qualitatively support the same conclusion (i.e. evidence as insensitive, or as supporting H_0 , or as supporting H_1), notated as: $RR_{B>6} [x_1, x_2]$ or $RR_{B<1/6} [x_1, x_2]$ or $RR_{1/6>B>6} [x_1, x_2]$ where x_1 is the smallest SD that gives the same conclusion and x_2 is the largest (see Dienes, 2019).

4.2.2.1. Data pre-processing

In order to account for potential response bias, type I d' (a Signal Detection Theoretic measure of sensitivity to signal versus noise; Stanislaw & Todorov, 1999) was computed for both PDT2 and the main conditioning task. Type I d' can be used to index awareness level corresponding to the objective threshold, where chance performance corresponds to lack of awareness, regardless of confidence or subjective awareness reports. Note that this measure is used here following the procedure of Pessiglione et al (2008). For the PDT, correct same/different responses were treated as hits, and incorrect responses as false alarms. In the conditioning task, Go responses to rewarding cues were treated as hits, and Go responses to punishing cues as false alarms. Go responses to neutral cues were discounted, as participants are expected to respond arbitrarily to them due to their null outcome.

4.2.2.2. Awareness check

At the group level, d' scores for the PDT2 were entered into a one-way t-test against 0, which indicates no ability to discriminate the stimuli (no sensitivity between signal versus noise, akin to chance performance). A Bayes Factor (B) was computed for the difference, with the predictions of H_1 (awareness is present) modelled as a half-normal distribution centred on 0, with an SD equal to a d' of 1 (the average expected effect size corresponding to 70% hit rate (accuracy) and 30% false alarms, an estimate of above-chance and below-ceiling performance).

The results indicate that participants were able to successfully discriminate the stimuli, with the average d' significantly greater than zero ($M = 0.946$, $SE = 0.197$, $p < 0.001$, $B_{H(0,1)} = 25554$, $RR[0.07, 303.5]$).

4.2.2.3. Main conditioning task

The d' scores for the conditioning task were entered into a one-way t-tests against 0, indicating lack of discrimination between the cues, and consequently, lack of learning. B was computed for the difference, with the predictions of H_1 (learning is present) modelled as a half-normal distribution centred on 0, with an SD equal to 0.7 (the expected effect size if learning is present, derived from Pessiglione et al., 2008).

The results indicate that participants were able to successfully learn, with the average d' significantly greater than zero ($M = 1.793$, $SE = 0.301$, $p < 0.001$, $B_{H(0,0.7)} = 2514517$, $RR[0.081, 633.5]$).




Figure 15: Learning (d') plotted against objective level of awareness (d') obtained in the supraliminal pilot study. Ribbon represents a 95% confidence interval. $N = 25$.

4.2.3. Pilot: Discussion

The purpose of the pilot was to establish a rough relation between the level of awareness as measured by the awareness measure by Pessiglione et al (2008), and the level of learning it can support. In the pilot, the mean awareness was $d' = 0.9$, and the mean learning was $d' = 1.793$. These are all the crucial facts we need. If both these measures result from the influence of the same knowledge base, namely conscious perception, then as conscious perception goes to zero, both should also go to zero (Dienes, 2015). Thus, on a plot of awareness against learning (Figure 15), a line from the point given by the two means going to (0,0) gives a rough estimate of the relation that should be obtained between awareness and learning, assuming it is linear. While there are uncertainties in both the estimates and their linearity, we only need a rough estimate, as we will model uncertainty around this estimate, and robustness regions will be provided. Now we are in a position to proceed with the replication.

The theory that the Pessiglione et al. (2008) method produces unconscious learning involves two predictions: 1) participants will perform at chance on the awareness measure; and 2) participants will show conditioning. These are the crucial tests we will consider below. 1) might be regarded as an outcome neutral test in order for the paradigm to be relevant for showing unconscious learning. From the point of view of a replication, however, it constitutes a crucial test of whether the procedure does result in stimuli being subliminal.

4.3. Experiment 1: Direct replication

Experiment 1 aims to directly replicate unconscious instrumental conditioning found in Pessiglione et al. (2008). For this reason, all methods will be in keeping with those employed in that original study. The original frequentist analyses will be supplemented with Bayes Factors, in order to disambiguate potentially non-significant results as either indicating support for the null hypothesis, or indicating insensitive data.

4.3.1. Method

4.3.1.1. Participants

Participants will be recruited at the University of Sussex. Sample size will be determined with the Bayesian Stopping Rule, using previously obtained effect sizes as empirical priors, or cease at 200 participants if the result in the awareness check remains insensitive (see 3.2.2. *Planned Analyses* for detail). In keeping with the original study, participants will be told they will be reimbursed with their earnings from the task, but at the end this will be rounded to a fixed amount of £6. Ethical approval was granted by the School of Psychology ethics committee at the University of Sussex, and the study will be conducted in accordance with the Declaration of Helsinki.

4.3.1.2. Stimuli and materials

All stimuli and materials used will be equivalent to those reported in the original study. The stimuli will be 15 randomly chosen characters from Agathodaimon font, presented in white typeface on a black screen, with a size of 240 by 180 pixels. For each participant, 3 will be randomly assigned to PDT1, 3 to PDT2, and 9 to the main task, with 1 rewarding, 1 punishing, and 1 neutral cue in each of the 3 blocks. Two black-and-white visual noise masks of the same size as the stimuli will be generated by scrambling one character image into 30 by 30 pixel squares. The same two masks will be used for all participants in the same fashion (one preceding and one following the target stimulus). The outcome images will be a circled £1 coin image for reward, a crossed-out £1 image for punishment, and a greyed-out coin for neutral.

The task will be programmed using Matlab 2018b (MathWorks, 2018), running Psychophysics Toolbox (Brainard, 1997), and presented on a Dell LCD monitor with a 60Hz refresh rate (manufactured in 2006 to approximate the screen technology used in the original experiment by Pessiglione et al). Responses will be collected with a standard keyboard.

4.3.1.3. Procedure

Perceptual discrimination task 1.

Participants will be seated with their chin on a chin rest placed at 50cm distance from the screen. Each session will commence with a PDT used to determine the individual cue display duration. The duration will be either (33 or 50ms); the largest for which they show chance-level ($\leq 50\%$) discrimination performance. In the task, participants will be shown two cues, each forward-backward masked (67ms), separated by a 3s interval indicated with a fixation cross, following the method of Pessiglione et al (2008). Following the display, they will be asked to report whether the cues presented were the same or different, using the arrow-keys. The task will consist of 2 blocks of 120 trials (with 60 same and 60 different trials in each, in a randomised order). The first block will be conducted with 50ms display duration of each cue. If discrimination accuracy at this stage is at chance (assessed with a chi-squared test for each participant), the task will be ended and a 50ms display duration adopted in the main task. If performance in the first block is above chance, the duration will be decreased to 33ms for the second block, and performance assessed again. If it is at chance at the end of the second block, the duration of 33ms will be adopted in the main task. Participants who remain above chance at 33ms will not be able to take part.

Main conditioning task.

The task will be identical to the original protocol and the pilot study, with the exception that the cues were presented subliminally, for the duration determined in the PDT1 (33 or 50ms). Participants will be asked to choose between making a response by pressing a spacebar (Go), or refraining from a response (NoGo), to masked cues. In each block, one cue will be paired with reward, one with punishment, and one with the neutral outcome. Hence, participants can choose either to take a “risky” action (where they might win £1, lose £1, or have a neutral outcome depending on the proceeding cue) or to refrain from acting and thus ensure a neutral outcome.

Each trial will begin with a fixation cross (500ms), followed by a forward mask (67ms), one of the target cues (determined subliminal display duration), and backward mask (67ms). Subsequently, a question mark will appear on the screen, indicating that the response may be made. Regardless of the response (Go or NoGo), the response window will remain open for 3000ms, and participants’ response will be collected at the end – Go if the spacebar is being pressed, and NoGo if it is released. Finally, the choice made (Go! or No!) will be displayed (500ms), followed by the outcome (reward, punishment, or neutral; 2000ms).

In order to counterbalance motor conditions, the 'risky' response will be pseudo-randomised to be Go for half the participants, and NoGo for the other half. There will be 3 blocks of 120 trials, with 40 trials of each type (rewarding, punishing, neutral).

Preference task.

Following each of the three conditioning blocks, the three cues used will be shown on the screen side by side, unmasked, in a randomised order. Participants will be asked to rate them in order of preference, from most (3) to least (1) liked.

Perceptual discrimination task 2.

A PDT with 120 trials (60 same, 60 different) and 3 new stimuli will be repeated at the end of the testing session. There will be no adjustments to the display duration, which will be kept at the level determined in PDT1. The task will allow to determine whether or not participants' cue awareness remained at chance level.

4.3.2. Planned Analysis

4.3.2.1. Data pre-processing

Identical to pilot study.

4.3.2.2. First crucial test: awareness check

Absence of awareness will be determined by assessing discrimination performance on the second perceptual discrimination task, indexed by type I d' scores (corresponding to the objective threshold of awareness). At the group level, d' scores will be entered into a one-way t -test against 0, which indicates no ability to discriminate between the stimuli (no sensitivity between signal versus noise). B will be computed on the obtained mean d' , with the H_1 (awareness present) modelled as a half-normal distribution with a mean of 0 and a SD equal to the value derived from the pilot study, following the regression method outlined by Dienes (2015, p.211-213). The mean learning d' (1.793) and the corresponding mean awareness d' (0.9) from the supraliminal pilot will be used to estimate the mean awareness d' expected from the obtained level of learning in the main subliminal conditioning task. This will be done using the regression line drawn between the supraliminal mean values and the point of origin (no learning, no awareness). The expected d' if awareness is present in the unconscious task (H_1) can then be derived from the learning d' value we actually obtain in this experiment. In line with the Bayesian Stopping Rule, data collection will continue until support for the H_0 at the group level is found

($B_{H(0,X)} < 1/6$), or cease at 200 in the event of the data remaining insensitive. The upper cap was determined with the Bayesian sample size estimation method (Dienes, 2015b), using the above regression method to derive the average awareness d' (0.37) expected from Pessiglione et al (2008) learning d' (0.7), and adjusting the awareness SE obtained in the pilot (0.2) in line with sample size increases. A robustness region will be reported, as described in the Pilot.

Following the original method, performance for every individual compared to chance (50% accuracy) will be assessed with chi-square tests. Participants who show significant above-chance performance will be excluded from further analysis, as well as those who explicitly report seeing the stimuli on-screen.

4.3.2.3. Second crucial test: main conditioning task

Presence of learning in the conditioning task will be assessed with d' scores. d' scores will be entered into a one-way t-tests against 0, indicating lack of discrimination between the cues, and consequently, lack of learning. B will be computed with H_1 modelled as a half-normal distribution with a mean of 0 and a SD equal to 0.7 (expected effect size if learning is present, derived from Pessiglione et al., 2008). Resulting $B_{H(0,0.7)} > 6$, will be taken as evidence of learning. $B_{H(0,0.7)} < 1/6$, will be taken as evidence for absence of learning. In the event of an insensitive result, data collection will cease at 170 participants (upper cap estimated in the same way as in section 3.2.2., using a learning d' of 0.7 as the expected effect size and learning SE of 0.3 obtained in the pilot). A robustness region will be reported, as described in the pilot.

4.4. Experiment 2

Experiment 2 will be conducted only in the event of Experiment 1 replicating the effect found in Pessiglione et al (2008). In light of the methodological and theoretical advances and debates in the field of unconscious learning (e.g. Dienes, 2015; Mudrik et al., 2014; Newell & Shanks, 2014), Experiment 2 will aim to replicate the result, introducing changes to the paradigm targeted at increasing the methodological rigour.

Firstly, in the original study, the measures in the awareness check and the learning task pertain to two different aspects of decision-making. The perceptual discrimination tasks (serving as threshold-setting and as awareness check) required a same/different perceptual judgment of 2 cues, separated by a 3s interval. In contrast, the main conditioning task required an approach/avoid response after a single stimulus. Hence, the measure used in the perceptual discrimination task reflected a different decision process than was required in the conditioning task, violating the relevance and sensitivity criteria (Berry & Dienes, 1993; Newell & Shanks,

2014). As such, the threshold-setting task will be amended to match the conditioning task more closely.

Secondly, the separate awareness check will be replaced with a trial-by-trial measure, allowing to access the information about participants' awareness in a more immediate fashion (Berry & Dienes, 1993; Newell & Shanks, 2014).

Thirdly, the original task design leaves open the possibility that participants might occasionally experience awareness of the stimuli in the learning task, which does not become apparent in the final PDT. This might occur either where the same brief moments of awareness do not reoccur in the PDT or where they are too infrequent to significantly influence the overall objective accuracy measure. Reliably excluding individual trials is impossible when only objective discrimination measures are collected. With this in mind, the trial-by-trial awareness check will also include confidence ratings, allowing to exclude trials where participants were subjectively aware. Because the initial staircase should sensitively settle on sub-threshold conditions, a trial-by-trial check should only elicit a small number of aware trials. While post-hoc trial exclusion of conscious trials can lead to regression to the mean (Shanks, 2017), its effect is negligible if the majority of trials are unconscious. We examined this assumption by modelling worst-case scenarios for presence of conscious trials at different proportions of observed unconscious trials, at different error rates (see Supplementary material). This allows to determine the maximum percentage of conscious trials which could inadvertently contribute to unconscious knowledge. Thus, if we observe 80% of unconscious trials (leaving room for error in the remaining 0-20%), the maximum proportion of conscious trials possibly contained within our observed unconscious trials is 1.59%. This would be the percentage of our conscious knowledge potentially accounting for the learning effect. Using the observed d' found in the pilot study of 1.8, we find that the maximum influence from conscious knowledge where 80% of responses are attributed to unconscious responding is $d' = 0.03$. We consider this negligible and as such will adopt a strategy whereby provided the proportions of responses attributed to conscious responding does not exceed 20%, our exclusion criteria will be applied (see section 4.2.2.).

Finally, the forward and backward masks will be generated afresh on each trial by randomly scrambling a black-and-white noise image. The use of different masks on each trial reduces the likelihood of participants building erroneous associations from possible salient repetitive features of the masks (or of some stimulus-mask combinations).

As in Experiment 1, original frequentist analyses will be supplemented with Bayes Factors.

4.4.1. Method

4.4.1.1. Participants

Participants will be recruited at the University of Sussex. Sample size will be determined with the Bayesian Stopping Rule, using previously obtained effect sizes as empirical priors, or cease at 170 participants should the result remain insensitive (see 4.2.2. *Planned Analyses* for detail). In keeping with the original study, participants will be told they will be reimbursed with their earnings from the task, but at the end this will be rounded to a fixed amount of £6. Ethical approval was granted by the School of Psychology ethics committee at the University of Sussex, and the study will be conducted in accordance with the Declaration of Helsinki.

4.4.1.2. Stimuli and Materials

The stimuli will be 11 characters from Agathodaimon font, chosen pseudo-randomly to ensure six symmetrical and seven asymmetrical characters. All will be presented in white typeface on a black screen, with a size of 240 by 180 pixels. For each participant, two characters will be randomly assigned to PDT1 (with one symmetrical and one asymmetrical character). The nine remaining stimuli will be pseudo-randomly assigned to the main task, with one rewarding, one punishing, and one neutral cue in each of the three blocks, such that each block contains both symmetrical and asymmetrical cues. Both the forward and backward masks will be generated afresh on each trial by randomly scrambling a 240 by 180 pixels black-and-white noise image in blocks of 3x3 pixels.

4.4.1.3. Procedure

Perceptual discrimination task 1.

Each session will commence with a PDT, allowing to determine the individual cue display duration. The duration will be either 33 or 50ms, the largest for which results in chance-level ($\leq 50\%$) discrimination performance. In this task (in contrast to Experiment 1), participants will be shown a fixation cross (500ms), followed by a single cue (display duration starting at 50ms), forward-backward masked (67ms) cue. Following the display, they will be asked to report whether the cue presented was symmetrical or asymmetrical, using the arrow-keys. As in Experiment 1, the task will consist of 2 blocks of 120 trials (with 60 symmetrical and 60 asymmetrical trials, in a randomised order, in each block). The first block will be conducted with 50ms display duration of each cue. If discrimination accuracy at this stage is at chance (assessed with a chi-squared test for each participant), the task will be ended and a 50ms display duration adopted in the main task. If performance in the first block is above chance, the duration will be

decreased to 33ms for the second block, and performance assessed again. If it is at chance at the end of the second block, the duration of 33ms will be adopted in the main task. Participants who remain above chance at 33ms will not be able to take part.

Main conditioning task.

The conditioning task will be identical to Experiment 1, with the exception that an awareness check will be added at the end of each trial. Following feedback presentation (reward/punishment), participants will be asked to report if the masked cue was symmetrical or asymmetrical, using the arrow keys. Next, they will be asked to report their confidence in that judgment on a binary scale (between ‘some confidence’ and ‘total guess’). There will be 3 blocks of 120 trials, with 40 trials of each type (rewarding, punishing, neutral), in a randomised order. Prior to beginning, participants will be explicitly instructed that the symmetry judgments are not related to the rewarding/punishing outcomes. They will also be shown a different pair of example cues to illustrate what is meant by symmetry.

Preference task.

Identical to Experiment 1.

4.4.2. Planned Analysis

4.4.2.1. Data pre-processing

Identical to Experiment 1.

4.4.2.2. Crucial test: main conditioning task

Individual trials where participants make a correct symmetry judgment with confidence will be marked as ‘aware’ trials and excluded. In cases where exclusions exceed 20% of trials, the entire participant will be excluded from analysis. The remaining trials will be analysed with type I d' in a manner identical to Experiment 1. B will be computed with H_1 modelled as a half-normal distribution with a mean of 0 and a SD equal to 0.7 (expected effect size if learning is present, derived from the original study). Resulting $B_{H(0,0.7)} > 6$, will be taken as evidence of learning. $B_{H(0,0.7)} < 1/6$, will be taken as evidence for absence of learning. In the event of an insensitive result, data collection will cease at 170 participants (upper cap estimated in the same way as in section 3.2.3., using a learning d' of 0.7 as the expected effect size and learning SE of 0.3 obtained in the pilot). A robustness region will be reported.

For reasons given by the simulations in the Supplementary material, if more than 20% of participants have to be excluded due to having more than 20% of conscious trials, we will use an interval null hypothesis, with the interval being from 0 to the maximum d' that could be contributed by regression to the mean, as determined by the exact formulae in the Supplementary materials. We will calculate the contribution to d' from excluding trials, then from excluding participants, and add the contributions together. The model of H1 will remain the same. (This entails an overlap in the effects predicted by H1 and H0, implying a maximum expected Bayes factor given by the ratio of marginal likelihoods of H1 to H0 at the population effect size).

4.5. Supplementary material

Calculating the maximum potential influence of regression towards the mean resulting from post-hoc trial or participant exclusion.

Our paradigm proposes to use a trial-by-trial test of awareness in order to permit the exclusion of any trials where participants report having conscious knowledge of the stimuli. This approach has a number of key advantages over alternative approaches such as examining the mean accuracy in a separate test block. Firstly, it maximises the sensitivity of analysis by reducing unnecessary whole participant exclusions i.e. individual conscious trials can be excluded rather than all of a participant's responses excluded based on their mean accuracy. Secondly, the ability to exclude conscious trials in this way permits us to target an exposure duration that is as close as possible to the conscious threshold rather than having to deliberately aim substantially below that threshold in order to avoid the possibility of a small number of conscious trials contaminating the analysis. Finally, if a person's threshold is subject to variation over time, only a trial-by-trial analysis as proposed can reliably ensure that the effects of such momentary variations in conscious awareness are mitigated.

While employing a trial-by-trial awareness check has obvious advantages, any post-hoc exclusion of conscious trials or conscious subjects from the full sample brings a potential influence of regression to the mean (Shanks, 2017). In essence, while excluding the reported conscious trials will give the impression of leaving a set of purely unconscious trials for analysis, any noise in the measure of awareness (participants' reports of conscious perception) will mean that the selected subset may contain some proportion of conscious trials that were (due to measurement error) mislabelled as unconscious. Here we conduct an analysis that establishes the maximum potential influence that regression to the mean could have on our analyses for a given proportion of trials or participants excluded and a given error rate.

Our paradigm tests participants' ability to discriminate consciously whether a given shape is symmetrical or asymmetrical. For each symmetry-judgement, participants are required to report either that "I am guessing" or "I have some confidence". Any accuracy apparent in responses attributed to guessing is assumed to be unconscious and any apparent in responses attributed confidence is taken to be conscious. Because the awareness judgement is a two-alternative forced choice, the conscious plus unconscious attributions amount to 100% of responses. Crucially however, the observed proportions of each will differ from the true proportions according to the degree of error in the awareness report i.e. the number of trials where the true state is unconscious and yet reported with some confidence, or the true state is

conscious but reported as guessing. For simplicity, we assume that there is a fixed amount of error in this measure and that it is the same for both conscious and unconscious attributions. It follows that the observed proportion of unconscious trials (those attributed to guessing) will be the proportion of trials that are actually unconscious * the proportion of judgements without error + the proportion of trials that are conscious * the proportion of judgements made with error (misattributions). Note, that the maximum error is bounded by the proportion of responses attributed to each of the unconscious and conscious categories, e.g. if 90% of responses are attributed to unconscious judgement then the maximum error is 10% as this would account for all the conscious responses (10%). Adopting the following abbreviations, the formulae for observed proportion of unconscious trials can be written as below:

obsUC = The observed proportion of unconscious trials (trials attributed to guessing)

trueUC = The true proportion of unconscious trials

error = The proportion of judgments misattributed due to error

$$obsUC = (1 - error) * trueUC + error * (1 - trueUC)$$

This simplifies:

$$obsUC = trueUC + error * (1 - 2 * trueUC)$$

$$obsUC = error + trueUC * (1 - 2 * error)$$

From this, we can derive a formula for *trueUC*:

$$trueUC = \frac{(obsUC - error)}{(1 - 2 * error)}$$

And a formula for the proportion of conscious contamination of observed unconscious:

$$conscious\ contamination = \frac{error * (1 - trueUC)}{obsUC}$$

For example, if 90% of responses are attributed to unconscious responding (*obsUC* = 90%) and 10% to conscious it follows that the maximum error is 10% i.e. if 100% of responses are truly unconscious then 10% error would account for all the observed conscious responses. The true error can be anywhere between 0% and 10%. Applying the above formulae to each

possible error rate (0-10%), we find that an error of 5% gives the maximum conscious contamination, namely 0.31%. That is, where 90% of observed trials are attributed to unconscious responding and all conscious trials are excluded the maximum extent to which regression to the mean could influence the observed learning is 0.31% of the conscious learning effect. In the conscious pilot, the d' for learning was 1.8. Hence, the maximum effect of regression to the mean where 90% of responses are attributed to unconscious responding is $1.8 * 0.31\%$, $d' = 0.01$.

Supplementary table 1 shows the percentage of conscious contamination for each level of error for observed unconscious responding proportions of 95%, 90%, 85%, 80%, 75%, and 70%. Consulting this table and using the observed d' in the pilot study of 1.8 we find that the maximum influence from conscious knowledge where 80% of responses are attributed to unconscious responding is $d' = 0.03$. We feel that this is negligible and as such will adopt a strategy whereby provided the proportion of responses attributed to conscious responding does not exceed 20%, our exclusion criteria will be applied. Where participants make greater than 20% conscious attributions those individuals will be excluded as conscious. Precisely the same logic assessing the potential regression to the mean effect at the participant level applies. Hence, provided not more than 20% of participants make greater than 20% conscious attributions, we can be confident that the observed learning effect is not influenced by regression to the mean by more than a d' of 0.03. In the unlikely event that greater than 20% of participants make greater than 20% of conscious attributions then we will adopt an interval null in the Bayes analysis to account for this. This will modify the Bayes analysis only to the extent that the point null hypothesis will be replaced by an interval null modelled as a uniform from zero to the maximum influence of conscious contamination.

Table S4.1. Percentage of conscious contamination for each level of error for observed unconscious responding proportions of 95%, 90%, 85%, 80%, 75%, and 70%, computed with the formulae described in text. Frames highlight the maximum proportion of conscious contamination for the corresponding proportion of unconscious responding.

Observed unconscious	95%															
error	0%	1%	2%	3%	4%	5%										
true percentage unconscious	95.0%	95.9%	96.9%	97.9%	98.9%	100.0%										
conscious contamination	0.00%	0.04%	0.07%	0.07%	0.05%	0.00%										
Observed unconscious	90%															
error	0%	1%	2%	3%	4%	5%	6%	7%	8%	9%	10%					
true percentage unconscious	90.0%	90.8%	91.7%	92.6%	93.5%	94.4%	95.5%	96.5%	97.6%	98.8%	100.0%					
conscious contamination	0.00%	0.10%	0.19%	0.25%	0.29%	0.31%	0.30%	0.27%	0.21%	0.12%	0.00%					
Observed unconscious	85%															
error	0%	1%	2%	3%	4%	5%	6%	7%	8%	9%	10%	11%	12%	13%	14%	15%
true percentage unconscious	85.0%	85.7%	86.5%	87.2%	88.0%	88.9%	89.8%	90.7%	91.7%	92.7%	93.8%	94.9%	96.1%	97.3%	98.6%	100.0%
conscious contamination	0.00%	0.17%	0.32%	0.45%	0.56%	0.65%	0.72%	0.77%	0.78%	0.77%	0.74%	0.66%	0.56%	0.41%	0.23%	0.00%
Observed unconscious	80%															
error	0%	1%	2%	3%	4%	5%	6%	7%	8%	9%	10%	11%	12%	13%	14%	15%
true percentage unconscious	80.0%	80.6%	81.3%	81.9%	82.6%	83.3%	84.1%	84.9%	85.7%	86.6%	87.5%	88.5%	89.5%	90.5%	91.7%	92.9%
conscious contamination	0.00%	0.24%	0.47%	0.68%	0.87%	1.04%	1.19%	1.32%	1.43%	1.51%	1.56%	1.59%	1.58%	1.54%	1.46%	1.34%
error	16%	17%	18%	19%	20%											
true percentage unconscious	94.1%	95.5%	96.9%	98.4%	100.0%											
conscious contamination	1.18%	0.97%	0.70%	0.38%	0.00%											

Table S4.1. *continued*

Observed unconscious	75%															
error	0%	1%	2%	3%	4%	5%	6%	7%	8%	9%	10%	11%	12%	13%	14%	15%
true percentage unconscious	75.0%	75.5%	76.0%	76.6%	77.2%	77.8%	78.4%	79.1%	79.8%	80.5%	81.3%	82.1%	82.9%	83.8%	84.7%	85.7%
conscious contamination	0.00%	0.33%	0.64%	0.94%	1.22%	1.48%	1.73%	1.95%	2.16%	2.34%	2.50%	2.63%	2.74%	2.81%	2.85%	2.86%
error	16%	17%	18%	19%	20%	21%	22%	23%	24%	25%						
true percentage unconscious	86.8%	87.9%	89.1%	90.3%	91.7%	93.1%	94.6%	96.3%	98.1%	100.0%						
conscious contamination	2.82%	2.75%	2.63%	2.45%	2.22%	1.93%	1.57%	1.14%	0.62%	0.00%						
Observed unconscious	70%															
error	0%	1%	2%	3%	4%	5%	6%	7%	8%	9%	10%	11%	12%	13%	14%	15%
true percentage unconscious	70.0%	70.4%	70.8%	71.3%	71.7%	72.2%	72.7%	73.3%	73.8%	74.4%	75.0%	75.6%	76.3%	77.0%	77.8%	78.6%
conscious contamination	0.00%	0.42%	0.83%	1.23%	1.61%	1.98%	2.34%	2.67%	2.99%	3.29%	3.57%	3.83%	4.06%	4.27%	4.44%	4.59%
error	16%	17%	18%	19%	20%	21%	22%	23%	24%	25%	26%	27%	28%	29%	30%	
true percentage unconscious	79.4%	80.3%	81.3%	82.3%	83.3%	84.5%	85.7%	87.0%	88.5%	90.0%	91.7%	93.5%	95.5%	97.6%	100.0%	
conscious contamination	4.71%	4.78%	4.82%	4.82%	4.76%	4.66%	4.49%	4.26%	3.96%	3.57%	3.10%	2.52%	1.82%	0.99%	0.00%	

5.

INSTRUMENTAL CONDITIONING REQUIRES CONSCIOUS AWARENESS

Abstract

Instrumental conditioning is a crucial substrate of adaptive behaviour, allowing individuals to selectively interact with the stimuli in their environment to maximise benefit and minimise harm. The extent to which complex forms of learning, such as instrumental conditioning, are possible without conscious awareness is a topic of substantial importance, and ongoing debate. In light of recent theoretical and empirical contributions casting doubt on the early demonstrations of unconscious instrumental conditioning, we revisit the question of its feasibility in two modes of conditioning. In Experiment 1, we used trace conditioning, following a prominent paradigm (Pessiglione et al., 2008) and enhancing its sensitivity. Success in this task requires participants to learn to approach reward-predictive stimuli and avoid punishment-predictive stimuli through monetary reinforcement. All stimuli were rendered unconscious using forward-backward masking. In Experiment 2, we used delay conditioning to shorten the stimulus-outcome delay, retaining the structure of the original task but presenting the stimuli under continuous flash suppression to allow for an overlap of the stimulus, action, and outcome, as well as replacing monetary reinforcement with primary appetitive reinforcement. In both experiments, we found evidence for absence of unconscious instrumental conditioning, showing that participants were unable to learn to adjust their behaviour to approach positive stimuli and avoid negative ones. This result is consistent with evidence that unconscious stimuli fail to bring about long-term behavioural adaptations, and provides empirical evidence to support theoretical proposals that consciousness might be necessary for adaptive behaviour, where selective action is required.

5.1. Introduction

The extent to which learning can proceed without conscious awareness has been an extensively studied question. While there have been successful demonstrations that simple associations can be learned without conscious awareness of the stimuli, even between different modalities (e.g. Scott, Samaha, Chrisley, & Dienes, 2018), the possibility of successful unconscious learning at increasing levels of complexity is unclear. One example of a complex learning process which has previously been suggested to proceed without cue awareness is instrumental conditioning, where agents learn the contingencies between stimuli in the world and their outcomes, and to selectively act upon them (e.g. approach/avoid; Mastropasqua & Turatto, 2015; Pessiglione et al., 2008). However, emerging evidence suggests that more complex forms of learning, including instrumental conditioning (Reber et al., 2018), contingency learning (Travers et al., 2018), and fear conditioning (Mertens & Engelhard, 2020) may not be possible unconsciously, as originally thought. Adding fuel to the debate are developments in experimental and statistical methods used to reliably assess absence of awareness (Dienes, 2015a; Rothkirch & Hesselmann, 2017; Shanks, 2017). With this in mind, we revisit the topic of unconscious instrumental conditioning. Instrumental learning is critical to adaptive behaviour (both ontogenetically and phylogenetically). As such, the extent to which it is feasible without conscious awareness is a question of considerable importance. In what follows, we attempt to replicate the original findings of Pessiglione *et al* (2008), employing more sensitive methods and a more robust statistical approach, in two different paradigms: a trace conditioning task using monetary reinforcements, following the original task (Experiment 1), and a delay conditioning task with appetitive and aversive primary reinforcements (Experiment 2).

Past research into unconscious learning suggests that simple forms of associative learning, ranging from classical conditioning, to emotional, visuospatial, or multisensory learning can proceed without conscious awareness of the stimuli (Clark & Squire, 1998; Faivre, Mudrik, Schwartz, & Koch, 2014a; Knight et al., 2003; Lin & He, 2009; Olsson & Phelps, 2004; Rosenthal et al., 2016, 2010; Scott et al., 2018; Seitz et al., 2009a). Learning is usually indexed with presence or absence of the conditioned response, or with reaction times (e.g. where shorter reaction times to congruent vs incongruent pairings are indicative of learning). Unaware learning (as well as priming) is typically observed when stimuli are presented at short spatiotemporal intervals, or concurrently (van Gaal et al., 2012). For example, unaware classical conditioning was demonstrated in delay scenarios (where stimuli to be associated overlap temporally), but not in trace scenarios (where they are separated by an interval; Clark & Squire, 1998).

In contrast, instrumental conditioning can be classified as a more complex process. Agents must not only learn the associations between different stimuli, or stimuli and their outcomes, but also deploy action selectively (e.g. approach or avoid), and adapt their behaviour in the long-term. This involves integrating information over a long temporal scale and distinct modalities, as involved in processing the visual input, extracting its predictive value, deploying a selective response, processing the reinforcement, and comparing the expected outcome with the actual outcome in order to update the representations of stimulus value. As such, it is a considerably more complex process than the aforementioned simpler forms of associative learning, which do not require selective decisions on whether to act or not, or a behavioural adaptation from trial to trial.

Past theoretical and empirical accounts of conscious versus unconscious processing suggest that such an increased level of complexity should require conscious access. A number of theoretical contributions propose that consciousness is related to long-lasting, long-range connections between distinct brain regions, supporting recurrent information integration across distinct cognitive modules (Dehaene & Naccache, 2001; Dehaene & Changeux, 2011; Dehaene, Charles, King, & Marti, 2014; Lamme, 2006; Mudrik, Faivre, & Koch, 2014). As such, low-level or short-range (spatial or temporal) information integration might be possible without conscious awareness, but consciousness might be necessary at increased levels of complexity – for example, semantic knowledge, complex visual processing, decision-making and problem-solving (Baars, 2002), all of which involve integration of information across longer spatiotemporal intervals or larger spatial distance. Indeed, neuroimaging evidence has found conscious processing to be characterised by global, long-range spread of activity, in contrast to more localised, shorter-range projections when processing is unconscious (Baars, 2002; Baars, Ramsøy, & Laureys, 2003; Dehaene et al., 2014, 2001; Dehaene & Naccache, 2001; Melloni et al., 2007).

Another important consideration is related to the methodology of assessing absence of consciousness during task performance. Although it is a frequent practice in unconscious processing research to infer absence of awareness when a behavioural measure (e.g. conditioning, priming, etc.) is above chance, while an independent measure of awareness is not significantly different from chance (e.g. a separate unconscious discrimination task), this approach has drawn criticism (Dienes, 2015a; Vadillo et al., 2016). Finding that performance on a separate awareness check does not differ from chance is not enough to infer true absence of awareness – a non-significant result cannot disambiguate between absence of an effect (i.e. support for the null hypothesis) and absence of evidence for an effect (i.e. insensitive data). In a meta-analysis, Vadillo, Konstantindis and Shanks (2016) argue that the seemingly chance-

performance on this type of awareness test is more likely to reflect a false negative, especially in low-powered studies with small samples. This approach, also adopted in the Pessiglione *et al* study, can be rectified in two ways. One is to ensure that the methods to assess awareness are relevant and sensitive (Berry & Dienes, 1993; Shanks, 2017; Shanks & St. John, 1994). This could be achieved by, for instance, a closer similarity between the awareness test and the measure of interest. The second is to apply statistical methods, such as the Bayes factor, which allows to determine whether a null result indicates support for the null (e.g. absence of awareness) over the alternative hypothesis (presence of awareness), or whether the data are insensitive (Dienes, 2014, 2016; Sand & Nilsson, 2016). Both of those approaches will be used in the present paper to ensure true absence of awareness.

Instrumental conditioning is a crucial substrate of adaptive behaviour. Understanding whether it is achievable without conscious awareness is therefore of vital importance for theoretical and practical research investigating the boundaries of unconscious processing. With recent research casting doubt on the feasibility of more complex forms of unconscious learning, as well as the methodological and statistical artifacts which can contribute to falsely assuming absence of awareness, it is imperative to rigorously address the question of the feasibility of unconscious instrumental conditioning. Here, we present two experimental attempts to conceptually replicate Pessiglione and colleagues' unconscious instrumental conditioning task. The original task constituted trace conditioning – participants learned the association between the subliminally presented stimulus and its outcome (presented supraliminally as monetary reinforcement), with an intervening gap of 3 seconds, during which action could have been executed (approach) or not (avoid). Experiment 1 is closely based on this paradigm, similarly adopting trace conditioning, and introducing minor methodological changes to enhance the sensitivity of the task. Experiment 2 seeks to evaluate whether unconscious instrumental conditioning can be achieved under conditions previously found to be more favourable for unconscious learning, namely where there is no temporal delay between the stimulus and outcome. Hence, Experiment 2 is a delay instrumental conditioning task, retaining the general structure of the original task, but ensuring that stimulus presentation overlaps with both the response, and the delivery of the reinforcement. To allow for the overlap of all three (which is not possible in just the visual domain), we replaced the secondary visual (monetary) reinforcement with primary reinforcement (appetitive and aversive tastes), which have been previously shown to be highly effective in achieving learning in humans (Birbaumer, Klucken, Angrilli, Andreatta, & Pauli, 2015; Martin-Soelch, Linthicum, & Ernst, 2007; Parkinson, Cardinal, & Everitt, 2000). The strong spatiotemporal overlap was designed in order to maximise the chance of observing conditioning. Notably, primary reinforcements might also offer a stronger

incentive to learn. In both experiments, we used Bayes Factors to determine whether the results sensitively support the null hypothesis (i.e. absence of learning) or the alternative hypothesis (presence of learning).

5.2. Experiment 1

Experiment 1 sought to conceptually replicate the unconscious instrumental conditioning task (Pessiglione *et al.*, 2008), where participants learned to associate predictive cues with reward and punishment, and to adjust their behaviour accordingly by responding (Go) only to rewarding cues, and refraining from a response (NoGo) to punishing cues. Here, we retained the structure of the task, and introduced a number of changes to increase the sensitivity of the paradigm. Firstly, we used a staircased perceptual discrimination task, allowing to determine the sub-threshold stimulus presentation duration for each participant individually (in contrast to keeping the duration at 33 or 50ms, as in Pessiglione *et al.*). We exploited modern screen technologies with higher refresh rates, allowing to present stimuli closer to participants' actual thresholds. Secondly, stimuli were presented in lower contrast, in order to maximise the potential exposure durations without being consciously perceived (following Scott *et al.*, 2018). Thirdly, we disposed of the neutral cue, which was not predictive of any outcome and used in the original experiment primarily for the purpose of comparisons in functional neuroimaging.

Finally, we used a trial-by-trial awareness check, allowing to monitor the levels of awareness on an ongoing basis, and analyse only the trials which were genuinely outside of conscious awareness. In the original Pessiglione *et al.* study, awareness was assessed with a separate measure, and participants' performance on the task of interest is considered unaware if they perform at chance in the separate awareness check (this has also been a standard approach in the field). However, this does not account for the possibility of transient, undetected moments of awareness during the task being misattributed to unconscious knowledge. As such, a trial-by-trial check allows us to measure awareness in a more immediate fashion (Berry & Dienes, 1993; Newell & Shanks, 2013). Because our staircase should sensitively settle on sub-threshold conditions for each participant, the number of conscious trials to be excluded should be small. While excluding individual aware trials (or participants) can lead to regression to the mean (Shanks, 2017), our analyses have demonstrated its effect to be negligible if only a small proportion of trials or participants are excluded. Here, we will adopt a threshold of 25%, which corresponds to a maximum conscious contamination of a type I d' of 0.05 (see supplementary material in Skora, Livermore, Dienes, Seth, & Scott (*accepted*) for a method to determine the maximum proportion of conscious knowledge inadvertently

contributing to observed unconscious performance at different proportions of trials classified as unconscious).

5.2.1. Method

5.2.1.1. Participants

46 participants (15 males, 1 unknown) with a mean age of 23.7 years ($SD = 9.84$) were recruited for participation via the University of Sussex online recruitment system. Sample size was determined using the Bayesian Stopping Rule, with data collection continuing until a sensitive result was obtained in the main (conditioning) task (see *Analyses and Results* for detail). All participated in exchange for course credit. All reported having normal or corrected-to-normal vision, and no current or history of cardiac or neurological illness. Ethical approval was granted by the School of Psychology ethics committee at the University of Sussex, and the study was conducted in accordance with the Declaration of Helsinki.

5.2.1.2. Stimuli and Materials

The experiment was implemented in Matlab 2017b (MathWorks, 2017), running Psychophysics Toolbox (Brainard, 1997). All stimuli were presented on a Samsung 2233RZ LCD monitor (1680 by 1050 pixels) with a refresh rate of 120Hz, with the aim of ensuring fast and precise stimulus presentation in line with previous recommendations (Wang & Nikolic, 2011).

The target stimuli included 12 neutral symbols (6 vertically symmetrical, 6 asymmetrical) obtained from Agathodaimon font in the main task, and two circular shapes in the perceptual discrimination task used for threshold finding (see Appendix 1). All were 180x180 pixels in size, and presented in light grey (RGB: 217 217 217) on white background. The stimuli were forward and backward masked with black and white noise masks, also 180x180 pixels in size, with block size of 3x3 pixels. The forward and backward masks were different, but kept consistent between blocks and participants. Low contrast cues and the type of mask were deliberately chosen in order to increase the duration of presentation without conscious awareness, following Scott, Samaha, Chrisley and Dienes (2018).

5.2.1.3. Procedure

Threshold setting.

Participants were seated with their chin on a chin rest placed at 45cm distance from the screen. Each session began with the threshold of visual awareness determined individually for each participant using a masked perceptual discrimination task. Each trial began with a fixation

cross (500ms), followed by a mask (300ms), a target cue (either a symmetrical or an asymmetrical circular shape, starting at 600ms), and another mask (300ms). After each sequence, participants were asked to determine whether the target cue was symmetrical or asymmetrical by pressing corresponding arrows. Next, they were asked to assess whether they had any confidence in their judgment, or if they were guessing, also using corresponding arrows (following Scott et al., 2018). They were explicitly instructed to indicate 'some confidence' if they had any degree of confidence, and 'total guess' only if they felt they did not see the cue and were responding randomly. Each time a correct response was made with confidence, the display duration of the target cue was reduced by 50ms on the following trial. When a duration of 100ms was reached, or the first guess response was made, the display duration returned to the previous level (+50ms), and subsequently reduced in 8.35ms steps on the following trials, corresponding to a single screen refresh duration for a 120Hz screen. A reduction in exposure duration continued to be made after each non-guess response but not after guess responses. This process continued until participants indicated guessing on six consecutive trials, regardless of the accuracy of responses. The cue display duration on those trials was set as their individual unconscious threshold. The average display duration was 208ms ($SD = 64ms$).

Conditioning task.

The main task was adapted from the subliminal instrumental conditioning task approximating a deterministic instrumental learning scenario used previously (Pessiglione et al., 2008), in which participants learn to approach (Go response) or avoid (NoGo) the presented stimuli. Here, each trial consisted of a fixation cross (500ms), mask (300ms), target cue (display duration determined in the perceptual discrimination task), and mask (300ms), followed by a decision prompt in the form of a question mark, during which participants had 2 seconds to make a response (see *Fig.16.* for task chronology). While the question mark was presented, participants could decide if they want to 'approach' the cue by pressing the spacebar (Go) or 'avoid' the cue by refraining from pressing (NoGo). Pressing the space bar (Go) indicated a decision to take the risk, following which the participant could win 1 token (golden token displayed on the screen) or lose 1 token (a red cross over the golden token displayed), depending on the type of cue presented between the masks. Not pressing the space bar (NoGo) indicated a safe choice, which always resulted in a null outcome (greyed-out token). Participants were instructed that they should follow their instincts or gut feelings in making the decisions, as they should not be consciously aware of the type of cue presented.

Following the feedback presentation, participants were asked to report whether they thought the cue was vertically symmetrical or asymmetrical, and their confidence in that

judgment on a binary scale (some confidence or total guess). Both responses were made with the arrow keys before proceeding to the next trial. Following 3 correct and confident responses (indicating awareness), display duration was reduced by another single screen refresh (8.35ms). Participants were explicitly instructed that cue symmetry was unrelated to its rewarding or punishing nature. They were also shown examples of vertical symmetrical and asymmetrical stimuli on a different sample of images.

The task contained 2 blocks of 100 trials, each with 50 rewarding and 50 punishing trials in randomised order. For each block, two target cues (one symmetrical, one asymmetrical) from the pool of 12 were randomly assigned to be rewarding or punishing, without replacement, in order to ensure each block contained a novel pair.

Note that the task was conducted as a control condition of a larger experiment (Skora, 2020; doctoral thesis, also <https://osf.io/b34uv/>), not reported here, where 4 additional blocks were conducted with an added manipulation (600 trials in total). The control blocks reported here are fully independent from the manipulation.

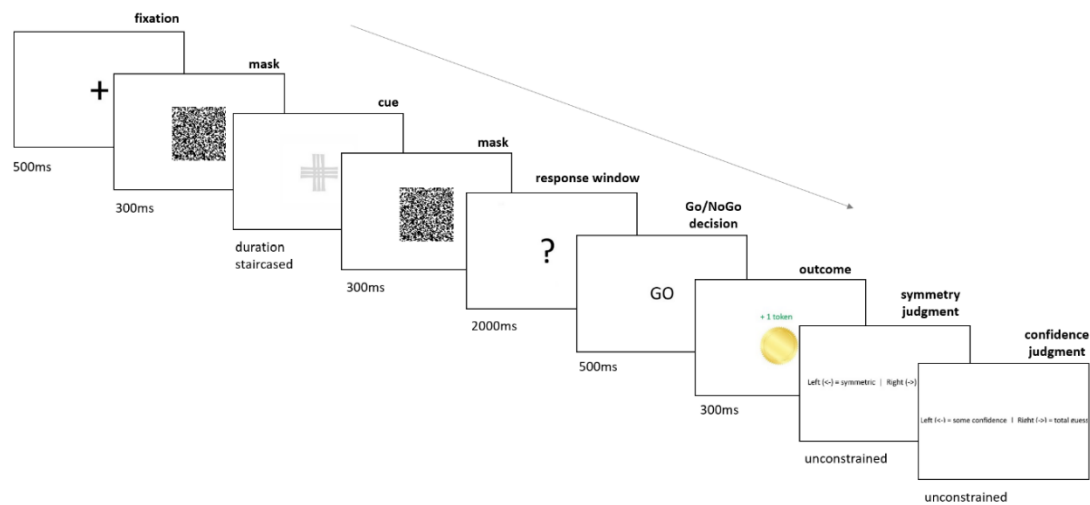


Figure 16. A single trial sequence. Participants are presented with a single cue (rewarding or punishing), rendered unconscious with forward-backward masking. After cue presentation, they are requested to make an approach (Go) or avoid (NoGo) response, following their instincts. In the event of a Go response, the response window closes immediately, but stays open for 2 seconds if no response is recorded. Here, a participant executed the Go response following a rewarding cue, and were rewarded with one token. Each trial ends with a cue symmetry judgment and confidence judgment to assess cue awareness.

5.2.2. Analyses and Results

5.2.2.1. Exclusion criteria

In order to ascertain that analyses are conducted only on trials that were truly unconscious, all individual trials where participants made a correct symmetry judgment with confidence were marked as aware and excluded (12% of all trials). Seven participants (15%) who were aware on more than 25% (50 of 200) of all trials were excluded from further analysis. Two participants were excluded due to failing to make any Go responses, yielding a final sample of 37 participants.

For the exploratory reaction time (RT) analysis, RTs under 100ms or over 2 standard deviations from individual means were excluded (8%). One participant who lost more than 25% of RT trials (50) to those criteria was excluded, yielding a sample of 36 participants for the RT analysis. Note that NoGo trials yielded no RTs, thus were not included in the analysis.

5.2.2.2. Evidence of learning: Performance in the conditioning task

Type I d' (a Signal Detection Theoretic measure of sensitivity to signal versus noise; Stanislaw & Todorow, 1999) was computed for the main conditioning task, treating Go responses to rewarding cues as Hits, and Go responses to punishing cues as False Alarms. The resulting measure of sensitivity can be taken as evidence of successful learning (i.e. discrimination between the cues) if it is significantly above 0.

At the group level, d' scores were entered into a one-sample t-test against 0, which indicates no ability to discriminate the stimuli (no sensitivity between signal versus noise, akin to chance performance). A Bayes Factor (B; Dienes, 2015, 2016) was computed for the difference, with a half-normal distribution, mean specified as 0, and d' of 0.7 as the SD of the mean (corresponding to the expected effect size, based on Pessiglione et al., 2008). In line with the Bayesian Stopping Rule (Dienes, 2015b), data collection continued until a sensitive result was found in support of either H_0 (absence of learning; by convention indicated by a B smaller than 0.3) or H_1 (presence of learning; indicated by a B larger than 3).

Total d' was not significantly different from 0 ($M = 0.02$, $SE = 0.05$; $t(36) = 0.46$, $p = 0.649$, $B_{H(0,0.7)} = 0.10$; see *Figure 17*). As such, we found evidence for the absence of unconscious instrumental conditioning.




Figure 17. Left: average type I d' with individual data points. **Right:** average proportions of correct responses (Go to rewarding and NoGo to punishing cues) across the length of the trial block. Ribbon represents a regression line (90% CI). Participants' choices throughout the duration of the block remained erratic and around chance level, failing to display any evidence of learning the correct decisions.

5.2.2.3. Exploratory analysis: Reaction times

An $RT_{\text{difference}}$ index was computed by subtracting RTs to rewarding cues from RTs to punishing cues. As such, positive values indicate that participants took a longer time to respond to punishing cues than to rewarding cues, in line with RT-oriented indicators of learning (e.g. Atas, Faivre, Timmermans, Cleeremans, & Kouider, 2014). Zero indicates that there was no difference between the two.

The index was entered into a one-sample t-test against 0. B was computed with a half-normal distribution, mean specified as 0, and a value of 34ms as the SD of the mean (obtained from a past study which found a RT difference in the absence of performance effects in a similar task (Atas et al., 2014)). The t-test revealed that $RT_{\text{difference}}$ was not significantly different from 0 ($M = 14.06$, $SE = 17.81$, $t(35) = 0.79$, $p = 0.435$, $B_{H(0,34)} = 0.90$).

5.2.3. Conclusions of Experiment 1

Experiment 1 attempted to conceptually replicate the unconscious instrumental (trace) conditioning task, used by Pessiglione *et al* (2008). The results indicate that no unconscious instrumental conditioning was achieved – when participants had no conscious awareness of the rewarding or punishing nature of the stimuli, they were unable to discriminate between them and adjust their behaviour accordingly (i.e. choose to Go or NoGo, respectively). Bayes Factors were used to supplement null hypothesis significance tests, allowing us to assert that the null

result obtained indicates absence of learning. This result is a failure to replicate the original Pessiglione et al (2008) finding.

In light of previous research and the theoretical contributions about the limits of unconscious processing (e.g. Mudrik et al., 2014), we considered it plausible that the trace conditioning paradigm used here might make it harder to achieve learning, even if it is possible without cue awareness. In the present task, participants had to learn from two events – the stimulus and its consequence – separated by an interval of up to 2 seconds, during which they may or may not execute an action. If the large temporal interval is preventing the unconsciously presented predictive cue from being integrated with subsequent information, shortening the delay might offer a solution. We address this possibility in Experiment 2 by amending the task to constitute delay, rather than trace, instrumental conditioning.

5.3. Experiment 2

Experiment 2 sought to replicate the unconscious instrumental conditioning task again, this time shortening the stimulus-outcome delay. To achieve this, we used a delay conditioning paradigm, while retaining as much of the structure of the original task as possible. As such, we employed the same Go/NoGo task, where participants learned to associate predictive cues with reward and punishment, and to adjust their behaviour accordingly by responding (Go) only to rewarding cues, and refraining from a response (NoGo) to punishing cues. In contrast to Experiment 1, the response and reinforcement delivery occurred during (rather than following) stimulus presentation, constituting delay conditioning. Prolonged stimulus presentation was enabled through continuous flash suppression (CFS; Tsuchiya & Koch, 2005). CFS uses interocular suppression to render the target stimulus unconscious by presenting it to the non-dominant eye, while simultaneously presenting a high-contrast flashing pattern (called the Mondrian pattern) to the dominant eye. Instead of monetary rewards, Experiment 2 used appetitive and aversive tastes as primary rewards, allowing for extended delivery to overlap with stimulus presentation. In a fashion identical to Experiment 1, a trial-by-trial awareness check was used, with the same exclusion threshold of a maximum of 25% of individual trials.

5.3.1. Method

5.3.1.1. Participants

39 participants (5 male) with a mean age of 21 years ($SD = 5.2$) were recruited for the study via the University of Sussex online recruitment system. Sample size was determined with the Bayesian Stopping Rule (see section 3.2.2. for detail). All participated in exchange for course

credit. Participants were required to have normal or corrected-to-normal binocular vision, no current or history of neurological illness and/or diabetes, and no allergies to gluten, lactose, milk or chocolate. Further exclusion criteria included dislike of chocolate milkshakes and/or smoking more than five cigarettes a day, as those factors might affect responsiveness to tastes. Participants were requested to refrain from eating and drinking (apart from water) for 2 hours prior to the experiment. Ethical approval was granted by the School of Psychology ethics committee at the University of Sussex, and the study was conducted in accordance with the Declaration of Helsinki. The study was pre-registered on Open Science Framework at <https://osf.io/jt8m9>.

5.3.1.2. Stimuli and Materials

The experiment was implemented in Matlab 2018b (MathWorks, 2018) with the Cogent toolbox (UCL LoN, 2003). All stimuli were presented on a Dell monitor (1280 by 1024 pixels), with a refresh rate of 60Hz⁷. Binocular presentation was ensured with a mirror stereoscope. Responses were collected with a standard keyboard.

The target stimuli were three geometrical shapes, including one symmetrical and two asymmetrical shapes (generated by randomly overlaid shapes; see Appendix 3). All stimuli were 120 by 120 pixels in size, presented in dark grey (RGB: 80,80,80) on lighter grey (RGB: 128,128,128) background.

The solution used as the appetitive stimulus was a chocolate milkshake (Shaken Udder Chocolush). The aversive stimulus was a mixture of food-grade quinine monohydrochloride dihydrate, at the FDA-approved maximum concentration (80mg/l; Sigma Aldrich), and standard table salt (NaCl; 12.5g/l). Custom-made syringe pumps were used to administer the solutions to participants' mouths through two clean silastic tubes (one per solution), inserted into the corners of the mouth and placed in the middle of the tongue. The pumps were triggered directly from the Matlab task script, delivering a 0.5ml bolus of the corresponding solution at a time.

5.3.1.3. Procedure

Conditioning task.

The conditioning task followed a similar Go/NoGo procedure as outlined in Experiment 1, amended for presentation under continuous flash suppression (CFS; Tsuchiya & Koch, 2005). During the unconscious presentation of each cue, participants were asked to decide if they want

⁷ A 120Hz screen, as used in Experiment for the purpose of sensitive threshold finding, was no longer necessary for a CFS setup.

to ‘approach’ the cue by pressing the spacebar (Go) or ‘avoid’ the cue by refraining from pressing (NoGo). Approaching the cue resulted in the delivery of a 0.5ml bolus of the corresponding solution – appetitive if they approached the rewarding cue, and aversive if they approached the punishing cue. A NoGo response was again considered a safe choice, which resulted in no taste outcome. In order to further incentivise learning and ensure participants do not default to the safe choice, each correct response additionally resulted in winning a token. Participants were instructed to maximise their token winnings for an extra prize draw entry.

The cues were presented unconsciously using CFS, with the cue itself presented to participants’ non-dominant eye, and a colourful Mondrian pattern flashing at the rate of 10Hz presented simultaneously to the dominant eye through a mirror stereoscope, placed atop a chinrest at a 50cm distance from the screen. Ocular dominance was established prior to beginning with a standard Miles test (Miles, 1930). The unconscious presentation of each cue started with a fade-up period of 500ms. A further 500ms after the cue reached full contrast, the fixation dot (overlaid on displays to both eyes) changed colour from black to white, indicating that response was required. The response window remained open for 1000ms. When the participant executed a Go response, a pump was triggered to deliver the corresponding taste outcome (appetitive or aversive, depending on the type of the cue). The taste delivery continued for 2000ms, while the cue remained on-screen (concurrently with the Mondrian pattern) for 1000ms, ensuring a full second of temporal overlap between the cue and the outcome (reward or punishment), constituting delay conditioning. See *Figure 18* for an illustration of the trial sequence.

Following the taste delivery, participants were asked to report whether the cue was symmetrical or asymmetrical, and their confidence in that judgment on a binary scale (some confidence or total guess). Both responses were made with the arrow keys before proceeding to the next trial.

The task consisted of 1 block of 100 trials, with a total of 50 rewarding and 50 punishing trials in a randomised order, as well as a short practice round. The task used two stimuli (one symmetrical, one asymmetrical), out of which one was randomly assigned to be appetitive, and one to be aversive, for each participant.




Figure 18. Single trial sequence in Experiment 2, illustrating the binocular presentation with the non-dominant eye (left) receiving the cue (predictive of reward or punishment), and the dominant eye receiving the Mondrian pattern. Following the cue reaching full contrast, participants are required to make an approach (Go) or avoid (NoGo) response, and are rewarded with 0.5ml of chocolate milkshake or punished with a 0.5ml of a quinine/saline solution accordingly. The taste delivery overlaps with cue presentation for 1000ms and continues for another 1000ms. Each trial ends with a cue symmetry judgment and confidence judgment to assess cue awareness on the trial.

Liking and familiarity ratings.

Following the conditioning task, the two cues used and the third, previously unseen cue were presented on the screen in a randomised order in a single column, alongside corresponding VAS scales. First, participants were asked to rate their liking of each stimulus, on a scale from 0 (not at all) to 10 (extremely). Next, they were asked to rate their familiarity with each, again on the same 0-10 scale. These were not analysed in the context of this study. Following the ratings, participants were thanked and debriefed.

5.3.2. Results

5.3.2.1. Exclusion criteria

The exclusion criteria were identical to Experiment 1. 31.8% of all trials were marked as aware and excluded. 20 (51%) participants who were aware on more than 25% (25) of all trials were classified as aware. The remaining 19 participants were found to be aware on 9.6% of all trials. The proportion of participants classified as aware on more than 25% of the trials was surprisingly high. We suspect that the length of stimulus presentation under CFS (2 seconds) may have made it easier for some participants to see the stimuli on some occasions. Once seen, the stimuli are typically easier to see on subsequent trials, which is what likely resulted in the high rate of aware trials. Because of the risk of regression to the mean following post-hoc selection of a large number of participants (Shanks, 2017), we found that the maximum proportion of affected trials could be 55%, corresponding to a d' of 0.99 (computed with the method used in Experiment 1). Excluding such a large proportion of participants could lead to a false-positive result (presence of learning), where some conscious knowledge is mistakenly included in the measure of unconscious knowledge, in this case overinflating the d' by up to 0.99. As we will show below, this possibility is not relevant here, because the result shows absence of learning (even after allowing for the contribution of conscious knowledge).

For maximum rigour and transparency, analysis was conducted on both the unaware sample as planned, and the full sample (those classified as aware and unaware), after excluding individual aware trials.

For the reaction time (RT) analysis, RTs under 100ms or over 2 standard deviations from individual means were excluded (4.49%). No participant lost more than 12% of their trials to this criterion. Again, NoGo trials were not included in the analysis.

5.3.2.2. Evidence of learning: Performance in the conditioning task

At the group level, d' scores were entered into a one-sample t-test against 0, which indicates no ability to discriminate between the stimuli (chance performance). B was computed for the difference, with a half-normal distribution, mean specified as 0, and d' of 0.7 as the SD of the mean (corresponding to the expected effect size from Pessiglione et al., 2008). In line with the Bayesian Stopping Rule, data collection continued until a sensitive result was found either in support of the H_0 (absence of unconscious learning; $B < 0.3$) or the H_1 (evidence for unconscious learning; $B > 3$).

For the sample with aware participants excluded ($N = 19$), total d' was not significantly different from 0 ($M = -0.10$, $SE = 0.10$; $t(18) = -1.00$, $p = 0.329$, $B_{H(0,0.7)} = 0.07$; see *Figure 19*). As such, we found evidence for the absence of unconscious instrumental conditioning. For the full sample (participants classified both as aware and unaware; $N = 39$), total d' was also not significantly different from 0, but the B was insensitive ($M = 0.13$, $SE = 0.14$; $t(38) = 0.97$, $p = 0.339$, $B_{H(0,0.7)} = 0.51$).

5.3.2.3. Exploratory analysis: Reaction times

The $RT_{\text{difference}}$ index was computed again by subtracting RTs to rewarding cues from RTs to punishing cues (with positive values indicating that participants took a longer time to respond to punishing than to rewarding cues), and entered into a one-sample t-test against 0. B was computed with a half-normal distribution, mean specified as 0, and a value of 34ms as the SD of the mean (obtained from a past study which found a RT difference in absence of performance effects in a similar task, Atas et al., 2014). For the sample with aware participants excluded ($N = 19$), the t-test revealed that $RT_{\text{difference}}$ was not significantly different from 0 ($M = 7.98\text{ms}$, $SE = 9.69$, $t(18) = 0.83$, $p = 0.420$, $B_{H(0,34)} = 0.59$). The same was true for the full sample ($N = 39$; $M = 2.84$, $SE = 15.69$, $t(38) = 0.18$, $p = 0.857$, $B_{H(0,34)} = 0.48$).

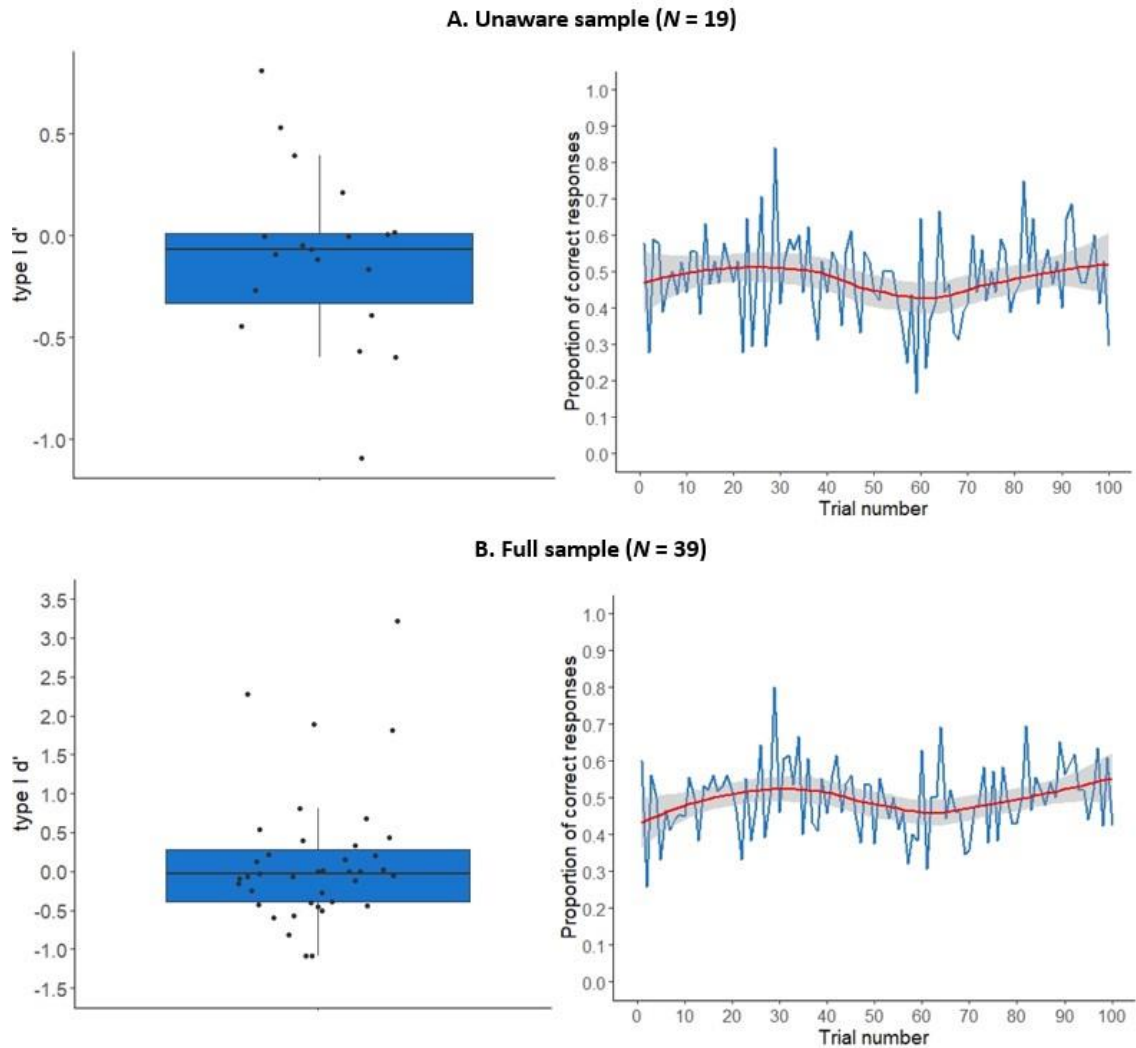


Figure 19. Panel A, left: average type I d' with individual data points (unaware sample, $N=19$). Panel A, right: average proportions of correct responses (Go to rewarding and NoGo to punishing cues) across the length of the trial block. Ribbon represents a regression line (90% CI). Panel B, left: average type I d' with individual data points (full sample, $N = 39$). Panel B, right: average proportions of correct responses (Go to rewarding and NoGo to punishing cues) across the length of the trial block. Ribbon represents a regression line (90% CI). Participants' choices throughout the duration of the block remained erratic and around chance level, failing to display any evidence of learning the correct decisions. Note that individual aware trials were excluded.

5.3.3. Conclusions of Experiment 2

Experiment 2 made a second attempt to replicate unconscious instrumental conditioning (Pessiglione et al., 2008), this time with a delay conditioning paradigm. The results demonstrate that no unconscious instrumental conditioning was achieved, even when the temporal delay between the presentation of the stimulus and its consequence overlapped in time. Again, when participants had no conscious awareness of the stimuli, they were unable to adjust their

responses accordingly (i.e. choose to Go in response to rewarding stimuli and NoGo to punishing stimuli). Bayes Factors allowed us to assert that the null result obtained indicates lack of learning. This result constitutes a failure to conceptually replicate unconscious instrumental conditioning.

5.4. General Discussion

The extent to which complex forms of learning, such as instrumental conditioning, are possible without conscious awareness is still an unresolved question. To our knowledge, evidence for unconscious instrumental learning has been limited (Pessiglione et al., 2008; Mastropasqua & Turatto, 2015) and disputed (Reber et al., 2018). Yet, instrumental behaviour is vital and one of the oldest forms of adaptive behaviour (both phylogenetically and ontogenetically). As such, we consider the extent to which it can or cannot proceed unconsciously an important question to settle. In the present paper, we revisited unconscious instrumental conditioning in two modes of conditioning: trace (Experiment 1) and delay (Experiment 2). We enhanced the sensitivity of the paradigms, and manipulated the temporal delay between stimulus and outcome – previously identified as a key aspect of unconscious learning.

Experiment 1 attempted to revisit the original Pessiglione *et al* (2008) finding obtained in an unconscious instrumental (trace) conditioning task. Our results show evidence for the absence of unconscious instrumental conditioning – when participants had no conscious awareness of the rewarding or punishing nature of the stimuli, they were unable to discriminate between them and adjust their behaviour accordingly (i.e. choose to Go or NoGo, respectively). Bayes Factors allow us to assert that the null result obtained indicates an absence of learning and not simply insensitivity. This result is a failure to replicate the original Pessiglione et al (2008) finding.

Following the failure to demonstrate unconscious learning, we hypothesised that the large temporal gap between the stimulus and its outcome might hinder the ability to learn. Past theoretical and empirical work has shown that unconscious associative learning or priming is more likely to occur at shorter temporal intervals (e.g. van Gaal et al., 2012), or when the events are overlapping in time (delay conditioning; Clark & Squire, 1998). Thus, Experiment 2 made a second attempt to investigate unconscious instrumental conditioning, this time with a delay conditioning paradigm. Again, the results show evidence for absence of unconscious instrumental conditioning, even when the stimulus, response and reinforcement temporally overlapped. As before, when participants had no conscious awareness of the stimuli, they failed to adjust their responses accordingly (i.e. choose to Go in response to rewarding stimuli and

NoGo to punishing stimuli). Bayes Factors again allowed us to assert that the null result obtained indicates absence of learning. As in Experiment 1, this result constitutes a failure to conceptually replicate unconscious instrumental conditioning.

Despite extensively demonstrated effects in simpler associative learning or priming tasks, it appears that unconscious instrumental learning cannot be achieved even if the stimulus-outcome delay is shortened to the point of total overlap. As such, perhaps the problem is not the length of the delay. Instead, it may in fact be the instrumental component of the task itself. Previous theoretical contributions have already suggested that conscious access might be necessary for higher-order processes, including selective decision-making (Baars, 2002; Dehaene & Changeux, 2011; Lamme, 2006). On those accounts, consciousness allows for information exchange between distinct cognitive modules through long-lasting, long-range recurrent interactions between brain areas. Hence, it might be key for flexible and lasting information processing strategies. Regardless of the temporal intervals between events, instrumental learning is a complex process, necessitating integration of information across distinct modalities involved in processing the visual input, extracting its predictive value from past interactions, deploying the response, processing the reinforcement, and comparing the expected outcome with actual outcome to update the representations of stimulus value. Such integration might simply not be possible when the stimulus is presented unconsciously, as it may not have the capacity to evoke such a broad range of activity. Corroborating this interpretation, past work has already pointed out that unconscious stimuli are less likely to evoke long-term behavioural adaptations than conscious stimuli (e.g. in post-error slowing or conflict adaptation; (de Lange, van Gaal, Lamme, & Dehaene, 2011; Kunde, Reuss, & Kiesel, 2012; Reber et al., 2018; Travers et al., 2018; van Gaal et al., 2012)). Instrumental learning might constitute a similar case.

If instrumental learning does indeed require conscious access to proceed, the earlier results supporting unconscious instrumental learning may have resulted from transient moments of awareness during the learning task contributing to the unconscious knowledge. Those may have been undetected due to the lack of immediate, sensitive measures (Berry & Dienes, 1993; Newell & Shanks, 2013; Shanks & St. John, 1994), such as trial-by-trial awareness checks. In the two experiments reported here, it was evident that most participants were aware of the nature of the stimuli on at least some of the trials, and that there are individual differences in those proportions. Capturing aware trials and excluding them ensured that the trials analysed were genuinely unconscious, and constituted a considerable improvement over previously employed methods. Crucially, it also avoids the potential for anything other than a negligible effect from regression to the mean (Shanks, 2017), the extent of which we computed for both experiments.

Nonetheless, it is worthwhile to address whether our methodology, although more rigorous, made it more difficult for learning to happen. While a trial-by-trial awareness check was essential to rigour, it introduced a small, variable time lag (around 1 second) between the trials. This might have further disrupted the process of temporal integration across trials, affecting the long-term consolidation of the stimulus-outcome associations learned within them. However, it is the integration of events within the trial that is vital to learning in the first place. As such, we consider any plausible interference of the inter-trial delay to be minor.

Another important consideration is the large number of participants displaying awareness of stimuli in Experiment 2. According to our criteria laid out in Experiment 1, we allowed a maximum of 25% of aware trials to be excluded per participant. This corresponds to the maximum level of conscious contamination of a type I d' of 0.05 (Skora et al., accepted). Yet, in Experiment 2 a half of the sample was found to be aware on more than 25% of the trials, and hence excluded. Exclusion of this surprisingly large sample poses a risk of regression to the mean, where the remaining sample might show a false-positive (evidence of learning) purely as a statistical artefact (Shanks, 2017). In the interest of maximum transparency, we analysed both samples (with individual aware trials excluded), showing that both exhibited absence of learning – thus rendering the risk of a false-positive on the remaining unaware sample irrelevant.

The surprisingly large proportion of participants showing awareness of the stimuli under CFS on over 25% of trials is likely to have stemmed from the long duration of each CFS trial (2 seconds), which was necessary to allow for cue presentation to temporally overlap with the response and the response-contingent taste reinforcement. With known individual differences in CFS (e.g. Gayet & Stein, 2017), it is plausible that some people may have seen the cue presented under interocular suppression at this duration easier than others. Once seen, it typically becomes easier for the stimulus to be perceived on subsequent occasions. This could be an issue for consideration for future research. Here, we were able to reliably account for that in the analyses presented.

To summarise, we used a trace conditioning paradigm (Experiment 1, following a method used by Pessiglione et al., 2008) and a delay conditioning paradigm (Experiment 2) to investigate the feasibility of unconscious instrumental conditioning with different temporal delays. We provide rigorous evidence demonstrating that instrumental conditioning cannot be achieved without stimulus awareness, corroborating the emerging evidence that complex types of learning, including instrumental conditioning, require conscious awareness. This provides strong support for the theoretical proposals that unconscious stimuli fail to bring about long-term behavioural adaptations, and that consciousness might be necessary for complex

processes, where selective action is required. As such, this result contributes to mapping out the boundaries of unconscious processing in adaptive behaviour.

6.

SENSORIMOTOR PREDICTIONS SHAPE REPORTED CONSCIOUS VISUAL EXPERIENCE

Abstract

Accounts of predictive processing propose that conscious experience is influenced not only by passive predictions about the world, but also by predictions encompassing how the world changes in relation to our actions – that is, on predictions about sensorimotor contingencies. We tested whether valid sensorimotor predictions, such as learned associations between stimuli and actions, shape reports about conscious visual experience. Two experiments used instrumental conditioning to build sensorimotor predictions, linking different stimuli with distinct actions. Conditioning was followed by a breaking continuous flash suppression (b-CFS) task, measuring the speed of reported breakthrough for different pairings between the stimuli and prepared actions; including those congruent and incongruent with the trained sensorimotor predictions. In Experiment 1, counterbalancing of the b-CFS task was achieved by repeating the same action within each block but having them differ across the two. Experiment 2 sought to increase the predictive salience of the actions by avoiding the repetition within blocks.

Mixed-effects modelling assessed the speed of reported conscious perception in the b-CFS task. In Experiment 1, breakthrough times were numerically shorter for congruent than incongruent pairings but Bayesian analysis supported the null hypothesis of no influence from the sensorimotor predictions. In Experiment 2, reported consciousness perception was significantly faster for congruent than for incongruent pairings. A meta-analytic Bayes factor combining the two experiments provided further confirmation of this effect. Hence, we provide evidence for a key implication of the action-oriented predictive processing approach to conscious perception, namely that sensorimotor predictions shape our conscious experience of the world.

6.1. Introduction

A growing body of experimental work, rooted in the predictive processing framework (Clark, 2013; Hohwy, 2013, 2020; Rao & Ballard, 1999), shows that perceptual experiences are influenced by beliefs or predictions about the world. Valid predictions have been shown to facilitate access to visual consciousness (Meijs, Slagter, De Lange, & Simon Van Gaal, 2018; Melloni, Schwiedrzik, Muller, Rodriguez, & Singer, 2011; Pinto, van Gaal, de Lange, Lamme, & Seth, 2015), reduce repetition suppression (Summerfield, Trittschuh, Monti, Mesulam, & Egner, 2008), improve metacognition (Sherman, Seth, Barrett, & Kanai, 2015), and aid interpretation under perceptual ambiguity (Aru, Rutiku, Wibral, Singer, & Melloni, 2016; Panichello et al., 2013).

Within the predictive processing framework, predictions are instantiated by probabilistic generative models, encoded in cortical hierarchies. Incoming sensory signals, such as visual input, are compared against descending predictions to give rise to prediction errors (PEs) at each hierarchical level of processing. PEs can then be suppressed by updating predictions (corresponding to perceptual inference). In this framework, conscious sensory experience has been proposed to reflect the perceptual prediction that best suppresses PEs across hierarchical levels (e.g. Hohwy, 2013; Seth et al., 2016). In other words, conscious experience is shaped by the posterior prediction that ‘best’ predicts the (hidden) causes of sensory signals. Hence, when perceptual predictions are valid, conscious access to the target stimulus should be facilitated, as reported in the above studies.

Recent accounts of predictive processing move beyond passive perception by putting *action* at the centre of the relationship between prediction and conscious experience. From this theoretical perspective, perceptual experience is influenced not only by passive predictions about the world, but more generally by predictions encompassing the coupling or contingency between actions and sensory signals – i.e., on predictions about sensorimotor contingencies (Clark, 2015a; O’Regan & Noë, 2001; Seth, 2014). According to this view, the predictions that best suppress PEs are not necessarily those which are the most veridical, but those which best support adaptive interactions with the world (Clark, 2015a, 2016; Seth, 2014, 2015; Tschantz, Seth, & Buckley, 2020). In this light, action emerges as not just an output, but as an integral part of our experience of the world. This perspective raises an important question: how is conscious perception shaped by the validity of predictions about sensorimotor contingencies, where such contingencies could reflect, for example, learned associations between a stimulus and an action?

Action is already known to modulate conscious perception in a number of ways. Voluntary action has been found to help disambiguate a bistable or otherwise ambiguous percept if the action is congruent with an aspect of that percept, e.g. when the direction of movement corresponds to the direction of moving dots (Beets et al., 2010; Di Pace & Saracini, 2014; Maruya et al., 2007; Mitsumatsu, 2009; Suzuki et al., 2019). A comparable effect has been found for proprioception, where correspondence between the target and proprioceptive information can bias or sharpen the visual percept (Butz et al., 2010; Salomon et al., 2013; Yon et al., 2018). For instance, access to visual awareness was shown to be faster for images of hands that matched the orientation of participants' own hands (Salomon et al., 2013). Crucially, none of these studies directly manipulated sensorimotor contingencies. As such, they cannot distinguish the effects of predictions about sensorimotor contingencies from the effects of congruency between stimuli and actions. This is a distinction of theoretical importance. It is the difference between examining the contribution to conscious experience made by maintaining a prediction about the action one will take in response to a stimulus (sensorimotor prediction), versus the contribution of, for example, a simple correspondence between the direction of an agent's action and the direction of some stimuli in the world (stimulus-action congruence).

We set out to test whether predictions about sensorimotor contingencies affect reportable conscious perception of visual stimuli. We developed a novel two-stage paradigm in which we operationalised sensorimotor contingencies as learned associations between visual stimuli and subsequent actions. We call these learned associations 'sensorimotor predictions'. We were then able to examine the effects of valid versus invalid sensorimotor predictions on speed of access to visual consciousness, quantified through the proxy of breakthrough time in continuous flash suppression (b-CFS; Jiang, Costello, & He, 2007; Tsuchiya & Koch, 2005); we return to the limitations of b-CFS for measuring the speed of conscious access later.

In the first stage of the paradigm, we leveraged instrumental conditioning to build sensorimotor predictions by linking distinct stimuli with specific actions. Two stimuli were arbitrarily associated with equally simple but distinguishable actions (an index finger or a little finger button press), and a third stimulus with no action. In the second stage, each stimulus was presented under CFS, where a dynamically flashing high-contrast pattern displayed to one eye is used to suppress visual awareness of the target stimulus displayed to the other eye, while participants prepared to respond with stimulus-contingent or non-contingent action. Thus, we test the extent to which maintaining the prediction of the relevant sensorimotor association (through response preparation) facilitates reported conscious access to the associated stimulus. Speed of access of each stimulus to visual consciousness was assessed using breakthrough time in CFS (breaking-CFS or b-CFS; Jiang, Costello, & He, 2007), measuring the time it takes for the

target to overcome interocular suppression and become consciously visible. If sensorimotor predictions affect conscious experience, then preparing an action conditioned on a specific relevant stimulus should engage a valid sensorimotor prediction, facilitating conscious experience of that stimulus and yielding faster reported breakthrough times.

Two experiments were conducted with the same conditioning task but with minor differences to the b-CFS task. In Experiment 1, the counterbalancing of action and stimuli was achieved at the block level by having the same prepared action for each trial within a b-CFS block and actions only differing between blocks. As each block contained one presentation of each of the four visual stimuli, this ensured that each action was paired with each stimulus on only one occasion. For example, if the action for block 1 was the index finger button press, participants would be using the index finger to respond when each of the four different visual stimuli broke through b-CFS. Then in block 2, they would be using their little finger to respond, again as each of the four different visual stimuli broke through b-CFS. However, requiring repetition of the action within a block may have limited the action's predictive salience. Experiment 2 avoided this repetition by varying the prepared action within each b-CFS block, while still ensuring counterbalancing of the stimulus-action pairing across the two blocks. A second minor variation made to the b-CFS task in Experiment 2 sought to separate two different contributions to response times. In Experiment 1 participants had to indicate the orientation of a single pixel line overlaid on the stimulus when it broke through suppression; the inclusion of this judgment seeks to reduce premature responses. In Experiment 2, this line orientation judgment was made after the initial reaction time response, thus avoiding adding noise to the measure of breakthrough time. Both Experiments were pre-registered on the Open Science Framework at <https://osf.io/ensba> (Experiment 1) and <https://osf.io/ez62m> (Experiment 2). All material, task code and data will be available at <https://osf.io/hpsju/>.

6.2. Experiment 1

6.2.1. Method

6.2.1.1. Participants

68 participants (11 male; M age = 20, SD = 2.38, range: 19-32) were recruited for the study via the University of Sussex online recruitment system, and an internal mailing list. All participated in exchange for course credit. Participants were required to have normal or corrected-to-normal vision, and no current or history of neurological illness. Ethical approval was granted by the Science and Technology Cross-School Research Ethics Committee at the University of Sussex, and the study was conducted in accordance with the Declaration of Helsinki.

6.2.1.2. Stimuli and Materials

The experiment was implemented in Matlab 2017b (MathWorks, 2017) with the Cogent2000 toolbox (UCL LoN, 2003). All stimuli were presented on a Dell monitor (1280 by 1024 pixels), with a refresh rate of 60Hz. Responses were collected with a standard keyboard.

The target stimuli included three sets of four cues. Each set contained four 90° rotations of the same symbol, a neutral, asymmetrical graphic, generated by randomly overlaid shapes (see *Fig 1.* for an example and Appendix 4 for the full set). All stimuli were 119 by 119 pixels in size (3.43° visual angle), presented in dark grey (RGB: 80,80,80) on lighter grey (RGB: 128,128,128) background. The Mondrian patterns for CFS were composed of coloured rectangles in a box 240 by 240 pixels in size (6.91° visual angle), flashing at a rate of 10Hz, and were dynamically created on each trial (see *Fig 20.*).

6.2.1.3. Procedure

Each participant completed two experimental conditions in a single experimental session, one where conditioning was completed consciously (where stimuli were clearly visible), and a second where conditioning was attempted subliminally (where stimuli were presented without conscious awareness). The conditions consisted of the same two main stages. Stage 1 was an instrumental conditioning task (2 blocks of 60 trials), and Stage 2 was a breaking-CFS task (2 blocks of 4 trials each). A brief set of practice trials was completed prior to Stage 1. Throughout both stages, participants were asked to look at the screen through a mirror stereoscope, fitted atop a chinrest at a 50cm distance from the screen. Ocular dominance was established prior to beginning the experiment with a standard Miles test (Miles, 1930).

Here, we report the conscious condition only, as we found strong evidence that the stimulus-action associations were not learnt in the unconscious condition (see Supplementary Material for the procedure and results of the unconscious instrumental conditioning task). Given that our primary interest here is to examine the effect that predictions about learned stimulus-action associations have on breaking-CFS, it would not be informative to analyse data from a condition where those associations were not acquired. The order of the conditions was randomised for each participant. The conditions were independent, and used different sets of stimuli (randomised).

Conditioning task.

The conditioning task was used to establish the sensorimotor (stimulus-action) associations. The task used three stimuli selected from the assigned set of four. One of the

stimuli (*cue A*) was paired with action A (an index finger button press), a second (*cue B*) with action B (a little finger button press), and a third (*no-action cue*) with no action. For practical reasons, the index finger button press was made on the left arrow, and the little finger button press on the right arrow. Note that while the keys being pressed differed, we consider the finger used to make the button press to be the conditioned action, not the button. Cues were randomly assigned to a given stimulus type for each participant. The order of presentation was randomised with an equal number of exposures to each cue occurring in each block of trials. Participants performed 120 conditioning trials in two blocks of 60, with a 1 minute break between the blocks.

In this task, both eyes were presented with the stimuli, i.e. there was no CFS and the stimuli were clearly visible. On each trial, stimulus presentation started with a fade-up period of 500ms. A further 500ms after the stimulus reached full contrast, the fixation dot (overlaid on top of the cue) changed colour, indicating that a response was required (see *Fig. 20* for a trial sequence). Fixation dot colours provided the distinction between action and no-action trials. For action cues, participants were instructed to respond with either action A or B (of their choosing), when prompted by the fixation dot changing colour from white to black. Following the action, positive or negative reinforcing feedback was delivered ('correct!' or 'wrong!' printed on the screen paired with a cash register or buzz sound, respectively) depending on the correspondence between the executed action and the cue presented, thus instantiating instrumental conditioning. For example, if *cue A* was presented and *action A* was executed, positive feedback was delivered, but if *cue A* was presented and *action B* was executed, negative feedback was delivered. For *no-action cues*, the fixation dot started as blue, and changed colour to red. This indicated to participants at the onset of the trial that no response will be required. In these no-action trials negative feedback was delivered if participants responded with either action A or B, and positive feedback if they correctly refrained from action. As such, the presented cue became associated with the 'no action' response.

On action trials, the cue disappeared as soon as participants made a response (a total of 1000ms + reaction time after cue onset). This ensured that the action was prepared and executed while the participant was exposed to the cue. On no-action trials, the cue disappeared as soon as the fixation dot changed which for these trials was set to be 1500ms after cue onset (500ms longer than in action trials). This additional 500ms sought to roughly equalise the total exposure duration for no-action trials, where there was no response time, and action trials where the exposure extended to include the participant's response time.

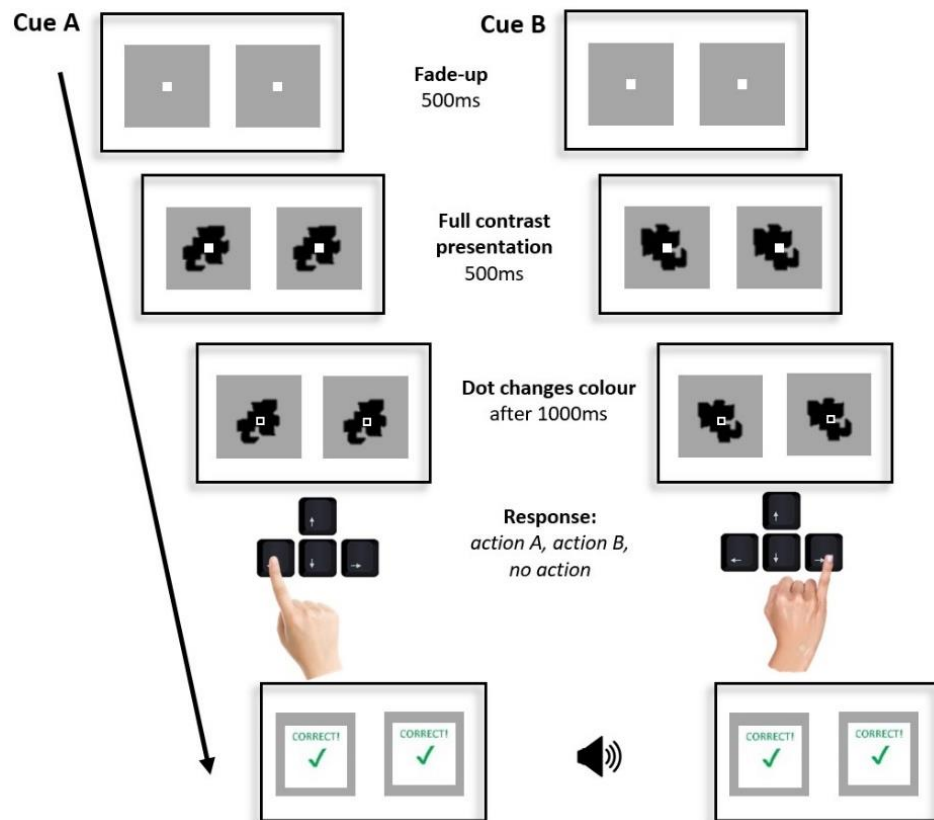


Figure 20. Instrumental conditioning task. Chronological screenshots depict a single trial sequence for the action cues – each panel shows the images shown to the left and right eyes. After 1000ms (including 500ms fade-up of the cue), the fixation dot changes colour from white to black (with a white border), and participants can execute the desired action (A, an index finger button press on the left arrow, or B, a little finger button press on the right arrow). For a no-action cue, the fixation dot would change colour from red to blue. Here, the action executed corresponded to the cue type, and the participant was rewarded with visual and auditory feedback.

Breakthrough task.

The conditioning task was immediately followed by a breaking-CFS task, using the set of cues from the conditioning task (*cue A*, *cue B*, *no-action cue*), as well as the fourth, novel cue, never previously seen (the last from the set). In the breakthrough task, the stimuli were presented under CFS, with the target cue presented to the non-dominant eye, while the dominant eye received a Mondrian pattern (see *Stimuli and materials*).

Each cue was presented with a randomly assigned horizontal or vertical line (1 pixel wide) overlaid on top of it. In order to quantify the time of breakthrough, participants were asked to report the orientation of the line using the “1” (horizontal) and “2” (vertical) buttons on the keypad, as soon as they were able to discriminate it. Note that they were required to report the orientation even for cues previously associated with ‘no-action’. The cue faded up

over 500ms, and remained on screen in full contrast until the response was made. See *Figure 21* for an illustration of the trial sequence.

The task was split into two blocks, with all four stimuli presented in each block in a randomised order, resulting in a total of 8 trials. When indicating the line orientation, participants were required to make their response with the same action for all the trials in a single block, i.e. using the index finger to press button “1” or “2”, or the little finger to press button “1” or “2” throughout. The assigned action (use of index finger or little finger) was randomised between blocks. This design ensured that each cue was matched with each action type once (e.g. action A is performed to indicate time of breakthrough for cue A, cue B, no-action cue, and the novel cue in one block, and action B in the other).




Figure 21. Breaking-CFS task. Screenshots depict a single trial sequence (identical for all cues). Following 500ms fade-up, the cue remained on-screen with a horizontal or vertical line overlaid on top of it. Participants were requested to make a response indicating the observed line orientation as soon as they could. This response was to be made with either their index finger or little finger depending on the action randomly assigned to that block. On a given trial, the cue presented could be congruent with the trained action (i.e. associated with it in the conditioning stage), incongruent with it, associated with no action, or novel (not associated with any action).

6.2.2. Data Analysis and Results

6.2.2.1. Bayes Factors

For hypothesis testing, Bayes Factors (Bs) will be reported alongside p -values for all comparisons. Bs can help to disambiguate non-significant results as either indicating support for the null hypothesis (H_0 , positing no effect), support for the alternative hypothesis (H_1 , which uses an estimated raw effect size as the standard deviation of its distribution), or indicating insensitive data (i.e. the data are not in favour of either H_0 or H_1 ; Dienes, 2014). By convention, Bs smaller than 1/3 indicate evidence for H_0 . Bs larger than 3 indicate evidence for H_1 . Bs between those values indicate insensitive data.

6.2.2.2. Conditioning task

Data pre-processing and exclusions.

All trials with RTs under 100ms (suggesting automatic, rather than deliberate, responding) and greater than 3SD from each subject's mean were excluded. This resulted in removal of 373 trials (4.57% of all trials). Subjects missing over 25% of trials were marked for exclusion. No such subjects were identified.

Accuracy data for the conditioning task was then converted into type I d' , in order to account for potential response bias. Type I d' is a signal detection-theoretic measure of sensitivity to signal versus noise (Stanislaw & Todorov, 1999), and was computed for each cue separately using the proportions of correct (correct action deployed for the corresponding cue, e.g. cue A – action A; Hits) and incorrect (wrong action deployed, e.g. cue A – action B; False Alarm) responses.

Evidence of conditioning.

In order to assess the presence of learning, one-sample t-tests were used to contrast the d' values for each cue type against 0 (indicating no ability to discern signal from noise). By proxy, a d' of 0 can be taken as an indicator that learning failed to take place. Bs were computed with H_1 modelled as a half-normal distribution centred on 0, with an SD equal to an approximate expected effect size of $d' = 1$ (corresponding to 70% hit rate). Analysis was conducted in R (R Core Team, 2018).

The average d' scores for all cues throughout the task were significantly above 0 (see Table 6), suggesting that learning of the cue-action association was successfully established.

Table 6. Mean type I d' and SE for each cue type (A, B, NA (no-action)) in both blocks, and total, in the conscious conditioning task. Stars indicate significant difference from 0 (*: $p < 0.05$, ***: $p < 0.001$). Cross indicates a sensitive B favouring H_1 (*: $B_{H(0,1)} > 3$). $N = 68$.

Cue	Conscious conditioning											
	Task total				Block 1				Block2			
	d'	SE	p	$B_{H(0,1)}$	d'	SE	p	$B_{H(0,1)}$	d'	SE	p	$B_{H(0,1)}$
A	1.58****	0.21	<0.001	>10 ¹⁰	1.01****	0.22	<0.001	>10 ¹⁰	2.46****	0.25	<0.001	>10 ¹⁰
B	1.50****	0.23	<0.001	>10 ¹⁰	0.67**	0.26	0.011	>10 ¹⁰	2.71****	0.27	<0.001	>10 ¹⁰
NA	4.46****	0.06	<0.001	Inf	4.43****	0.08	<0.001	Inf	4.53****	0.05	<0.001	Inf

In the conditioning stage, subjects were marked as ‘learners’ if their d' scores for both cues A and B in block 2 (where learning should be evident if it had taken place) were greater than 0. The no-action cue was not included in this criterion, as due to the obvious no-action requirement it yielded a nearly perfect accuracy regardless of the ability to learn the associations between the other cues and their outcomes. 55 (out of 68) subjects were identified as learners and included in the next analysis stage.

6.2.2.3. Breakthrough task

Data pre-processing and exclusions.

No trials with response times under 100ms were identified. Unlike the conditioning task, no upper cut-off on RTs was applied, as long response times were expected. All trials where subjects made an incorrect line discrimination (horizontal/vertical) were excluded in order to reduce premature responses and ensure only trials where participants paid attention were analysed (see Discussion for consideration of the related issue of accurately guessed responses). This resulted in the removal of 65 trials (11.96%). Five subjects with 50% (4) or more missing trials were removed. One extra trial was removed due to the participant failing to engage with the task, which resulted in a response time of over 10 minutes. One subject was removed from both conditions due to a disruption in the testing session. Both cases were noted by the experimenter in the session log. Only the subjects identified as learners in the conditioning stage were brought into the breakthrough time analysis stage. This resulted in a final sample of 49. We confirmed that, averaging across trial types, there was no significant difference in response times between actions A and B ($M_{\text{actionA}} = 5817.87$, $M_{\text{actionB}} = 5894.63\text{ms}$; $t(172) = -0.52$, $p = 0.607$).

Each breakthrough trial was given a label describing its cue-action congruence status (i.e. whether the action prepared to indicate breakthrough was congruent or incongruent with the cue, as established in the conditioning task). Action A – cue A and action B – cue B pairs were labelled as congruent pairs. Action A – cue B and action B – cue A were labelled as incongruent pairs. Cue C was always labelled as no-action, and the fourth, previously unseen cue, was labelled novel.

Breakthrough time results.

Due to the unevenly distributed missing values (i.e. exclusions due to incorrect line orientation judgments) across few data points per participant, the pre-registered analysis method (repeated-measures analysis of variance) was rendered inappropriate. An ANOVA excludes such cases listwise, resulting in excluded participants and reduced power. Given superior performance in treatment of repeated-measures data and data with unevenly distributed missing values, as well as superior ability to model repeated-measures, a generalised linear mixed model (GLMM) was fitted instead. Analysis was conducted using the lme4 package (Bates et al., 2015) in R (R Core Team, 2018).

The model included the raw times of breakthrough as the response variable, cue-action congruence status (4 levels: congruent, incongruent, no-action, novel) as a fixed effect, and subject-specific random intercepts⁸ (Matuschek et al., 2017; Singmann & Kellen, 2019). A gamma distribution of the response variable, with an identity link function, was specified in order to approximate the nature of response time data without the need for transformations (Lo & Andrews, 2015). The model was fitted by maximum likelihood estimation. All following comparisons were conducted on that model.

The GLMM revealed a significant main effect of congruence status on time of breakthrough ($\chi^2(3) = 12.52, p = 0.006$; see *Table 7* for regression coefficients). Subsequent pairwise comparisons on estimated means (Tukey-adjusted for multiple comparisons; see *Table 8*) showed that only the novel, previously unseen cue ($M = 6207\text{ms}$, $SE = 64.39$) resulted in significantly shorter breakthrough time than both congruent ($M = 6311\text{ms}$, $SE = 50.09$) and incongruent cues ($M = 6356\text{ms}$, $SE = 62.44$). Despite a marginally shorter breakthrough time for congruent than incongruent cues, given the adopted priors and the size of the observed effect, the Bayes factor indicates strong evidence against a genuine difference. For B calculation, H_1 was modelled as a normal distribution centred on 0, with an SD equal to an estimated effect size of 774ms. This estimate was derived from the observed difference between rewarded and unrewarded cues in a similar b-CFS task conducted earlier by some of the authors (Scott et al., *in preparation*).

⁸ In R notation, the fixed and random effects of the model were specified as: breakthrough time ~ congruence index + (1|subjectID).

Table 7. Regression estimates from the GLMM. Congruent cue-action status serves as reference point. Stars indicate significant difference from the intercept (*: $p < 0.05$, **: $p < 0.001$). $N = 55$.

	Estimate	SE	t	p (> z)
Intercept (congruent)	6311.12	50.09	126.01	< 0.001 **
Incongruent	45.05	30.62	1.47	0.141
No-action	-30.92	36.12	-0.86	0.392
Novel	-103.30	33.93	-3.04	0.002 *

Table 8. Pairwise comparisons (Tukey-adjusted) of breakthrough time means, estimated in the GLMM. Star indicates a significant difference (*: $p < 0.05$). Cross indicates a sensitive B favouring H_1 (+: $B_{N(0,774)} > 3$). Tilde indicates a sensitive B favouring H_0 (~: $B_{N(0,774)} < 0.3$).

Congruence status contrast	Estimated mean difference (ms)	SE	df	z-ratio	p	$B_{N(0,774)}$
congruent-incongruent	-45.06	30.62	Inf	-1.472	0.459	0.12 ~
congruent-no-action	30.92	36.12	Inf	0.86	0.828	0.07 ~
congruent-novel	103.30	33.93	Inf	3.04	0.013*	4.48 +
incongruent-no-action	75.97	48.42	Inf	1.57	0.397	0.21 ~
incongruent-novel	148.35	46.24	Inf	3.20	0.008*	10.07 +
no-action-novel	72.38	52.48	Inf	1.38	0.512	0.17 ~

6.2.3. Conclusions of Experiment 1

In Experiment 1, we investigated whether valid sensorimotor predictions, built through instrumental conditioning, can affect conscious experience, operationalised in terms of reportable access to consciousness in breaking interocular suppression. If sensorimotor predictions shape conscious experience, we hypothesised that congruency between the cue and the prepared action would result in the cue breaking through CFS faster.

While the data showed, numerically, marginally shorter breakthrough times for congruent than for incongruent pairs, given the adopted priors and the size of the observed difference the evidence was in favour of the null hypothesis. Specifically, preparing an action congruent with a specific cue (i.e. one previously conditioned with that action) while it was presented under interocular suppression does not shorten the suppression duration relative to preparing an action incongruent with the cue. Therefore, these data do not support the hypothesis that a valid prediction of a cue-action association speeds up reported access to consciousness of the target.

These findings are counter to the results of previous research which, while not directly manipulating sensorimotor predictions as we did here, did show modulation of perception by action in line with action-congruent percepts (e.g. Beets et al., 2010; Maruya, Yang, Blake, 2007; Mitsumatsu, 2009). While it is possible that our result reflects a genuine absence of influence arising from sensorimotor predictions, it is also possible that the paradigm may have limited the

predictive salience of the prepared action. While requiring the same action on each trial within a block ensured counterbalancing of the action-preparation and stimulus congruency, it may also have caused participants to deploy action in an automatic, rather than voluntary, goal-oriented manner. Indeed, previous research has suggested that the stimulus-action contingencies should adaptively reflect a goal or objective in order to facilitate interactions with the world (Di Pace & Saracini, 2014; Hommel et al., 2001; Mitsumatsu, 2009; Prinz, 2003; Seth, 2014; Seth et al., 2016; Wohlschläger, 2000). The requirement for repetitive action may have inadvertently eliminated this important aspect of the behaviour.

A second potentially confounding influence arises from the timing of the requirement to report the line orientation. Participants were required to press the key corresponding to a vertical or horizontal line overlaid on the stimulus as soon as they began to see the stimulus break through CFS. It is plausible that the decision time relating to the orientation judgment added noise to the measure of breakthrough time. We designed a second experiment, Experiment 2, to address both these issues.

6.3. Experiment 2

Experiment 2 introduced two changes to the breakthrough task described in Experiment 1. Because keeping the action requirement consistent across the entire block may have resulted in participants executing the action in an automatic manner, thus reducing the predictive salience of the prepared action, we varied the action requirement from trial to trial. We also requested the participants to make a single response (with the corresponding action) to indicate that they can see the stimulus, followed by a line orientation judgment performed independently with the other hand. All other task parameters remained the same.

6.3.1. Method

6.3.1.1. Participants

65 participants (18 males; M age = 20.78, SD = 4.31, range = 18-47) were recruited for the study via the University of Sussex online recruitment system, and an internal mailing list. All participation criteria were identical to Experiment 1. Data for one participant was unusable due to software malfunction, resulting in a sample of 64.

6.3.1.2. Stimuli and Materials

All stimuli and materials were identical to Experiment 1.

6.3.1.3. Procedure

The procedure was identical to the conscious conditioning task of Experiment 1 with the exception of the minor changes made to the breakthrough task outlined below.

Conditioning task.

Identical to Experiment 1.

Breakthrough task.

The conditioning task was immediately followed by the b-CFS task, using the same set of cues (*cue A*, *cue B*, *no-action cue*), as well as fourth, novel cue, never previously seen. The task parameters remained the same as in Experiment 1. In contrast to Experiment 1, the response requirement (action A or B) was no longer repeated without variation within each block. Instead, participants were instructed at trial onset which action would be required to make the response (in a randomised order, but counterbalanced such that each of the four cues is matched with each of the two actions once, resulting in eight trials). This change was considered to be important in order to maintain the predictive salience of the action.

Additionally, in order to eliminate the potentially confounding influence of line orientation judgment on pure breakthrough time, these two responses were separated. Participants were required to respond with the instructed action (A or B) using the return key as soon as they saw the image break through (as opposed to responding as soon as they were able to discriminate the line orientation). After the initial response both the image and the Mondrian pattern disappeared and participants were required to indicate the perceived line orientation; this was done using the “1” (horizontal) and “2” (vertical) buttons at the top of the keyboard using their left hand.

6.3.2. Results

6.3.2.1. Conditioning task

Data pre-processing and exclusions.

Pre-processing and exclusion procedures were identical to Experiment 1. All trials with RTs under 100ms and under or over 3SD from each subject’s mean were excluded. This resulted in removal of 225 trials (2.9% of all trials). Subjects missing over 25% of trials were marked for exclusion. No such subjects were identified.

Evidence of conditioning.

As in Experiment 1, one-sample t-tests were used to contrast the d' values for each cue type against 0 (indicating no sensitivity to signal versus noise). By proxy, a d' of 0 can be taken as an indicator that learning failed to take place. Bs were computed with H_1 modelled as a half-normal distribution centred on 0, with an SD equal to an estimated expected effect size of $d' = 1$ (corresponding to 70% hit rate).

The d' scores for all cues in both blocks of the conditioning task, and in total, were significantly above 0 (see Table 9), suggesting that learning of the cue-action association was again successfully established.

Table 9. Mean type I d' and SE for each cue type (A, B, NA (no-action)) in both blocks, and total, in the conscious conditioning task. Stars indicate significant difference from 0 (***: $p < 0.001$). Cross indicates a sensitive B favouring H_1 (*: $B_{H(0,1)} > 3$). N = 64.

Cue	Conscious conditioning											
	Task total				Block 1				Block2			
	d'	SE	p	$B_{H(0,1)}$	d'	SE	p	$B_{H(0,1)}$	d'	SE	p	$B_{H(0,1)}$
A	1.54****	0.16	<0.001	>10 ¹⁰	0.90****	0.20	<0.001	6632	2.61****	0.20	<0.001	>10 ¹⁰
B	1.64****	0.19	<0.001	>10 ¹⁰	0.75****	0.24	<0.001	47.17	3.04****	0.23	<0.001	>10 ¹⁰
NA	4.54****	0.07	<0.001	Inf	4.54****	0.08	<0.001	Inf	4.61****	0.07	<0.001	Inf

Subjects were marked as learners if their d' scores for both cues A and B in block 2 (where learning should be evident if it had taken place) were greater than 0. The *no-action cue* was not included in this criterion, as due to the obvious no-action requirement it yielded a nearly perfect accuracy regardless of the ability to learn the associations between the other cues and their outcomes. 58 subjects were identified as learners and included in the next analysis stage.

6.3.2.2. Breakthrough task

Data pre-processing and exclusions.

Pre-processing and exclusion procedures were identical to Experiment 1. No trials with response times under 100ms were identified. Again, no upper cut-off on RTs was applied, as long response times were expected. All trials where subjects made an incorrect line discrimination were excluded. This resulted in the removal of 46 trials (9%). Three subjects with 50% (4) or more missing trials were removed. This resulted in the final sample of 55 subjects. Outlier analysis identified 5 extremely outlying trials, which were also removed. Again we confirmed that, averaging across trial types, there was no significant difference in response times between actions A and B ($M_{\text{actionA}} = 6610.41\text{ms}$, $M_{\text{actionB}} = 6217.86\text{ms}$; $t(192) = 0.36$, $p = 0.715$).

Each breakthrough trial was given a label describing its cue-action congruence status following the identical procedure as employed in Experiment 1.

Breakthrough time results.

The model and analysis performed was identical to that in Experiment 1. A GLMM was fitted to the data, including the raw times of breakthrough as the response variable. Model specification included the cue-action congruence status (4 levels: congruent, incongruent, no-action, novel) as a fixed effect, and subject-specific random intercepts, and was fit with a gamma distribution with an identity link function (Lo & Andrews, 2015). The model was fitted by maximum likelihood estimation. All following comparisons were conducted on that model.

As in Experiment 1, the GLMM revealed a significant main effect of congruence status on time of breakthrough ($\chi^2(3) = 158.28, p < 0.001$; see *Table 10* for regression coefficients). Subsequent pairwise comparisons on estimated means (Tukey-adjusted for multiple comparisons; see *Table 11*) showed that cue-action congruent breakthrough times ($M = 6516\text{ms}$, $SE = 30.3$) were significantly shorter than incongruent breakthrough times ($M = 6809\text{ms}$, $SE = 44.6$), mirroring the direction of effect found in Experiment 1. Crucially, however, in this instance the size of the difference was substantial, with the Bayes factor providing strong evidence in favour of H_1 . Cue-action congruent breakthrough times were also shorter than no-action breakthrough times ($M = 6691\text{ms}$, $SE = 39.6$), but significantly longer than breakthrough times for novel cues ($M = 6340\text{ms}$, $SE = 41.9$; see *Table 11* for pairwise comparisons). In addition, the cue-action incongruent breakthrough times were significantly longer than no-action breakthrough times. Novel cues resulted in the shortest breakthrough times.

The statistics obtained in pairwise comparisons for Experiment 1 and 2 were then combined in order to calculate a meta-analytic Bayes factor, a single Bayes factor indicating evidence for H_0 or H_1 in a group of studies. Posterior parameters were computed for each pairwise comparison using the estimated mean differences and standard errors from Experiment 1 as priors, and the estimated mean differences and standard errors from Experiment 2 as likelihoods (following the method from Dienes, 2014). The resulting mean and the standard deviation of the posterior distribution were then used in the meta-B calculation in a manner identical to regular B calculation, where H_1 was modelled as a normal distribution centred on 0, with an SD equal to an estimated effect size of 774ms. See *Table 11* for the meta-Bs for each cue-action pairing.

Table 10. Regression estimates from the GLMM. Congruent cue-action status serves as reference point. Stars indicate significant difference from the intercept (*: $p < 0.05$, **: $p < 0.001$). $N = 49$. N of observations: 412.

	Estimate	SE	<i>t</i>	<i>p</i> ($> z $)
Intercept (congruent)	6516.27	30.34	214.80	<0.001***
Incongruent	292.94	34.47	8.50	<0.001***
No-action	174.73	23.52	7.43	<0.001***
Novel	-176.47	28.88	-6.111	<0.001***

Table 11. Pairwise comparisons (Tukey-adjusted) of breakthrough time means, estimated in the GLMM. Star indicates a significant difference (*: $p < 0.05$, ***: $p < 0.001$). Cross indicates a sensitive B favouring H_1 (+: $B_{N(0,774)} > 3$). The final column presents meta-analytic Bayes factors, obtained by pooling the parameters from Experiments 1 and 2.

Congruence status contrast	Estimated mean difference (ms)	SE	df	z-ratio	<i>p</i>	$B_{N(0,774)}$	meta-B $N(0,774)$
congruent-incongruent	-293	34.5	Inf	-8.50	<0.001***	>10 ¹⁰⁺	>10 ⁷⁺
congruent-no-action	-175	23.5	Inf	-7.43	<0.001***	>10 ¹⁰⁺	>10 ⁷⁺
congruent-novel	176	28.9	Inf	6.11	<0.001***	>10 ⁶⁺	>10 ⁶⁺
incongruent-no-action	118	41.5	Inf	2.85	0.023 *	3.13 ⁺	6.63 ⁺
incongruent-novel	469	44.6	Inf	10.53	<0.001***	>10 ¹⁰⁺	>10 ¹⁰⁺
no-action-novel	351	38.7	Inf	9.06	<0.001***	>10 ¹⁰⁺	>10 ¹⁰⁺

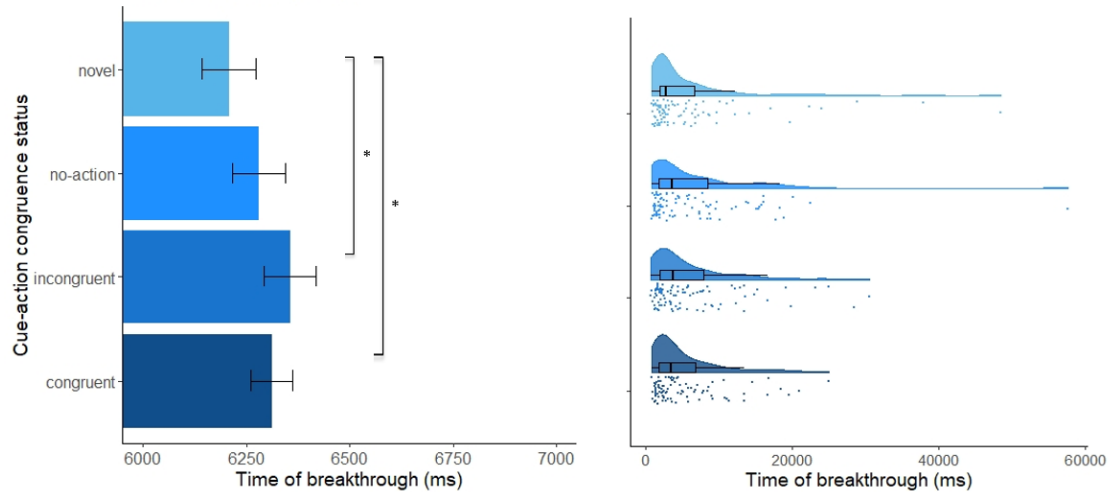
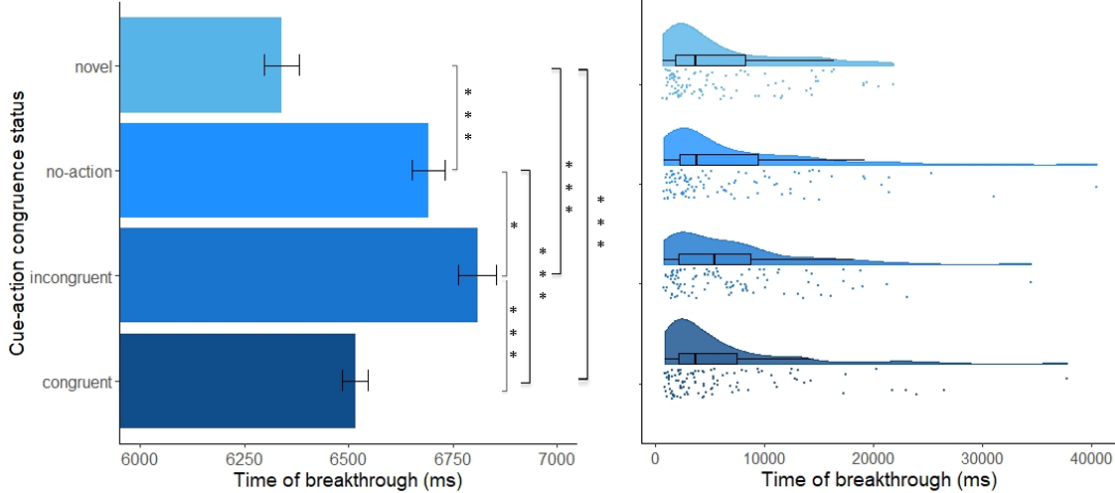
EXPERIMENT 1:**EXPERIMENT 2:**

Figure 22. Left panels: Mean reported times of breakthrough (ms) by cue-action congruence status in Experiment 1 (top; $N = 49$) and Experiment 2 (bottom; $N = 55$), estimated by the GLMM (± 1 SEM). Stars indicate $p < 0.05$. **Right panels:** Raw distributions of breakthrough times with boxplots in Experiment 1 and 2.

6.3.3. Conclusions of Experiment 2

In Experiment 2, we continued to investigate whether sensorimotor predictions, built through instrumental conditioning, can affect conscious experience, as measured by reported access to consciousness in breaking interocular suppression. We amended the b-CFS task used in Experiment 1 in order to address limitations in the initial design. If sensorimotor predictions shape reported conscious experience, we predicted that congruency between the prepared action and the cue would cause the cue to break through CFS faster.

The results show that a valid sensorimotor prediction does affect the speed of reported access to consciousness. Preparing an action that was congruent with the cue it was conditioned on resulted in significantly faster breakthrough than an incongruent action. This finding is further supported by the meta-analytic Bayes factor analysis, which pooled the results from Experiments 1 and 2. Together the data provides strong support for the hypothesis that conscious experience is affected by sensorimotor predictions – if the executed action is congruent with the cue-action predictive model, reported access to consciousness is facilitated. While the central hypothesis was supported, the shortest breakthrough time was seen for the novel cue. While not in conflict with the key finding, the framework of action-oriented predictive processing does not offer a ready explanation for this pattern. We return to this issue in the general discussion.

6.4. General Discussion

Predictive processing proposes that conscious experience is shaped or constituted by the predictive model that best explains the incoming sensory input. Action-oriented interpretations expand this notion, proposing that perceptual predictions should encompass sensorimotor contingencies – i.e., predictions of the relationships between an agent’s actions and changes in the world (Clark, 2015, 2016; Seth, 2014, 2015). We sought to empirically test whether conscious experience is affected by valid predictions about sensorimotor (stimulus-action) associations – which we here call ‘sensorimotor predictions’.

We developed a novel paradigm in which we leveraged instrumental conditioning to build arbitrary sensorimotor associations, linking distinct cues with specific actions or with no action. Conditioning was followed by a breaking-CFS task, where a high-contrast pattern was used to suppress visual awareness of the target cues while participants prepared to respond with cue-contingent or non-contingent actions. Through the proxy of breaking interocular suppression, this paradigm allowed us to test the extent to which maintaining the relevant sensorimotor prediction through action preparation facilitates reported conscious access to the associated stimulus, relative to non-contingent stimuli, as well as stimuli associated with no action, and novel stimuli that had not undergone any conditioning exposure.

In directly manipulating sensorimotor predictions, the present study extends previous research in a specific way. Prior investigations examining the modulation of conscious perception by action have suffered from the limitation that observed effects could potentially arise from the actions themselves resulting in response congruency or attentional cueing (e.g. the direction of hand movement facilitating the perception of the stimulus which moves in the

corresponding direction; Beets et al., 2010). In the present study, we observe the effect of manipulating the sensorimotor prediction itself prior to the action being taken, and hence limit these potentially confounding effects.

Experiment 1 failed to find a relationship between sensorimotor predictions and breakthrough times. This finding went against our initial hypothesis and conflicted with previous evidence showing modulation of conscious perception by action (albeit not reflecting sensorimotor predictions; Beets et al., 2010; di Pace & Saracini, 2014; Maruya, Yang & Blake, 2007; Mitsumatsu, 2009). However, the design of Experiment 1 likely limited the predictive salience of the prepared action. Requiring the same action on each trial within a block (targeted at counterbalancing the stimulus-action congruency) may have caused participants to deploy action in an automatic, rather than dynamic, goal-oriented fashion that diminished the relative importance of action related predictions.

In Experiment 2, we addressed this limitation by varying the action requirement on each trial. We also reduced the potential for measurement noise in the breakthrough judgment by separating the response indicating the moment of conscious breakthrough from the response indicating the line orientation. While both experiments delivered numerically shorter breakthrough times for congruent versus incongruent cue-action pairings, the modifications designed to improve the predictive salience of the action resulted in considerably larger differences. In Experiment 2, Bayes factor analysis showed strong evidence in favour of a genuine difference based on congruency. Breakthrough time for the no-action cue was also longer than for the action-congruent cue. This is again consistent with our hypothesis, because executing either action A or B would engage the corresponding prediction related to perception of cue A or B, effectively rendering the no-action cue incongruent with either. Importantly, a meta-analytic Bayes factor analysis, where the results from Experiments 1 and 2 were pooled to obtain a single Bayes factor for each comparison, show strong support for the alternative hypothesis (i.e., congruent actions accelerate conscious access) across the two experiments.

A surprising finding in both Experiments was that the novel cue was significantly faster to break through than both the congruent and incongruent cue-pairings. While this result is contrary to some previous research examining the effects of expectation on access to awareness (e.g. Pinto et al., 2015), it is consistent with reports of greater attentional capture for the first unannounced presentation of novel cues (Al-Aidroos, Guo, & Pratt, 2010; Becker & Horstmann, 2011; Ernst, Becker, & Horstmann, 2020; cf. Meijs, Klaassen, Bokeria, Van Gaal, & de Lange, 2018). Indeed, there is an active debate in the predictive processing literature about whether predictions or prediction errors should dominate perceptual content and perceptual access, with the balance proposed to be modulated by attention (Feldman & Friston, 2010; Hohwy,

2012). Future extensions of our paradigm may usefully target this specific question (Press, Kok, & Yon, 2020).

It is noteworthy that in the present task we operationalised sensorimotor contingencies as learned arbitrary associations between visual stimuli and subsequent actions. We called those associations ‘sensorimotor predictions’, and manipulated their predictive validity to evaluate reported conscious access. While this mapping constitutes a kind of contingency, one might note that it reverses the direction of what is typically understood as sensorimotor contingencies, i.e. how the world changes in response to our actions (e.g. Seth, 2014). Nor do our learned associations correspond to temporally extended relationships between actions and sensory signals, such as when an object is rotated (e.g. Suzuki et al, 2019). As such, the sensorimotor predictions used here do not correspond to fully-fledged sensorimotor contingencies, but rather to simpler sensorimotor mappings. Nonetheless, the present paradigm could be fruitfully extended to investigate sensorimotor contingencies in the more traditional sense.

An important criticism of b-CFS that should be noted is its interpretation in relation to “conscious access”. In b-CFS, the moment understood as the moment of access to consciousness is in fact the moment participants *report* having conscious awareness of the stimulus. Those events may or may not be the same. For example, participants might have rudimentary awareness of the stimulus but delay reporting it (e.g. due to low confidence). As such, our paradigm and many similar paradigms use b-CFS as a proxy measure of speed of access to consciousness, while recognizing this caveat.

Another general criticism of b-CFS is that the time of breakthrough measure may be affected by other processes or biases (Stein & Sterzer, 2014; Sterzer et al., 2014). A challenge pertaining to all b-CFS experiments is the difficulty of disentangling the speed of access to conscious (reportable) awareness from the time needed to prepare the response – the measured time of breakthrough inherently includes the time taken to respond, which could vary between stimuli or types of response. In the present experiment, two different responses (index and little finger button presses) were a core feature of our design – indeed, we argued that the preparation of a congruent versus incongruent action is what causes the stimulus to break through faster, in line with the idea that valid sensorimotor predictions shape conscious experience. Importantly, however, average response times (irrespective of congruence) did not differ between the two actions, allowing us to rule out the possibility that the any underlying difference in time taken to respond with either action contributed to the result. Nonetheless, this does not preclude the possibility of post-perceptual biases affecting the total response time, such as stimulus features slowing down the action. We attempted to guard against this by ensuring the stimuli were as similar as possible. We used four 90° rotations of the same shape,

and (in Experiment 2) disentangled the response indicating the stimulus breakthrough from the response indicating line orientation. Still, however, the extent to which post-perceptual biases affect b-CFS remains a topic deserving further study.

Our b-CFS task used only one presentation per cue-action pairing (e.g. action A was paired once with cue A, cue B, no-action cue, and the novel cue). This design was adopted because we felt that to best address the question of how sensorimotor predictions affect conscious experience, we should focus on the very first conscious experience of the action-associated stimulus. While some previous b-CFS experiments have used multiple trials (e.g. Salomon et al., 2013), repeated exposure itself tends to reduce breakthrough times, which could reduce the ability to observe the main effect of interest. Adopting a single trial per cue-action pairing avoids this issue albeit at the cost of the reduced statistical power resulting from a smaller number of trials. Fortunately, our mixed-effects model proved robust to the small number of observations per participant and converged without issue. In addition, the Bayes Factor calculation ensured that the evidence for or against the differences in breakthrough times was sensitive.

The inclusion of a line orientation judgment was intended to reduce the occurrence of premature responding (i.e. responding prior to conscious perception), and to identify instances where participants' attention had been suboptimal. We excluded all such trials from analysis. However, theoretically this still leaves open the possibility that the analysed trials contained some instances where participants responded prematurely but correctly guessed the line orientation; any such trials would contaminate the RTs as a pure indicator of conscious breakthrough. Given that the line orientation decisions were not accompanied by confidence judgments, it not possible to directly identify accurately guessed responses. However, if guessing of the line orientation was a common occurrence, then we would reasonably expect the accuracy of the line orientation judgements to be substantially reduced. Given the observed accuracy was 88 and 91% in Experiment 1 and 2 respectively, we feel reassured that the potential contribution from accurate guesses is likely to have been negligible.

In conclusion, in two experiments we investigated the effect of valid sensorimotor predictions on conscious experience, measured through breaking interocular suppression. The combined data from the two experiments provides strong evidence that preparing a cue-congruent action results in more rapid reported conscious perception of the suppressed stimuli. This provides evidence for a key theoretical implication of the action-oriented predictive processing approach to conscious perception, namely, that sensorimotor predictions shape conscious experience of the world.

6.5. Supplementary material

Unconscious condition of Experiment 1.

6.5.1. Method

6.5.1.1. Participants

As reported in Experiment 1 described in the main paper.

6.5.1.2. Stimuli

As reported in Experiment 1 described in the main paper.

6.5.1.3. Procedure

Unconscious conditioning.

The conditioning task was used to establish the sensorimotor (stimulus-action) associations. The task used stimuli selected from the assigned set of four. One of the stimuli (*cue A*) was paired with action A (an index finger button press), a second (*cue B*) with action B (a little finger button press), and a third (*no-action cue*) with no action. For practical reasons, the index finger button press was made on the left arrow, and the little finger button press on the right arrow. Note, while the keys being pressed differed, we consider the finger used to make the button press to be the conditioned action. Cues were randomly assigned to a given type for each participant. The order of presentation was randomised.

In the unconscious condition, stimuli were presented under CFS (Tsuchiya & Koch, 2005), in a mirror stereoscope setup. The dominant eye received a Mondrian pattern composed of coloured rectangles, flashing at a rate of 10Hz. The non-dominant eye received the target stimulus. Ocular dominance was established prior to beginning with a standard Miles test (Miles, 1930). Three cues selected from the assigned set of four were presented subliminally using CFS, in a randomised order.

The unconscious presentation of each cue started with a fade-up period of 500ms. Further 500ms after the cue reached full contrast, the fixation dot (overlaid on displays to both eyes) changed colour, indicating that a response was required (see *Figure 1* for an illustration of trial sequence). Fixation dot colours provided the distinction between action and no-action trials. For action cues, participants were instructed to respond with either action A or B (of their choosing), when prompted by the fixation dot changing colour from white to black. Following the action, positive or negative reinforcing feedback ('correct!' or 'wrong!' printed on the

screen, and a cash register or buzz sound, respectively) was delivered, depending on the correspondence of the executed action to the cue presented, thus instantiating instrumental conditioning (e.g. if *cue A* is presented and *action A* is executed, positive feedback is delivered, but if *action B* is executed, negative feedback is delivered). For non-action cues, the fixation dot started as blue and changed colour to red. This indicated to participants at the onset of the trial that no response will be required. As such, the subliminally presented cue became associated with the ‘no action’ response. Participants performed 120 conditioning trials in two blocks of 60, with a 1 minute break between the blocks.

On action cue trials, the cue then disappeared as soon as the participant made a response (1000ms + reaction time after cue onset). This ensured that the action was prepared and executed while the participant was exposed to the cue, but exposure was not prolonged unnecessarily following the response. On no-action trials, the cue disappeared as soon as the fixation dot changed (1500ms after cue onset). This additional 500ms sought to roughly equalise the total exposure duration for no-action trials, where there was no response time, and action trials where the exposure extended to include the participant’s response time.

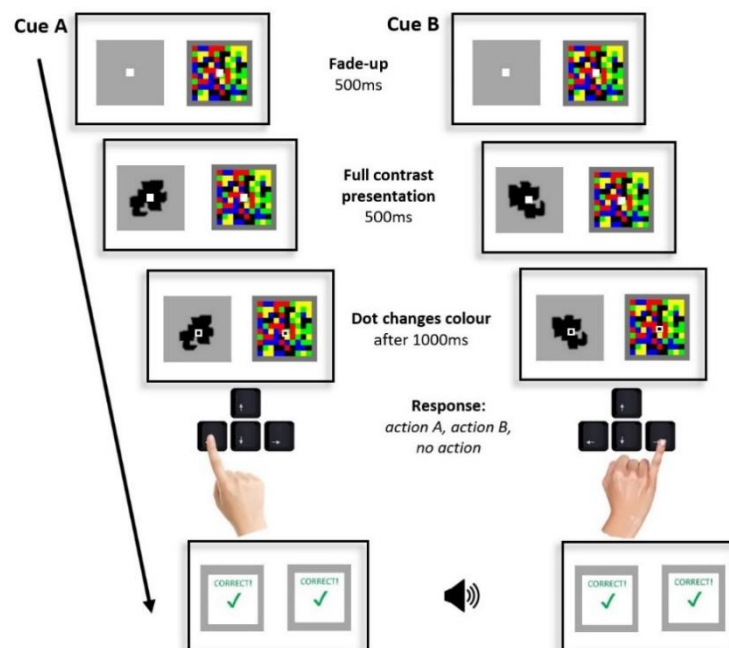


Figure S6.1. Unconscious instrumental conditioning task. Chronological screenshots depict a single trial sequence for two action cues – each panel shows the images shown to the left and right eyes. After 1000ms (including 500ms fade-up of the cue), the fixation dot changes colour from white to black (with a white border), participants can execute the desired action (A, an index finger button press on the left arrow, or B, a little finger button press on the right arrow). For a no-action cue, the fixation dot would change colour from red to blue. Here, the action executed corresponded to the cue type, and the participant was rewarded with visual and auditory feedback.

Breakthrough task.

The breaking CFS task in the unconscious condition was identical to the breaking CFS task in the conscious condition, as described in Experiment 1 in the main paper.




Figure S6.2. Breaking-CFS task. Screenshots depict a single trial sequence (identical for all cues). Following 500ms fade-up, the cue remained on-screen with a horizontal or vertical line overlaid on top of it. Participants were requested to make a response indicating the observed line orientation as soon as they could. This response was to be made with either their index finger or little finger depending on the action randomly assigned to that block. On a given trial, the cue presented could be congruent with the trained action (i.e. associated with it in the conditioning stage), incongruent with it, associated with no action, or novel (not associated with any action).

Awareness check.

The session ended with an awareness check, which sought to verify that the cues presented under CFS in the unconscious conditioning task were outside of conscious awareness. The task included one block of 60 trials, and was identical to the unconscious conditioning task, with the exception of the cues, which were two distinct symbols: one black rectangle, and one irregular shape resembling the shapes used in the main task. The symbols were randomly allocated to be rewarding or punishing for each participant, and their order was randomised. There was no no-action cue. Following the feedback, participants were asked to judge if the cue presented under CFS was symmetrical (rectangle) or asymmetrical (the other cue). Next, they were asked to assess their confidence on a binary scale (*some confidence vs total guess*). Accuracy, confidence, and their correlations were computed to assess participant's awareness of the cues during the task. Upon completion, participants were thanked and debriefed.

6.5.2. Results

6.5.2.1. Conditioning task

Data pre-processing and exclusions.

All trials with RTs under 100ms (suggesting automatic, rather than deliberate, responding) and under/over 3SD from each subject's mean were excluded. This resulted in removal of 113 trials (1.38%) in the unconscious conditioning task. Subjects missing over 25% of trials were marked for exclusion. No such subjects were identified.

Subjects who showed visual awareness of the cues, assessed by above-chance judgment accuracy and a positive correlation with confidence in the awareness check, were excluded from further analysis. Those who scored below chance on accuracy, and those who were above chance but showed no correlation (0 or below) were retained. This resulted in the exclusion of 37 out of 68 subjects in the unconscious task (31 participants remaining).

Evidence of conditioning.

In order to assess the presence of learning one-sample t-tests were used to contrast the d' values for each cue type against 0 (indicating no sensitivity to signal versus noise). By proxy, a d' of 0 can be taken as an indicator that learning failed to take place. Bayes Factors were computed with H_1 modelled as a half-normal distribution centred on 0, with an SD equal to approximate expected effect size of $d' = 0.7$ for the unconscious condition, following past research examining subliminal learning (Pessiglione et al., 2008). Analysis was conducted in R (R Core Team, 2018).

In unconscious conditioning, only the *no-action cue* was significantly above 0 (see *Table S6.1*). For the action-related cues, average d' for the length of the task was not significantly different from 0, with evidence favouring the null hypothesis, suggesting that subjects were unable to discriminate between the cues, and thus failed to learn the cue-action association. Performance on *cue A* in block 2 was an exception, with an insensitive B. On average, this result demonstrates that unconscious instrumental conditioning was not achieved in the present paradigm.

Table S6.1. Mean type I d' and SE for each cue type (A, B, NA (no-action)) in both blocks, and total, in the unconscious conditioning task. Stars indicate significant difference from 0 (*: $p < 0.05$, **: $p < 0.001$). Cross indicates a sensitive B favouring H_1 (+: $B_{H(0,0.7)} > 3$). Tilde indicates a sensitive BF favouring H_0 (~: $B_{H(0,0.7)} < 0.3$). $N = 31$.

Cue	Task total				Block 1				Block2			
	d'	SE	p	B	d'	SE	p	B	d'	SE	p	B
A	-0.06	0.21	0.80	0.23	-0.25~	0.20	0.23	0.13	0.19	0.28	0.51	0.66
B	-0.32*~	0.17	0.07	0.08	-0.44*~	0.19	0.03	0.08	-0.20~	0.21	0.344	0.16
NA	4.44***	0.08	<0.001	Inf	4.56***	0.06	<0.001	Inf	4.35***	0.11	<0.001	Inf

In the conditioning stage, subjects were marked as 'learners' if their d' scores for both cues A and B in block 2 (where learning should be evident if it had taken place) was greater than 0. The *no-action cue* was not included in this criterion – due to the obvious no-action requirement, it yielded a nearly perfect accuracy for all subjects in both conditioning tasks, regardless of their ability to learn the associations between the other cues and their outcomes.

12 (out of the 31 determined as unaware) subjects were identified as 'learners'. However, given that learning was absent at the group level, the learners were not passed on to the next analysis stage due to the risks for robustness of the analysis. Firstly, post-hoc selection of the aware subjects from the sample risks obtaining a false-positive through regression to the mean (Shanks, 2017). Secondly, a small sample size was likely to result in insufficient power. As such, analysis for the unconscious conditioning condition of the task was terminated after the conditioning stage. Since we were interested in predictions about learned stimulus-action associations (sensorimotor predictions) on breaking-CFS, it is not informative to analyse data from a condition where these associations were not learned.

7.

GENERAL DISCUSSION

7.1. Summary of findings – revisiting the main aims

The present thesis set out to address three specific aims concerning human instrumental learning. This section will re-state those aims, first set in the theoretical overview, and summarise the contributions of the empirical chapters (2-6) to addressing them.

1. *Investigate the role of interoceptive information in instrumental learning. The extent to which instrumental learning is affected by the salience of bodily information (i.e. interoceptive precision) can shed light on the importance of bodily signals in adaptive processes, especially in primitive scenarios, where learning might occur without the influence of higher-order cognition.*

A vital part of successful adaptive behaviour is monitoring performance for errors. Error-related autonomic activity, such as error-related cardiac deceleration, was proposed to reflect the performance monitoring mechanism (Hajcak, McDonald, et al., 2003), akin to the orienting response (Sokolov, 1963). However, the extent to which such activity is directly driving adaptive behaviour is unknown.

Chapter 2 attempted to investigate this question through the role of cardiac precision in an unconscious instrumental conditioning task, a simple form of adaptive learning. In the task, participants learned to approach reward-predictive subliminal stimuli, and avoid punishment-predictive ones. Performance feedback was always presented consciously following their action. The precision of error-related cardiac information was amplified and disrupted through real-time auditory cardiac feedback, played as participants learned. Precision was amplified through synchronous cardiac feedback (an auditory stream played in rhythm of participants' actual heartbeat), and disrupted with an asynchronous stream. If cardiac precision plays a functional role in learning, amplifying it should improve learning, while disrupting it should disrupt learning. Unfortunately, conclusions about the effect of the manipulation on learning were hindered by the strong evidence for an absence of unconscious instrumental learning at baseline.

Chapter 3 thus attempted to replicate this result, with an added measure of cardiac activity throughout the learning task. It investigated whether error-related cardiac deceleration, an autonomic marker of performance, might be present even in the absence of any behavioural manifestations of learning. The results again provide strong evidence for the absence of unconscious instrumental learning. Crucially, error-related heart rate deceleration (i.e. deceleration in response to performance feedback following an error in behaviour such as

approaching a punishing stimulus), which occurs reliably if performance feedback is meaningful, was also absent. This shows that unconscious instrumental learning fails to produce the autonomic markers of performance monitoring, supporting the claim that unconscious instrumental learning may be infeasible.

2. Investigate the role of consciousness in instrumental learning. The extent to which instrumental learning is possible in simple scenarios – such as without conscious awareness of the stimuli – is itself an unsettled case. Aside from the bodily influences, learning may require conscious awareness to be able to drive motivated instrumental responses.

Instrumental learning is a fundamental aspect of successful adaptive behaviour, allowing agents to learn the associations between stimuli and outcomes, and to act upon them to achieve reward and avoid punishment. As such, the extent to which this requires conscious awareness to proceed is of considerable theoretical importance for the field of consciousness research. While some research found successful instrumental learning in the absence of stimulus awareness (Mastropasqua & Turatto, 2015; Pessiglione et al., 2008), other work has cast doubt on this finding (Reber et al., 2018).

While Chapters 2 and 3 were not inherently designed to probe this question, their results show strong evidence against instrumental learning in the absence of stimulus awareness. Experiments presented in both those chapters employed the same core task as the most prominent paradigm in the field (Pessiglione et al., 2008), with enhanced task sensitivity and statistical methods.

Thus, Chapters 4 and 5 were designed to specifically probe that surprising result. Chapter 4 is a Stage 1 Registered Replication Report, aimed at rigorously replicating the aforementioned paradigm. Unfortunately, due to external circumstances, data collection was impossible during the period of registration. Chapter 5 is a completed conceptual replication, again showing strong evidence for absence of unconscious instrumental conditioning – when the stimuli was hidden from conscious awareness, participants were unable to learn to approach reward-predictive stimuli and avoid punishment-predictive stimuli. This was robustly shown across two different modes of conditioning: trace (where the stimulus, response, and outcome followed one another), and delay (where the stimulus, response, and outcome overlapped temporally).

These results provide strong evidence that instrumental learning, a fundamental aspect of selective, adaptive behaviour, cannot proceed in absence of stimulus awareness.

3. *Investigate the role of instrumental action in conscious experience. According to action-oriented predictive approaches to consciousness, it is the prediction of our interactions with the world – the sensorimotor contingencies – that directly shapes conscious experience. Yet, direct investigations of this hypothesis are still lacking.*

The final aim shifted focus, attempting to investigate the role instrumental action might play in conscious experience of the world. Action-oriented approaches to perception and cognition propose that action is a vital, inseparable aspect of our experience of the world, rather than a mere form of response (Clark, 2015a).

Chapter 6 was a test of an important implication of the action-oriented predictive processing approach to consciousness, namely that conscious experience is shaped by prediction of sensorimotor contingencies. In two experiments, sensorimotor contingencies were operationalised as learned instrumental associations between stimuli and different actions. Reported access to consciousness was then measured with breaking interocular suppression in CFS (bCFS). The results demonstrate that learned instrumental (stimulus-action) associations, constituting sensorimotor predictions, speed up reported access to consciousness for predicted stimuli as opposed to unpredicted ones. This provides support for the assertion that prediction of our instrumental interactions with the world affects conscious perceptual experience.

7.2. The role of the body in instrumental learning

As reviewed in the introduction, there is extensive theoretical research pointing to the importance of bodily information in adaptive behaviour. Under this perspective, the current and expected condition of the body informs value judgments (e.g. whether stimuli in the world are rewarding or punishing to the organism), and drives motivation and invigoration of goal-directed behaviour, targeted at obtaining rewarding outcomes and avoiding negative outcomes. Empirically, behavioural and neuroimaging evidence provides support for this view, showing interconnectedness of networks related to interoception, value computation, reward, and decision-making (Gu & Fitzgerald, 2014). Another source of support comes from the evidence for autonomic markers of performance monitoring (Hajcak, McDonald, et al., 2003), akin to the much-documented orienting response, characterised by reorienting the organism's internal resources to deal with new environmental demands (Pavlov, 1927; Sokolov, 1963). Those converging lines of evidence opened an important question: to what extent is performance-related bodily (in this case, cardiac) information *functionally* important in driving adaptive behaviour, such as instrumental learning?

Chapter 2 of this thesis came short of answering that question due to strong evidence against the presence of unconscious instrumental learning at baseline. Consequently, learning was unaffected by the cardiac precision manipulation. Chapter 3 investigated this question in more detail through a measure of cardiac activity throughout the learning process. The physiological results suggest that the heart failed to differentiate between rewarding and punishing performance feedback, contrary to past evidence showing that the heart decelerates more in response to punishing feedback (indicative of an error) than to rewarding feedback (Crone et al., 2005b, 2003; van der Veen et al., 2004a). Instead, cardiac deceleration was evident for both types of feedback. Failure to observe a differentiated response to rewarding and punishing feedback suggests that it was not informative – it did not reflect a mismatch between the expected outcome, predicted by the cue, and the actual outcome. This, in turn, suggests that the unconscious cue was not processed to an extent permitting its integration with feedback. That is, cue identity was not predictive of the expected outcome, so any feedback signifying actual outcome was rendered meaningless and uninformative, as the performance monitoring mechanism was unable relate the two.

Despite evidence for absence of learning, it is also worth noting that the current cardiac feedback manipulation technique might not be robust enough for the purpose of manipulating interoceptive precision. Presenting an auditory stream corresponding to participants' own heartbeats is undoubtedly an exteroceptive signal, and there is still limited evidence regarding its processing in the brain. However, recent EEG evidence demonstrates that the brain does process external synchronous cardiac feedback in a manner similar to other self-generated sounds, suggesting that it can be differentiated from cardiac-asynchronous sounds (van Elk et al., 2014). Elsewhere, unexpected omissions from a synchronous, but not asynchronous, auditory stream were found to elicit a surprise response (Pfeiffer & De Lucia, 2017). Still, the proposal that such manipulation is in fact affecting precision of the cardiac channel, as formalised in the interoceptive inference framework (Petzschner, 2017; Petzschner et al., 2019), remains theoretical. As such, different manipulations of cardiac precision should be explored and tested to fully understand its role in other processes, as well as direct manipulations of the interoceptive signal itself. One such manipulation could be Vagus Nerve Stimulation (Paciorek & Skora, 2020), which was recently found to improve interoceptive accuracy (Villani et al., 2019).

While the setback encountered in the form of absence of instrumental learning prevented any conclusions about the role of cardiac precision in learning, it suggests that perhaps there is another, equally (or more) vital ingredient for adaptive behaviour: consciousness. Past theoretical contributions suggested that information about the body's condition is indispensable to the organisms' survival, rendering it the lowest-level substrate,

almost a foundation of subsequent processing, including consciousness and the sense of self (Damasio, 1994, 1999). Indeed, some authors argue that the body constitutes the ‘first-prior’ (Allen & Tsakiris, 2018), or a frame of reference for subsequent processes and behaviour, providing the sense of subjective “I” or “me” (Azzalini, Rebollo, & Tallon-Baudry, 2019; Tallon-Baudry & Park, 2014). Inference on bodily processes has also been proposed to be the basis of the sense of conscious self (Seth & Tsakiris, 2018) and emotion (Seth, 2013). Consequently, the body and consciousness seem tightly intertwined. With the imperative to keep the organism alive through reward pursuit, this relationship is particularly evident in adaptive behaviour. However, the exact mechanism of this relationship is still a matter of dispute, and is more readily accounted for by certain theoretical positions (e.g. embodied accounts of predictive processing or the FEP) than others. Further research is necessary to chart the details of the functional brain-body connection, and to provide further evidence for or against the theoretical implications.

7.3. The role of consciousness in instrumental learning

The requirement for consciousness for processes of varying degrees of complexity has been an extensively studied question, both theoretically and empirically. In the context of learning especially, it has been found that simpler forms of learning may be possible without conscious awareness of the stimuli, and / or of the stimulus-outcome contingencies, including classical conditioning, to emotional, visuospatial, or multisensory learning (Clark & Squire, 1998; Faivre, Mudrik, Schwartz, & Koch, 2014a; Knight et al., 2003; Lin & He, 2009; Olsson & Phelps, 2004; Rosenthal et al., 2016, 2010; Scott et al., 2018; Seitz et al., 2009a). However, more complex forms of learning, including instrumental learning, may require conscious awareness (Reber et al., 2018; Travers et al., 2018).

A fundamental reliance on conscious awareness is indeed the conclusion supported by the results of Chapters 2, 3, and 5 of this thesis (as well as the data presented in the supplementary material of Chapter 6). Those three chapters used variants of an unconscious instrumental conditioning task, where participants learned to approach reward-predictive stimuli and avoid punishment-predictive stimuli, rendered unconscious with masking or CFS. In all three cases, the results demonstrate strong evidence for the absence of unconscious instrumental learning – participants were unable to learn those associations and adjust their behaviour accordingly. Those results support the view that consciousness may be a necessary component for higher-level processing, including problem solving, decision-making, and behavioural adaptations (Baars, 2002; de Lange, van Gaal, Lamme, & Dehaene, 2011; van Gaal et al., 2012).

Instrumental learning can be classified as a complex adaptive process, requiring an agent to not only learn the associations between different stimuli, or stimuli and their outcomes, but also deploy action selectively (e.g. approach or avoid), and adapt their behaviour in the long-term. This involves integrating information over a long temporal scale and distinct modalities, as involved in processing the visual input, extracting its predictive value, deploying a selective response, processing the reinforcement, and comparing the expected outcome with the actual outcome in order to update the representations of stimulus value. As such, it is a considerably more complex process than the simpler forms of associative learning (e.g. stimulus-stimulus associations), which do not require selective decisions on whether to act or not, or a behavioural adaptation from trial to trial.

Indeed, as outlined in the theoretical overview, past theoretical and empirical accounts of conscious versus unconscious processing suggest that such an increased level of complexity should require conscious access. A number of theoretical contributions propose that consciousness is related to long-lasting, long-range connections between distinct brain regions, supporting recurrent information integration across distinct cognitive modules (Dehaene & Naccache, 2001; Dehaene & Changeux, 2011; Dehaene, Charles, King, & Marti, 2014; Lamme, 2006; Mudrik, Faivre, & Koch, 2014). As such, low-level or short-range (spatial or temporal) information integration might be possible without conscious awareness, but consciousness might be necessary at increased levels of complexity, requiring integration of information across longer spatiotemporal intervals or larger spatial distance, as in the case of instrumental learning.

Nonetheless, the exact dividing line between conscious and unconscious processing remains a topic of ongoing debate. It has been shown that certain aspects of cognitive and behavioural control can still take place unconsciously. Unconsciously presented primes were found to trigger response inhibition (Van Gaal, Ridderinkhof, Fahrenfort, Scholte, & Lamme, 2008; van Gaal, Ridderinkhof, Scholte, & Lamme, 2010), although the effect is less pronounced than for conscious primes, and occurs only in a decision-making setting (Chiu & Aron, 2014). Similarly, unconsciously presented rewards can enhance invigoration of behaviour (Bijleveld, Custers, & Aarts, 2010; Pessiglione et al., 2007), but, interestingly, they fail to engage the striatum and connected areas involved in motivated behaviour (including the motor cortex and superior temporal gyrus; Bijleveld et al., 2014). The error-monitoring network also appears to be active when errors are committed without awareness, with the ACC responding to aware and unaware errors alike (Hester, Foxe, Molholm, Shpaner, & Garavan, 2005). Similar results were obtained for the ERN, but not the Pe (Nieuwenhuis, Ridderinkhof, Blom, Band, & Kok, 2001; note that error awareness is different from stimulus awareness).

Together, those findings suggest that aspects of instrumental performance may proceed without conscious awareness. However, none of those studies investigated the entire process of instrumental learning. As such, it could be the case that an agent requires conscious awareness for novel associations to be learned and selectively acted upon, but once the associations are learned the brain can detect stimulus value and invigorate or monitor behaviour without conscious awareness. Indeed, there is some evidence that prior conscious exposure facilitates unconscious multimodal learning (Faivre et al., 2014b). Whether unconscious instrumental learning can be aided in a similar way could be an avenue for future research.

It is also noteworthy that the field of unconscious processing is inherently difficult to study, and fraught with methodological and statistical issues. Perhaps the most serious challenge lies in ensuring the genuine absence of awareness of the desired targets. It has been proposed that awareness test should be sensitive, relevant, and immediate (Berry & Dienes, 1993; Shanks, 2017; Shanks & St. John, 1994) – for example, testing for awareness immediately after the performance measure of interest, rather than in a separate task. However, a considerable portion of research, especially older studies, fail to provide appropriate awareness checks. If awareness is not properly controlled for, one cannot preclude the possibility that rudimentary awareness was in fact present throughout the task of interest and may have contaminated the putatively “unconscious” effect. Certain methodological and statistical advances have been developed to address this issue. One approach is, as mentioned, to measure presence of (or the level of) awareness on an ongoing, trial-by-trial basis, allowing to exclude any aware trials. This approach has been used in the present thesis while ensuring that individual trial exclusions did not lead to regression to the mean (Chapters 4 and 5).

However, the problem of conscious “contamination” becomes even more serious in unconscious learning research. Learning is an incremental process, so transient awareness of the target stimulus can not only contaminate the perception at that one moment, but possibly also affect its subsequent perception. In that sense, learning from a stimulus once consciously seen might not be entirely unconscious anymore. Nonetheless, this is a speculative consideration – more research is needed to address this issue and its implication for task design.

Finally, while the evidence presented in this thesis overwhelmingly supports the genuine absence of unconscious instrumental learning, one cannot ignore the possibility that there may be other ways to probe this phenomenon. Type I d' , the Signal Detection Theoretic measure used here, is superior to pure accuracy measures, but it fails to take into account the progression of learning. As such, more sensitive metrics, such as RL models, might illuminate the finer-

grained detail of an unconscious instrumental learning process⁹. It is also possible that unconscious instrumental learning may require a more sensitive task – for example, one including longer learning periods or more repetitions. This is also a matter for future research.

7.4. The role of instrumental action in shaping conscious perception

Action and conscious experience have long been proposed to be closely intertwined (James, 1891). Under the action-oriented accounts, our actions are not merely responses to the events in the world – instead, they directly shape how we perceive our environment. This relationship has been formalised by action-oriented predictive processing approaches to consciousness. Under predictive processing, our conscious experience is shaped by the predictive model that best explains the causes of sensory inputs (i.e. the best posterior prediction). Action-oriented accounts extend this view by proposing that perceptual experience is influenced not only by passive predictions about the world, but more generally by predictions encompassing the coupling or contingency between actions and sensory signals – i.e., on predictions about sensorimotor contingencies (Clark, 2015b; Seth, 2014). If that assumption is correct, then holding valid sensorimotor predictions should directly shape access to consciousness in line with those predictions.

This is indeed the pattern observed in Chapter 6 of this thesis. Stimuli which were congruent with the sensorimotor prediction instantiated through training an instrumental stimulus-action relationship showed facilitated access to reported conscious awareness, relative to stimuli which were incongruent with the prediction. Stimuli unrelated to any instrumental action behaved similarly to incongruent stimuli. Interestingly, stimuli which were completely novel – that is, never previously experienced in relation to an action or no-action – showed the fastest access to consciousness. This result suggests that while sensorimotor predictions have a direct impact on conscious experience, action might not be the dominant factor. A novel cue could be motivationally important or life-threatening, thus requiring priority processing to enable an immediate reaction. This effect also points to an open question in the predictive processing literature – the roles of predictions versus PEs in shaping perception, and the potential roles of expected precisions (proposed to be mediated by attention).

⁹ A type of RL model, Q-learning (Watkins & Dayan, 1992), was in fact cautiously applied to the learning data in Chapter 2. The learning rates estimated by the model suggested that participants failed to learn the stimulus values, and their instrumental behaviour remained erratic, rather than guided by stimulus value. This data was not included in the empirical chapters as it was considered a side-project, and a work in progress besides – the range of models and parameters which can be applied is huge, and can often considerably affect the result.

The result of Chapter 6 provides an important extension to previous findings in this line of research, which showed that action can influence perception, for example by biasing the perceived direction of stimulus motion in line with the direction of movement (e.g. Maruya, Yang, & Blake, 2007; Mitsumatsu, 2009; Suzuki, Schwartzman, Augusto, & Seth, 2019). Crucially, however, the paradigm presented in this thesis goes beyond studying a match in stimulus and percept characteristics (such as a match in directions) to evaluate the role of holding stimulus-action *predictions* on perceptual experience. This allows to directly investigate the mechanism behind conscious experience, as postulated by predictive processing.

The results showing changes in access to consciousness in line with sensorimotor predictions speak to the proposals that conscious experience is organised so as to support streamlined, dynamic interactions with the world. In this sense, conscious experience does not represent the world in the most veridical way, but rather in a way which is optimal for action (Clark, 2016; Seth, 2014, 2015; Tschantz, Seth, & Buckley, 2020). Indeed, it has been long suggested that the role of conscious experience is to provide a unified subjective frame of reference that enables efficient interaction with the world (Land, 2012; Merker, 2013).

Nonetheless, conscious experience is a complex topic to operationalise in research. Firstly, there is no agreed-upon definition of consciousness – the nature of consciousness differs based on the approach taken. For example, the Global Neuronal Workspace hypothesis might see consciousness as a result of global broadcasting of information in the brain (Dehaene & Naccache, 2001), but under the Integrated Information Theory consciousness is a property of information integration in any system (Tononi, 2004, 2012). Consequently, measures of consciousness also differ based on the definition used. Secondly, consciousness has many facets – one may think of levels of consciousness (e.g. from absence of consciousness, through being just awake, to full alertness), contents of consciousness (e.g. the phenomenal experience), or the subjective confidence in one's conscious experience. Operationalising consciousness based on the chosen aspect of consciousness results in a host of different metrics and measures (Overgaard, 2015; Persuh, 2017; Seth, Dienes, Cleeremans, Overgaard, & Pessoa, 2008). The method used here, breaking CFS, can only address the speed of access to conscious, reportable awareness (in other words, the speed of reported conscious detection of a stimulus under interocular suppression). Although a common measure, it cannot address consciousness comprehensively (indeed, it is likely that no single measure can). As such, future research is needed to assess the role of action in the other aspects of consciousness.

7.5. Conclusions

Instrumental learning is a fundamental substrate of adaptive behaviour, allowing humans and other animals to interact with their environments to promote survival. The present work investigated human instrumental learning from three perspectives: focusing on the roles of bodily information, consciousness, and action in shaping instrumental learning and the associated experience of interacting with the world. While adaptive behaviour – including instrumental learning – is a vast and complex phenomenon, the present thesis contributed to charting the potential influences from the body (on the outside, and on the inside) and consciousness on adaptive learning processes.

The main finding of the empirical research contained in the present thesis is that conscious access is a key ingredient for instrumental learning to proceed successfully. Across multiple experiments, the results show that instrumental learning cannot proceed in the absence of conscious awareness of the stimuli. Unfortunately, this finding prevents drawing conclusions about the functional role of cardiac precision in unconscious instrumental learning, one of the three main aims of this project. Nonetheless, this result may inform further research into the intersection and interrelationships between the body and consciousness in adaptive behaviour. Finally, while showing that consciousness may be needed for instrumental behaviour, we also show that instrumental behaviour further shapes conscious experience of the world in line with it.

Together, those findings point to a tight, bidirectional relationship between consciousness and the body – on the outside, and on the inside – where the two are not independent of each other, but rather a part of the same package shaping the agents' interactions with their environments. Consequently, this work extends the current understanding of instrumental learning as a fundamental component of adaptive behaviour, showing that conscious access is required to drive adaptive interactions with the world, and to further shape our conscious experience.

Reference list

- Adams, R. A., Shipp, S., & Friston, K. J. (2013). Predictions not commands: active inference in the motor system. *Brain Structure and Function*, 218, 611–643. Retrieved from <http://www.fil.ion.ucl.ac.uk/~karl/Predictions not commands - active inference in the motor system.pdf>
- Adams, R. A., Stephan, K. E., Brown, H. R., Frith, C. D., & Friston, K. J. (2013). The Computational Anatomy of Psychosis. *Frontiers in Psychiatry*, 4, 47. <https://doi.org/10.3389/fpsy.2013.00047>
- Al-Aidroos, N., Guo, R. M., & Pratt, J. (2010). You can't stop new motion: Attentional capture despite a control set for colour. *Visual Cognition*, 18(6), 859–880. <https://doi.org/10.1080/13506280903343085>
- Allen, M., & Friston, K. J. (2016). From cognitivism to autopoiesis: towards a computational framework for the embodied mind. *Synthese*. <https://doi.org/10.1007/s11229-016-1288-5>
- Allen, M., & Tsakiris, M. (2018). The body as first prior: Interoceptive predictive processing and the primacy of self-models. In M. Tsakiris & H. De Preester (Eds.), *The Interoceptive Mind: from Homeostasis to Awareness* (pp. 27–45). Oxford, UK: Oxford University Press.
- Apps, M. A. J., & Tsakiris, M. (2014). The free-energy self: A predictive coding account of self-recognition. *Neuroscience and Biobehavioral Reviews*. <https://doi.org/10.1016/j.neubiorev.2013.01.029>
- Aru, J., Rutiku, R., Wibral, M., Singer, W., & Melloni, L. (2016). Early effects of previous experience on conscious perception. *Neuroscience of Consciousness*, 2, 1–10. <https://doi.org/10.1093/nc/niw004>
- Aspell, J. E., Heydrich, L., Marillier, G., Lavanchy, T., Herbelin, B., & Blanke, O. (2013). Turning Body and Self Inside Out: Visualized Heartbeats Alter Bodily Self-Consciousness and Tactile Perception. *Psychological Science*, 24(12), 2445–2453. <https://doi.org/10.1177/0956797613498395>
- Atas, A., Faivre, N., Timmermans, B., Cleeremans, A., & Kouider, S. (2014). Nonconscious learning from crowded sequences. *Psychological Science*, 25(1), 113–119. <https://doi.org/10.1177/0956797613499591>
- Axelrod, V., & Rees, G. (2014). Conscious awareness is required for holistic face processing. *Consciousness and Cognition*, 27, 233–245. <https://doi.org/10.1016/j.concog.2014.05.004>
- Azzalini, D., Rebollo, I., & Tallon-Baudry, C. (2019). Visceral Signals Shape Brain Dynamics and Cognition. *Trends in Cognitive Sciences*. <https://doi.org/10.1016/J.TICS.2019.03.007>

- Baars, B. J. (2002). The conscious access hypothesis: origins and recent evidence. *Trends in Cognitive Sciences*, 6(1), 47–52. [https://doi.org/10.1016/S1364-6613\(00\)01819-2](https://doi.org/10.1016/S1364-6613(00)01819-2)
- Baars, B. J., Ramsøy, T. Z., & Laureys, S. (2003). Brain, conscious experience and the observing self. *Trends in Neurosciences*, 26(12), 671–675. <https://doi.org/10.1016/j.tins.2003.09.015>
- Balleine, B., Ball, J., & Dickinson, A. (1994). Benzodiazepine-induced outcome revaluation and the motivational control of instrumental action in rats. *Behavioral Neuroscience*, 108(3), 573–589. <https://doi.org/10.1037/0735-7044.108.3.573>
- Balleine, B., Davies, A., & Dickinson, A. (1995). Cholecystokinin Attenuates Incentive Learning in Rats. *Behavioral Neuroscience*, 109(2), 312–319. <https://doi.org/10.1037/0735-7044.109.2.312>
- Balleine, B. W. (2011). Sensation, Incentive Learning, and the Motivational Control of Goal-Directed Action. In J. A. Gottfried (Ed.), *Neurobiology of Sensation and Reward* (pp. 287–310). Boca Raton, FL: CRC Press/Taylor & Francis. Retrieved from <http://oreilly.com/catalog/errata.csp?isbn=9781449340377>
- Balleine, B. W., Delgado, M. R., & Hikosaka, O. (2007). The Role of the Dorsal Striatum in Reward and Decision-Making. *Journal of Neuroscience*, 27, 8161–8165. <https://doi.org/10.1523/JNEUROSCI.1554-07.2007>
- Balleine, B. W., & Dickinson, A. (1998a). Goal-directed instrumental action: contingency and incentive learning and their cortical substrates. *Neuropharmacology*, 37(4–5), 407–419. [https://doi.org/10.1016/S0028-3908\(98\)00033-1](https://doi.org/10.1016/S0028-3908(98)00033-1)
- Balleine, B. W., & Dickinson, A. (1998b). Goal-directed instrumental action: contingency and incentive learning and their cortical substrates. *Neuropharmacology*, 37, 407–419. Retrieved from https://ac.els-cdn.com/S0028390898000331/1-s2.0-S0028390898000331-main.pdf?_tid=d6c235fe-b0e8-11e7-8317-00000aab0f02&acdnat=1507990154_9f524296100ca81a42a6af6a5da07d0d
- Balleine, B. W., & O'Doherty, J. P. (2010). Human and Rodent Homologies in Action Control: Corticostriatal Determinants of Goal-Directed and Habitual Action. *Neuropsychopharmacology Reviews*, 35, 48–69. <https://doi.org/10.1038/npp.2009.131>
- Balleine, B. W., & Ostlund, S. B. (2007). Still at the choice-point: Action selection and initiation in instrumental conditioning. *Annals of the New York Academy of Sciences*, 1104, 147–171. <https://doi.org/10.1196/annals.1390.006>
- Baltieri, M., & Buckley, C. L. (2017). An active inference implementation of phototaxis. In *Proceedings of the ECAL* (pp. 36–43). Lyon, France.
- Banellis, L., & Cruse, D. (2020). *Skipping a beat: heartbeat-evoked potentials reflect predictions*

during interoceptive-exteroceptive integration.

<https://doi.org/10.1101/2020.04.21.053173>

Barrett, L. F. (2016). The theory of constructed emotion: an active inference account of interoception and categorization. *Social Cognitive and Affective Neuroscience*.

<https://doi.org/10.1093/scan/nsw154>

Barrett, L. F., & Bar, M. (2009). See it with feeling: affective predictions during object perception. *Philosophical Transactions of the Royal Society B*, 364, 1325–1334.

<https://doi.org/10.1098/rstb.2008.0312>

Barrett, L. F., Quigley, K. S., & Hamilton, P. (2016). An active inference theory of allostasis and interoception in depression. *Philosophical Transactions of the Royal Society B*, 371, 1–17.

<https://doi.org/10.1098/rstb.2016.0011>

Barrett, L. F., & Simmons, K. W. (2015). Interoceptive predictions in the brain. *Nature Reviews Neuroscience*, 16, 419–429. <https://doi.org/10.1038/nrn3950>

Bates, D., Mächler, M., Bolker, B., & Walker, S. (2015). Fitting Linear Mixed-Effects Models Using **lme4**. *Journal of Statistical Software*, 67(1). <https://doi.org/10.18637/jss.v067.i01>

Bechara, A., Damasio, H., Tranel, D., & Damasio, A. R. (1997). Deciding advantageously before knowing the advantageous strategy. *Science*, 275(94), 2193–2195. Retrieved from <http://science.sciencemag.org/content/sci/275/5304/2193.full.pdf>

Bechara, A. (2004). The role of emotion in decision-making: Evidence from neurological patients with orbitofrontal damage. *Brain and Cognition*, 55(30), 40.

<https://doi.org/10.1016/j.bandc.2003.04.001>

Bechara, A., Damasio, A. R., Damasio, H., & Anderson, S. W. (1994). Insensitivity to future consequences following damage to human prefrontal cortex. *Cognition*, 50(1–3), 7–15.

[https://doi.org/10.1016/0010-0277\(94\)90018-3](https://doi.org/10.1016/0010-0277(94)90018-3)

Bechara, A., Tranel, D., & Damasio, H. (2000). *Characterization of the decision-making deficit of patients with ventromedial prefrontal cortex lesions*. *Brain* (Vol. 123).

Bechara, A., Tranel, D., Damasio, H., & Damasio, A. R. (1996). Failure to Respond Autonomically to Anticipated Future Outcomes Following Damage to Prefrontal Cortex. *Cerebral Cortex*, 6(2), 215–225. Retrieved from <https://academic.oup.com/cercor/article-abstract/6/2/215/338675>

Bechara, Antoine, Damasio, H., Damasio, A. R., & Lee, G. P. (1999). Different Contributions of the Human Amygdala and Ventromedial Prefrontal Cortex to Decision-Making. *Journal of Neuroscience*, 19, 5473–5481.

Becker, S. I., & Horstmann, G. (2011). Novelty and saliency in attentional capture by unannounced motion singletons. *Acta Psychologica*, 136(3), 290–299.

<https://doi.org/10.1016/j.actpsy.2010.12.002>

- Beets, I. A. M., 't Hart, B. M., Rösler, F., Henriques, D. Y. P., Einhäuser, W., & Fiehler, K. (2010). Online action-to-perception transfer: Only percept-dependent action affects perception. *Vision Research*, 50, 2633–2641. <https://doi.org/10.1016/j.visres.2010.10.004>
- Benarroch, E. E. (1993). The Central Autonomic Network: Functional organization, dysfunction, and perspective. *Mayo Clinic Proceedings*, 68(10), 988–1001. [https://doi.org/10.1016/S0025-6196\(12\)62272-1](https://doi.org/10.1016/S0025-6196(12)62272-1)
- Benarroch, E. E. (2004). Autonomy of Central Autonomic Areas. In D. Robertson, I. Biaggioni, G. Burnstock, & P. Low (Eds.), *Primer on the Autonomic Nervous System* (2nd Editio, Vol. 2 ed, pp. 17–19).
- Berntson, G. G., & Cacioppo, J. T. (2007). Integrative Physiology: Homeostasis, Allostasis, and the Orchestration of Systemic Physiology. In *Handbook of Psychophysiology* (pp. 433–452). <https://doi.org/10.1017/CBO9780511546396>
- Berntson, G. G., Cacioppo, J. T., & Bosch, J. A. (2016). From Homeostasis to Allodynamic Regulation. In J. T. Cacioppo, L. G. Tassinary, & G. G. Berntson (Eds.), *Handbook of Psychophysiology* (pp. 401–426). Cambridge, UK: Cambridge University Press. <https://doi.org/10.1017/9781107415782.018>
- Berntson, G. G., Cacioppo, J. T., & Quigley, K. S. (1993). Respiratory sinus arrhythmia: Autonomic origins, physiological mechanisms, and psychophysiological implications. *Psychophysiology*, 30(2), 183–196. <https://doi.org/10.1111/j.1469-8986.1993.tb01731.x>
- Berntson, G. G., Gianaros, P. J., & Tsakiris, M. (2018). Interoception and the autonomic nervous system: Bottom-up meets top-down. In Manos Tsakiris & H. De Preester (Eds.), *The Interoceptive Mind: from Homeostasis to Awareness* (pp. 3–25). New York: Oxford University Press.
- Berridge, K. C. (2000). Reward learning: Reinforcement, incentives, and expectations. *Psychology of Learning and Motivation*, 40, 223–278. [https://doi.org/10.1016/S0079-7421\(00\)80022-5](https://doi.org/10.1016/S0079-7421(00)80022-5)
- Berry, D. C., & Dienes, Z. (1993). *Implicit Learning: Theoretical and Empirical Issues*. Hove, UK: Lawrence Erlbaum.
- Bijleveld, E., Custers, R., & Aarts, H. (2010). Unconscious reward cues increase invested effort, but do not change speed-accuracy tradeoffs. *Cognition*, (115), 330–335. <https://doi.org/10.1016/j.cognition.2009.12.012>
- Bijleveld, E., Custers, R., & Aarts, H. (2012a). Adaptive reward pursuit: How effort requirements affect unconscious reward responses and conscious reward decisions. *Journal of Experimental Psychology: General*, 141, 728–742. <https://doi.org/10.1037/a0027615>

- Bijleveld, E., Custers, R., & Aarts, H. (2012b). Human reward pursuit: From rudimentary to higher-level functions. *Current Directions in Psychological Science*, (21), 194–199. <https://doi.org/10.1177/0963721412438463>
- Bijleveld, E., Custers, R., Van der Stigchel, S., Aarts, H., Pas, P., & Vink, M. (2014). Distinct neural responses to conscious versus unconscious monetary reward cues. *Human Brain Mapping*, 35, 5578–5586. <https://doi.org/10.1002/hbm.22571>
- Birbaumer, N., Klucken, T., Angrilli, A., Andreatta, M., & Pauli, P. (2015). Appetitive vs. Aversive conditioning in humans. *Front. Behav. Neurosci*, 9, 128. <https://doi.org/10.3389/fnbeh.2015.00128>
- Blake, R., & Fox, R. (1974). Adaptation to invisible gratings and the site of binocular rivalry suppression. *Nature*, 249(5456), 488–490. <https://doi.org/10.1038/249488a0>
- Blessing, W. (1997). *The Lower Brainstem and Bodily Homeostasis*. *Annals of Neurology* (Vol. 43). Oxford, UK: Oxford University Press. <https://doi.org/10.1002/ana.410430624>
- Brainard, D. H. (1997). The Psychophysics Toolbox. *Spatial Vision*, 10(4), 433–436. Retrieved from <http://www.ncbi.nlm.nih.gov/pubmed/9176952>
- Breitmeyer, B. G. (2015). Psychophysical “blinding” methods reveal a functional hierarchy of unconscious visual processing. *Consciousness and Cognition*, 35, 234–250. <https://doi.org/10.1016/j.concog.2015.01.012>
- Brener, J., & Ring, C. (2016). Towards a psychophysics of interoceptive processes: the measurement of heartbeat detection. *Philosophical Transactions of the Royal Society B*, 371, 1–9. <https://doi.org/10.1098/rstb.2016.0015>
- Buser, P., & Rougeul, A. (1961). Somatic functions of the nervous system. *Annual Review of Physiology*, 23, 387–418. Retrieved from www.annualreviews.org
- Butz, M. V, Thomaschke, R., Linhardt, M. J., & Herbolt, O. (2010). Remapping motion across modalities: tactile rotations influence visual motion judgments. *Experimental Brain Research*, 207, 1–11. <https://doi.org/10.1007/s00221-010-2420-2>
- Cameron, O. G. (2002). *Visceral sensory neuroscience: Interoception*. Oxford, UK: Oxford University Press.
- Campbell, M. C., Stout, J. C., & Finn, P. R. (2004). Reduced autonomic responsiveness to gambling task losses in Huntington’s disease. *Journal of the International Neuropsychological Society*, (10), 239–245. <https://doi.org/10.1017/S1355617704102105>
- Cannon, W. (1915). *Bodily changes in pain, hunger, fear and rage: An account of recent researches into the function of emotional excitement*. D Appleton & Company.
- Cannon, W. B. (1929). Organisation for physiological homeostasis. *Physiological Reviews*, IX(3), 399–431.

- Carter, S., & Smith Pasqualini, M. (2004). Stronger autonomic response accompanies better learning: A test of Damasio's somatic marker hypothesis. *Cognition and Emotion*, 18, 901–911. Retrieved from <http://www.tandfonline.com/doi/pdf/10.1080/02699930341000338>
- Cechetti, D. F., & Shoemaker, J. K. (2009). Functional neuroanatomy of autonomic regulation. *NeuroImage*, 47(3), 795–803. <https://doi.org/10.1016/j.neuroimage.2009.05.024>
- Cheesman, J., & Merikle, P. M. (1984). Priming with and without awareness. *Perception & Psychophysics*, 36(4), 387–395. <https://doi.org/10.3758/BF03202793>
- Cheesman, J., & Merikle, P. M. (1986). Distinguishing conscious from unconscious perceptual processes. *Canadian Journal of Psychology*, 40(4), 343–367. <https://doi.org/10.1037/h0080103>
- Chiu, Y.-C., & Aron, A. R. (2014). Unconsciously triggered response inhibition requires an executive setting. *Journal of Experimental Psychology: General*, 143, 56–61. <https://doi.org/10.1037/a0031497>
- Clark, A. (1998). *Being there. Putting brain, body, and world together again*. Cambridge, MA: MIT Press.
- Clark, A. (2013). Whatever next? Predictive brains, situated agents, and the future of cognitive science. *Behavioral and Brain Sciences*, 36, 181–253. <https://doi.org/10.1017/S0140525X12000477>
- Clark, A. (2015a). Embodied Prediction. *Open MIND*, 7, 1–21. <https://doi.org/10.15502/9783958570115>
- Clark, A. (2015b). Radical Predictive Processing. *The Southern Journal of Philosophy*, 53, 3–27. <https://doi.org/10.1111/sjp.12120>
- Clark, A. (2016). *Surfing Uncertainty: Prediction, action and the embodied mind*. Oxford, UK: Oxford University Press.
- Clark, R. E., & Squire, L. R. (1998). Classical conditioning and brain systems: The role of awareness. *Science*, 280, 77–81. Retrieved from <http://science.sciencemag.org/>
- Corcoran, A. W., & Hohwy, J. (2018). Allostasis, interoception, and the free energy principle: Feeling our way forward. In M. Tsakiris & H. De Preester (Eds.), *The Interoceptive Mind: From homeostasis to awareness* (pp. 272–292). Oxford, UK: Oxford University Press. <https://doi.org/10.1093/oso/9780198811930.001.0001>
- Craig, A. D. (2009). How do you feel - now? The anterior insula and human awareness. *Nature Reviews Neuroscience*, 10, 59–70. Retrieved from www.nature.com/reviews/neuro
- Craig, A. D. (2013, December). An interoceptive neuroanatomical perspective on feelings, energy, and effort. *Behavioral and Brain Sciences*. Cambridge University Press.

<https://doi.org/10.1017/S0140525X13001489>

Craig, A. D. (2014). *How do you feel? An interoceptive moment with your neurobiological self*. Princeton, NJ: Princeton University Press.

Critchley, H. D. (2005). Neural Mechanisms of Autonomic, Affective, and Cognitive Integration. *J. Comp. Neurol*, 493, 154–166. <https://doi.org/10.1002/cne.20749>

Critchley, H. D., Elliott, R., Mathias, C. J., & Dolan, R. J. (2000). *Neural Activity Relating to Generation and Representation of Galvanic Skin Conductance Responses: A Functional Magnetic Resonance Imaging Study*.

Critchley, H. D., & Garfinkel, S. N. (2017). Interoception and emotion. *Current Opinion in Psychology*, 17, 7–14. <https://doi.org/10.1016/j.copsyc.2017.04.020>

Critchley, H. D., & Garfinkel, S. N. (2018). The influence of physiological signals on cognition. *Current Opinion in Behavioral Sciences*, 19, 13–18. <https://doi.org/10.1016/j.cobeha.2017.08.014>

Critchley, H. D., & Harrison, N. A. (2013, February 20). Visceral Influences on Brain and Behavior. *Neuron*. Elsevier. <https://doi.org/10.1016/j.neuron.2013.02.008>

Critchley, H. D., & Rolls, E. T. (1996). *Olfactory Neuronal Responses in the Primate Orbitofrontal Cortex: Analysis in an Olfactory Discrimination Task*. *JOURNAL OF NEUROPHYSIOLOGY* (Vol. 75).

Critchley, H. D., Tang, J., Glaser, D., Butterworth, B., & Dolan, R. J. (2005). Anterior cingulate activity during error and autonomic response. *NeuroImage*, 27(4), 885–895. <https://doi.org/10.1016/j.neuroimage.2005.05.047>

Crone, E. A., Bunge, S. A., de Klerk, P., & van der Molen, M. W. (2005a). Cardiac concomitants of performance monitoring: Context dependence and individual differences. *Cognitive Brain Research*, 23(1), 93–106. <https://doi.org/10.1016/J.COGBRAINRES.2005.01.009>

Crone, E. A., Bunge, S. A., de Klerk, P., & van der Molen, M. W. (2005b). Cardiac concomitants of performance monitoring: Context dependence and individual differences. *Cognitive Brain Research*, 23(1), 93–106. <https://doi.org/10.1016/j.cogbrainres.2005.01.009>

Crone, E. A., Somsen, R. J. M., Beek, B. Van, & Van Der Molena, M. W. (2004a). Heart rate and skin conductance analysis of antecedents and consequences of decision making. *Psychophysiology*, 41(4), 531–540. <https://doi.org/10.1111/j.1469-8986.2004.00197.x>

Crone, E. A., Somsen, R. J. M., Beek, B. Van, & Van Der Molena, M. W. (2004b). Heart rate and skin conductance analysis of antecedents and consequences of decision making. *Psychophysiology*, 41(4), 531–540. <https://doi.org/10.1111/j.1469-8986.2004.00197.x>

Crone, E. A., van der Veen, F. M., van der Molen, M. W., Somsen, R. J. M., van Beek, B., & Jennings, J. R. (2003). Cardiac concomitants of feedback processing. *Biological*

- Psychology*, 64(1–2), 143–156. [https://doi.org/10.1016/S0301-0511\(03\)00106-6](https://doi.org/10.1016/S0301-0511(03)00106-6)
- Damasio, A. (1994). *Descartes' error*. New York, USA: Vintage.
- Damasio, A. (1999). *The feeling of what happens: Body, emotion, and the making of consciousness*. London, UK: Vintage.
- Damasio, A. (2004). *Looking for Spinoza*. London, UK: Vintage.
- Damasio, A. R. (1999). *The Feeling of What Happens: Body and emotion in the making of consciousness*. New York: Harcourt Brace & Company.
- Daw, N. D., Niv, Y., & Dayan, P. (2005). Uncertainty-based competition between prefrontal and dorsolateral striatal systems for behavioral control. *Nature Neuroscience*, 8(12), 1704–1711. <https://doi.org/10.1038/nn1560>
- Day, T. A. (2005). Defining stress as a prelude to mapping its neurocircuitry: No help from allostasis. <https://doi.org/10.1016/j.pnpbp.2005.08.005>
- Dayan, P., & Daw, N. D. (2008). Decision theory, reinforcement learning, and the brain. *Cognitive, Affective, & Behavioral Neuroscience*, 8(4), 429–453. <https://doi.org/10.3758/CABN.8.4.429>
- Dayan, Peter, & Balleine, B. W. (2002). Reward, motivation, and Reinforcement Learning. *Neuron*, 36, 285–298. [https://doi.org/10.1016/S0896-6273\(02\)00963-7](https://doi.org/10.1016/S0896-6273(02)00963-7)
- Dayan, Peter, & Berridge, K. C. (2014). Model-based and model-free Pavlovian reward learning: Revaluation, revision, and revelation. <https://doi.org/10.3758/s13415-014-0277-8>
- Dayan, Peter, & Niv, Y. (2008). Reinforcement learning: The Good, The Bad and The Ugly. *Current Opinion in Neurobiology*, 18(2), 185–196. <https://doi.org/10.1016/j.conb.2008.08.003>
- de Lange, F. P., van Gaal, S., Lamme, V. A. F., & Dehaene, S. (2011). How awareness changes the relative weights of evidence during human decision-making. *PLoS Biology*, 9(11). <https://doi.org/10.1371/journal.pbio.1001203>
- Dehaene, S., & Naccache, L. (2001). Towards a cognitive neuroscience of consciousness: Basic evidence and a workspace framework. *Cognition*, 79, 1–37. Retrieved from https://www.jsmf.org/meetings/2003/nov/Dehaene_Cognition_2001.pdf
- Dehaene, S., & Changeux, J.-P. (2011). Experimental and Theoretical Approaches to Conscious Processing. *Neuron*, 70(2), 200–227. <https://doi.org/10.1016/J.NEURON.2011.03.018>
- Dehaene, Stanislas, Charles, L., King, J.-R., & Marti, S. (2014). Toward a computational theory of conscious processing. *Current Opinion in Neurobiology*, 25, 76–84. <https://doi.org/10.1016/j.conb.2013.12.005>
- Dehaene, Stanislas, & Naccache, L. (2001). Towards a cognitive neuroscience of consciousness: basic evidence and a workspace framework. *Cognition*, 79, 1–37. Retrieved from

www.elsevier.com/locate/cognit

- Dehaene, Stanislas, Naccache, L., Cohen, L., Bihan, D. Le, Mangin, J.-F., Poline, J.-B., & Rivière, D. (2001). Cerebral mechanisms of word masking and unconscious repetition priming. *Nature Neuroscience*, 4(7), 752–758. <https://doi.org/10.1038/89551>
- Dehaene, Stanislas, Sergent, C., & Changeux, J.-P. (2003). A neuronal network model linking subjective reports and objective physiological data during conscious perception. *Proceedings of the National Academy of Sciences*, 100(14), 8520–8525. <https://doi.org/10.1073/pnas.1332574100>
- Desmedt, O., Luminet, O., & Corneille, O. (2018). The heartbeat counting task largely involves non-interoceptive processes: Evidence from both the original and an adapted counting task. *Biological Psychology*, 138, 185–188. <https://doi.org/10.1016/j.biopsycho.2018.09.004>
- Devinsky, O., Morrell, M. J., & Vogt, B. A. (1995). *Contributions of anterior cingulate cortex to behaviour*. *Brain* (Vol. 118). Oxford University Press. Retrieved from <http://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.667.1375&rep=rep1&type=pdf>
- Dewey, J. (1896). The reflex arc concept in psychology. *Psychological Review*, 3, 357–370.
- Dezfouli, A., & Balleine, B. W. (2012). Habits, action sequences and reinforcement learning. *European Journal of Neuroscience*, 35(7), 1036–1051. <https://doi.org/10.1111/j.1460-9568.2012.08050.x>
- Di Pace, E., & Saracini, C. (2014). Action Imitation Changes Perceptual Alternations in Binocular Rivalry. *PLoS ONE*, 9(5). <https://doi.org/10.1371/journal.pone>
- Dickinson, A. (1994). Instrumental conditioning. In *Animal learning and cognition* (pp. 45–78). https://doi.org/10.1007/978-3-540-68706-1_343
- Dickinson, A., & Balleine, B. (1994). Motivational control of goal-directed action. *Animal Learning & Behavior*, 22(1), 1–18. <https://doi.org/10.3758/BF03199951>
- Dickinson, A., & Dawson, G. R. (1987). *Pavlovian Processes in the Motivational Control of Instrumental Performance*. *The Quarterly Journal of Experimental Psychology* (Vol. 39). Retrieved from <https://journals.sagepub.com/doi/pdf/10.1080/14640748708402264>
- Dienes, Z. (2004). Assumptions of Subjective Measures of Unconscious Mental States. Higher Order Thoughts and Bias. *Journal of Consciousness Studies*, 9, 25–45.
- Dienes, Z. (2007). Subjective measures of unconscious knowledge. *Progress in Brain Research*, 168, 49–269. [https://doi.org/10.1016/S0079-6123\(07\)68005-4](https://doi.org/10.1016/S0079-6123(07)68005-4)
- Dienes, Z. (2014). Using Bayes to get the most out of non-significant results, 5(781), 1–17. <https://doi.org/10.3389/fpsyg.2014.00781>

- Dienes, Z. (2015a). How bayesian statistics are needed to determine whether mental states are unconscious. In M. Overgaard (Ed.), *Behavioral Methods in Consciousness Research* (pp. 199–220). Oxford: Oxford University Press.
- Dienes, Z. (2015b). How many participants might I need? - YouTube. Retrieved November 14, 2019, from https://www.youtube.com/watch?v=10Lsm_o_GRg
- Dienes, Z. (2016). How Bayes factors change scientific practice. *Journal of Mathematical Psychology*, 72, 78–89. <https://doi.org/10.1016/j.jmp.2015.10.003>
- Dienes, Z. (2019). How do I know what my theory predicts? *Advances in Methods and Practices in Psychological Science*, 1–14. <https://doi.org/10.1177/2515245919876960>
- Dienes, Z., & McLatchie, N. (2017). Four reasons to prefer Bayesian analyses over significance testing. *Psychonomic Bulletin & Review*, 25(1), 207–218. <https://doi.org/10.3758/s13423-017-1266-z>
- Dienes, Zoltan. (2014). Using Bayes to get the most out of non-significant results. *Frontiers in Psychology*, 5, 781. <https://doi.org/10.3389/fpsyg.2014.00781>
- Dienes, Zoltan, & Seth, A. K. (2010). Measuring any conscious content versus measuring the relevant conscious content: Comment on Sandberg et al. *Consciousness and Cognition*, 19(19), 1079–1080. <https://doi.org/10.1016/j.concog.2010.03.009>
- Dogge, M., Gayet, S., Custers, R., & Aarts, H. (2018). The influence of action-effect anticipation on bistable perception: differences between onset rivalry and ambiguous motion. *Neuroscience of Consciousness*, 4(1), 1–12. <https://doi.org/10.1093/nc/niy004>
- Dolega, K. (2017). Moderate predictive processing. *Philosophy and Predictive Processing*, (March), 154–172. <https://doi.org/10.15502/9783958573116>
- Draine, S. C., & Greenwald, A. G. (1998). Replicable Unconscious Semantic Priming. *Journal of Experimental Psychology: General*, 127(3), 286–303.
- Dunn, B. D., Dalgleish, T., & Lawrence, A. D. (2006). The somatic marker hypothesis: A critical evaluation. *Neuroscience and Biobehavioral Reviews*, 30, 239–271. <https://doi.org/10.1016/j.neubiorev.2005.07.001>
- Dunn, B. D., Galton, H. C., Morgan, R., Evans, D., Oliver, C., Meyer, M., ... Dalgleish, T. (2010). Listening to your heart: How interoception shapes emotion experience and intuitive decision making. *Psychological Science*, 21(12), 1835–1844. <https://doi.org/10.1177/0956797610389191>
- Dupoux, E., Gardelle, V. de, & Kouider, S. (2008). Subliminal speech perception and auditory streaming. *Cognition*, 109(2), 267–273. <https://doi.org/10.1016/j.cognition.2008.06.012>
- Duschek, S., Werner, N. S., & Reyes del Paso, G. A. (2013). The behavioral impact of baroreflex function: A review. *Psychophysiology*, 50, 1183–1193.

<https://doi.org/10.1111/psyp.12136>

Dworkin, B. R. (1981). The role of instrumental learning in the organization and maintenance of physiological control mechanisms. *Brain and Behaviour*, 17, 169–176.

<https://doi.org/10.1016/b978-0-08-027338-9.50029-3>

Elliott, R., Friston, K. J., & Dolan, R. J. (2000). *Dissociable Neural Responses in Human Reward Systems*.

Engel, A. K., Friston, K. J., & Kragic, D. (2015). *The pragmatic turn: Toward action-oriented views in cognitive science*. Cambridge, MA: MIT Press.

Ernst, D., Becker, S., & Horstmann, G. (2020). Novelty competes with saliency for attention. *Vision Research*, 168, 42–52. <https://doi.org/10.1016/j.visres.2020.01.004>

Faivre, N., Mudrik, L., Schwartz, N., & Koch, C. (2014a). Multisensory integration in complete unawareness: Evidence from audiovisual congruency priming. *Psychological Science*, 25(11). <https://doi.org/10.1177/0956797614547916>

Faivre, N., Mudrik, L., Schwartz, N., & Koch, C. (2014b). Multisensory intergration in complete unawareness: Evidence from audiovisual congruency priming. *Psychological Science*, 25, 2006–2016. <https://doi.org/10.1177/0956797614547916>

Falkenstein, M., Hoormann, J., Christ, S., & Hohnsbein, J. (2000). ERP components on reaction errors and their functional significance: a tutorial. *Biological Psychology*, 51(2–3), 87–107. [https://doi.org/10.1016/S0301-0511\(99\)00031-9](https://doi.org/10.1016/S0301-0511(99)00031-9)

Falkenstein, M., Hoormann, J., & Hohnsbein, J. (1999). ERP components in Go/Nogo tasks and their relation to inhibition. *Acta Psychologica*, 101(2–3), 267–291. [https://doi.org/10.1016/S0001-6918\(99\)00008-6](https://doi.org/10.1016/S0001-6918(99)00008-6)

Faul, F., Erdfelder, E., Lang, A.-G., & Buchner, A. (2007). G*Power 3: a flexible statistical power analysis program for the social, behavioral, and biomedical sciences. *Behavior Research Methods*, 39(2), 175–191. <https://doi.org/10.3758/bf03193146>

Feldman, H., & Friston, K. J. (2010). Attention, uncertainty, and free-energy. *Frontiers in Human Neuroscience*, 4(215), 1. <https://doi.org/10.3389/fnhum.2010.00215>

Fodor, J. (1983). *The modularity of mind: An essay on faculty psychology*. Cambridge, MA: MIT Press.

Forkmann, T., Scherer, A., Meessen, J., Michal, M., Schächinger, H., Vögele, C., & Schulz, A. (2016). Making sense of what you sense: Disentangling interoceptive awareness, sensibility and accuracy. *International Journal of Psychophysiology*, 109, 71–80. <https://doi.org/10.1016/j.ijpsycho.2016.09.019>

Fox, J., & Weisberg, S. (2011). *An R companion to applied regression*. Thousand Oaks, California: SAGE Publications.

- Friston, K. (2009). The free-energy principle: a rough guide to the brain? *Trends in Cognitive Sciences*, 13, 293–301. <https://doi.org/10.1016/j.tics.2009.04.005>
- Friston, K. (2011). What Is Optimal about Motor Control? *Neuron*, 72, 488–498. <https://doi.org/10.1016/j.neuron.2011.10.018>
- Friston, K. J. (2010). The free energy principle: A unified brain theory? *Nature Reviews: Neuroscience*, 11, 127–138. <https://doi.org/10.1038/nrn2787>
- Friston, K. J., Adams, R. A., Perrinet, L., & Breakspear, M. (2012). Perceptions as Hypotheses: Saccades as Experiments. *Frontiers in Psychology*, 3(151), 1–20. <https://doi.org/10.3389/fpsyg.2012.00151>
- Friston, K. J., Daunizeau, J., Kilner, J., Kiebel, S. J., Friston, K. J., Daunizeau, J., ... Kiebel, S. J. (2010). Action and behavior: a free-energy formulation. *Biological Cybernetics*, 102, 227–260. <https://doi.org/10.1007/s00422-010-0364-z>
- Friston, K., & Kiebel, S. (2009). Predictive coding under the free-energy principle. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 364, 1211–1221. <https://doi.org/10.1098/rstb.2008.0300>
- Friston, Karl, Rigoli, F., Ognibene, D., Mathys, C., Fitzgerald, T., & Pezzulo, G. (2015). Active inference and epistemic value. *Cognitive Neuroscience*, 6(4), 187–224. <https://doi.org/10.1080/17588928.2015.1020053>
- Gaillard, R., Cohen, L., Adam, C., Clemenceau, S., Hasboun, D., Baulac, M., ... Naccache, L. (2007). Subliminal words durably affect neuronal activity. *Neuroreport*, 18(15), 1527–1531. <https://doi.org/10.1097/WNR.0b013e3282f0b6cd>
- Galvez-Pol, A., McConnell, R., & Kilner, J. (2018). Voluntary visual search is modulated by the cardiac cycle. <https://doi.org/10.1101/405902>
- García-Cordero, I., Esteves, S., Mikulan, E. P., Hesse, E., Baglivo, F. H., Silva, W., ... Sedeño, L. (2017). Attention, in and Out: Scalp-Level and Intracranial EEG Correlates of Interoception and Exteroception. *Frontiers in Neuroscience*, 11, 411. <https://doi.org/10.3389/fnins.2017.00411>
- Garfinkel, S. N., & Critchley, H. D. (2016). Threat and the Body: How the Heart Supports Fear Processing. *Trends in Cognitive Sciences*, 20, 34–46. <https://doi.org/10.1016/j.tics.2015.10.005>
- Garfinkel, S. N., Minati, L., Gray, M. A., Seth, A. K., Dolan, R. J., & Critchley, H. D. (2014). Fear from the Heart: Sensitivity to Fear Stimuli Depends on Individual Heartbeats. *Journal of Neuroscience*, 34, 6573–6582. <https://doi.org/10.1523/JNEUROSCI.3507-13.2014>
- Garfinkel, S. N., Seth, A. K., Barrett, A. B., Suzuki, K., & Critchley, H. D. (2015). Knowing your own heart: Distinguishing interoceptive accuracy from interoceptive awareness.

- Biological Psychology*. <https://doi.org/10.1016/j.biopsycho.2014.11.004>
- Gayet, S., & Stein, T. (2017). Between-Subject Variability in the Breaking Continuous Flash Suppression Paradigm: Potential Causes, Consequences, and Solutions. *Frontiers in Psychology*, 8, 437. <https://doi.org/10.3389/fpsyg.2017.00437>
- Gendron, M., & Barrett, L. F. (2009). Reconstructing the Past: A Century of Ideas About Emotion in Psychology. *Emotion Review*, 1((4)), 316–339. <https://doi.org/10.1177/1754073909338877>
- Gentsch, A., Sel, A., Marshall, A. C., & Schütz-Bosbach, S. (2019). Affective interoceptive inference: Evidence from heart-beat evoked brain potentials. *Human Brain Mapping*, 40(1), 20–33. <https://doi.org/10.1002/hbm.24352>
- Gibson, J. J. (1979). *The Ecological Approach to Visual Perception* (2015th ed.). New York: Psychology Press.
- Gottfried, J. A., O'Doherty, J., & Dolan, R. J. (2003). Encoding predictive reward value in human amygdala and orbitofrontal cortex. *Science*, 301(5636), 1104–1107. <https://doi.org/10.1126/science.1087919>
- Gray, M. A., Harrison, N. A., Wiens, S., & Critchley, H. D. (2007). Modulation of Emotional Appraisal by False Physiological Feedback during fMRI. *PLoS ONE*, 2(6). <https://doi.org/10.1371/journal.pone.0000546>
- Groen, Y., Wijers, A. A., Mulder, L. J. M., Minderaa, R. B., & Althaus, M. (2007). Physiological correlates of learning by performance feedback in children: a study of EEG event-related potentials and evoked heart rate. *Biological Psychology*, 76(3), 174–187. <https://doi.org/10.1016/J.BIOPSYCHO.2007.07.006>
- Gu, X., & Fitzgerald, T. H. B. (2014). Interoceptive inference: homeostasis and decision-making. *Trends in Cognitive Sciences*, 18(6), 269–270. <https://doi.org/10.1016/j.tics.2014.02.001>
- Gu, X., Hof, P. R., Friston, K. J., & Fan, J. (2013). Anterior insular cortex and emotional awareness. *Journal of Comparative Neurology*, 521(15), 3371–3388. <https://doi.org/10.1002/cne.23368>
- Hajcak, G., McDonald, N., & Simons, R. F. (2003). To err is autonomic: Error-related brain potentials, ANS activity, and post-error compensatory behavior. *Psychophysiology*, 40, 895–903. <https://doi.org/10.1111/1469-8986.00107>
- Hajcak, G., McDonald, N., & Simons, R. F. (2003). To err is autonomic: Error-related brain potentials, ANS activity, and post-error compensatory behavior. *Psychophysiology*, 40(6), 895–903. <https://doi.org/10.1111/1469-8986.00107>
- Hajcak, G., McDonald, N., & Simons, R. F. (2004). Error-related psychophysiology and negative affect. *Brain and Cognition*, 56(2), 189–197.

- <https://doi.org/10.1016/J.BANDC.2003.11.001>
- Hassanpour, M. S., Simmons, W. K., Feinstein, J. S., Luo, Q., Lapidus, R. C., Bodurka, J., ... Khalsa, S. S. (2018). The insular cortex dynamically maps changes in cardiorespiratory interoception. *Neuropsychopharmacology*, 43(2), 426–434.
<https://doi.org/10.1038/npp.2017.154>
- Hauser, T. U., Iannaccone, R., Stämpfli, P., Drechsler, R., Brandeis, D., Walitza, S., & Brem, S. (2014). The feedback-related negativity (FRN) revisited: New insights into the localization, meaning and network organization. *NeuroImage*, 84, 159–168.
<https://doi.org/10.1016/J.NEUROIMAGE.2013.08.028>
- Helmholtz, H. (1867). *Concerning the perceptions in general*. New York, NY: Dover. Retrieved from <https://psycnet.apa.org/buy/2006-10213-027>
- Hester, R., Foxe, J., Molholm, S., Shpaner, M., & Garavan, H. (2005). Neural correlates of error detection with and without awareness. *NeuroImage*, 27, 602–608.
<https://doi.org/10.1016/j.neuroimage.2005.04.035>
- Hohwy, J. (2012). Attention and Conscious Perception in the Hypothesis Testing Brain. *Frontiers in Psychology*, 3(APR), 96. <https://doi.org/10.3389/fpsyg.2012.00096>
- Hohwy, J. (2013). *The Predictive Mind*. Oxford, UK: Oxford University Press.
- Hohwy, J. (2020). New directions in predictive processing. *Mind & Language*, 35, 209–233.
<https://doi.org/10.1111/mila.12281>
- Holroyd, C. B., & Coles, M. G. H. (2002). The neural basis of human error processing: Reinforcement learning, dopamine, and the error-related negativity. *Psychological Review*, 109(4), 679–709. <https://doi.org/10.1037/0033-295X.109.4.679>
- Hommel, B., Muesseler, J., Aschersleben, G., & Prinz, W. (2001). The Theory of Event Coding (TEC): A framework for perception and action planning. *Behavioral and Brain Sciences*, 24, 849–937. Retrieved from www.mpipf-muenchen.mpg.de/~prinz
- Hull, C. L. (1943). *Principles of behavior*. New York: Appleton-Century-Crofts. Retrieved from <https://pdfs.semanticscholar.org/37f2/16ef09af4b813bb6041e0e9ff77c12c0b32a.pdf>
- Hull, C. L. (1952). *A behavior system; an introduction to behavior theory concerning the individual organism*. Yale University Press. Retrieved from <https://psycnet.apa.org/record/1953-03120-000>
- Hulme, O. J., Morville, T., & Gutkin, B. (2019). Neurocomputational theories of homeostatic control. *Physics of Life Reviews*, 31, 214–232. <https://doi.org/10.1016/j.plrev.2019.07.005>
- James, W. (1884). What is an emotion? *Mind*, os-IX, 188–205. Retrieved from <https://doi.org/10.1093/mind/os-IX.34.188>
- James, W. (1891). *The principles of psychology*. London: Macmillan. Retrieved from

https://openlibrary.org/books/OL14035625M/The_principles_of_psychology

James, W. (2001). *Psychology: The Briefer Course*. Courier Dover Publications (Original work published 1892).

Jiang, Y., Costello, P., & He, S. (2007). Processing of Invisible Stimuli Advantage of Upright Faces and Recognizable Words in Overcoming Interocular Suppression. *Psychological Science*, 18, 349–355.

Kanai, R., Komura, Y., Shipp, S., Friston, K., Komura, Y., Shipp, S., & Friston, K. (2015). Cerebral hierarchies : predictive processing , precision and the pulvinar. *Philosophical Transactions of the Royal Society of London. Series B, Biological Sciences*, 370.

<https://doi.org/10.1098/rstb.2014.0169>

Kandasamy, N., Garfinkel, S. N., Page, L., Hardy, B., Critchley, H. D., Gurnell, M., & Coates, J. M. (2016). Interoceptive Ability Predicts Survival on a London Trading Floor. *Scientific Reports*, 6(1), 1–7. <https://doi.org/10.1038/srep32986>

Kastner, L., Kube, J., Villringer, A., & Neumann, J. (2017). Cardiac Concomitants of Feedback and Prediction Error Processing in Reinforcement Learning. *Frontiers in Neuroscience / Www.Frontiersin.Org*, 11, 598. <https://doi.org/10.3389/fnins.2017.00598>

Katkin, E. S., Blascovich, J., & Goldband, S. (1981). Empirical assessment of visceral self-perception: Individual and sex differences in the acquisition of heartbeat discrimination. *Journal of Personality and Social Psychology*, 40(6), 1095–1101. <https://doi.org/10.1037/0022-3514.40.6.1095>

Keramati, M., & Gutkin, B. (2011). A Reinforcement Learning Theory for Homeostatic Regulation.

Keramati, M., & Gutkin, B. (2014). Homeostatic reinforcement learning for integrating reward collection and physiological stability. *ELife*, 3(e04811), 1–26. <https://doi.org/10.7554/eLife.04811>

Kern, M., Aertsen, A., Schulze-Bonhage, A., & Ball, T. (2013). Heart cycle-related effects on event-related potentials, spectral power changes, and connectivity patterns in the human ECoG. *NeuroImage*, 81, 178–190. <https://doi.org/10.1016/j.neuroimage.2013.05.042>

Khalsa, S. S., Adolphs, R., Cameron, O. G., Critchley, H. D., Davenport, P. W., Feinstein, J. S., ... Paulus, M. P. (2017). Interoception and Mental Health: A Roadmap. *Biological Psychiatry: Cognitive Neuroscience and Neuroimaging*. <https://doi.org/10.1016/j.bpsc.2017.12.004>

Kim, H., Shimojo, S., & O'Doherty, J. P. (2011). Overlapping responses for the expectation of juice and money rewards in human ventromedial prefrontal cortex. *Cerebral Cortex*. <https://doi.org/10.1093/cercor/bhq145>

King, A. B., Menon, R. S., Hachinski, V., & Cechetto, D. F. (1999). Human forebrain activation by

- visceral stimuli. *Journal of Comparative Neurology*, 413(4), 572–582.
[https://doi.org/10.1002/\(SICI\)1096-9861\(19991101\)413:4<572::AID-CNE6>3.0.CO;2-S](https://doi.org/10.1002/(SICI)1096-9861(19991101)413:4<572::AID-CNE6>3.0.CO;2-S)
- Kleckner, I. R., Zhang, J., Touroutoglou, A., Chanes, L., Xia, C., Simmons, W. K., ... Feldman Barrett, L. (2017). Evidence for a large-scale brain system supporting allostasis and interoception in humans. *Nature Human Behaviour*, 1(0069), 1–14.
<https://doi.org/10.1038/s41562-017-0069>
- Knight, D. C., Nguyen, H. T., & Bandettini, P. A. (2003). Expression of conditional fear with and without awareness. *Proceedings of the National Academy of Sciences*, 100, 15270–15283.
- Korn, C. W., & Bach, D. R. (2015). Maintaining Homeostasis by Decision-Making. *PLOS Computational Biology*, 11(5), e1004301. <https://doi.org/10.1371/journal.pcbi.1004301>
- Kouider, S., & Dehaene, S. (2007). Levels of processing during non-conscious perception: a critical review of visual masking. *Philosophical Transactions of the Royal Society of London. Series B, Biological Sciences*, 362(1481), 857–875.
<https://doi.org/10.1098/rstb.2007.2093>
- Kouider, S., & Dupoux, E. (2001). A functional disconnection between spoken and visual word recognition: evidence from unconscious priming. *Cognition*, 82(1), B35–B49.
[https://doi.org/10.1016/S0010-0277\(01\)00152-4](https://doi.org/10.1016/S0010-0277(01)00152-4)
- Kunde, W., Reuss, H., & Kiesel, A. (2012). Consciousness and cognitive control. *Advances in Cognitive Psychology*, 8(1), 9–18. <https://doi.org/10.5709/acp-0097-x>
- Kunzendorf, S., Klotzsche, F., Akbal, M., Villringer, A., Ohl, S., & Chael Gaebler, M.-3. (2019). Active vision varies across the cardiac cycle 1 2. *Psychophysiology*, 56(5), e13322.
<https://doi.org/10.1101/283838>
- Laborde, S., Mosley, E., & Thayer, J. F. (2017, February 20). Heart rate variability and cardiac vagal tone in psychophysiological research - Recommendations for experiment planning, data analysis, and data reporting. *Frontiers in Psychology*. Frontiers Research Foundation.
<https://doi.org/10.3389/fpsyg.2017.00213>
- Lamme, V. A. F. (2006). Towards a true neural stance on consciousness. *Trends in Cognitive Sciences*, 10(11), 494–501. <https://doi.org/10.1016/j.tics.2006.09.001>
- Land, M. F. (2012). Minireview The Operation of the Visual System in Relation to Action. *Current Biology*, 22, 811–817. <https://doi.org/10.1016/j.cub.2012.06.049>
- Lange, C. G. (1885). The mechanism of emotions. In B. Rand (Ed.), *The classical psychologists* (pp. 672–684). Copenhagen (Original work published 1885, Om Sindsbevaegelser et Psyko-Fysiologisk Studie).
- Lau, H. C., & Passingham, R. E. (2006). Relative blindsight in normal observers and the neural correlate of visual consciousness. *Proceedings of the National Academy of Sciences of the*

- United States of America*, 103(49), 18763–18768.
<https://doi.org/10.1073/pnas.0607716103>
- Lauwereyns, J. (2012). *Brain and the Gaze: On the Active Boundaries of Vision* - Jan Lauwereyns - *Google Books*. Cambridge, MA: MIT Press. Retrieved from
https://books.google.co.uk/books/about/Brain_and_the_Gaze.html?id=NrirB_KZ4uoC&printsec=frontcover&source=kp_read_button&redir_esc=y#v=onepage&q&f=false
- Lenggenhager, B., Azevedo, R. T., Mancini, A., & Aglioti, S. M. (2013). Listening to your heart and feeling yourself: effects of exposure to interoceptive signals during the ultimatum game. *Experimental Brain Research Brain Res*, 230, 233–241.
<https://doi.org/10.1007/s00221-013-3647-5>
- Lin, Z., & He, S. (2009). Seeing the invisible: The scope and limits of unconscious processing in binocular rivalry. *Progress in Neurobiology*, 87(4), 195–211.
<https://doi.org/10.1016/J.PNEUROBIO.2008.09.002>
- Lindley, D. V. (1957). A Statistical Paradox. *Biometrika*, 44(1/2), 187.
<https://doi.org/10.2307/2333251>
- Lo, S., & Andrews, S. (2015). To transform or not to transform: using generalized linear mixed models to analyse reaction time data. *Frontiers in Psychology*, 6, 1171.
<https://doi.org/10.3389/fpsyg.2015.01171>
- Lovero, K. L., Simmons, A. N., Aron, J. L., & Paulus, M. P. (2009). Anterior Insular Cortex Anticipates Impending Stimulus Significance. *NeuroImage*, 45(3), 976–983.
- Lowe, R., Morse, A., & Ziemke, T. (2017). Predictive Regulation: Allostasis, Behavioural Flexibility and Fear Learning. In J. Vallverdú, M. Mazzara, M. Talanov, S. Distefano, & R. Lowe (Eds.), *Advanced research on biologically inspired cognitive architectures*. Hershey, PA: IGI Global.
- Łukowska, M., Sznajder, M., & Wierzchoń, M. (2018). Error-related cardiac response as information for visibility judgements. *Scientific Reports*, 8(1).
<https://doi.org/10.1038/s41598-018-19144-0>
- Maia, T. V., & McClelland, J. L. (2004). A reexamination of the evidence for the somatic marker hypothesis: What participants really know in the Iowa gambling task. *Proceedings of the National Academy of Sciences*, 101, 16075–16080. Retrieved from
<http://www.pnas.org/content/101/45/16075.full.pdf>
- Maia, T. V., & McClelland, J. L. (2005). The somatic marker hypothesis: still many questions but no answers. *Trends in Cognitive Sciences*, 9, 162–164.
<https://doi.org/10.1016/j.tics.2005.02.002>
- Martin-Soelch, C., Linthicum, J., & Ernst, M. (2007). Appetitive conditioning: neural bases and

- implications for psychopathology. *Neuroscience and Biobehavioral Reviews*, 31, 426–440.
<https://doi.org/10.1016/j.neubiorev.2006.11.002>
- Maruya, K., Yang, E., & Blake, R. (2007). Voluntary Action Influences Visual Competition. *Psychological Science*, 18, 1090–1098. Retrieved from
<https://pdfs.semanticscholar.org/7861/100714c1b9644e69e186d68a0265d1cb7a13.pdf>
- Mastropasqua, T., & Turatto, M. (2015). Attention is necessary for subliminal instrumental conditioning. *Scientific Reports*, 5(1290), 1–7. <https://doi.org/10.1038/srep12920>
- Maturana, H. R., & Varela, F. J. (1980). *Autopoiesis and Cognition: The Realization of the Living*. Dordrecht: Holland: D. Reidel Publishing Company.
- Matuschek, H., Kliegl, R., Vasishth, S., Baayen, H., & Bates, D. (2017). Balancing Type I error and power in linear mixed models. *Journal of Memory and Language*, 94, 305–315.
<https://doi.org/10.1016/J.JML.2017.01.001>
- McEwen, B. S. (1998). Stress, Adaptation, and Disease: Allostasis and Allostatic Load. *Annals of the New York Academy of Sciences*, 840(1), 33–44. <https://doi.org/10.1111/j.1749-6632.1998.tb09546.x>
- Meijs, E. L., Klaassen, F. H., Bokeria, L., Van Gaal, S., & de Lange, F. P. (2018). Cue predictability does not modulate bottom-up attentional capture. *Royal Society Open Science*, 5, 18054. <https://doi.org/10.1098/rsos.180524>
- Meijs, E. L., Slagter, H. A., De Lange, F. P., & Simon Van Gaal, X. (2018). Behavioral/Cognitive Dynamic Interactions between Top-Down Expectations and Conscious Awareness. <https://doi.org/10.1523/JNEUROSCI.1952-17.2017>
- Melloni, L., Molina, C., Pena, M., Torres, D., Singer, W., & Rodriguez, E. (2007). Synchronization of Neural Activity across Cortical Areas Correlates with Conscious Perception. *Journal of Neuroscience*, 27(11), 2858–2865. <https://doi.org/10.1523/JNEUROSCI.4623-06.2007>
- Melloni, L., Schwiedrzik, C. M., Muller, N., Rodriguez, E., & Singer, W. (2011). Expectations Change the Signatures and Timing of Electrophysiological Correlates of Perceptual Awareness. *Journal of Neuroscience*, 31, 1386–1396.
<https://doi.org/10.1523/JNEUROSCI.4570-10.2011>
- Merikle, P. M., & Daneman, M. (2000). Conscious vs unconscious perception. In M. S. Gazzaniga (Ed.), *The new cognitive neurosciences*. Cambridge, MA: MIT Press.
- Merker, B. (2013). The efference cascade, consciousness, and its self: naturalizing the first person pivot of action control. <https://doi.org/10.3389/fpsyg.2013.00501>
- Mertens, G., & Engelhard, I. M. (2020). A systematic review and meta-analysis of the evidence for unaware fear conditioning. *Neuroscience & Biobehavioral Reviews*, 108, 254–268.
<https://doi.org/10.1016/j.neubiorev.2019.11.012>

- Mies, G. W., Van der Veen, F. M., Tulen, J. H. M., Hengeveld, M. W., & Van der Molen, M. W. (2011). Cardiac and Electrophysiological Responses to Valid and Invalid Feedback in a Time-Estimation Task. *Journal of Psychophysiology*, 25(3), 131–142.
<https://doi.org/10.1027/0269-8803/a000049>
- Mitsumatsu, H. (2009). Voluntary action affects perception of bistable motion display. *Perception*, 38, 1522–1535. <https://doi.org/10.1068/p6298>
- Montoya, P., Schandry, R., & Müller, A. (1993). *Heartbeat evoked potentials (HEP) : topography and influence of cardiac awareness and focus of attention. Electroencephalography and Clinical Neurophysiology* (Vol. 88). Retrieved from https://ac.els-cdn.com/0168559793900016/1-s2.0-0168559793900016-main.pdf?_tid=5c6b7c97-74c5-4a9c-b13f-15c179e16ab9&acdnat=1548764889_39f72b91702c368ece40ba4a111ab888
- Moors, A. (2009). Theories of emotion causation: A review. *Cognition & Emotion*, 23(4), 625–662. <https://doi.org/10.1080/02699930802645739>
- Morris, J. S., Ohman, A., & Dolan, R. J. (1998). Conscious and unconscious emotional learning in the human amygdala. *Nature*, 393, 467–470. Retrieved from <https://www.nature.com/nature/journal/v393/n6684/pdf/393467a0.pdf>
- Motyka, P., Grund, M., Forschack, N., Al, E., Villringer, A., & Gaebler, M. (2019). Interactions between cardiac activity and conscious somatosensory perception. *Psychophysiology*, e13424. <https://doi.org/10.1111/psyp.13424>
- Mudrik, L., Faivre, N., & Koch, C. (2014). Information integration without awareness. *Trends in Cognitive Sciences*, 18, 488–496. <https://doi.org/10.1016/j.tics.2014.04.009>
- Müller, S. V., Möller, J., Rodriguez-Fornells, A., & Münte, T. F. (2005). Brain potentials related to self-generated and external information used for performance monitoring. *Clinical Neurophysiology*, 116(1), 63–74. <https://doi.org/10.1016/j.clinph.2004.07.009>
- Newell, B. R., & Shanks, D. R. (2013). Unconscious Influences on Decision Making : A Critical review. *Behavioural and Brain Sciences*, 37(1), 1–19. <https://doi.org/10.1111/j.1365-2141.1970.tb01443.x>
- Nieuwenhuis, S., Ridderinkhof, K. R., Blom, J., Band, G. P. H., & Kok, A. (2001). Error-related brain potentials are differentially related to awareness of response errors: Evidence from an antisaccade task. *Psychophysiology*, 38, 752–760. Retrieved from http://sandernieuwenhuis.nl/pdfs/nieuwenhuis_et_al_2001.pdf
- O'Doherty, J. (2004). Reward representations and reward-related learning in the human brain: Insights from neuroimaging. *Current Opinion in Neurobiology*, 14, 769–776.
<https://doi.org/10.1016/j.conb.2004.10.016>
- O'Doherty, J., Deichmann, R., Critchley, H. D., & Dolan, R. J. (2002). Neural responses during

- anticipation of a primary taste reward. *Neuron*, 33, 815–826.
- O'Doherty, J., Kringelbach, M. L., Rolls, E. T., Hornak, J., & Andrews, C. (2001). Abstract reward and punishment representations in the human orbitofrontal cortex. *Nature Neuroscience*, 4(1), 95–102. <https://doi.org/10.1038/82959>
- O'Regan J.K, & Noë, A. (2001). A sensorimotor account of vision and visual consci.pdf. *Behavioral and Brain Sciences*, 24, 939–1031.
- Olsson, A., & Phelps, E. A. (2004). Learned Fear of “Unseen” Faces After Pavlovian, Observational, and Instructed Fear. *Psychological Science*, 15, 822–828. Retrieved from <http://journals.sagepub.com/doi/pdf/10.1111/j.0956-7976.2004.00762.x>
- Oppenheimer, S., & Cechetto, D. (2016). The insular cortex and the regulation of cardiac function. *Comprehensive Physiology*, 6(2), 1081–1133. <https://doi.org/10.1002/cphy.c140076>
- Overbeek, T. J. M., Nieuwenhuis, S., & Ridderinkhof, K. R. (2005). Dissociable Components of Error Processing. *Journal of Psychophysiology*, 19(4), 319–329. <https://doi.org/10.1027/0269-8803.19.4.319>
- Overgaard, Morten. (2015). *Behavioral methods in consciousness research*.
- Owens, A. P., Allen, M., Ondobaka, S., & Friston, K. J. (2018). Interoceptive inference: From computational neuroscience to clinic. *Neuroscience & Biobehavioral Reviews*, 90, 174–183. <https://doi.org/10.1016/J.NEUBIOREV.2018.04.017>
- Owens, A. P., Friston, K. J., Low, D. A., Mathias, C. J., Critchley, H. D., & Owens, A. (2018). Investigating the relationship between cardiac interoception and autonomic cardiac control using a predictive coding framework. *Autonomic Neuroscience: Basic and Clinical*, 210, 65–71. Retrieved from [http://sro.sussex.ac.uk/73135/1/Investigating the relationship between cardiac interoception and autonomic cardiac control using a predictive coding framework.pdf](http://sro.sussex.ac.uk/73135/1/Investigating%20the%20relationship%20between%20cardiac%20interoception%20and%20autonomic%20cardiac%20control%20using%20a%20predictive%20coding%20framework.pdf)
- Oya, H., Adolphs, R., Kawasaki, H., Bechara, A., Damasio, A., Iii, M. A. H., & Raichle, M. E. (2005). *Electrophysiological correlates of reward prediction error recorded in the human prefrontal cortex*. Retrieved from www.pnas.org/cgi/doi/10.1073/pnas.0500899102
- Paciorek, A., & Skora, L. (2020). Vagus Nerve Stimulation as a Gateway to Interoception. *Frontiers in Psychology*, 11, 1659. <https://doi.org/10.3389/FPSYG.2020.01659>
- Panichello, M. F., Cheung, O. S., Bar, M., Muckli, L., Seth, A. K., & Gordon, R. D. (2013). Predictive feedback and conscious visual experience. <https://doi.org/10.3389/fpsyg.2012.00620>
- Parkinson, J. A., Cardinal, R. N., & Everitt, B. J. (2000). Limbic cortical-ventral striatal systems underlying appetitive conditioning. *Progress in Brain Research*, 126, 263–285.

[https://doi.org/10.1016/S0079-6123\(00\)26019-6](https://doi.org/10.1016/S0079-6123(00)26019-6)

- Pavlov, I. (1928). *Conditioned reflexes: An investigation of the physiological activity of the cerebral cortex*. Oxford: Oxford University Press: Humphrey Milford.
- Pavlov, I. P. (1927). Conditioned reflexes. An investigation of the physiological activity of the cerebral cortex. Translated and edited by E. V. ANREP. *Conditioned Reflexes. An Investigation of the Physiological Activity of the Cerebral Cortex. Translated and Edited by E. V. ANREP*, xv+427p. 18 fig.--xv+427p. Retrieved from <http://s-f-walker.org.uk/pubsebooks/pdfs/Conditioned-Reflexes-Pavlov.pdf>
- Peirce, C. S., & Jastrow, J. (1886). On small differences of sensation. *Memoirs of the National Academy of Sciences*, 3, 73–83.
- Persuh, M. (2017). Measuring Perceptual Consciousness. *Frontiers in Psychology*, 8, 2320. <https://doi.org/10.3389/fpsyg.2017.02320>
- Pessiglione, M., Schmidt, L., Draganski, B., Kalisch, R., Lau, H., Dolan, R. J., & Frith, C. D. (2007). How the Brain Translates Money into Force: A Neuroimaging Study of Subliminal Motivation. *Science*. <https://doi.org/10.1126/science.1140459>
- Pessiglione, Mathias, Petrovic, P., Daunizeau, J., Palminteri, S., Dolan, R. J., & Frith, C. D. (2008). Subliminal Instrumental Conditioning Demonstrated in the Human Brain. *Neuron*, 59, 561–567. <https://doi.org/10.1016/j.neuron.2008.07.005>
- Petzschner, F. (2017). Interoceptive Predictive Coding of Heartbeats. Beijing, China: Association for the Scientific Study of Consciousness (ASSC).
- Petzschner, F., Weber, L. A., Wellstein, K. V., Paolini, G., Do, C. T., & Stephan, K. E. (2019). Focus of attention modulates the heartbeat evoked potential. *NeuroImage*, 186, 595–606. <https://doi.org/10.1016/J.NEUROIMAGE.2018.11.037>
- Pezzulo, G. (2014). Why do you fear the bogeyman? An embodied predictive coding model of perceptual inference. *Cognitive, Affective, and Behavioral Neuroscience*, 14, 902–911. <https://doi.org/10.3758/s13415-013-0227-x>
- Pezzulo, G., Barca, L., & Friston, K. (2015). Active inference and cognitive-emotional interactions in the brain. *Behavioral and Brain Sciences*, 38.
- Pezzulo, Giovanni, Rigoli, F., & Friston, K. (2015). Active Inference, homeostatic regulation and adaptive behavioural control. *Progress in Neurobiology*, 134, 17–35. <https://doi.org/10.1016/j.pneurobio.2015.09.001>
- Pfeiffer, C., & De Lucia, M. (2017). Cardio-audio synchronization drives neural surprise response. *Scientific Reports*, 7(1). <https://doi.org/10.1038/s41598-017-13861-8>
- Pinto, Y., van Gaal, S., de Lange, F. P., Lamme, V. A. F., & Seth, A. K. (2015). Expectations accelerate entry of visual stimuli into awareness. *Journal of Vision*, 15, 13.

<https://doi.org/10.1167/15.8.13>

- Ploghaus, A., Tracey, I., Gati, J. S., Clare, S., Menon, R. S., Matthews, P. M., & Rawlins, J. N. P. (1999). Dissociating pain from its anticipation in the human brain. *Science*, 284(5422), 1979–1981. <https://doi.org/10.1126/science.284.5422.1979>
- Pollatos, O., Herbert, B. M., Matthias, E., & Schandry, R. (2007). Heart rate response after emotional picture presentation is modulated by interoceptive awareness. *International Journal of Psychophysiology*, 63, 117–124. <https://doi.org/10.1016/j.ijpsycho.2006.09.003>
- Poppa, T., & Bechara, A. (2018). The somatic marker hypothesis: revisiting the role of the “body-loop” in decision-making. *Current Opinion in Behavioral Sciences*, 19, 61–66. <https://doi.org/10.1016/j.cobeha.2017.10.007>
- Power, M. L. (2004). Commentary. Viability as Opposed to Stability: An Evolutionary Perspective on Physiological Regulation. In J. Schulkin (Ed.), *Allostasis, Homeostasis, and the Costs of Physiological Adaptation* (pp. 343–364). Cambridge, UK: Cambridge University Press. <https://doi.org/10.1017/CBO9781316257081.012>
- Pramme, L., Larra, M. F., Schachinger, H., & Frings, C. (2016). Cardiac cycle time effects on selection efficiency in vision. *Psychophysiology*, 53, 1702–1711. <https://doi.org/10.1111/psyp.12728>
- Press, C., Kok, P., & Yon, D. (2020). The Perceptual Prediction Paradox. *Trends in Cognitive Sciences*, 24(1), 13–24. <https://doi.org/10.1016/j.tics.2019.11.003>
- Prinz, W. (2003). Experimental Approaches to Action. In J. Roessler & N. Eilan (Eds.), *Agency and Self-Awareness: Issues in Philosophy and Psychology* (pp. 165–187). New York: Clarendon Press.
- Quadt, L., Critchley, H. D., & Garfinkel, S. N. (2018a). Interoception and emotion: Shared mechanisms and clinical implications. In M. Tsakiris & H. De Preester (Eds.), *The Interoceptive Mind: from Homeostasis to Awareness* (p. 123). Oxford, UK: Oxford University Press.
- Quadt, L., Critchley, H. D., & Garfinkel, S. N. (2018b). The neurobiology of interoception in health and disease. *Ann. N.Y. Acad. Sci.*, 1428, 112–128. <https://doi.org/10.1111/nyas.13915>
- Quintana, D. S., Heathers, J. A. J., Kemp, A., Tarvainen, M., & Billman, G. E. (2014). Considerations in the assessment of heart rate variability in biobehavioral research. <https://doi.org/10.3389/fpsyg.2014.00805>
- Rae, C. L., Botan, V. E., Gould van Praag, C. D., Herman, A. M., Nyssönen, J. A. K., Watson, D. R., ... Critchley, H. D. (2018). Response inhibition on the stop signal task improves during

- cardiac contraction. *Scientific Reports*, 8(1), 9136. <https://doi.org/10.1038/s41598-018-27513-y>
- Raio, C. M., Carmel, D., Carrasco, M., & Phelps, E. A. (2012). Nonconscious fear is quickly acquired but swiftly forgotten. *Current Biology*, 22(12), R477–R479. <https://doi.org/10.1016/j.cub.2012.04.023>
- Ramsay, D. S., & Woods, S. C. (2014). Clarifying the Roles of Homeostasis and Allostasis in Physiological Regulation. *Psychological Review*, 121(2), 225–247. <https://doi.org/10.1037/a0035942>
- Rangel, A., Camerer, C., & Montague, P. R. (2008). A framework for studying the neurobiology of value-based decision making. *Nature Reviews Neuroscience*, 9, 545–556. <https://doi.org/10.1038/nrn2357>
- Rao, R. P. N., & Ballard, D. H. (1999). Predictive coding in the visual cortex: a functional interpretation of some extra-classical receptive-field effects. *Nature Neuroscience*, 2, 79–87. Retrieved from <http://neurosci.nature.com>
- Reber, T. P., Luechinger, R., Boesiger, P., & Henke, K. (2012). Unconscious relational inference recruits the hippocampus. *Journal of Neuroscience*, 32(18), 6138–6148. <https://doi.org/10.1523/JNEUROSCI.5639-11.2012>
- Reber, T. P., Samimizad, B., & Mormann, F. (2018). Cue discriminability predicts instrumental conditioning. *Consciousness and Cognition*, 61, 49–60. <https://doi.org/10.1016/J.CONCOG.2018.03.013>
- Rebollo, I., Devauchelle, A. D., Béranger, B., & Tallon-Baudry, C. (2018). Stomach-brain synchrony reveals a novel, delayed-connectivity resting-state network in humans. *ELife*, 7, 1–25. <https://doi.org/10.7554/eLife.33321>
- Rebuschat, P. (2013). Measuring Implicit and Explicit Knowledge in Second Language Research. *Language Learning*, 63(3), 595–626. <https://doi.org/10.1111/lang.12010>
- Reimann, M., & Bechara, A. (2010). The somatic marker framework as a neurological theory of decision-making: Review, conceptual comparisons, and future neuroeconomics research. *Journal of Economic Psychology*, 31(5), 767–776. <https://doi.org/10.1016/j.joep.2010.03.002>
- Rescorla, R. A. (1988). Pavlovian conditioning. It's not what you think it is. *American Psychologist*, 43(3), 151–160. <https://doi.org/10.1016/j.neuroimage.2007.11.042>
- Ridderinkhof, K. R., Nieuwenhuis, S., & Bashore, T. R. (2003). Errors are foreshadowed in brain potentials associated with action monitoring in cingulate cortex in humans. *Neuroscience Letters*, 348(1), 1–4. [https://doi.org/10.1016/S0304-3940\(03\)00566-4](https://doi.org/10.1016/S0304-3940(03)00566-4)
- Ring, C., & Brener, J. (2018). Heartbeat counting is unrelated to heartbeat detection: A

- comparison of methods to quantify interoception. *Psychophysiology*, 55(9), 1–10.
<https://doi.org/10.1111/psyp.13084>
- Ring, C., Brener, J., Knapp, K., & Mailloux, J. (2015). Effects of heartbeat feedback on beliefs about heart rate and heartbeat counting: A cautionary tale about interoceptive awareness. *Biological Psychology*, 104, 193–198.
<https://doi.org/10.1016/j.biopsycho.2014.12.010>
- Rolls, E. T. (2000). The orbitofrontal cortex and reward. *Cerebral Cortex*, 10, 284–294.
- Rosen, J. B., & Schulkin, J. (2004). Adaptive Fear, Allostasis, and the Pathology of Anxiety and Depression. In J. Schulkin (Ed.), *Allostasis, Homeostasis, and the Costs of Physiological Adaptation* (pp. 164–227). Cambridge, UK: Cambridge University Press.
<https://doi.org/10.1017/CBO9781316257081.009>
- Rosenthal, C. R., Andrews, S. K., Antoniadis, C. A., Kennard, C., Soto Correspondence, D., & Soto, D. (2016). Learning and Recognition of a Non-conscious Sequence of Events in Human Primary Visual Cortex. *Current Biology*, 26, 834–841.
<https://doi.org/10.1016/j.cub.2016.01.040>
- Rosenthal, C. R., Kennard, C., & Soto, D. (2010). Visuospatial Sequence Learning without Seeing. *PLoS ONE*, 5(7), e11906. <https://doi.org/10.1371/journal.pone.0011906>
- Rothkirch, M., & Hesselmann, G. (2017). What we talk about when we talk about unconscious processing - A plea for best practices. *Frontiers in Psychology*, 8(835), 1–6.
<https://doi.org/10.3389/fpsyg.2017.00835>
- Rushworth, M. F. S., Kolling, N., Sallet, J., & Mars, R. B. (2012). Valuation and decision-making in frontal cortex: One or many serial or parallel systems? *Current Opinion in Neurobiology*. <https://doi.org/10.1016/j.conb.2012.04.011>
- Rushworth, M. F. S., Noonan, M. A. P., Boorman, E. D., Walton, M. E., & Behrens, T. E. (2011). Frontal Cortex and Reward-Guided Learning and Decision-Making. *Neuron*.
<https://doi.org/10.1016/j.neuron.2011.05.014>
- Russell, J. A. (2003). Core Affect and the Psychological Construction of Emotion. *Psychological Review*, 110(1), 145–172. <https://doi.org/10.1037/0033-295X.110.1.145>
- Salomon, R., Lim, M., Herbelin, B., Hesselmann, G., & Blanke, O. (2013). Posing for awareness: Proprioception modulates access to visual consciousness in a continuous flash suppression task. *Journal of Vision*, 13(7), 2–2. <https://doi.org/10.1167/13.7.2>
- Salomon, Roy, Ronchi, R., Dönn, J., Bello-Ruiz, J., Herbelin, B., Martet, R., ... Blanke, O. (2016). The Insula Mediates Access to Awareness of Visual Stimuli Presented Synchronously to the Heartbeat. *Journal of Neuroscience*, 36, 5115–5127.
<https://doi.org/10.1523/JNEUROSCI.4262-15.2016>

- Sand, A., & Nilsson, M. E. (2016). Subliminal or not? Comparing null-hypothesis and Bayesian methods for testing subliminal priming. *Consciousness and Cognition*, 44, 29–40.
<https://doi.org/10.1016/J.CONCOG.2016.06.012>
- Saper, C. B. (2002). The Central Autonomic Nervous System: Conscious Visceral Perception and Autonomic Pattern Generation. *Annual Review of Neuroscience*, 25(1), 433–469.
<https://doi.org/10.1146/annurev.neuro.25.032502.111311>
- Schachter, S., & Singer, J. (1962). Cognitive, social, and physiological determinants of emotional state. *Psychological Review*, 69(5), 379–399.
<https://doi.org/10.1037/h0046234>
- Schandry, R. (1981). Heart Beat Perception and Emotional Experience. *Psychophysiology*, 18(4), 483–488. <https://doi.org/10.1111/j.1469-8986.1981.tb02486.x>
- Schulkin, J. (2003a). Allostasis: A neural behavioral perspective. *Hormones and Behavior*, 43(1), 21–27. [https://doi.org/10.1016/S0018-506X\(02\)00035-1](https://doi.org/10.1016/S0018-506X(02)00035-1)
- Schulkin, J. (2003b). *Rethinking Homeostasis Allostatic Regulation in Physiology and Pathophysiology*. Cambridge, MA: MIT Press.
- Schulkin, J. (2011). Social allostasis: anticipatory regulation of the internal milieu. *Frontiers in Evolutionary Neuroscience*, 2(111), 1–15. <https://doi.org/10.3389/fnevo.2010.00111>
- Scott, R. B., Samaha, J., Chrisley, R., & Dienes, Z. (2018). Prevailing theories of consciousness are challenged by novel cross-modal associations acquired between subliminal stimuli. *Cognition*, 175, 169–185. <https://doi.org/10.1016/j.cognition.2018.02.008>
- Seitz, A. R., Kim, D., & Watanabe, T. (2009a). Rewards Evoke Learning of Unconsciously Processed Visual Stimuli in Adult Humans. *Neuron*, 61, 700–707.
<https://doi.org/10.1016/j.neuron.2009.01.016>
- Seitz, A. R., Kim, D., & Watanabe, T. (2009b). Rewards Evoke Learning of Unconsciously Processed Visual Stimuli in Adult Humans. *Neuron*, 61(5), 700–707.
<https://doi.org/10.1016/j.neuron.2009.01.016>
- Seth, A. K. (2013). Interoceptive inference, emotion, and the embodied self. *Trends in Cognitive Sciences*, 17, 565–573. <https://doi.org/10.1016/j.tics.2013.09.007>
- Seth, A. K. (2014). A predictive processing theory of sensorimotor contingencies: Explaining the puzzle of perceptual presence and its absence in synesthesia. *Cognitive Neuroscience*, 5(2), 97–118. <https://doi.org/10.1080/17588928.2013.877880>
- Seth, A. K. (2015). The Cybernetic Bayesian Brain. In T. Metzinger & J. M. Windt (Eds.), *Open MIND* (pp. 1–24). Frankfurt am Main: MIND Group.
<https://doi.org/10.15502/9783958570108>
- Seth, A. K., & Critchley, H. D. (2013). Extending predictive processing to the body: Emotion as

- interoceptive inference. *Behavioral and Brain Sciences*.
<https://doi.org/10.1017/S0140525X12002270>
- Seth, A. K., Dienes, Z., Cleeremans, A., Overgaard, M., & Pessoa, L. (2008). Measuring consciousness: relating behavioural and neurophysiological approaches. *Trends in Cognitive Sciences*, 12(8), 314–321. <https://doi.org/10.1016/j.tics.2008.04.008>
- Seth, A. K., & Friston, K. J. (2016). Active interoceptive inference and the emotional brain. *Philosophical Transactions of the Royal Society B*, 371, 20160007.
<https://doi.org/10.1098/rstb.2016.0007>
- Seth, A. K., Suzuki, K., & Critchley, H. D. (2012). An interoceptive predictive coding model of conscious presence. *Frontiers in Psychology*, 2, 395.
<https://doi.org/10.3389/fpsyg.2011.00395>
- Seth, A. K., & Tsakiris, M. (2018, November 1). Being a Beast Machine: The Somatic Basis of Selfhood. *Trends in Cognitive Sciences*. Elsevier Ltd.
<https://doi.org/10.1016/j.tics.2018.08.008>
- Seth, A. K., Verschure, P. F. M. J., Blanke, O., Butz, M. V, Ford, J. M., Frith, C. D., ... Regan, J. K. O. (2016). Action-Oriented Understanding of Consciousness and the Structure of Experience. In A. K. Engel, K. J. Friston, & D. Kragic (Eds.), *The Pragmatic Turn: Toward Action-Oriented Views in Cognitive Science* (pp. 261–281). Cambridge, MA: MIT Press.
- Shanks, D. R. (2017). Regressive research: The pitfalls of post hoc data selection in the study of unconscious mental processes. *Psychonomic Bulletin & Review*, 24(3), 752–775.
<https://doi.org/10.3758/s13423-016-1170-y>
- Shanks, D. R., & St. John, M. F. (1994). Characteristics of dissociable human learning systems. *Behavioural and Brain Sciences*, 17, 367–447. Retrieved from
https://www.cambridge.org/core/services/aop-cambridge-core/content/view/98272161869C71FE8FDF29145DF79342/S0140525X00035032a.pdf/characteristics_of_dissociable_human_learning_systems.pdf
- Sherman, M. T., Seth, A. K., Barrett, A. B., & Kanai, R. (2015). Prior expectations facilitate metacognition for perceptual decision. *Consciousness and Cognition*, 35, 53–65.
<https://doi.org/10.1016/j.concog.2015.04.015>
- Sherrington, C. S. (1948). *Integrative action of the nervous system*. Cambridge, UK: Cambridge University Press.
- Singer, T., Critchley, H. D., & Preuschoff, K. (2009). A common role of insula in feelings, empathy and uncertainty. *Trends in Cognitive Sciences*, 13(8), 334–340.
<https://doi.org/10.1016/j.tics.2009.05.001>
- Singmann, H., & Kellen, D. (2019). An introduction to mixed models for experimental

- psychology. In D. Spieler & E. Schumacher (Eds.), *New methods in neuroscience and cognitive psychology*. New York, NY: Routledge. Retrieved from http://singmann.org/download/publications/singmann_kellen-introduction-mixed-models.pdf
- Skinner, B. F. (1932). On the rate of formation of a conditioned reflex. *Journal of General Psychology*, 7(2), 274–286. <https://doi.org/10.1080/00221309.1932.9918467>
- Skora, L. I., Livermore, J. J. A., Dienes, Z., Seth, A. K., & Scott, R. B. (n.d.). Feasibility of Unconscious Instrumental Conditioning: A Registered Replication. *Cortex*. <https://doi.org/10.31234/OSF.IO/P9DGN>
- Sokol-Hessner, P., Hartley, C. A., Hamilton, J. R., & Phelps, E. A. (2015). Interoceptive ability predicts aversion to losses. *Cognition and Emotion*, 29, 695–701. <https://doi.org/10.1080/02699931.2014.925426>
- Sokolov, E. N. (1963). Higher nervous functions: The orienting reflex. *Annual Review of Physiology*, 25, 545–580. Retrieved from www.annualreviews.org
- Somsen, R. J. M., Van Der Molen, M. W., Jennings, J. R., & Van Beek, B. (2000). Wisconsin Card Sorting in adolescents: Analysis of performance, response times and heart rate. *Acta Psychologica*, 104(2), 227–257. [https://doi.org/10.1016/S0001-6918\(00\)00030-5](https://doi.org/10.1016/S0001-6918(00)00030-5)
- Staddon, J. E. R., & Cerutti, D. T. (2003). Operant Conditioning. *Annual Review of Psychology*, 54, 115–144. <https://doi.org/10.1016/B978-0-12-375000-6.00262-7>
- Stanislaw, H., & Todorov, N. (1999). Calculation of signal detection theory measures. *Behavior Research Methods, Instruments, & Computers*, 31(1), 137–149. Retrieved from <http://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.453.5684&rep=rep1&type=pdf>
- Stein, T., & Sterzer, P. (2014). Unconscious processing under interocular suppression: getting the right measure. *Frontiers in Psychology*, 5(387), 1–5. <https://doi.org/10.3389/fpsyg.2014.00387>
- Sterling, P. (2004). Principles of allostasis: Optimal design, predictive regulation, pathophysiology, and rational therapeutics. In J. Schulkin (Ed.), *Allostasis, Homeostasis, and the Costs of Physiological Adaptation* (pp. 17–64). Cambridge, UK: Cambridge University Press. <https://doi.org/10.1017/CBO9781316257081.004>
- Sterling, P. (2012). Allostasis: A model of predictive regulation. *Physiology and Behavior*, 106(1), 5–15. <https://doi.org/10.1016/j.physbeh.2011.06.004>
- Sterling, P., & Eyer, J. (1988). Allostasis: A new paradigm to explain arousal pathology. In S. Fisher & J. Reason (Eds.), *Handbook of life stress, cognition and health* (pp. 629–649). John Wiley & Sons. Retrieved from <https://psycnet.apa.org/record/1988-98352-034>

- Sterzer, P., Stein, T., Ludwig, K., Rothkirch, M., Hesselmann, G., Dubois, J., ... Yuval-Greenberg, S. (2014). Neural processing of visual information under interocular suppression: a critical review. *Frontiers in Psychology*, 5(453), 1–12. <https://doi.org/10.3389/fpsyg.2014.00453>
- Strigo, I. A., & Craig, A. D. (2016). Interoception, homeostatic emotions and sympathovagal balance. *Philosophical Transactions of the Royal Society B*, 317, 20160010. <https://doi.org/10.6084/m9>
- Summerfield, C., Trittschuh, E. H., Monti, J. M., Mesulam, M.-M., & Egner, T. (2008). Neural repetition suppression reflects fulfilled perceptual expectations. *Nature Neuroscience*, 11(9), 1004–1006. <https://doi.org/10.1038/nn.2163>
- Sutton, R. S., & Barto, A. G. (1998). *Introduction to Reinforcement Learning*. Cambridge, MA: MIT Press.
- Suzuki, A., Hirota, A., Takasawa, N., & Shigemasa, K. (2003). Application of the somatic marker hypothesis to individual differences in decision making. *Biological Psychology*, 65(1), 81–88. [https://doi.org/10.1016/S0301-0511\(03\)00093-0](https://doi.org/10.1016/S0301-0511(03)00093-0)
- Suzuki, K., Schwartzman, D. J., Augusto, R., & Seth, A. K. (2019). Sensorimotor contingency modulates breakthrough of virtual 3D objects during a breaking continuous flash suppression paradigm. *Cognition*, 187, 95–107. <https://doi.org/10.1016/j.cognition.2019.03.003>
- Suzuki, Keisuke, Garfinkel, S. N., Critchley, H. D., & Seth, A. K. (2013). Multisensory integration across exteroceptive and interoceptive domains modulates self-experience in the rubber-hand illusion. *Neuropsychologia*, 51, 2909–2917. <https://doi.org/10.1016/j.neuropsychologia.2013.08.014>
- Tallon-Baudry, C., & Park, H.-D. (2014). The neural subjective frame: from bodily signals to perceptual consciousness. *Philosophical Transactions of the Royal Society B*, 369, 20130208. <https://doi.org/10.1098/rstb.2013.0208>
- Thayer, J. F., & Lane, R. D. (2000). A model of neurovisceral integration in emotion regulation and dysregulation. *Journal of Affective Disorders*, 61(3), 201–216. [https://doi.org/10.1016/S0165-0327\(00\)00338-4](https://doi.org/10.1016/S0165-0327(00)00338-4)
- Thayer, J. F., Loerbroks, A., & Sternberg, E. M. (2011, September 30). Inflammation and cardiorespiratory control: The role of the vagus nerve. *Respiratory Physiology and Neurobiology*. <https://doi.org/10.1016/j.resp.2011.05.016>
- Thorndike, E. L. (1898). Animal intelligence: An experimental study of the associative processes in animals. *The Psychological Review: Monograph Supplements*, 2(4), i–109. <https://doi.org/10.1037/h0092987>
- Thorndike, E. L. (1911). *Animal intelligence*. New York: Macmillan.

- Timmermans, B., & Cleeremans, A. (2015). How can we measure awareness? An overview of current methods. In *Behavioral Methods in Consciousness Research* (pp. 21–46). Oxford University Press. <https://doi.org/10.1093/acprof:oso/9780199688890.003.0003>
- Tononi, G. (2004). An information integration theory of consciousness. *BMC Neuroscience*, 5(1), 42. <https://doi.org/10.1186/1471-2202-5-42>
- Tononi, G. (2012). Integrated information theory of consciousness: an updated account. *Archives Italiennes de Biologie*, 150, 293–329. Retrieved from <http://www.architalbiol.org/aib/article/viewFile/15056/23165867>
- Travers, E., Frith, C. D., & Shea, N. (2018). Learning rapidly about the relevance of visual cues requires conscious awareness. *Quarterly Journal of Experimental Psychology*, 71(8), 1698–1713. <https://doi.org/10.1080/17470218.2017.1373834>
- Tsakiris, Manos, & Critchley, H. (2016, November 19). Interoception beyond homeostasis: Affect, cognition and mental health. *Philosophical Transactions of the Royal Society B: Biological Sciences*. Royal Society of London. <https://doi.org/10.1098/rstb.2016.0002>
- Tsakiris, Manos, Tajadura-Jiménez, A., & Costantini, M. (2011). Just a heartbeat away from one's body: interoceptive sensitivity predicts malleability of body-representations. *Proceedings of the Royal Society B*, 278, 2470–2467. <https://doi.org/10.1098/rspb.2010.2547>
- Tschantz, A., Seth, A. K., & Buckley, C. L. (2019). Learning action-oriented models through active inference. *ArXiv*. <https://doi.org/10.1101/764969>
- Tschantz, A., Seth, A. K., & Buckley, C. L. (2020). Learning action-oriented models through active inference. *PLOS Computational Biology*, 16(4), E1007805. <https://doi.org/10.1371/journal.pcbi.1007805>
- Tsuchiya, N., & Koch, C. (2005). Continuous flash suppression reduces negative afterimages. *Nature Neuroscience*, 8(8), 1096–1101. <https://doi.org/10.1038/nn1500>
- Ullsperger, M., Danielmeier, C., & Jocham, G. (2014). Neurophysiology of Performance Monitoring and Adaptive Behavior. *Physiological Review*, 94, 35–79. <https://doi.org/10.1152/physrev.00041.2012.-Successful>
- Ullsperger, M., & Von Cramon, D. Y. (2004, November). Decision making, performance and outcome monitoring in frontal cortical areas. *Nature Neuroscience*. Nature Publishing Group. <https://doi.org/10.1038/nn1104-1173>
- Vadillo, M. A., Konstantinidis, E., & Shanks, D. R. (2016). Underpowered samples, false negatives, and unconscious learning. *Psychonomic Bulletin & Review*, 23(1), 87–102. <https://doi.org/10.3758/s13423-015-0892-6>
- Valera, F., Thompson, E., & Rosch, E. (1992). *The embodied mind: Cognitive science and human*

- experience*. Cambridge, MA: MIT Press.
- Van Den Bussche, E., Van Den Noortgate, W., & Reynvoet, B. (2009). Mechanisms of Masked Priming: A Meta-Analysis. *Psychological Bulletin*, 135, 452–477.
<https://doi.org/10.1037/a0015329>
- van der Veen, F. M., van der Molen, M. W., Crone, E. A., & Jennings, J. R. (2004a). Phasic heart rate responses to performance feedback in a time production task: effects of information versus valence. *Biological Psychology*, 65(2), 147–161.
<https://doi.org/10.1016/j.biopsycho.2003.07.003>
- van der Veen, F. M., van der Molen, M. W., Crone, E. A., & Jennings, J. R. (2004b). Phasic heart rate responses to performance feedback in a time production task: effects of information versus valence. *Biological Psychology*, 65(2), 147–161.
<https://doi.org/10.1016/J.BIOPSYCHO.2003.07.003>
- van Elk, M., Lenggenhager, B., Heydrich, L., & Blanke, O. (2014). Suppression of the auditory N1-component for heartbeat-related sounds reflects interoceptive predictive coding. *Biological Psychology*, 99, 172–182. <https://doi.org/10.1016/J.BIOPSYCHO.2014.03.004>
- van Gaal, S., de Lange, F. P., & Cohen, M. X. (2012, May 7). The role of consciousness in cognitive control and decision making. *Frontiers in Human Neuroscience*. Frontiers Media S. A. <https://doi.org/10.3389/fnhum.2012.00121>
- Van Gaal, S., Ridderinkhof, K. R., Fahrenfort, J. J., Scholte, H. S., & Lamme, V. A. F. (2008). Frontal cortex mediates unconsciously triggered inhibitory control. *Journal of Neuroscience*, 28(32), 8053–8062. <https://doi.org/10.1523/JNEUROSCI.1278-08.2008>
- van Gaal, S. van, Ridderinkhof, K. R., Scholte, H. S., & Lamme, V. A. F. (2010). Unconscious Activation of the Prefrontal No-Go Network. *Journal of Neuroscience*, 30(11), 4143–4150.
<https://doi.org/10.1523/JNEUROSCI.2992-09.2010>
- Villani, V., Tsakiris, M., & Azevedo, R. T. (2019). Transcutaneous vagus nerve stimulation improves interoceptive accuracy. *Neuropsychologia*, 134.
<https://doi.org/10.1016/j.neuropsychologia.2019.107201>
- Wagenmakers, E.-J., Marsman, M., Jamil, T., Ly, A., Verhagen, J., Love, J., ... Morey, R. D. (2017). Bayesian inference for psychology. Part I: Theoretical advantages and practical ramifications. In S. O. Lilienfeld & I. Waldman (Eds.), *Psychological Science Under Scrutiny: Recent Challenges and Proposed* (pp. 123–138). John Wiley & Sons.
<https://doi.org/10.3758/s13423-017-1343-3>
- Wang, P., & Nikolić, D. (2011). An LCD monitor with sufficiently precise timing for research in vision. *Frontiers in Human Neuroscience*, (AUGUST).
<https://doi.org/10.3389/fnhum.2011.00085>

- Wang, X., Wu, Q., Egan, L., Gu, X., Liu, P., Gu, H., ... Fan, J. (2019). Anterior insular cortex plays a critical role in interoceptive attention. *ELife*, 8(e42265), 1–31.
<https://doi.org/10.7554/eLife.42265.001>
- Watkins, C., & Dayan, P. (1992). Technical Note: Q-learning. *Machine Learning*, 8, 279–292.
- Werner, N. S., Jung, K., Duschek, S., & Schandry, R. (2009). Enhanced cardiac perception is associated with benefits in decision-making. *Psychophysiology*, 46(6), 1123–1129.
<https://doi.org/10.1111/j.1469-8986.2009.00855.x>
- Werner, N. S., Schweitzer, N., Meindl, T., Duschek, S., Kambeitz, J., & Schandry, R. (2013). Interoceptive awareness moderates neural activity during decision-making. *Biological Psychology*, 94(3), 498–506. <https://doi.org/10.1016/j.biopsycho.2013.09.002>
- Wessel, J. R., Danielmeier, C., & Ullsperger, M. (2011). Error awareness revisited: Accumulation of multimodal evidence from central and autonomic nervous systems. *Journal of Cognitive Neuroscience*, 23(10), 3021–3036. <https://doi.org/10.1162/jocn.2011.21635>
- Whitehead, W. E., Drescher, V. M., Heiman, P., & Blackwell, B. (1977). *Relation of Heart Rate Control to Heartbeat Perception 1,2. Biofeedback and Self-Regulation* (Vol. 2). Retrieved from <https://link.springer.com/content/pdf/10.1007%2FBF00998623.pdf>
- Wohlschläger, A. (2000). Visual motion priming by invisible actions. *Vision Research*, 40, 925–930.
- Yon, D., Gilbert, S. J., De Lange, F. P., & Press, C. (2018). Action sharpens sensory representations of expected outcomes. *Nature Communications*, 9(4288), 1–8.
<https://doi.org/10.1038/s41467-018-06752-7>
- Zamariola, G., Maurage, P., Luminet, O., & Corneille, O. (2018). Interoceptive accuracy scores from the heartbeat counting task are problematic: Evidence from simple bivariate correlations. *Biological Psychology*, 137, 12–17.
<https://doi.org/10.1016/j.biopsycho.2018.06.006>
- Zhou, W., Jiang, Y., He, S., & Chen, D. (2010). Olfaction modulates visual perception in binocular rivalry. *Current Biology*, 20(15), 1356–1358.
<https://doi.org/10.1016/j.cub.2010.05.059>

Appendices

Appendix 1: Stimuli used in Chapters 2, 3, 5 (Exp.1)

Main task:



Threshold-finding:




Example masks (note that masks were generated afresh on each trial):



Appendix 2: Stimuli used in Chapter 4 (pilot)

Main task and threshold-finding:



Masks:




Appendix 3: Stimuli used in Chapter 5 (Exp. 2)



Appendix 4: Stimuli used in Chapter 6

Main sets:



Demo set:

