Probing Visual Consciousness with Transcranial Magnetic Stimulation

Christopher. P. G. Allen

PhD Thesis

2012

Cardiff University

School of Psychology

Declaration

This work has not been submitted in substance for any other degree or award at this or any other university or place of learning, nor is being submitted concurrently in candidature for any degree or other award.

Signed …………………………………………… (candidate)   Date …………………………

STATEMENT 1

This thesis is being submitted in partial fulfillment of the requirements for the degree of PhD

Signed …………………………………………… (candidate)   Date …………………………

STATEMENT 2

This thesis is the result of my own independent work/investigation, except where otherwise stated.

Other sources are acknowledged by explicit references. The views expressed are my own.

Signed …………………………………………… (candidate)   Date …………………………

STATEMENT 3

I hereby give consent for my thesis, if accepted, to be available for photocopying and for inter-library loan, and for the title and summary to be made available to outside organisations.

Signed …………………………………………… (candidate)   Date …………………………

STATEMENT 4: PREVIOUSLY APPROVED BAR ON ACCESS

I hereby give consent for my thesis, if accepted, to be available for photocopying and for inter-library loans **after expiry of a bar on access previously approved by the Academic Standards & Quality Committee.**

Signed …………………………………………… (candidate)   Date …………………………

Acknowledgments

Contents

Glossary

BIP	Blindsight Inducing Pulse
BrC	Measure of response criteria in conscious detection
cTBS	Continuous Theta Burst Stimulation
ERD	Event Related Desynchronisation
ERF	Event Related Field
GABA	Gamma-Aminobutyric Acid
GLX	Glutamate and glutamine
GW	Global Workspace
HOT	Higher Order Thought
iTBS	Intermittent Theta Burst Stimulation
LGN	Lateral Geniculate Nucleus
MEG	Magnetoencephalography
MRS	Magnetic Resonance Spectroscopy
NAA	N-Acetyl Aspartate
NCC	Neural Correlate of Consciousness
PcU	Measure of proportion correct when reportedly 'unseen'
Ppm	Parts Per Million of the proton frequency
PrC	Measure of conscious detection
PT	Phosphene Threshold
R/M	Retinotectal and Magnocellular pathways
rTMS	Repetitive Transcranial Magnetic Stimulation
SC	Superior Colliculus
SDT	Signal Detection Theory
SNR	Signal-to-Noise Ration
TMS	Transcranial Magnetic Stimulation

"When there is name-and-form, consciousness comes to be; consciousness has name-and-form as its condition… it does not go further back than name-and-form."

(Samyutta Nikaya (Connected Discourses) 12:65, pp. 602. Cf. D. ii 32, pp.211, circa 500 BC, Siddhartha Gautama (the Historical Buddha), translated by Stephen Batchelor)

Chapter 1.
Introduction to experimental design

1.1 Overview

This thesis explores the effects of transcranial magnetic stimulation (TMS) on conscious perception and visual processing. Chapter 1 addresses issues of experimental design. Two broad classes of TMS intervention were used and are reported in separate chapters. Chapter 2 involves repetitive 'off-line' TMS combined with neuroimaging techniques. Chapter 3 employs 'on-line' TMS applied with temporal specificity to track the passage of information through early visual cortex. Chapter 4 is a general discussion primarily concerned with the issues encountered experiments oriented towards consciousness.

The area of interest in this research is conscious processing. One approach to this field is to contrast clear reports of awareness with perception explicitly lacking in consciousness. The neurological phenomenon of *blindsight* exemplifies this contrast. Blindsight can follow damage to the primary visual cortex and is expressed by subjects insisting that they are not consciously aware of stimuli while nevertheless being able to guess the stimulus characteristics above chance. Similar effects have previously been reported following TMS of early visual cortex in healthy observers. Like clinical blindsight, TMS-induced blindsight straddles the boundary between consciousness and perception lacking in consciousness.

Most of the experiments reported in this thesis exploited a similar behavioural paradigm in which subjects were presented with simple arrow stimuli and asked several questions. One set of questions was a forced choice discrimination decision which allowed for the tracking of perception independently of conscious criteria. Other questions, such as 'Did you consciously see the target?' were employed to ascertain measures of the subjects' conscious experience of the stimuli. As with clinical blindsight, independent changes in these measures illustrate the borders that surround conscious awareness, revealing the neuronal

processes that contribute to conscious processing in contrast to those responsible for perception lacking in consciousness.

The effects of off-line repetitive TMS can last for approximately one hour or longer and are explored in Chapter 2. This series of experiments initially aimed to use the suppressive after-effect of TMS to study the behavioural consequences of reducing cortical excitability in areas where lesions can cause blindsight. However, contrary to expectations, a repetitive TMS protocol that suppresses cortical excitability was found to enhance rather than impair a measure of subjects' conscious experience. This finding contradicted the hypothesis and warranted further investigation. The experimental configuration and initial results were therefore replicated (2.6.1) in conjunction with magnetoencephalography (MEG) and other experimental manipulations (2.6). In addition, magnetic resonance spectroscopy (MRS) was employed to measure corresponding biochemical changes induced by TMS (2.5). The overall conclusion of the off-line experiments was that repetitive TMS increased conscious detection of stimuli, and that this intervention potentiated mechanisms of gating through inhibition.

Chapter 3 reports a series of experiments involving on-line (or event-related) occipital TMS, in which single pulses or short bursts of stimulation are delivered at varying intervals relative to visual stimuli. This approach allowed for mapping the time course and role of activity that is causally constitutive for visual perception and awareness. Visual stimulus characteristics were also manipulated to permit conclusions concerning the role of specific anatomical pathways in perception and awareness. Two main on-line studies experiments were conducted: a between-subjects study (3.3) and a single-subject case study (3.4-3.7). The single-subject case study included three independent experiments (3.5-3.7) that considered the role of different phases of occipital activity and different anatomical pathways in perception. The between-subjects experiment employed similar interventions to those of the single-subject study but with lower resolution. Overall, the on-line TMS experiments suggest that conscious awareness of stimuli depends on visual cortical activity from approximately 100ms after stimulus onset, and that it may benefit from input through retinotectal and/or magnocellular pathways during the very early and very late stages of processing. In contrast the preserved 'unseen' abilities during TMS-induced blindsight did not appear to rely on such input, contrary to previous suggestions, and may instead depend on geniculostriate activity and activity subsequent to the 100ms epoch that is central to conscious processing.

In sum, the experiments described here explored the neural basis of consciousness. By analysing changes in reported awareness of stimuli following different experimental interventions, it has been possible to link various neuronal substrates to distinct aspects of visual processing.

1.2 General background

1.2.1 Philosophical context

This section will explore some of the background considerations and aims that informed the experimental design.

The philosophical strategy adopted for this research has been that of Neurophenomenology (Lutz, 2003; Varela, 1996; Varela, Thompson, & Rosch, 1991). This framework seeks to explore the coupling between physiology and phenomenology[1]. Subjects' phenomenological experiences of experiments were quantified predominantly through psychophysical methods, and related to physiological changes assessed through the use of neuro-interventional and neuroimaging techniques.

The principal technique used in all experiments was TMS. This method perturbs neuronal activity and can also have phenomenological consequences (e.g. Boyer, Harrison, & Ro, 2005). TMS thus has a unique advantage over the more traditional or common cognitive neuroscience techniques: it allows us to move beyond the correlation of brain-behaviour states to assess their causal relationship (Campbell, 2007; Chambers & Mattingley, 2005; Walsh & Cowey, 2000).

---

[1] Phenomenology can be understood as the study of the first person subjective element of consciousness and experience. It is commonly linked with the 'what it is like' (Nagel, 1974) aspect to experience and the notion of 'qualia' (Lewis, 1929). Historically it derives from the work of continental philosophers such as Husserl and Merleau-Ponty, but it can also be attributed to the investigations of Hume and eastern traditions.

The investigation of differences between states of consciousness - that is, differences in content - appears to be a perfectly valid approach to the investigation of consciousness (e.g. Lutz, Lachaux, Martinerie, & Varela, 2002) but it is not the one adopted here (some of these alternative approaches are considered in Chapter 4). Rather, my approach has been to consider the presence versus lack of subjective awareness. This can be thought of as a border that surrounds conscious experience. This framework conveys an understanding of what is involved in conscious experience in contrast to processes closely related but not involved, thus drawing what Dehaene calls a 'negative picture' (pp13 Dehaene & Changeux, 2004). This contrast between perception and explicit conscious awareness is also discussed by Dretske (2000) as one of the few viable directions for investigations that isolate, and therefore operationally define, consciousness.

There is scope within this contrast between conscious processing and perception lacking consciousness to look at differences that are uninformative in terms of the conscious experiences in question. For example, the processes of the circulatory system are necessary for, and coupled to, conscious experience, but the contrasts between a functioning and a non-functioning heart are not particularly informative about consciousness. Likewise, inflammatory responses due to skin contact with poisonous plants can be described as a form of perception, but again the utility of this as a contrast state would be uninformative due to vast differences between the inflammatory system and the systems that support consciousness (Dretske, 2000). There is, therefore, a need to be as 'minimal' (Chalmers, 2000) or tight around this border as possible.

In order to frame an optimally informative contrast, we require both states of conscious experience and states in which all things are equal except for the conscious experience. Blindsight (Sanders, Warrington, Marshall, & Wieskrantz, 1974) is deemed to be just such a state (Holt, 2003). Subjects with blindsight are able to respond veridically to the environment (in a manner that in other circumstances might be credited with conscious awareness) but insist that they are not aware of the environmental stimuli being presented. Hence, what they lack is specifically conscious experience (Holt, 2003). This is why blindsight, and in particular, TMS-induced blindsight (e.g. Boyer, et al., 2005) forms the focus of this thesis.

1.2.2 Blindsight

Blindsight classically arises following lesions to primary visual cortices and has been defined as "visual capacity in a field defect in the absence of acknowledged awareness" (pp166 Weiskrantz, 1986). Blindsight patients demonstrate perception insofar as their performance in forced choice tasks is greater than expected by chance, yet at the same time will not acknowledge awareness of the stimuli having been presented, often insisting that they are guessing (Weiskrantz, 1986). These demonstrable abilities, in the presence of a negative report of awareness, include localisation, orientation, colour, and lexical discrimination (Cowey, 2010; Weiskrantz, 1986). Residual abilities in blindsight are most often demonstrated using forced choice tasks (e.g. Sanders, et al., 1974). However, blindsight has also been shown *indirectly* where a cue presented within a subject's 'blind field' may not be acknowledged but still affects other motor responses (e.g. Kentridge, Heywood, & Weiskrantz, 1999). Blindsight therefore corresponds to a dissociation or disjunction between conscious detection and perceptual abilities (Weiskrantz, 1986, 2001, 2009). Although Weiskrantz was the principal investigator on the paper that coined the term (Sanders, et al., 1974), and remains a central figure in blindsight literature, it is worth noting that a dissociation between subjects reporting not having 'seen' stimuli despite above-chance performance was first documented over 100 years ago (Sidis, 1898).

Blindsight is a controversial phenomenon; its very existence has been questioned (Campion, Latto, & Smith, 1983) and its implications and interpretations are diverse (compare Block, 1995; Chalmers, 1996; Dennett, 1991). Some of these issues are dealt with in the general discussion (Chapter 4). For now, it need only be acknowledged that blindsight illustrates a condition in which subjective conscious awareness is suppressed while perceptual performance without acknowledged awareness is relatively preserved; this is a relatively uncontroversial claim around which the current project was conceived. Experiments were therefore required to be capable of tracking subjects' awareness of stimuli and, independently, subjects' perceptual abilities when they reported a lack of awareness. Hence experiments should be capable of illustrating the dissociations between conscious experience and response abilities, i.e. blindsight type effects. As in previous studies, the general rationale for using blindsight in the current investigation of consciousness is that it

offers a set of contrast states where all things are equal except, specifically, conscious awareness.

1.2.3 TMS-induced blindsight

During TMS a strong current (~1.8 kV, 7kA) briefly (~250 μs) is passed through a coil to induce a time-varying magnetic field (~2 tesla) (Magstim Company Limited, Carmarthenshire; stimulator specifications; Wagner, Rushmore, Eden, & Valero-Cabre, 2009). When applied over the head this time-varying magnetic field induces an electrical field in the underlying cortex, resulting in the activation of nerve cells (Barker, Jalinous, & Freeston, 1985).

Under certain circumstances, disruption to ongoing activity caused by TMS has been seen as analogous to lesions (Pascual-Leone, Walsh, & Rothwell, 2000). Given the existence of lesion-based blindsight, TMS-induced blindsight was therefore a natural progression for experimentation. Although electromagnetic disruption of visual perception has been known about for over 100 years (d'Arsonval, 1896) it was not until 1989 that such disruption was clearly linked with visual cortical processes (Amassian et al., 1989). In 2004 the first study used the term blindsight to describe a disruption of visual areas caused by TMS (Ro, Shelton, Lee, & Chang, 2004). Ro and colleagues used TMS to reduce subjects' reported awareness of centrally presented distracters and showed that the distracter, even though it was reportedly beyond awareness, increased response times during a saccadic but not manual response task (Ro, et al., 2004). The distracter was thus described as being rendered 'unaware' by the TMS, even though it still had a measurable perceptual effect – a 'blindsight type' effect.

It is possible to argue that oculomotor mechanisms are somewhat lower down the cognitive scale from the perceptual abilities common in the most influential demonstrations of classic blindsight (Weiskrantz, 1986), due to their governance, at least in part, being reflexive (de No, 1933). Hence, influence over saccades might be expected to function independently of effects on conscious awareness and therefore might not be considered relevant to blindsight. However, Ro and colleagues went on to demonstrate different aspects

of TMS induced blindsight, including those involving higher, more complex, perceptual abilities such as colour and orientation discrimination (Boyer, et al., 2005).

Using a variety of tasks, investigators have been able to demonstrate that different perceptual abilities can be dissociated from conscious awareness. This has involved the application of TMS to attenuate measures of awareness relative to measures of perceptual ability. For example, Jolij et al (2005) claimed to uncover 'affective blindsight' by showing that occipital TMS applied 110ms after target onset impaired reportedly conscious localisation of targets without degrading the emotional content of targets (cartoon faces). Others have shown that deflections in grasping action can result from distracters for which awareness is suppressed by occipital TMS (Christensen, Kristiansen, Rowe, & Nielsen, 2008).

Apart from its reversible nature, a great advantage of TMS-induced blindsight over classic blindsight is the opportunity it affords to explore the time course of the dissociation between conscious and unconscious vision. Pulses can be applied and activity disrupted at specific times relative to visual stimuli (on-line), and periods of activity may therefore be revealed which can be related to phenomenal and behavioural consequences. TMS-induced blindsight thus contributes to a wider literature where TMS is used to elucidate the timing and role of visual cortical activity (see section 3.1.2). One of the main aims of the work described here has been to disrupt activity in a temporally precise manner and relate this to the contrast involving consciousness. Similar reasoning and experiments have led to the proposal that there is an early feedforward sweep of activity that is 'unconscious' and preserved in classic blindsight, which is then followed by a recurrent phase of processing that confers conscious awareness (Lamme, 2001, 2006b; Lamme & Roelfsema, 2000). Probing this theory was a major aim of the experiments undertaken here, particularly the 'on-line' experiments. The research which inspired, and is therefore most relevant to, the experiments described here are previous demonstrations of TMS induced blindsight (Boyer, et al., 2005; Christensen, et al., 2008; Jolij & Lamme, 2005; Ro, et al., 2004), all of which involve on-line TMS. These are considered in greater depth in Chapter 3.

The off-line TMS experiments (Chapter 2) focused instead on the neuronal correlates of changes in conscious awareness following repetitive TMS. In particular, these experiments

studied the oscillatory and neurochemical changes that might govern the conscious detection of visual stimuli.

1.2.4 Types of blindsight and the pathways involved

Blindsight has been divided into sub-classes. The most commonly discussed differentiation is between type I and type II (Weiskrantz, 1997). Type I corresponds to an absolute absence of awareness, whereas type II demonstrates a markedly reduced form of awareness. Type II subjects are not reflectively aware of the precise content of a presented stimulus yet can guess as to its character and had a minimal sense that something was presented to them (as in 'agnosopsia', not knowing what one sees Danckert & Rossetti, 2005). A similar phenomenon that has been considered akin to type II blindsight is Riddoch syndrome, where a subject may be aware of moving but not stationary visual stimuli in the blind field. This ability has been related to pathways that bypass the primary visual cortex and directly innervate the motion sensitive area V5/middle temporal cortex (MT+) (Zeki & Ffytche, 1998).

Danckert and Rossetti (2005) offer a further distinction between 'action' vs 'attentional' blindsight (as well as agnosopsia). However, this differentiation does not seem to have been widely adopted, perhaps because it attempts to reflect differences in underlying neuronal pathways that remain unproven: action blindsight is assumed to involve direct innovation of posterior parietal cortex, whereas attention blindsight is assumed to rely on extrastriate visual areas, including V5/MT+. The distinction may also blur the categorisation between blindsight and widely accepted neurological disorders, such as neglect and visual agnosia.

In terms of the nature of classic blindsight (type I and II), a frequently overlooked point is that subjects with blindsight express their condition in only a portion of their visual field; they are often normally sighted across the bulk of their field. The straightforward possibility of light being scattered from retinal locations corresponding to the blind field into their sighted field has been raised as an explanation that trivialises blindsight (Campion, et al., 1983). Proponents of blindsight have gone to great lengths in order to demonstrate that

this is not the case (e.g. Kentridge, Heywood, & Weiskrantz, 1997). Another aspect, again often overlooked, is the role of the intact visual cortex and plasticity in supporting residual guessing abilities, although this has been suggested (Campion, et al., 1983 commentary by Bach-y-Rita, P.)

In addition to debates over the nature and interpretation of blindsight, the physiology that facilitates the residual abilities is one of the most keenly debated aspects in the literature (Cowey, 2010), driven by compelling evidence for opposing points of view (cf. Leh, Mullen, & Ptito, 2006; Schmid et al., 2010). There are broadly two theories or positions for the pathways that facilitate blindsight type abilities. Some advocate a role for spared projections from the lateral geniculate nucleus (LGN) (Schmid, et al., 2010) which mainly projects to the primary visual cortex and is consistent with there being 'spared islands' of primary visual cortex that facilitate residual perception (Radoeva, Prasad, Brainard, & Aguirre, 2008). An opposing view is that the residual abilities of blindsight are supported by pathways that mainly bypass the striate cortex, primarily via the retinotectal route that includes the superior colliculus (e.g. Leh et al 2006, 2009). In terms of classic blindsight, the majority of studies favour the latter interpretation, although the matter is by no means settled (see section 3.1.3, and figure 3.1.3.1). Moreover, in terms of the relatively new phenomenon of TMS-induced blindsight, the question of which pathway supports these perceptual abilities is still very much open. It could be the case that the pathways which support the positive abilities seen in classic blindsight are not the same as those found to mediate abilities seen in TMS-induced blindsight, or possibly that no single pathway underlies the spared abilities of blindsight. Rather, the pathways that support the spared abilities of blindsight may be dependent on the task used to quantify the residual abilities and method of neurological intervention.

These accounts were contrasted in Chapter 3 (see section 3.1.3), exploiting the observation that the retinotectal and magnocellular pathways of the LGN do not appear to receive colour opponent input from short wave cones of the retina (Sumner, Adamjee, & Mollon, 2002). A key aim of these on-line experiments was to adjudicate between competing theories of blindsight by combining time-locked TMS with stimuli that were calibrated to activate or bypass specific anatomical pathways.

1.2.5 Summary of background and aims

The broadest aim of this research has been to study the relationship between physiological and phenomenological measures (Varela, 1996), with a view to elucidating neural mechanisms of visual awareness in the human brain. In doing so the experiments focused on TMS-induced blindsight to investigate the mechanisms of consciousness and phenomenology through a contrast between explicit conscious awareness and states specifically lacking in consciousness. The next section focuses on establishing a reliable method for deriving phenomenologically-related measures that allow this contrast. Chapters 2 and 3 address the physiological side of these relationships. Specifically, Chapter 2 makes use of several imaging and interventional modalities to probe the physiological mechanisms that may underlie modulation of conscious awareness, with a particular focus on biochemical and oscillatory effects of TMS. Chapter 3 aims to elucidate the time course and pathways of early visual processing that support conscious awareness and perception lacking in consciousness.

1.3 Tasks, measures and stimuli

This section provides a general overview of the stimuli and behavioural measures employed in this project. It focuses on the stimuli and tasks originally developed and deployed in the off-line TMS experiments and the reasons for their use. Issues arising and subsequent adaptations are covered in the discussion of the off-line experiments (2.8) and introduction to the on-line experiments (3.2).

Fred Dretske has pointed to the utility of comparing at least two tasks which give access to the subjective experience of subjects, in contrast with those that probe raw perceptual abilities (Dretske, 2000). This is also what is required to demonstrate the dissociation of blindsight (Weiskrantz, 1986) and a principle upon which these experiments were designed.

There are many perceptual responses to stimuli that have the potential to reveal blindsight type abilities. Here an explicit question, 'Did you consciously see the stimulus?'

was posed in conjunction with a non-speeded forced choice discrimination judgement about the stimulus. Subjects' forced choice discrimination ability, specifically when they reported *not* having seen the stimulus, served as a measure of their perceptual ability, explicitly lacking in conscious awareness. This design was intended to reflect the principal components of the classic demonstrations of blindsight by Weiskrantz and colleagues (1986). It was also partly inspired by the Perceptual Awareness Scale (PAS) of Overgaard and colleagues, which has been used to compare the confidence subjects place in a judgment with their objective performance (Overgaard, 2003; Ramsoy. T.Z., 2004).

Reaction times, manual response effects in terms of speeded cueing, saccadic effects and indirect measures in general are also capable of demonstrating blindsight type effects (e.g. Leh, Mullen, et al., 2006; Ro, et al., 2004; Weiskrantz, 1996). However, these were not used in the studies presented here for the following three related reasons. First, there are statistical considerations: forced choice discrimination performance can be more straightforwardly equated with detection performance than can effects on reaction times or movement deflections. Thus, dissociations between two forced choice measures are easier to interpret and are less likely to be conflated by differential sensitivity of very different measures to cortical interference. Second, and more importantly, if blindsight is to be used to reveal consciousness through a comparison with perceptual states lacking in awareness, then the differences between the perceptual and awareness probes should be minimised in all other respects (Chalmers, 2000; Marcel, 1983). Intuitively and phenomenally, detection tasks have more in common with overt forced choice discrimination tasks than with cueing tasks and other more indirect measures. Relatedly, it seems entirely possible that the neuronal architecture underlying cueing might be quite different from that underlying perceptual abilities involved in higher perceptual discriminations; thus the use of cueing as a contrast state might reveal something quite different, and possibly less informative. It is also possible that abilities demonstrated with different tasks (cueing vs. forced choice discrimination) could reflect different types of blindsight (e.g. attentional vs. action). Third, one of the questions we wished to pose was in relation to the role of the superior colliculus, so we required a probe that could operate via this route but did not necessarily require such input under normal conditions. If we had chosen a probe that had previously been shown to depend critically upon the superior colliculus under normal conditions, such as saccadic dynamics (Lee, Rohrer, & Sparks, 1988; Ro, et al., 2004), then the capacity of experiments to reveal

whether or not residual abilities of blindsight depend on the superior colliculus would be compromised.

1.3.1 Stimuli and questions

The visual stimuli were arrows (see figure 1.3.1.1). At least two tasks or probes are required to track the borders between conscious perception and perception that is independent of, or lacking in, awareness (Dretske, 2000). One task is required to be capable of tracking subjective phenomenal experience of presented stimuli. The second task must be capable of illustrating perception and of being modulated independently of the first task. Two questions were therefore applied to each trial: i) 'Was the arrow pointing left or right?' (prompted by the display of 'L/R') – a forced choice discrimination judgement that is capable of revealing perception even when the subject reports not being aware of the stimulus (Kolb & Braun, 1995); and ii) 'Did you consciously see the arrow? Yes or No' (prompted by 'Y/N') – a question designed to track the subject's reflective experience of the stimulus[2].

Figure 1.3.1.1. A left arrow, approximately as presented to subjects. The surrounding bars are the luminance noise, which alternated every 20ms with a range of 10cd/m[2].

---

[2] To minimise the influence of the two questions upon one another, the ('L/R') discrimination question was posed prior to the detection ('Y/N') question. This was because, intuitively, the reflection upon the detection question could affect the subject's response to the discrimination judgment through the adjustment of criteria, whereas the effect of the discrimination task had upon the detection judgement was likely to be less influential. Therefore the discrimination question was presented first.

Arrows were chosen because their presence can be detected and they can be discriminated between. Several other features of arrows also encouraged their use. Most importantly they have particularly strong cueing properties (e.g. Ristic & Kingstone, 2006). This meant that they were thought likely to elicit above chance responses in the forced choice tasks at or below thresholds of awareness, as was confirmed in preliminary piloting. They also offered the possibility of examining the cueing or indirect elements of blindsight abilities, which are an ongoing aspect of this series of experiments not covered in this thesis.

Arrows were sometimes presented (stimulus-present) and sometimes not (stimulus-absent) in order to facilitate the use of Signal Detection Theory (SDT section 1.3.2) in relation to the conscious detection question. The luminance of the arrow was titrated for each subject so that, in the absence of TMS, it was detected on approximately half of stimulus-present trials. This level of detection was deemed low enough to be susceptible to TMS, yet was high enough to be capable of demonstrating a range of response changes.

A key feature of the stimuli was the luminance noise, which alternated at candela intervals over a range of ~10 cd/m^2 with the same periodicity as the stimulus presentation (20ms). Noise is a required feature of the 's-cone' stimuli used in Chapter 3 to isolate the collicular pathways (see section 3.1.3 and 3.2.1). Because this manipulation was a possibility in all experiments the noise was present during all stimuli presentations (although it was only exploited in the on-line experiments). An additional advantage of using a noisy background for stimuli is that it might increase 'false alarm rates', which aided the application of SDT to the detection task (see below 1.3.2.1). It is worth noting that the level of luminance of the arrows (or saturation of s-cone / purple colour in the case of the pathway-specific stimuli) was always greater than the range of the noise. Thus stimulus-present conditions always differed physically from stimulus-absent conditions. Noise was presented beyond the edge of the central region where the arrows were presented (see figure 1.3.1.1); this allowed the arrow to be presented at different locations within the noise (a manipulation not used in the final experiments). The special distribution of the noise also reduced the illusory appearance of arrows within the noise during stimulus-absent conditions, an effect which became apparent during preliminary piloting and development of the stimuli.

As with the choice of probes / questions, the intention in the design of the stimuli was to develop a task that required perception with a degree of abstraction or higher order

processing. If responses to stimuli are reflexive then it may not be appropriate to use them in the construction of a contrast state to reveal consciousness (Dretske, 2000). It was felt that the detection and discrimination of arrows required a sufficient level of abstraction (syntax) to draw on mechanisms of cognitive control and thus minimise the contribution of reflexive processes to performance.

Non-lateralised foveal stimuli allowed for foveally targeted TMS, avoiding potential inconsistencies in the lateralisation of TMS's effects (e.g. Corthout, Hallett, & Cowey, 2002). Additionally 'S-cones' (see section 3.1.3) are distributed more densely and consistently within the fovea than at the periphery (Sumner, 2000), additionally suggesting foveal presentation.

1.3.2 Derivation of measures

As mentioned, in order to illustrate blindsight type effects we require a measure of conscious detection and perceptual ability revealed by a discrimination judgement. This section describes the rationale and methodology for deriving these measures.

1.3.2.1 Detection

A major criticism of approaches that attempt to tackle measures of awareness, and phenomenologically led approaches in general, is that a subject's response criteria might fluctuate independently of changes in the conscious experience (e.g. Campion, et al., 1983; Eriksen, 1960; Nisbett, 1977). This concern of 'response bias' will be returned to in greater depth in Chapter 4. For now, it is sufficient to note that the application of Signal Detection Theory (SDT)(see Green & Swets, 1966; Macmillan & Creelman, 1990) to the 'Did you consciously see?' question allows us to compute measures of bias, and of conscious sensitivity largely independent of bias. In this context, the measure of sensitivity (e.g. d', Tanner & Swets, 1954) reflects the difference between subjects' propensity to say they saw the stimulus when it was present, and their propensity to say they saw it when it was absent. That is, it denotes the subject's acknowledged detection of, or explicit sensitivity to, external stimuli. Bias measures (e.g. β, Green & Swets, 1966) are expressions of the propensity of subjects to say they saw the stimuli independently of what was presented. The application of

such measures to blindsight type phenomena was first explored in the original case study by Weiskrantz (1986).

In SDT a positive response to a presented stimulus is a 'Hit', while a positive response when the stimulus is absent is a 'False Alarm'. If a stimulus is present but reported as unseen then the trial is classed as a 'Miss', whereas the same negative response to an absent stimulus is a 'Correct Rejection'. This stimulus-response mapping can be represented as follows:

Table 1.3.2.1.1. SDT classification used in the derivation of measures of conscious detection and bias in detection.

	Hit	False Alarm	Correct Rejection	Miss
Stimulus	Present	Absent	Absent	Present
Response	'Yes'	'Yes'	'No'	'No'

In the current studies, the occurrence of these conditions and responses was computed into hit rates and false alarm rates, forming the basic units from which the measures of sensitivity and bias were derived. Hit Rates (HR eq1.1) are the number of hits divided by the number of stimulus-present trials. False Alarms Rates (FAR eq1.2) are the number of false alarms divided by the number of stimulus-absent trials. There are however several different sensitivity and bias measures available to the experimenter; these form the focus of the next section.

$$HR = \frac{Hits}{(Hits + Misses)}$$

eq1.1

$$FAR = \frac{False\ Alarms}{(False\ Alarms + Correct\ Rejections)}$$

eq1.2

1.3.2.2 Which SDT measures?

Early piloting revealed that the conscious detection question, as opposed to the forced choice discrimination question, was subject to bias. This bias manifested conservatively as a preponderance toward saying 'No, I have not seen it', leading to low false alarm rates. Subjects often adopt conservative criteria in comparable 'yes/no' tasks (Macmillan & Creelman, 1990). The cause and interpretation of this phenomenon is considered in greater depth in Chapter 4. Here, the criterion has consequences for the choice of which SDT measures to apply.

Parametric versions of SDT such as the classic d', C or β involve z normalising HR and FAR (eq 1.3 to 1.5). If either of these rates is 1 or 0 then this is not possible. So where FAR is 0, classic SDT cannot be straightforwardly applied, which may often be the case in the context of such conservative criteria. Some transforms allow for extreme responses, including the 1/2n adjustment. The 1/2n calculation assumes that when extreme responses are found (e.g. FAR=0) the number of false alarms actually lies between 0 and 1 over the number of trials, i.e. 1/2n (Hautus, 1995; Macmillan & Kaplan, 1985).

$$d' \ = \ z(HR) - \ z(FAR)$$

eq1.3

$$c \ = \ -0.5(z(HR) \ + \ z(FAR))$$

eq1.4

$$\beta \ = \ e^{\frac{z^2(FAR) - z^2(HR)}{2}}$$

eq1.5

Formulae taken from Wright, Horry, & Skagerberg (2009)

An additional concern is the central assumption of classic SDT that the distributions of signal (from which hits result) and noise (from which false alarms result) are both normal and equal in their variance (pp140 Stanislaw & Todorov, 1999). The low incidence of false alarms suggested that this assumption may not have been valid, as did likely imbalances of trial numbers in the stimulus-present vs. stimulus-absent conditions (e.g. owing to trials being

excluded owing to TMS-induced blinks or different classes of stimuli). For these reasons a non-parametric form of SDT was employed throughout these experiments (Corwin, 1994).

Although A' and B" (Pollack & Norman, 1964; Stanislaw & Todorov, 1999 ; eq 1.6 and 1.7) are non-parametric SDT measures and thus not subject to parametric assumptions of d' and β, they nevertheless cope poorly with zero FARs or maximal HRs, again because such levels of detection would require division by 0 (see eq 1.6 and eq 1.7). For this reason, the simpler measures of Pr and Br (Corwin, 1994, eq 1.8. and 1.9) were used here as they are appropriate for low FARs, make fewer statistical assumptions, and have a far greater tolerance to imbalances in the ratio of stimulus-present to stimulus-absent trials (Corwin, 1994). Here Pr was used as an alternative to d' that is range-compatible with the proportion of correct 'unseen' discriminations in the forced choice task (1.3.2.3). The measure Br was adopted for bias, being approximately equivalent to β or C.

if HR≥FAR

$$A' = 0.5 + \frac{((HR - FAR)(1 + HR - FAR))}{(4HR(1 - FAR))}$$

else

$$A' = 0.5 - \frac{((FAR - HR)(1 + FAR - HR))}{4FAR(1 - HR)}$$

eq 1.6

if HR≥FAR

$$B'' = \frac{(HR(1 - HR) - FAR(1 - FAR))}{(HR(1 - HR) + FAR(1 - FAR))}$$

else

$$B'' = \frac{(FAR(1 - FAR) - HR(1 - HR))}{(FAR(1 - FAR) + HR(1 - HR))}$$

eq 1.7

$$Pr = HR - FAR$$

eq 1.8

$$Br = \frac{FAR}{(1 - (HR - FAR))}$$

eq 1.9

Eq 1.6 and 1.7 taken from Snodgrass & Corwin, (1988). Eq 1.8 and 1.9 taken from Corwin, (1994).

1.3.2.3 Discrimination

In deciding which discrimination judgments to incorporate in the measures of capacity which may persist when consciousness is suppressed, there two options are available. The first approach is to analyse all discrimination judgments; the alternative is to focus exclusively only on discrimination performance when subjects report the target as 'unseen'. Both approaches have been previously adopted. Classic blindsight studies (e.g. Weiskrantz, 1986) often opt for the use of all forced choice discrimination trials. However, since subjects rarely, if ever, report awareness in classic blindsight, the two measures would be expected to produce equivalent results. Studies of subliminal processing also often use all available trials to demonstrate perceptual abilities but may also assess the absence of awareness separately (e.g. Sumner, Tsai, Yu, & Nachev, 2006, for review see Kouider & Dehaene, 2007). The use of all discrimination trials has also been applied to the closely related, yet distinct, contrast between 'objective' and 'subjective' measures of detection (e.g. de Graaf, Goebel, & Sack, 2012); this distinction is discussed in section 4.4 and appendix A6. TMS-induced blindsight studies, which contrast conscious awareness with perception lacking in awareness, tend to use 'unseen' trials only in the demonstration of preserved abilities (Boyer, et al., 2005; Ro, et al., 2004).

Discrimination ability on reportedly 'seen' trials is arguably uninformative because subjects are invariably at ceiling; mean percentage correct for 'seen' discrimination in the first off-line TMS experiment (section 2.1) was 97.7% ± 0.53% (SD). This positive aspect of a report is unsurprising – if a subject sees an arrow and is aware of the experience, they tend correctly to judge its direction. Moreover, if the measure of discrimination includes conditions in which subjects are aware of the stimulus, then a genuine dissociation between detection and discrimination would be obscured because 'seeing' the stimulus more (or less) often would coincide with an increase (or decrease) in discrimination ability. Thus, the key measure of perception *independent* of consciousness is required to be discrimination on 'unseen' trials. In addition, to ensure that discrimination performance indexes perceptual

processing, reportedly 'unseen' performance must remain above chance in the control conditions (as was confirmed for these stimuli during pilot experiments). For these reasons discrimination responses were categorised according to whether or not they occurred on trials when the subject reported having seen the stimulus ('seen' or 'unseen'). 'Seen' trials were then removed from the analysis to produce a measure of 'unseen' discrimination.

The categorisation of the discrimination judgment according to the detection question has consequences for inferential statistical power: within any set number of trials, the greater the proportion of conscious detections then the lower the proportion of 'unseen' discriminations. This discrepancy between the measures can often be overcome by increasing the numbers of trials, although this adjustment is subject to limitations (as discussed in Chapter 2). Statistical tests can also be applied to confirm that the two measures are appropriately comparable in power. One such approach is to compare the variance of the two measures based on these trial numbers (Keppel, 1982). In particular, Keppel recommends that the ratio of standard error estimates should not exceed 3:1 if the measures are to be considered comparable. This approach was used in the on-line experiments in particular (e.g. see section 3.3).

With these considerations in mind, the number of trials per cell of the study design was determined by calculating the number of trials required to demonstrate above chance performance on 'unseen' trials. To determine this, a power analysis was applied based on a χ^2 distribution of performance on 'unseen' trials. As can be seen in figure 1.3.2.3.1, the benefit in terms of the likelihood of demonstrating above chance performance, conferred by adding additional trials, began to plateau from ~25 trials. Subjects were thresholded to detection levels at which they reported not having seen the arrow on approximately 25% of all trials. Therefore, the number of trials required for any one experimental condition was approximately 100 trials, which was usually rounded up to ~120 trials.
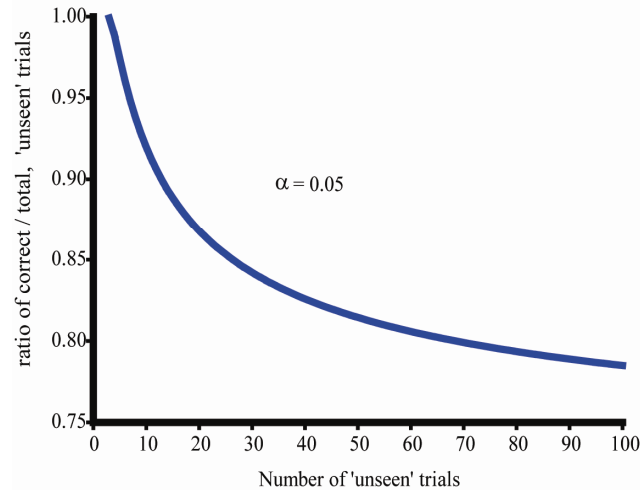
Figure 1.3.2.3.1 Illustration of the performance (proportion correct ratio) required to demonstrate above chance performance ($\alpha = 0.05$), given the number of 'unseen' trials, using a χ^2 test.

SDT can also be applied to the forced choice discrimination questions (Macmillan & Creelman, 1990) by calculating detection sensitivity to left arrows and bias towards saying 'left', or likewise for right arrows. For example, in terms of left arrows, the number of hits, false alarms, correct rejections and misses can be calculated according to Table 1.3.2.3.1 (below). These calculations were only used to ensure that subjects were performing the task as instructed (i.e. not simply giving the same response on all trials during which they felt they could not see the arrow). The measure used to quantify 'unseen' abilities in these experiments (PcU) was the proportion-correct performance of the 'left/right' decision, when subjects reported not having seen the arrow.

Table 1.3.2.3.1 Table of SDT classification used to ensure subjects performance.

	Hit	False alarm	Correct Rejection	Miss
Response	'Left'	'Left'	'Right'	'Right'
Stimulus	Left	Right	Right	Left

It is worth noting that describing these abilities as *unseen* in the literal sense is potentially misleading and can be considered incorrect because the fact that subjects perform above chance means they still *see* the stimuli (Dretske, 2000). This is why, throughout the thesis, **'unseen'** is used as the bracketing or epoché, which qualifies the term as referring to

the subjective state of the subject (Dapraz, 1999; Dennett, 1991; Lutz, 2003; Varela, et al., 1991). That is, objectively the stimuli are seen / perceived, but in terms of the subjective 1[st] person perspective they are 'unseen'. Bracketing moves the subjective report into the 3[rd] person domain. There is also the possibility of describing these abilities in terms of unconscious or subliminal processing, which is perfectly acceptable (see Dehaene, Changeux, Naccache, Sackur, & Sergent, 2006 for synopsis), but describing them as lacking in consciousness and reportedly 'unseen' is less controversial and captures the important elements of the task performance.

1.3.3 What do we define as blindsight?

What pattern of responses should we consider corresponds to TMS-induced blindsight? In classic lesion-based blindsight, where the suppression of conscious awareness is essentially permanent and absolute, subjects' conscious detection of stimuli will always approximate to a d' (eq 1.3) or Pr (eq 1.8) of 0, and their forced choice discrimination ability (which demonstrates blindsight) will dissociate and be greater than 0 (Cowey, 2010; Weiskrantz, 1986). TMS-induced blindsight, on the other hand, is a relative state revealed by the difference between active and control conditions. As such, the way in which blindsight manifests upon the measure of conscious detection is one of differential suppression between experimental conditions, with reportedly 'unseen' discrimination remaining above chance and relatively unaffected by TMS (Weiskrantz, 2001, 2009).

The possibility of reducing conscious detection to a d' or Pr of 0 is clearly preferable in terms of interpretation, as it allows greater confidence that the concurrent 'unseen' ability approaches a complete lack of awareness (Eriksen, 1960). However, such a conservative definition of blindsight is neither necessary nor practical: a dissociable effect of TMS on the two measures can be demonstrated independently of whether detection is fully abolished. Furthermore, simply reporting the stimuli as 'unseen' indicates that subjects were *less* consciously aware of the stimuli than when 'seen', highlighting the relative difference upon which the experiments are based. Finally, to my knowledge, none of the previous demonstrations of TMS-induced blindsight have reported complete elimination of conscious

detection (Boyer, et al., 2005; Christensen, et al., 2008; Jolij & Lamme, 2005; Ro, et al., 2004).

Demonstrations of blindsight in 'normal observers' are often based on masking paradigms which dissociate consciousness and performance (e.g. Lau & Passingham, 2006). However, whether or not this dissociation is complete and should be termed blindsight is controversial (for contrasting interpretations of data compare (Kolb & Braun, 1995; Lau & Passingham, 2006) to (Morgan, Mason, & Solomon, 1997; Robichaud & Stelmach, 2003)). Independently of whether or not these experiments failed or succeeded in demonstrating blindsight in 'normal observers', they all have a common approach; they do not seek to obtain complete suppressions of conscious detection type measures (Kolb & Braun, 1995; Lau & Passingham, 2006; Morgan, et al., 1997; Robichaud & Stelmach, 2003), but rather seek to demonstrate blindsight through a relative dissociation between subjective detection and accuracy performance under a task condition such as masking or rivalry.

Defining TMS-induced blindsight as a complete absence of conscious detection (i.e. $Pr \sim 0$ criteria) does not take account of baseline levels of performance. That is, if the levels of performance in the baseline condition were relatively high, then even large suppressions specific to conscious awareness might not be classed as blindsight. On the other hand, were baseline performance to be lower, much smaller effects would be classed as blindsight. Hence what is classed as blindsight might only depend upon the baseline condition. I therefore adopted the relative functional definition of blindsight as a statistically significant drop in the measure of conscious detection, while under the same experimental conditions, subjects' reportedly 'unseen' performance was unaffected or differentially affected, but above chance.

Chapter 2. Off-line experiments

The effects of continuous Theta Burst Stimulation on gating consciousness: a combined TMS, MEG and MRS study.

2.1 Overview

These experiments originally aimed to demonstrate blindsight type effects following a repetitive TMS (rTMS) protocol that is thought to reduce cortical excitability, in areas similar to those affected by lesions in classic blindsight. However, results indicated that rTMS had the *opposite* effect to that predicted: conscious detection increased (section 2.2). This counterintuitive finding therefore merited replication and further investigation. Subsequent experiments included the same behavioural and TMS paradigm in combination with magnetoencephalographic (MEG) recording and magnetic resonance spectroscopy (MRS) to investigate the electrophysiological and neurochemical basis of the observed effects.

This chapter begins with the report of the initial behavioural experiment (2A section 2.2) which is followed by a preliminary discussion of the effects on vision (2.3). Subsequent experiments to uncover the basis of these behavioural effects are then described (sections 2.4-6). Collectively, these comprised three experiments: a study of how rTMS influences the excitability of the visual cortex, (2B section 2.4), an MRS experiment to study how rTMS influences neurochemistry (2C section 2.5) and an MEG experiment (2D section 2.6), that includes a behavioural replication of the initial finding (2.6.1). These are then followed by a section-specific general discussion (2.7) that encompasses all of the off-line TMS experiments. Finally, I discuss the limitations of the behavioural paradigm (2.8), which informed modifications to the design in the subsequent on-line TMS experiments (Chapter 3). Additional analyses and relationships between measures are reported in appendix A1-3.

2.2 Behavioural experiment (2A)

2.2.1 Introduction

Blindsight has been one of the most informative conditions in recent investigations of consciousness (e.g. Block, 1995; Dehaene & Naccache, 2001; Lamme, 2001; Lau & Passingham, 2006). As previously described (1.2.2) patients with blindsight typically present with lesions to their early visual cortex and do not acknowledge conscious awareness of stimuli presented, yet can demonstrate residual perceptual abilities in forced choice tasks (Sanders, et al., 1974; Weiskrantz, 1986). Blindsight thus corresponds to a lack of conscious awareness that dissociates from residual perceptual abilities, highlighting the role of the affected brain regions in conscious processing.

Previous demonstrations of TMS-induced blindsight are based upon the application of single- or double-pulse TMS, which have been used to disrupt conscious processing through disruption of early visual areas at various times relative to stimulus onset (e.g. Boyer, et al., 2005). Continuous theta burst stimulation (cTBS) is an alternative 'off-line' repetitive TMS protocol that is thought to reduce cortical excitability and alter GABAergic function for approximately one hour (Franca, Koch, Mochizuki, Huang, & Rothwell, 2006; Huang, Edwards, Rounis, Bhatia, & Rothwell, 2005; Stagg et al., 2009). Following application of cTBS to motor areas the TMS intensity required to elicit a motor twitch for the same motor area, has been shown to increase (Huang, et al., 2005).

Here we sought to simulate an occipital lesion using cTBS within a behavioural paradigm capable of demonstrating blindsight-type effects. To this end we presented subjects with arrows and asked them in which direction the arrows were pointing and whether or not they were consciously aware of having seen the arrows (figure 2.2.1.1). From this was derived a measure of subjects' conscious detection of the arrows (PrC, section 1.3) and their forced choice discrimination ability, specifically, on trials where they reported not having seen the arrow (PcU). A blindsight-type effect, resulting from the cTBS-induced reduction in cortical excitability, was expected to manifest as a reduction in conscious detection while leaving reportedly 'unseen' discrimination ability above chance and relatively intact.

Using a related cTBS protocol over the dorsolateral prefrontal cortex, Rounis et al (2010) demonstrated the impairment of meta-cognitive confidence judgments (Rounis, Maniscalco, Rothwell, Passingham, & Lau, 2010). In addition, a suppression of similar capacities has been reported in a conference abstract when applying TBS to visual cortical areas (see Rahnev, Bahdo, Munneke, de Lange, & Lau, 2010).
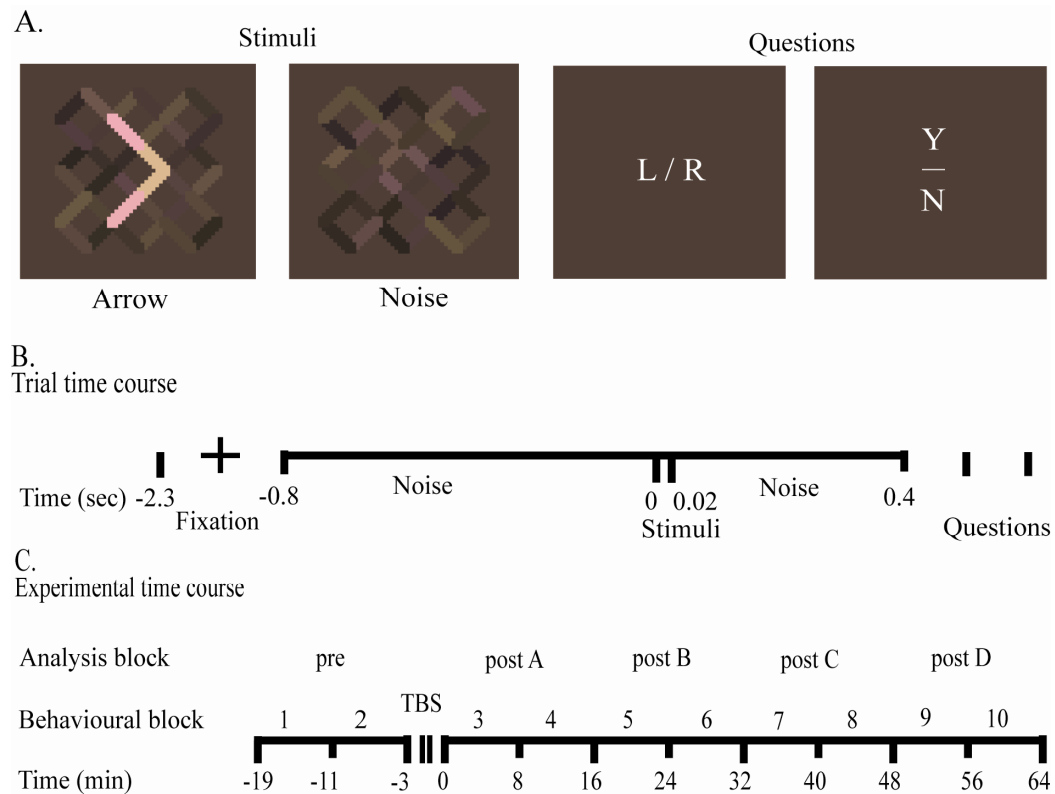
Figure 2.2.1.1. Experimental Design.

A. Example of arrow stimuli, noise (stimulus-absent) and the questions. The questions presented on every trial were 'Was the arrow pointing left or right?' denoted by 'L/R' and 'Did you see the arrow? Yes or No' denoted by 'Y/N'.

B. Time course of each trial. Fixation was followed by noise alternating at 50Hz with a stimulus frame (20ms) displayed at 800ms on half of the trials. Responses to questions followed after a further 400ms of noise and were not speeded. Questions commenced with the 'L/R?' decision.

C. Time course of the experiment. Behavioural and MEG acquisition blocks of eight minutes (Experiment 2A and 2D) were collapsed into 16-minute analysis blocks, to align with the acquisition of MRS (Experiment 2C) and phosphene threshold data

(Experiment 2B) acquisitions. Pre-TBS blocks were used to baseline the data (see section 2.2.2 for details). Active and control TMS were applied in separate sessions.

2.2.2 Methods

Sixteen neurologically healthy subjects participated in the first experiment, for which they were paid £10 per hour (5 female and 11 male, aged 22 to 35 mean 28). I was one of the subjects (see section 4.8.2). All provided informed consent and were screened for medical contraindications to TMS, including personal or family history of epilepsy (Maizey, In Press). This research was approved by the Ethics Committee at Cardiff University School of Psychology.

Visual stimuli were presented using a Cambridge Research Systems (CRS) Visage and Real Time Sequencer system on a Matlab platform, via a Mitsubishi Diamond Pro 2070sb monitor, refreshing at 100Hz, which was degaussed and regularly gamma-corrected. Infra-red eye tracking was carried out using a CRS 250Hz eye-tracker. Eye movements and pupil diameter were recorded. Two subjects were excluded from this analysis of eye-tracking data owing to failure of the eye-tracker to record reliable pupillometry data. Trials in which the subject blinked during the stimuli presentation, identified by a vertical shift in signal followed by a transitory loss of signal coincident with the stimuli presentation, were removed from the analysis. Across all subjects, this criterion resulted in the exclusion of 72 trials from a possible 25,600.

The arrow target stimulus was a 20ms increase in luminance amongst luminance noise (figure 2.2.1.1). Noise started 800ms prior to the target and continued for 400ms afterwards against a black background, alternating every 20ms within a range of 17.5 to 32.5 cd/m^2. Left and right arrows were presented singularly and in equal proportions, with equivalent luminance at the fixation point subtending the vertical meridian. The noise occupied 1.97° × 1.97° of the visual angle and the target arrow subtended 0.90° × 1.34°. On half the trials (stimulus-absent condition) a noise frame supplemented the target stimulus. On each trial, subjects were first asked in which direction the arrow was pointing (L/R?) and then whether or not they were consciously aware of having seen the arrow (Y/N?) (figure 2.2.1.1). From these questions three measures were derived: (a) their conscious detection of the arrow (PrC);

(b) their discrimination ability specifically when they reported not having seen the arrow (PcU); and (c) their response bias (BrC, see statistical methods below).

The luminance of the target was titrated in a separate session prior to experimental days, so that subjects were consciously detecting the target at Pr = 0.6. Following a period of familiarisation with the task and stimuli, this involved adjusting the luminance of the stimuli in 20-trial runs of 80-trial blocks and repeating for approximately 15 blocks. This produced a psychophysical function to which a linear or sigmoidal curve was fitted (depending on goodness of fit), which was then solved for the threshold value. A block consisting of 80 trials at the derived luminance value were then completed to confirm the threshold values, and small adjustments were made as required if performance exceeded a tolerance of ± 0.15 Pr units. If adjustments were made, blocks were repeated at the new luminance value. For experimental sessions, if the performance over baseline 'pre' TBS blocks (see figure 2.2.1.1.C) exceeded the tolerance, similar small adjustments were made to the luminance of the stimuli and the blocks repeated to maintain performance. If performance exceeded the criteria after 4 blocks then the experimental session was abandoned in order to avoid additional fatigue effects, and the session rescheduled.

Blocks consisted of 80 trials in a randomised order. Ten blocks were undertaken in each session, including two blocks prior to the application of the TBS to provide a baseline, and eight blocks following. Eight minutes were allocated for each block; as blocks took approximately 5 minutes to complete, a short rest period was included. The first block started 1 minute 20 seconds after TBS had ended, such that data was collected up to 66 minutes after the TBS had started. After every trial, subjects were asked two questions which were explained to them prior to testing. The two questions were: 'Was the arrow pointing Left or Right?' (the discrimination question for which they were prompted 'L/R'), and 'Did you consciously see the arrow?' (the detection question, for which they were prompted 'Y/N'). Questions were posed in that order to reduce the impact the later question might have upon the former. Responses were recorded via key press on a standard keyboard (response mapping is the same as in the participant inscriptions for the on-line studies: appendix A4).

A round coil (Magstim High Power 90mm Coil and Magstim biphasic Ripid2 stimulator) was used, delivering 600 pulses over 40 seconds (following Huang, et al., 2005) at

an intensity of 80% of the subjects' resting motor threshold[3]; this corresponded to a mean TMS intensity of 41.7% stimulator output (± 3.6% SD). Motor threshold was established using the observation of movement method, and was calculated as the average across hemispheres (Stokes et al., 2005; Varnava, Stokes, & Chambers, 2011). Positioning was achieved using a miniBIRD system (Ascension Technology Corp) in conjunction with MRIcro, MRIreg software (Chris Rorden) and structural MRI scans (Rorden & Brett, 2000). TMS was targeted at the striate cortex (V1); that is, the closest scalp coordinate to the mid-hemispheric termination of the left and right calcarine sulci, calculated with Surface Extraction Tool software (Version 2, Chris Chambers and Mark Stokes). However, because of the relatively diffuse effect of a round coil and consistent with the work of Thielscher et al (2010), the direct effect of this intervention in terms of magnetic induction can only definitively be said to be across occipital cortices (and not selectively V1). In the active cTBS condition the rim of the coil was positioned 2cm below the closest scalp coordinate to V1, which centred the rim over V1, with the handle pointing upward and side 'B' facing away from the subject. The position was closely matched in sham (control) conditions except that the coil was oriented horizontally and a 9mm spacer was placed between scalp and coil to replicate the contact artefact. The two conditions (active and sham) occurred on different days, the order of which was counter balanced between subjects. This protocol was used in all subsequent applications of TBS.

Statistical methods.

Behavioural blocks of 8 minutes were concatenated into blocks of 16 minutes to improve the power / reliability of individual data points and to align the behavioural experiment with subsequent MRS acquisitions (see Experiment 2C, section 2.5).

The behavioural measures were obtained via the application of non-parametric SDT to the subjects' responses (Corwin, 1994; Green & Swets, 1966; Macmillan & Creelman, 1990). Non-parametric SDT was used because of imbalances in stimulus conditions and response profiles, which violated the assumptions of classic SDT (Corwin, 1994). The measure of subjects' conscious detection (PrC) was taken as the veridical acknowledged sensitivity to the

---

[3] The motor thresholds used were collected by members of Cardiff University TMS-group according to the methods set out in (Stokes et al., 2005; Varnava, Stokes, & Chambers, 2011) and the Cardiff University TMS lab manual. Contributors to the motor threshold database therefore may include, myself, Alice Varnava, Chris Chambers, Frederick Verbruggen, Leah Maizey, Mark Stokes, Martynas Dervinis, Michail Kozlov and Rachel Adams.

arrow's presence. The application of SDT was therefore as follows: acknowledged awareness in the presence of an arrow was a 'hit'; acknowledged awareness when no arrow was presented was a 'false alarm'; denial of awareness when no arrow was presented was a 'correct rejection'; and denial of awareness when there was an arrow present was a 'miss'. The same allocations of SDT classes were used in the measure of response bias in conscious detection (BrC). The measure of 'unseen' discrimination ability was derived from the subject's left/right proportion correct performance, only when subjects reported not having seen the arrow (PcU). Bias in left/right discrimination was also calculated (for left bias) but this measure was only used to ensure that subjects were performing the task as instructed.

Effects of cTBS were assessed using repeated measures ANOVAs for each measure in the analysis. The dependent variable therefore consisted of the measure of interest for each post TBS block baselined to pre TBS levels. The factors for the ANOVA were TMS site (2 levels, active vs. control) and time since the TMS (4 levels, 2-18, 18-34, 34-50, 50-66 minutes). The experimental hypotheses are represented by the factor of TMS site. Of secondary concern was the development of these effects over the course of the experiment (TMS site × TMS time) and the purely time-dependent effects. For this reason the effects of TMS 'site' are reported in the first instance. Greenhouse-Geisser corrected p values are reported. The analyses were applied to measures separately as they informed independent hypotheses. All subsequent analyses of other dependent variables conform to this structure, unless otherwise stated.

'Unseen' performance was compared to chance using single sample t-tests applied to each block and appropriate Holm-Bonferroni corrections applied for the multiple blocks / comparisons (Holm, 1979).

Outliers were identified and excluded on the basis of Chauvenet's criterion (Taylor, 1997), applied to the dependent variable capable of representing a TMS dependent effect. Specifically, a vector was constructed that included active cTBS minus control (sham) cTBS, averaged across post-TBS blocks and further subtracted from the pre-TBS baseline. This provided a single measure per subject representing the distribution in effect sizes following cTBS. If the probability of any subject's data within this group vector, times the number of samples in that group, was less than 0.5 then the subject's data was excluded from the analysis of that measure. This resulted in the exclusion of a single subject's data from the analyses involving PrC, BrC and PcU.

2.2.3 Results

To reiterate, the experimental hypothesis was that conscious detection should be suppressed by the application of the cTBS to the occipital lobe; in other words, there should be a reduction in PrC following active TMS compared to control TMS. If concurrent 'unseen' abilities remain above chance but unaffected by disruption, this would fit the expected profile of a blindsight type effect. The central measures for this experiment were therefore conscious detection and 'unseen' discrimination. Both dependent variables were initially baselined to pre-TBS levels within session, thus taking into account day-to-day differences in performance.

Contrary to the hypothesis, conscious detection of stimuli *increased* following application of cTBS relative to sham (site effect: $F_{(1,14)}$=14.015, p=0.002, figure 2.2.3.1). The ability to discriminate the arrow's direction on trials when the subject reported not having seen the arrow, remained above chance over all experimental blocks (all $t_{(15)}$>7.170, all p<0.001, relative to a initial Holm-Bonferroni corrected alpha level of 0.005). This 'unseen' discrimination ability appeared to be unaffected by the application of the cTBS (site effect: $F_{(1,14)}$ =0.041, p=0.843, figure 2.2.3.2). The effect, being specific to the measure of conscious detection, in contrast to above chance 'unseen' discrimination ability, is of the same class as TMS-induced blindsight (e.g. Ro, et al., 2004) but is in the opposite direction: subjects were more aware of the arrow stimuli following a neuronally suppressive intervention. This was the central finding of Experiment 2A.
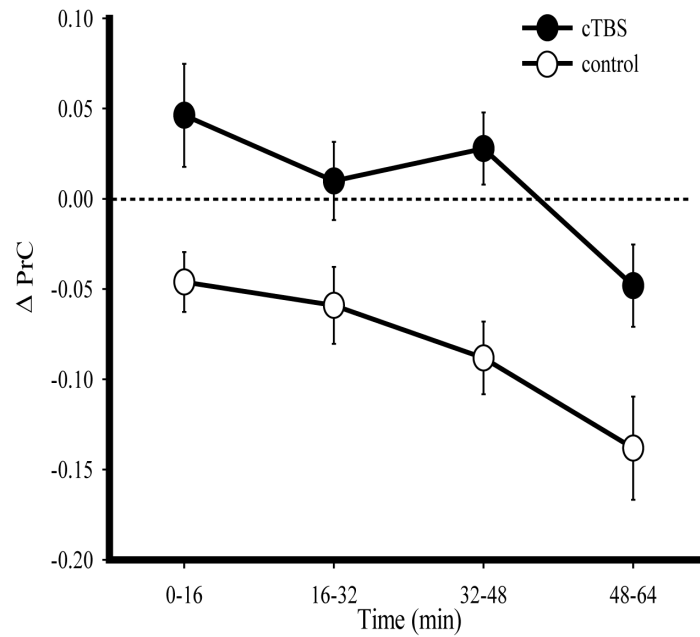
Figure 2.2.3.1. Effects of cortical stimulation on conscious detection (PrC) in Experiment 2A. The dependent variable is the change in PrC from the pre-TBS baseline following cTBS and control (sham) stimulation, illustrating the increase in PrC following cTBS relative to control. Error bars are within-subject standard error (Loftus & Masson, 1994; Masson & Loftus, 2003). Time corresponds to the 4 analysis blocks collected after the TBS was applied. All subsequent line plots in this chapter conform to this structure.
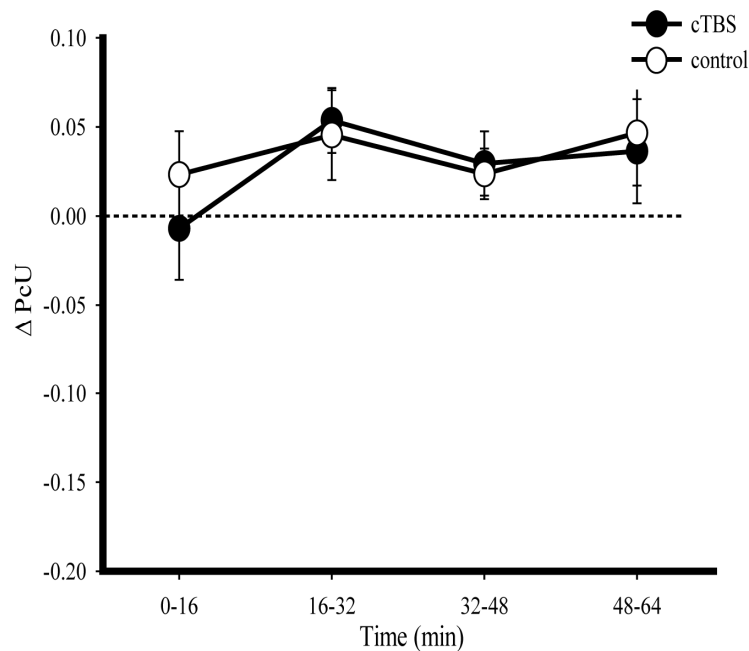
Figure 2.2.3.2. Change in proportion correct in 'unseen' discrimination (PcU) in Experiment 2A from pre-TBS baseline following cTBS and control (sham) stimulation.

The critical effect in terms of the hypothesis was the main effect of TMS site. However it is possible that this effect could evolve over the course of the experiment as illustrated by a potential site × time interaction. Since the effects of cTBS are thought to last for approximately the full course of the experiment, no predictions were cast in terms of the evolution of TMS over time and no corresponding interactions were observed on the main measures (PrC site × time interaction: $F_{(3,42)} = 0.369$, p=0.755; PcU site × time interaction: $F_{(3,42)} = 0.352$, p=0.729). However, following both control and active stimulation, a significant decline in conscious detection performance throughout the course of the experiment was observed (PrC time effect: $F_{(3,42)} = 6.511$, p=0.002). No time-dependent change was observed for 'unseen' discrimination (PcU time effect: $F_{(3,42)} = 1.738$, p=0.195). The drop in detection over the course of the experiment suggests fatigue, as did the discursive reports made by the subjects following the experiment. Consistent with this interpretation, which does not inform the main experimental hypothesis, are changes observed upon other measures associated with fatigue detailed in section 2.6.2 and appendix A2.

Asking subjects whether or not they were subjectively aware of stimuli, together with knowledge of stimulus presence/absence, enabled us to look at another secondary aspect of the subject's experience of the task, in addition to the contrast between conscious detection and perception lacking in awareness. That is, their response criteria or bias, quantified here as BrC (Corwin, 1994). Interpretations of what changes in bias refer to are diverse and controversial (cf. Campion, et al., 1983; Weiskrantz, 2001, discussed further in section 4.3 to 4.4). However, it seems to have been generally accepted that removing bias from measures of conscious detection makes changes in measures simpler to interpret, as was implemented here for the PrC measure.

Subjects reported seeing the stimuli more often irrespective of stimulus presence/absence following active TMS relative to sham (TMS site effect on the measure of bias in conscious detection: $F_{(1,14)} = 6.772$, p=0.021, see figure 2.2.3.3). There was also a general tendency to respond positively less often towards the end of the experiment (time effect: $F_{(3,42)} = 5.585$, p=0.015), again suggestive of fatigue. This change did not interact with the TMS dependent effect (site × time interaction: $F_{(3,42)} = 1.386$, p=0.267). One subject,

expressing a shift toward reporting 'no' more often following the application of the TMS, was excluded as an outlier from this analysis according to Chauvenet's criterion.
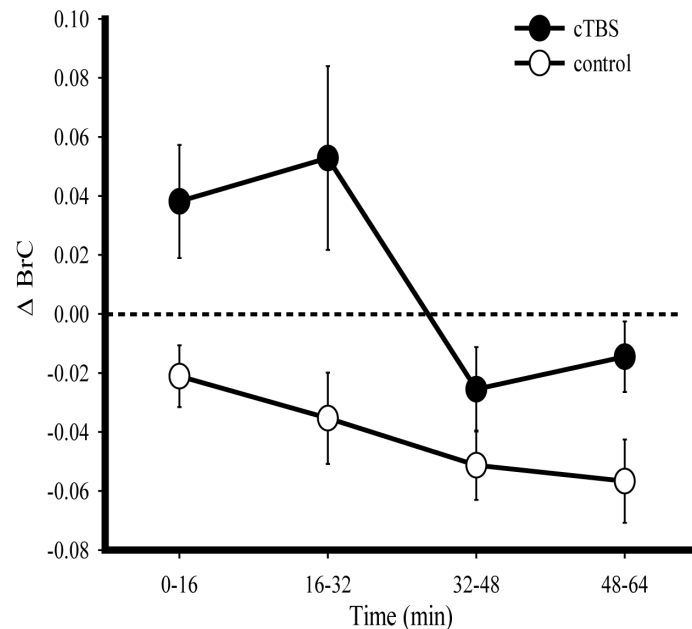
Figure 2.2.3.3. Change in response criteria / bias in conscious report (BrC), from pre-TBS baseline following cTBS and control (sham) stimulation.

2.3 Preliminary discussion and introduction to subsequent experiments

Contrary to the initial hypothesis, conscious detection of the stimuli increased following the application of cTBS relative to the sham control condition. 'Unseen' discrimination remained above chance but appeared to be unaffected by the TMS. The increase in conscious detection as a result of suppression was counterintuitive and thus warranted replication and further exploration (sections 2.4-6).

A common assumption is that increased activity in sensory cortical areas indicates increased conscious representations (e.g. Polonsky, Blake, Braun, & Heeger, 2000). However, it is possible that of equal importance is the negative aspect, that is, active

suppression of sensory representations relative to others is crucial for consciousness. That is, the quieting down or active inhibition of neuronal processes may be conducive to optimal detection. We hypothesised that this negative aspect of selection in the determination or gating of conscious percepts might explain the increase in detection, whilst being consistent with cTBS being a suppressive protocol. This is henceforth denoted as the *gating-by-inhibition* hypothesis. This hypothesis therefore became the focus of the subsequent investigation, in which we also replicated the original effect and explored alternative explanations, such as increased arousal and increased cortical excitability.

We applied the behavioural and TMS paradigms in the context of magnetoencephalographic (MEG) recording (Experiment 2D). Of particular interest were changes in oscillatory activity in the alpha band (6-12Hz). The alpha band has been associated with the active suppression of task irrelevant and superfluous activity (Foxe & Snyder, 2011; Jensen & Mazaheri, 2010; Varela, Toro, John, & Schwartz, 1981). The principal induced oscillatory response in the alpha band is the Event Related Desynchronisation (ERD; Pfurtscheller, Neuper, & Mohl, 1994) which can be seen as representative of a gating mechanism (Pfurtscheller, 1992): it is a shift from a synchronised state, where suppression is imposed by the alpha oscillation, to a less synchronised state following the presentation of stimuli. The gating by inhibition hypothesis therefore predicts that the ERD might be facilitated following the application of cTBS (section 2.6.4).

Active suppression almost certainly involves the principle inhibitory neurotransmitter Gamma-Aminobutyric Acid (GABA). Magnetic resonance spectroscopy (MRS) can quantify *in vivo* GABA concentration and has previously been used to demonstrate that levels of GABA increase following the application of cTBS to motor areas (Stagg, et al., 2009). Replication of this functional change in GABA over occipital areas would be consistent both with cTBS having a neuronally suppressive effect and with the gating by inhibition hypothesis for the increase in conscious detection (Experiment 2C section 2.5).

The ERD and GABA measures were therefore designed to provide evidence which informed the gating by inhibition hypothesis. Three further independent hypotheses were probed with additional measures. First, replication in the MEG allowed us to probe the possibility that cTBS might be effective in increasing neuronal responses to stimuli, rather than in suppressing neuronal noise (the *increased response* hypothesis). Later evoked responses were collected and have previously been associated with the processing of

conscious signals (Del Cul, Baillet, & Dehaene, 2007; Koivisto & Revonsuo, 2003; Melloni et al., 2007; Niedeggen, Wichmann, & Stoerig, 2001; Quiroga, Mukamel, Isham, Malach, & Fried, 2008).Therefore cTBS increasing late evoked responses might be taken as evidence for increased excitatory signals and responses, contrary to the gating hypothesis (detailed in section 2.6.3).

Another alternative hypothesis for the mechanism by which cTBS might be affecting conscious detection was the possibility that cTBS may increase general cortical excitability and possibly add noise to the affected regions. This could lead to an increased probability of any particular representation crossing a threshold for detection, as it would be boosted by the increased background activity. This would be a '*stochastic resonance*' type effect and has previously been suggested as the mechanism by which on-line TMS can enhance visual detection (Schwarzkopf, Silvanto, & Rees, 2011). Such an interpretation of increasing the likelihood of crossing a threshold for detection, irrespective of the stimulus condition, may also be consistent with the finding that response criteria was more liberal following the application of cTBS. In order to quantify levels of cortical excitability a separate experiment probed the effect of cTBS on the intensity of subsequent occipital TMS in eliciting visual perceptions known as phosphenes. The gating-by-inhibition hypothesis would predict an increased phosphene threshold (PT) following cTBS (Franca, et al., 2006) and a stochastic resonance or increased excitability based explanation might predict a reduction (Experiment 2B, section 2.4).

The final alternative explanation for the mechanism by which cTBS was affecting conscious detection was the possibility of it increasing levels of arousal. According to the *arousal* hypothesis, this could arise owing to the subject's expectations of cTBS and the generally ineffective blinding of active vs. sham conditions in TMS designs. Such a 'reactive' process (Ericsson, 2003, see section 4.3 for clarification) might operate independently of the direct neuronal effects of the TMS. We quantified arousal in both the original and replication of the behavioural effects through the concurrent measurement of pupil diameter commonly linked with autonomic arousal (Bradshaw, 1967). If the enhancement of conscious vision was caused by a general arousal effect then cTBS should increase pupil diameter relative to sham (section 2.6.2).

To anticipate, the ERD and GABA concentration were each potentiated following the application of the cTBS, whereas the three other hypotheses received no support.

Collectively, these experiments indicated that selection via inhibition is an important element in determining the conscious fate of representations.

2.4 Phosphene threshold experiment (2B)

2.4.1 Introduction

Phosphenes are percepts generated by activation of neurons in the visual cortex. To test the suppressive effect of cTBS at the neuronal level we used the stimulator intensity required to elicit phosphenes, known as phosphene threshold (PT), as a measure of intrinsic cortical excitability (Franca, et al., 2006). Although cTBS has been previously shown suppressive effects (Franca, et al., 2006; Huang, et al., 2005), it is possible that the cTBS applied here had the opposite effect; reversals of the expected effect (from suppression to facilitation) have been demonstrated when the levels of TMS used have been lowered (Abrahamyan, Clifford, Arabzadeh, & Harris, 2011). Here, the absolute stimulator output used was (42.42% of maximum ±7.35SD), which is less than those of previously reported suppressions by cTBS (45.68% of maximum ±10.88SD, Franca, et al., 2006).

2.4.2 Methods

Twelve subjects participated in Experiment 2B (aged 20 to 41, mean age=24, 7 females), three of whom also participated in behavioural experiments. The apparatus for this experiment is as described in the preceding section (2.2.2) unless otherwise stated.

The intensity of a single TMS pulse required to elicit a phosphene (PT) depends on the levels of cortical excitability within that region. The method used here resembled that of Franca et al., (2006). First, we assessed subjects' susceptibility to phosphenes within safety limits (160% of motor threshold, MT). The coil was initially positioned using the miniBird

system and tripod as in the main experiment, with single pulses applied at 120% of MT. If stimulation did not elicit phosphenes that the subject reported as being 'reasonably clear' the coil was moved until it did so, while minimising the distance to the original coil position. This location was recorded using a Brainsight system (Rogue Research Inc.) based on the subjects' anatomical MRI scans (collected on a 3T whole-body General Electric scanner). An approximate PT was obtained using an up-down staircase method, applying single pulses approximately every 5 seconds, starting at 50% of maximum output and adjusting TMS intensity in steps of 5%, then 2%, then 1%, so that subjects verbally reported seeing 5 phosphenes from 10 pulses. This level was then used as the basis of the more accurate thresholds where the number of reported phosphenes arising out of 10 pulses was recorded at -10, -5, 0, +5, +10 and +15% of the estimated PT. The orders of these sets of 10 pulses were randomised and the full range of intensities was repeated three times in separate runs separated by short breaks. The coil was repositioned at the start of each run. Averaging across runs yielded a function representing the number of phosphenes out of 10 over a range of intensities, to which a curve was fitted (sigmoid or linear depending on goodness of fit). Solving this curve for 5/10 phosphenes thus provided the PT.

These sets of three runs comprised a block of data, collected over 16 minutes, corresponding to the timing of the MRS acquisition (section 2.5.2) where one block was taken before the cTBS and three were acquired afterward. Three rather than four post-cTBS blocks were completed because, even with three post cTBS blocks, the quantity of TMS applied approached the limit permitted by the local ethics committee (based on Wassermann, 1998 and Maizey et al, In Press). Again, sham and active cTBS were applied on separate days, the order of which was counterbalanced across subjects. The mean intensity at which the cTBS was applied for this group was 42.4% ± 7.4 SD.

2.4.3 Results

The TMS intensity required to elicit phosphenes increased significantly following the application of cTBS relative to controls (site effect: $F_{(1,11)}$=5.395, p=0.040, figure. 2.4.3.1). No time dependent effects were observed (site × time interaction: $F_{(2,22)}$ =1.59, p=0.23, time effect: $F_{(2,22)}$ =1.59, p=0.23).
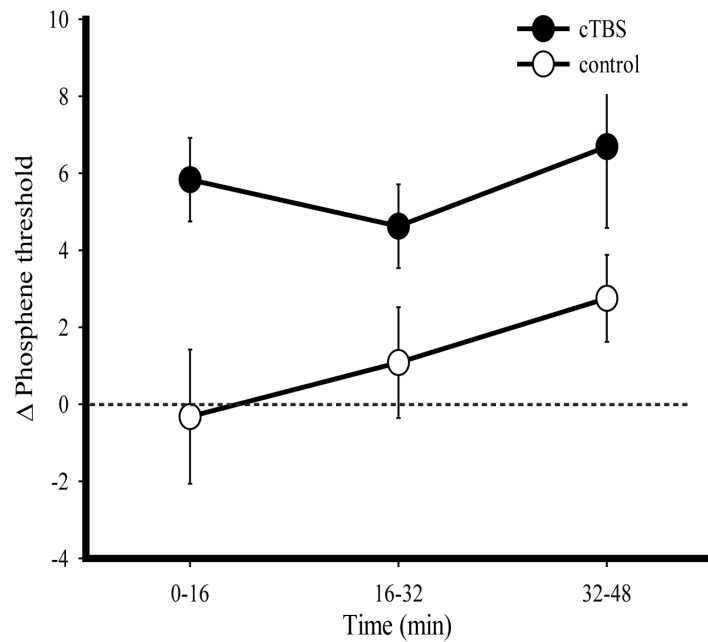
Figure 2.4.3.1. Change in phosphene threshold from pre-TBS levels, in Experiment 2B, following occipital cTBS and control stimulation.

These results are consistent with the expected inhibitory effect of cTBS, replicating previous observations (Franca, et al., 2006; Huang, et al., 2005). An explanation of the cTBS-induced enhancement of awareness in Experiment 2A according to increased excitability or stochastic resonance is therefore inconsistent with this data. The relationship between increased phosphene threshold and decreased threshold for conscious detection is discussed in section 2.7.

2.5 Magnetic Resonance Spectroscopy experiment (2C)

2.5.1 Introduction

Critical to the balance between neuronal excitation and suppression is the inhibitory neurotransmitter GABA. For example, many drugs used in the control of epilepsy target the GABAergic system, and its dysfunction appears to be the cause of some forms of epilepsy (Baulac et al., 2001). At the molecular / synaptic level, GABA increases chlorine flux across the neuronal membrane, hyperpolarising cells, making them less susceptible to excitatory input.

The GABA molecule ($H_3N^+CH_2CH_2CH_2COO^-$) has a magnetic resonance signature that can be measured using magnetic resonance spectroscopy (MRS). The pairs of carbon bound hydrogens occupy an almost (see methods) unique electromagnetic environment. The relaxation of their alignment with a strong magnetic field, following a radio frequency pulse, produces a response in the magnetic response spectra at 1.9, 2.3 and 3 ppm (parts per million of the proton frequency). The magnitude of the response, following editing (see methods), is proportional to levels of GABA concentration (see Puts & Edden, 2012).

This technique has revealed an increase in GABA concentration following cTBS of the motor cortex (Stagg, et al., 2009), raising the question of whether such effects are reproducible in non-motor cortical areas. Furthermore, the MRS applied to functional changes in GABA is a relatively new and unconfirmed approach, the implementation of which can differ widely between laboratories. For example, Stagg et al., calibrated their quantification of GABA according to *in situ* N-acetylaspartate levels, whereas here, as in many previous MRS studies (Puts & Edden, 2012), GABA concentration was normalised relative to water. Therefore, the reliability of MRS in terms of reproducibility of effects is central to the development and trustworthiness of the technique as a whole.

The measurement of GABA offers the opportunity to assess the brain's governance of the balance between excitation and inhibition. In addition to testing the replicability and generalizability of the findings by Stagg et al., the current experiment sought to contrast two competing hypotheses as to the effect cTBS might have. If the application of cTBS increases excitability, with the increase in conscious detection possibly operating through a stochastic resonance mechanism, then an overall decrease in occipital GABA concentration might be expected. In contrast, if the same TMS protocol that increased conscious detection also leads to an increase in GABA concentration, then the results would favour an explanation in terms of a potentiated inhibitory gating system.

2.5.2 Methods

Data was collected from 18 subjects (aged 21 to 40, mean 26, 7 female), of whom 12 participated in the behavioural experiment (2A), and 10 in the MEG experiment (2D).

MRS data were acquired on a GE 3 Tesla MRI scanner over two separate sessions corresponding to cTBS and sham control conditions, which were counterbalanced across subjects. The mean TMS intensity used was $40.8\% \pm 5.0$ SD.

Four MRS acquisitions were undertaken per session, the first prior to cTBS to provide a within-subject, within-session baseline. The three following cTBS ran consecutively, commencing 2 minutes 40 seconds after the start of the cTBS. Before the pre-cTBS MRS acquisition, subjects' T1-weighted anatomical scan (1 mm^3 isotropic resolution) were collected. This allowed positioning of the MRS voxel, which was a 3cm × 3cm × 3cm volume centred over V1 defined by the calcarine sulcus, as with TMS targeting. The voxel was then moved toward the centre of the head in order to avoid non-brain tissue, and dorsally, so that the lower edge followed the cortical surface and did not enter the cerebellum (figure 2.5.2.1). This voxel position was recorded relative to anatomical landmarks in three dimensions using a screen shot and repeated for subsequent acquisitions. The cTBS was performed in the MRI control room and the subject was then immediately transferred to the scanner. The first MRS acquisition commenced 3 minutes 40 seconds after the start of the TBS and was preceded by localisation and calibration scans (1 minute), allowing repositioning of the voxel. Each MRS

acquisition lasted 16 minutes and comprised a MEGA-PRESS (Mescher, Merkle, Kirsch, Garwood, & Gruetter, 1998) spectral sequence consisting of 512 transients (TE = 68ms, TR = 1800) and 8 unsuppressed water transients (PRESS). During the MEGA-PRESS sequence Gaussian GABA editing pulses (alternating at 1.9 and 7.5 ppm) were applied to the standard magnetic resonance spectra in order to isolate the spine coupling of the GABA molecule from the response due to glutamate, glutamine, creatine and N-acetyl aspartate, which lie at approximately the same position on the chemical shift spectra (Puts & Edden, 2012). During the MRS acquisitions all subjects watched the same film and did not perform any behavioural task.

GABA concentration was quantified using the Gannet analysis tool kit (Puts & Edden, 2012) applied to the MEGA-PRESS edited spectra (figure 2.5.2.2). This involved fitting a Gaussian curve to the GABA peak at 3ppm with a liner baseline component, where the area under the curve, relative to that of water, allowed quantification of GABA concentration in institutional units (see Mescher, et al., 1998).

Sessions containing clearly low quality data were repeated (2 sessions, owing to subject movement during acquisition). Fit quality was improved by manual phasing of the spectra where the experimenter was blind to the TMS condition. Because I was aware in some cases of which spectra corresponded to which condition, and because the manual phasing technique is developed with practice, it was completed by Dr C J Evans (MRI lab manager, CUBRIC).
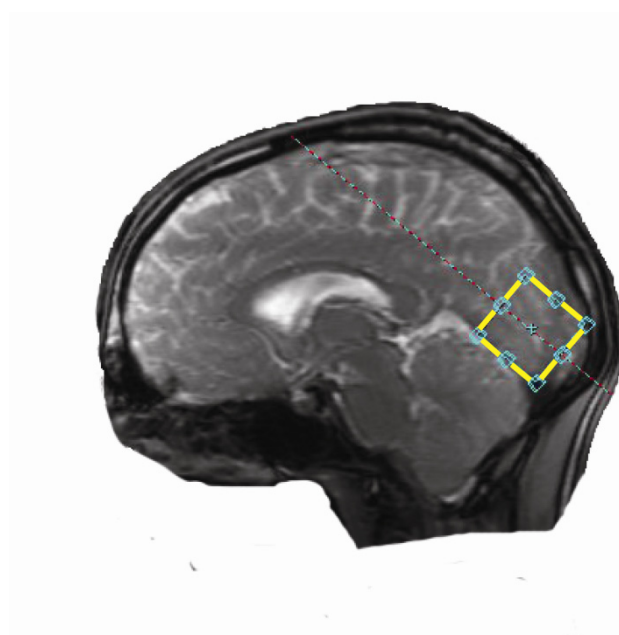
Figure 2.5.2.1. Illustration of the typical MRS voxel placement in Experiment 2C, as shown in the sagittal view of one participant.
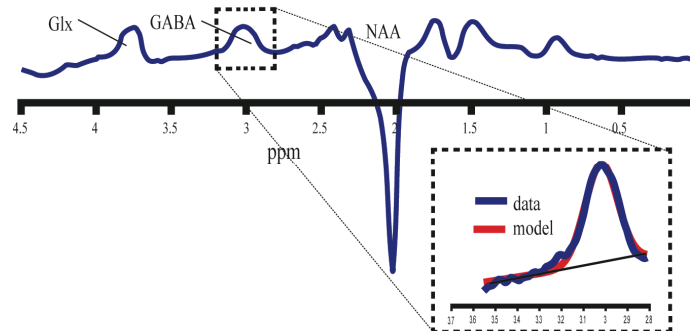
Figure 2.5.2.2. Illustration of model fitting applied to MEGA-PRESS edited spectra that allowed for quantification of GABA concentration. Units are parts per million (ppm) of proton frequency. Glx is the combined glutamate and glutamine peak. NAA is the peak caused by N-acetyl aspartate.

2.5.3 Results

Baseline-corrected GABA concentration in occipital cortex increased significantly following the application of occipital cTBS relative to sham cTBS (site effect: $F_{(1,16)}$=5.347, p= 0.034, figure 2.5.3.1). No time-dependent effects were observed (site × time interaction: $F_{(2,22)}$=0.416, p=0.603, time effect: $F_{(2,32)}$=0.347, p=0.656). One outlier was rejected from this data set on the basis of Chauvenet's criterion.
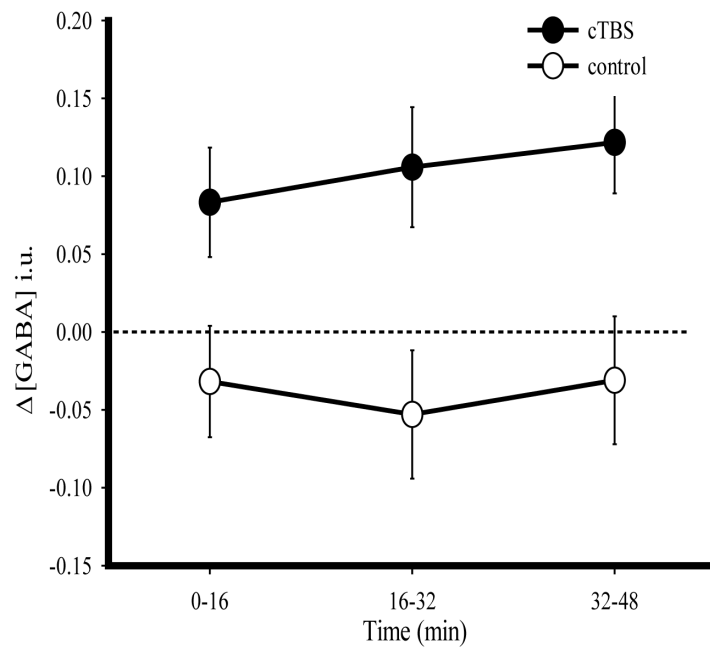
Figure 2.5.3.1. Change in GABA concentration in Experiment 2C. The ordinate indicates change in GABA in institutional units (i.u) relative to the pre-TBS baseline, plotted according to the TMS condition (occipital cTBS vs. Sham control) and time after stimulation (mins).

These findings replicate the increase in GABA concentration in motor cortex following cTBS of M1 (Stagg, et al., 2009) and, moreover, are consistent with the hypothesis that an increase in inhibitory processes may underlie the enhancement of conscious detection following occipital cTBS. As with Experiment 2B, the results run counter to the hypothesis that cTBS increases cortical excitability.

2.6 Magnetoencephalography experiment (2D)

This section includes four sub-sections that report the outcome of different questions addressed in Experiment 2D. The first section (2.6.1) describes the replication of the behavioural effects of cTBS, carried out inside the MEG scanner. The second section (2.6.2) analyses the eye tracking data with a view to probing the arousal hypothesis. The third and fourth sections concern the MEG data specifically, with section 2.6.3 addressing the evoked responses to stimuli that reflect of the increased signal hypothesis, section 2.6.4 reporting the induced oscillatory response in the alpha band associated with gating.

2.6.1 Behavioural replication

2.6.1.1 Methodological adaptations

Eighteen subjects (aged 20 to 32; mean age 26; 7 female) participated in Experiment 2D, four of whom had participated in the initial cTBS behavioural experiment (2A).

Subjects were thresholded to a level of $Pr = 0.5$ (as opposed to 0.6 in Experiment 2A) in order to optimise the sensitivity to detect both increases and decreases in detection ability. The ratio of stimuli-present to stimulus-absent was changed from 50:50 to 60:40 in favour of stimulus-present trials to optimise the power for the analysis of the stimulus-evoked electrophysiological response.

Responses were collected via a LUMItouch^{TM} response pad. Owing to the need to change rooms (TMS conducted outside the MEG magnetically sealed room) and the delay caused by the requirement to localise the subject's head before recording MEG data, 3 minutes were allocated between the termination of cTBS and commencement of the behavioural task. The mean TMS intensity was 42.1% stimulator output (± 4.6 SD).

A potential criticism of Experiment 2A (section 2.2) is that subjects were clearly aware of the difference between sham and cTBS conditions. The increase in conscious

detection might therefore have arisen due to a reactive response (Ericsson, 2003) by the subjects to the active TMS (see section 4.3 for further discussion). One alternative to sham TMS is the selection of an appropriate control site to replicate the auditory and tactile artefacts of stimulation. However, owing to the geometry of the round coil and the relatively broad distribution of the induced electric field in the brain, this approach was discounted in the current experiments: wherever the coil was placed, this relative lack of focality could be sufficient to disruption perceptual or cognitive processes.

Instead, we implemented intermittent theta burst stimulation (iTBS) as an alternative active control. During iTBS the same number of pulses is applied at the same intensity as cTBS but with a temporal profile that includes intermittent (8 second) gaps between continuous bursts of 2 second. In the motor cortex the physiological effect of iTBS opposes that of cTBS, reliably increasing cortical excitability (Huang, et al., 2005). For reasons that are not clear, this potentiating effect does not appear to be reproducible in the occipital cortex (Franca, et al., 2006). Nevertheless, evidence that iTBS produces either opposite or null effects on cortical excitability makes it an ideal occipital control condition to achieve effective participant blinding. The order of the different TBS conditions in Experiment 2D was thus counterbalanced across sessions between cTBS, iTBS and sham.

2.6.1.2 Behavioural results

The first question posed in the analysis of the behavioural results was whether iTBS is an appropriate control condition. As noted above, iTBS was implemented to address the concern that the effect of cTBS in Experiment 2A might be have resulted from subjects' expectations and/or non-specific effects of the active cTBS protocol. The central effect of interest, around which other analyses were based, was upon conscious detection. The question of whether iTBS was an appropriate baseline was therefore tested by comparing iTBS with sham performance on the measure of conscious detection. There was no evidence of a reliable difference between the effects of iTBS and those of sham on the critical central measure of conscious detection (TMS site effect iTBS vs. sham: $F_{(1,17)}=0.24$, p=0.63, site $\times$ time interaction: $F_{(3,51)}=1.24$, p=0.31). Because the application of iTBS did not discernibly differ from sham in terms of the central measure of interest, and to reduce the number of statistical

comparisons, all subsequent analyses in Experiment 2D collapsed the iTBS and sham conditions to yield a single control condition (also see section 4.2 and appendix A3).

The increase in conscious detection following cTBS relative to the control condition was replicated (site effect: $F_{(1,17)}$ =5.282, p=0.035, figure 2.6.1.2.1). Again 'unseen' discrimination ability appeared to be unaffected by cTBS (site effect: $F_{(1,17)}$=0.458, p=0.508, figure 2.6.1.2.2) while, crucially, remaining above chance (all $t_{(17)}$>6.747, all p<0.001).
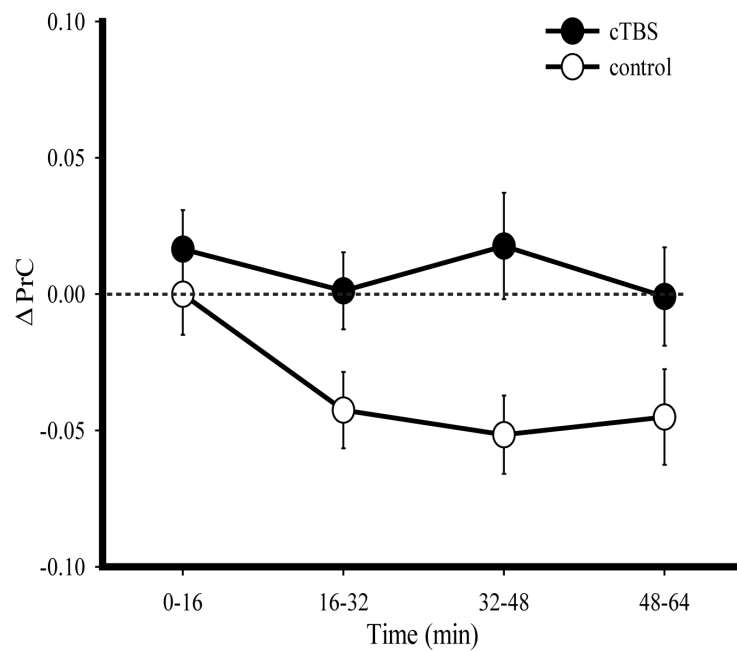
Figure 2.6.1.2.1 Change in the measure of conscious detection (PrC) from pre-TBS baseline following cTBS and control stimulation, in Experiment 2D.
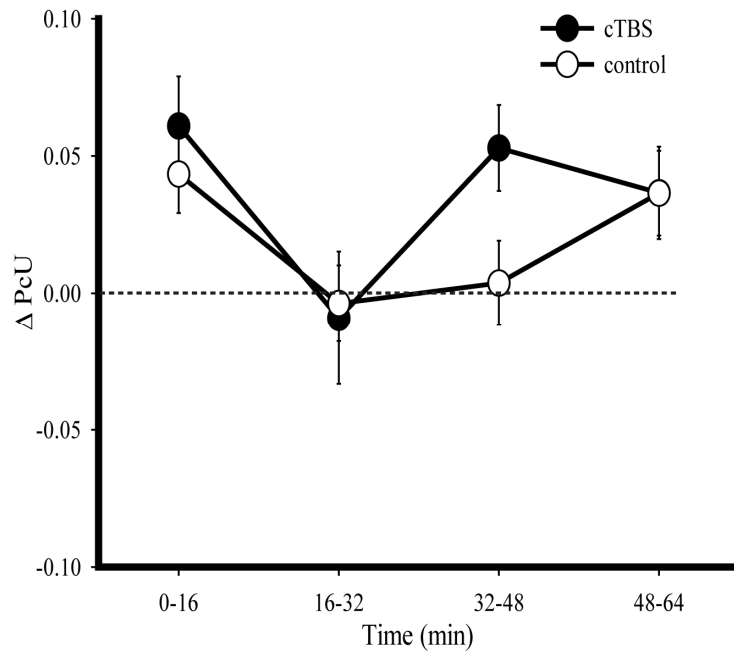
Figure 2.6.1.2.2 Change in 'unseen' discrimination measure (PcU) from pre-TBS baseline following cTBS and control stimulation, in Experiment 2D.

No interaction involving a change in conscious detection over the course of the experiment was observed (PrC site × time interaction: $F_{(3,51)}$=1.318, p=0.281), and, although a general decline in detection was observed throughout the experiment (as in Experiment 2A; see figure 2.2.3.1), this effect was not significant (PrC time effect: $F_{(3,51)}$ =1.557, p=0.221).

Changes in the measures of 'unseen' discrimination over the course of the experiment were observed (PcU time effect $F_{(3,51)}$=5.787, p=0.003). This effect did not appear to depend upon the presence of active TMS (PcU site × time interaction $F_{(3,51)}$=1.41, p=0.26) so does not relate to the experimental hypothesis. Given the 'u' shaped distribution of the change in PcU (see figure 2.6.1.2.2) it is possible to speculate that this effect might correspond to a combination of two effects: an initial fatigue effect, which the pupilometry data suggests had the greatest effect over the first two blocks (see figures 2.6.2.3.1 and 2.6.2.3.2), followed by a practice-driven improvement in performance towards the end of the sessions.
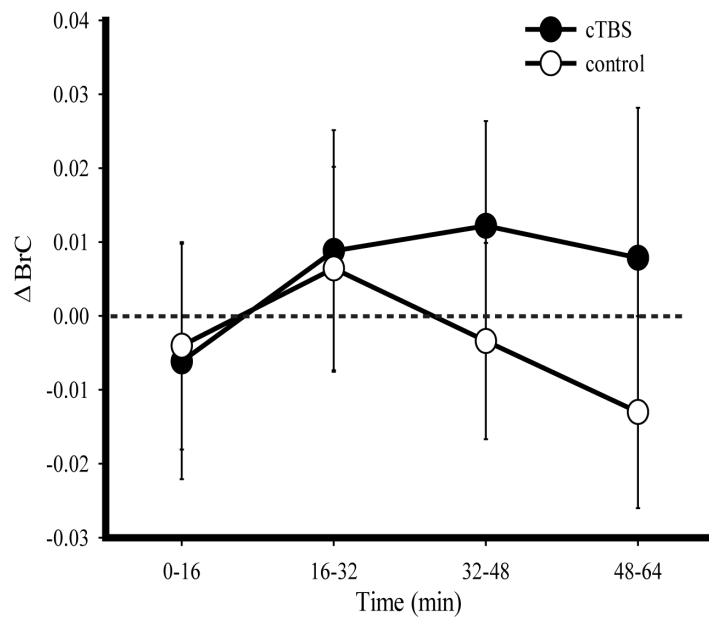
Figure 2.6.1.2.3. Change in response criteria / bias in conscious report (BrC) from pre-TBS baseline following cTBS and control stimulation, in Experiment 2D.

The effect of cTBS upon response bias observed in Experiment 2A did not replicate (BrC site effect: $F_{(1,17)} = 0.171$, p=0.706, BrC site × time interaction: $F_{(3,51)} = 0.347$, p=0.684, figure 2.2.6.3). Because the change in bias failed to replicate, further probing the possibility of an effect across experiments was justified. For subjects who participated in both experiments averages were drawn across repetitions. However, when the bias results from the original experiment were combined with the replication, no significant TMS effect was observed (site effect: $F_{(1,26)} = 2.430$, p=0.131). This lack of reproducibility indicates that little can be drawn from the data here in relation to response criteria/bias. A single subject's data was excluded from the analysis of BrC as an outlier, in both Experiment 2D and the combined analysis.

Overall, the central behavioural finding from Experiment 2D was that cTBS reliably increased conscious detection relative to the control condition, thus replicating the serendipitous effects observed in Experiment 2A.

2.6.2 Eye-tracking and arousal

2.6.2.1 Introduction

An alternative explanation for the effect of occipital cTBS on conscious detection is that it stems from enhanced arousal, potentially as part of a reaction (Ericsson, 2003) by subjects to active stimulation, and possibly independent of the direct neuronal effects of TMS. Pupil diameter is one of the most commonly used and reliable measures of autonomic arousal (Bradley, Miccoli, Escrig, & Lang, 2008; Bradshaw, 1967). Pupil diameter has previously been shown to increase following application of repetitive TMS, where it can be attributed to arousal \ sympathetic activation (Niehaus, Guldin, & Meyer, 2001).We measured pupil diameter in both the original and replication experiments and took this to quantify arousal. If arousal increased following the application of the active cTBS condition then we would expect an increase in pupil diameter relative to the control condition. The following section considers pupilometry data collected in Experiments 2A and 2D.

2.6.2.2 Methods

The original experiment (2A) made use of the CRS 250Hz chin mounted eye-tracker averaging pupil diameter collected on each trial into a block average. The MEG replication made use of a SensoMotoric Instruments infrared 50Hz eye tracking system mounted on a tripod. The replication did not divide the eye tracking data into trials, so averages were taken across blocks. This was because of the added time cost of relocating the subject into the magnetically sealed room following the TBS and head localisation procedures, preventing eye tracking calibration. This was also the reason for which the units of the replication are arbitrary rather than in mm and trials over which the subject blinked during stimuli presentation were not excluded in replication (the removal of such trials in the original experiment made no appreciable difference to the data and analysis). The data was filtered for losses of signal. Pupil diameter was averaged across blocks and the dependent variable was change in pupil diameter from pre-TBS baseline.

2.6.2.3 Results
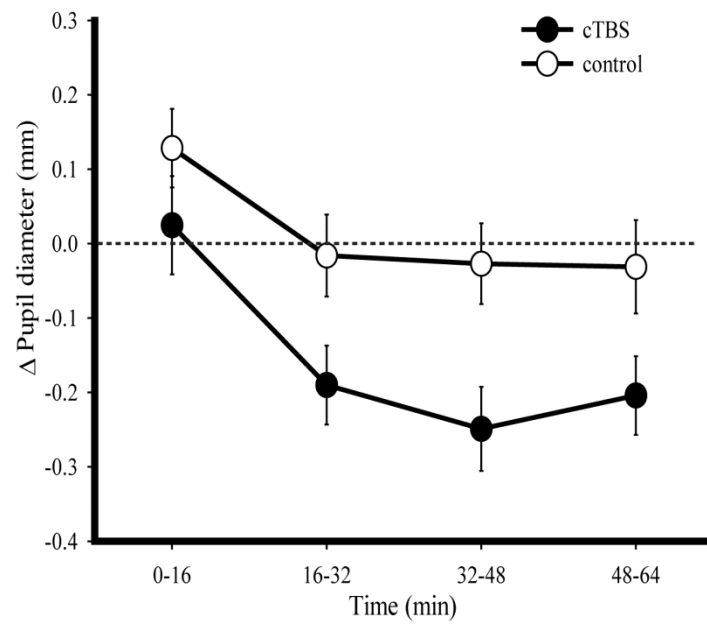
Figure 2.6.2.3.1. Change in pupil diameter in Experiment 2A, from pre-TBS baseline following cTBS and control stimulation.
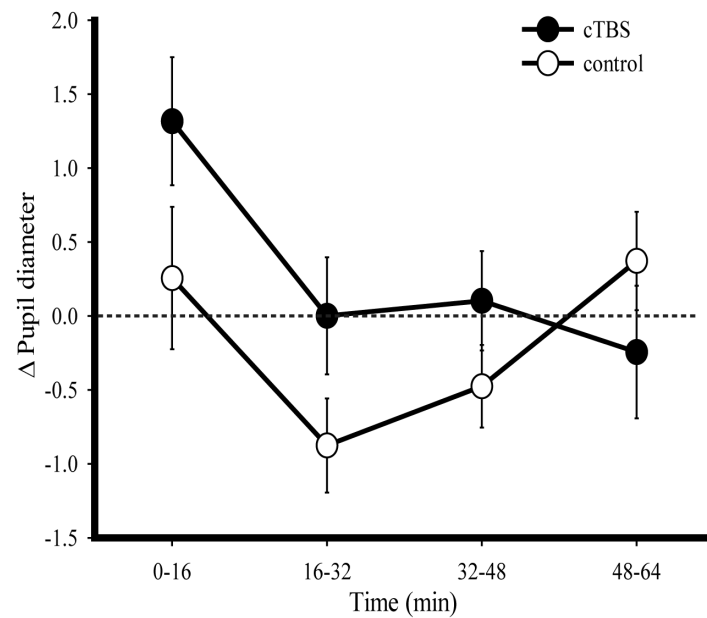
Figure 2.6.2.3.2. Change in pupil diameter in Experiment 2D, from pre-TBS baseline following cTBS and control stimulation. Units in replication are arbitrary (see methods)

In Experiment 2A, there was a trend towards a dissociation in the measure of pupil diameter between TMS conditions (site: $F_{(1,13)}$ =3.463, p=0.089, site × time interaction: $F_{(3,39)}$ =1.02, p=0.37) with arousal reduced following cTBS relative to sham (see figure 2.6.2.3.1). This direction of this trend suggests that the cTBS-induced increase in conscious detection in Experiment 2A is unlikely to have stemmed from increased arousal. In the replication experiment, average pupil diameter was instead greater following cTBS vs. control, although this difference was not statistically significant (site: $F_{(1,16)}$=1.124, p=0.305, site × time interaction: $F_{(3,48)}$=1.861, p=0.183, figure 2.6.2.3.2). A similar temporal profile in pupilometry was observed for both experiments: pupil diameter increased following the application of control and active TMS and then subsided throughout the course of the experiment, indicating its sensitivity to changes in arousal (original time effect on pupil diameter: $F_{(3,39)}$=9.71, p <0.001, and in replication: $F_{(3,48)}$=5.212, p=0.009, figure 2.6.2.3.1-2).

Two subjects were excluded from the analysis in Experiment 2A due to the eye tracker failing to acquire reliable data. One subject was excluded from the analysis of Experiment 2D as an outlier.

Because pupil diameter was not reliably affected by TMS, modulation of arousal is not a viable explanation for the effect of cTBS on conscious detection. Decreased arousal (fatigue) over the course of the experiment, may better explain many of the changes in the pupilometry data. Indeed, fatigue may also explain a number of other changes observed over the course of the experiments: there was a general increase in the power of alpha band oscillations throughout the course of the experiment (see appendix A2), which is consistent with increased fatigue (Boksem, Meijman, & Lorist, 2005). Also, there was a drop in performance in the conscious detection task, over the course of the study, irrespective of the TMS condition (section 2.2.3).

2.6.3 Evoked Magnetoencephalography

2.6.3.1 Introduction

Magnetoencephalography (MEG) offers the unique opportunity to observe a detailed and direct picture of how cTBS influences neural processing. The MEG detects electromagnetic changes emitted from the scalp, which are thought to originate predominantly from the summation of post-synaptic activity of pyramidal cells (Hamalain, 2002; Singh, 2006).

In addition to measuring the effect of cTBS on visual perception, Experiment 2D also tracked two aspects of the magnetic changes that have previously been associated with conscious processing. These are the evoked response to stimuli and the phase-independent oscillatory or induced changes (section 2.6.4).

The search for the neural correlates of consciousness (NCC; Crick, 1990) has become central to the scientific investigation of consciousness. Of the many correlates proposed, perhaps the most widely acknowledged NCC has been the correlation between conscious awareness and relatively late (>~100ms) cortical electromagnetic evoked responses to stimuli (Del Cul, et al., 2007; Kaernbach, Schroger, Jacobsen, & Roeber, 1999; Koivisto & Revonsuo, 2003; Quiroga, et al., 2008). One of the clearest demonstrations of this relationship and most relevant to the current studies is work carried out by Claire Sergent and colleagues using an attentional blink paradigm (Sergent, Baillet, & Dehaene, 2005). They subtracted the electroencephalographic (EEG) trace in the absence of stimuli from those collected in the presence of stimuli and showed that both reportedly 'seen' and 'unseen' trials resulted in comparable early (P1 and N1) evoked responses, whereas the later N3/4 (~300ms) components were potentiated when the subjects reported the stimuli as 'seen', hence correlating the late evoked responses with specifically conscious processing.

In contrast, there is some evidence that earlier potentials are also modulated according to the presence / absence of conscious processing (Kaernbach, et al., 1999; Marzi, Girelli, Miniussi, Smania, & Maravita, 2000; Schubert, Blankenburg, Lemm, Villringer, & Curio, 2006; Vuilleumier et al., 2001), although these studies also demonstrated later awareness-

dependent effects[4]. Therefore, there appears to be a broad consensus that beyond ~200ms the amplitude of evoked responses reflects the extent to which information is processed consciously (Babiloni, Vecchio, Miriello, Romani, & Rossini, 2006; Del Cul, et al., 2007; Kostandov & Arzumanov, 1977; Melloni, et al., 2007; Quiroga, et al., 2008). These processes may reflect the passage of conscious information in a recurrent occipital-frontal exchange (Del Cul, et al., 2007; Lamy, Salti, & Bar-Haim, 2009; Schubert, et al., 2006; Vuilleumier, et al., 2001).

The motivation for the quantification of late evoked responses applied here is in the service of the increased-response hypothesis. According this account, rather than suppressing superfluous representations, as posited by the gating-by-inhibition hypothesis, cTBS potentiated conscious representations directly. Late evoked responses may therefore be taken as representative of the conscious representation.

2.6.3.2 Methods

Figure 2.6.3.2.1. Illustration of example ERF produced from a single-subject, single analysis block, average across trials at the sensor level (peak channel selected)

---

[4] The study here was incapable of probing early evoked components as they were not clearly observable across the group. This is likely to be due to the stimuli being presented at peri-threshold levels, combined with their foveal presentation resulting in the cancellation of early evoked responses across the two hemifields (Luck, 2005).

changes from the -0.5 sec baseline. Also illustrated is the time course of the presentation of the task.

MEG acquisition

MEG was acquired on a 273-channel radial gradiometer system (CTF MEG, MEG International Services Ltd) sampled at 1200Hz, analysed as 3^{rd} order synthetic gradiometers (Vrba & Robinson, 2001). Data sets were collected in single 8-minute blocks with head localization procedures applied at the beginning and end of each block. Pairs of 8-minute blocks were then concatenated into single analysis blocks, resulting in 16-minute data sets that were consistent with the MRS acquisition. Trials were epoched from -2.3 to +1 second relative to the stimuli onset and band passed filtered with a 1-300Hz Butterworth filter. This resulted in 15 datasets for each subject (pre, post 1, post 2, post 3, post 4 × three TMS conditions of cTBS, iTBS and sham). Following DC offset the data were visually inspected and clearly corrupted data (e.g. from movement) were removed on a trial-by-trial basis. Together with other data loss, this resulted in a mean of 152 trials per data set (±13 trials SD). The two control TMS conditions (iTBS and sham) were processed separately and then collapsed prior to the application of inferential statistics.

Statistical analysis was conducted in sensor space and channels were clustered according to their CTF designation. All analyses were applied to the occipital / parietal cluster, unless otherwise stated, as these channels covered the region directly affected by the TMS.

Evoked response

The data was band pass filtered 1:40Hz (Butterworth). The baseline used was the mean field strength for each channel during the 500ms prior to the stimulus onset (see figure 2.6.3.2.1). To summate opposing evoked components, traces were squared and rooted. Evoked responses for each data set were measured by the peak amplitude of deflection from baselines applied to data sets averaged across stimulus-present trials. The peak was defined as the maximum amplitude of any channel in the cluster between 100ms and 500ms post stimulus.

This peak measure was taken to represent a combination of later M3 and M4 components which have been linked to conscious processing of a signal (Liddell, Williams, Rathjen, Shevrin, & Gordon, 2004; Sergent, et al., 2005). The mean latency for the peak was

320ms ($\pm$ 101ms SD), indicating that this methodology largely targeted the M3 component of the evoked response.

2.6.3.3 Results

Late evoked responses (Event Related Fields: ERFs) were reliably collected from stimulus-present trials but the peak amplitude did not appear to be affected by the application of cTBS (site effect: $F_{(1,16)} = 1.938$, $p=0.183$, figure 2.6.3.3.1-2). If peak amplitude had significantly increased then this would have lent weight to the alternative increased response hypothesis of the effect of cTBS on conscious detection. No time-dependent effects were observed (time: $F_{(3,48)} = 0.20$, $p=0.87$, site $\times$ time: $F_{(3,48)} = 0.38$, $p=0.75$). One outlier was excluded from this analysis.
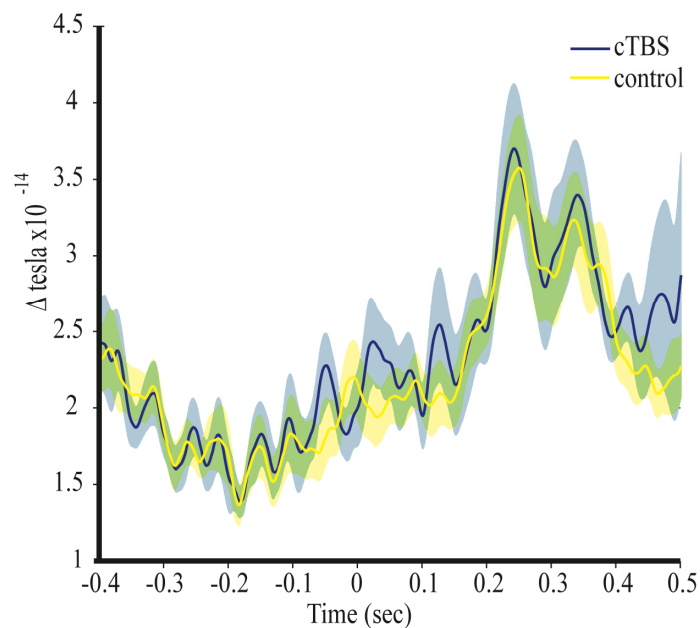
Figure 2.6.3.3.1. Group averaged evoked responses (ERF) following cTBS (blue) and control (yellow) stimulation, where shaded areas are one standard deviation across subjects. Plot derived from data averaged across post-TBS stimulus-present trials over occipital parietal clusters of channels. Δ refers to change from pre-stimulus baseline.
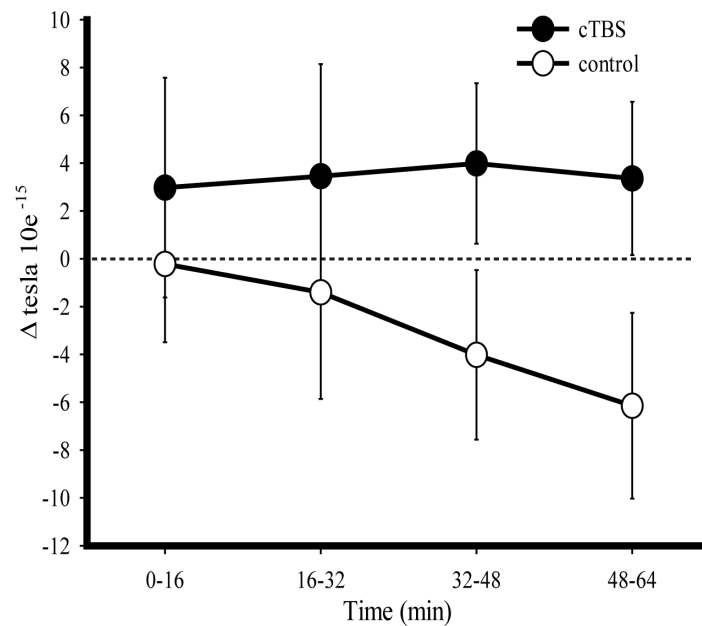
Figure 2.6.3.3.2. Change firm pre-TBS baseline in peak amplitude of evoked response following stimuli presentation for occipital / parietal channels. Active and control conditions are shown.

These results provides no conclusive evidence that occipital cTBS potentiates late evoked responses. Were there to have been such a demonstration, this would have provided evidence against the idea that the regulatory gating mechanisms, alone, were affected by the cTBS, favouring an explanation in terms of potentiation of conscious signals (the increased response hypothesis). However, it is worth emphasising that an absence of a statistically discernible effect should not be taken as evidence for or against any interpretation, especially given that there is a weak trend for increased evoked response following cTBS relative to controls as is apparent in figure 2.6.3.3.2. It is entirely possible that the measure of peak amplitude applied here was simply insensitive to the underlying changes, which, on the basis of the effect size in the measure of conscious detection (~8%[5]), may have been subtle.

---

[5] The mean difference between active and control conditions in the relevant MEG replication was 0.043 Pr units for the conscious detection measure. The maximum possible range of this measure was 0.5 Pr units. Therefore the effect size was roughly 8%.

2.6.4 Induced Magnetoencephalography

2.6.4.1 Introduction

Oscillations in magnetic activity at the scalp have been associated with a variety of brain processes, dissociable according to frequency. Higher frequencies, such as those in the gamma range (~30-100Hz) have been linked with the passage of explicit information throughout the brain (Crick, 1990, see appendix A2). Lower frequencies, particularly in the alpha band (~6-12Hz), been linked with the active suppression of superfluous information (Foxe & Snyder, 2011), gating (Jensen & Mazaheri, 2010) and idling (Pfurtscheller, 1992). Idling here refers to processes that occur during periods of inactivity or non-functionality. Alpha rhythms act to impose suppression in a temporal manner, such that suppression is maximal when the oscillation peaks; therefore it can be seen as segmenting information into discrete individual percepts (Varela, et al., 1981).

These attributes of alpha frequency oscillations clearly reflect the proposed gating-by-inhibition hypothesis. In particular, following the presentation of stimuli, there is a shift from a period when alpha activity is relatively high and suppression is imposed, to a less synchronised state following the appearance of the stimuli. This 'disinhibition' (Pfurtscheller, Schwarz, Pfurtscheller, & List, 1983; Romei, Gross, & Thut, 2010), operating through the induced[6] responses in the alpha band, is known as the event related desynchronisation (ERD, Pfurtscheller & Aranibar, 1979) and was the target of this experimental section.

A link between the alpha ERD and specifically conscious processing has been established (Pfurtscheller, et al., 1983). Event related synchronisation in response to visual

---

[6] Induced responses are changes in phase independent oscillatory activity following presentation of stimuli (Singh, 2006). There is a relationship between the evoked and induced responses: the evoked responses of the previous section summate activity over trials. Therefore, oscillations in phase with one another relative to or as a result of, the presentation of the stimuli will be expressed by the evoked responses. Because the late evoked responses (~300ms) were targeted in section 2.6.3, the measures used corresponded to phase locked activity in the delta-theta (1-6Hz) range (Yordanova, Kolev, & Polich, 2001). In contrast, induced responses to stimuli are the result of summation of oscillatory activity independent of phase and are applied to any frequency band prescribed, as it is performed via a Fourier (in this case Hilbert) transformation (Singh, 2006).

stimuli can be produced from comatose patients, whereas the extent to which the ERD is apparent follows recovery from such a state and appearance of conscious processing (Pfurtscheller, et al., 1983). In masking experiments it has also been shown that when subjects report stimuli as 'seen', the subsequent ERD over occipital and parietal channels is larger than when they report stimuli as 'not seen' (Babiloni, Vecchio, Bultrini, Luca Romani, & Rossini, 2006). According to the gating-by-inhibition account, the prediction here is that the ERD might be facilitated following the application of the cTBS.

2.6.4.2. Methods
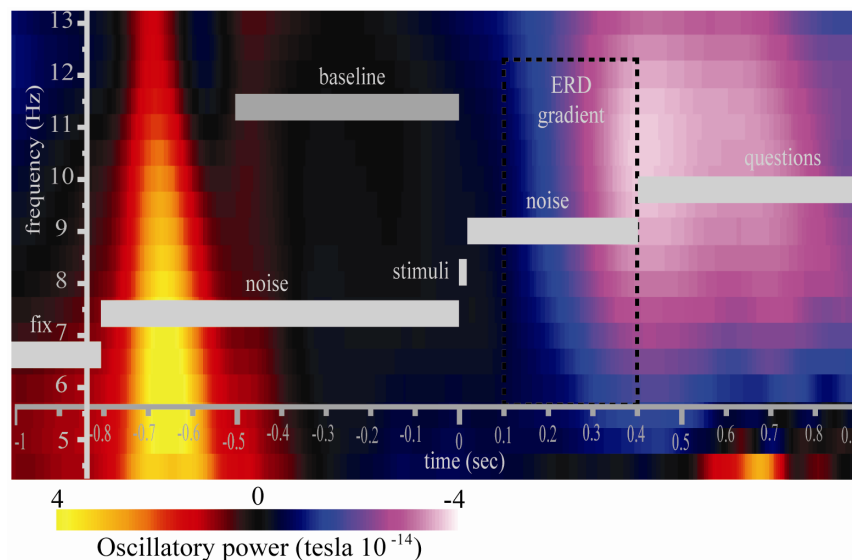
Figure 2.6.4.2.1. A typical time-frequency plot, illustrating event related desynchronisation (ERD) in the α band. Data is taken from group averaged post-TBS blocks, collapsed across TMS conditions. Power is derived from the application of a Hilbert transform to the data collected over occipital / parietal channels. Also illustrated is the time course of the presentation of the task, baseline period and period over which the ERD was quantified.

Acquisition parameters are as described in the evoked section (2.6.3.2). Oscillatory power within the alpha band (6-12Hz, Pfurtscheller, et al., 1994) was quantified by the application of a Hillbert transform (8Hz band width and 0.5Hz step size) applied to all trials in the data sets, down sampled to 600Hz. Calculated in sensor space, this resulted in a matrix for each data set of time (-2.3 to 1sec) × frequency (6-12Hz) × channel (85). From each

channel and frequency pair a baseline was taken from -500ms to the stimulus onset. This resulted in time × frequency plots for each data set of the induced oscillatory response (e.g. figure 2.6.4.2.1).

The ERD was quantified by the gradient of the desynchronisation between 100ms and 400ms following the stimuli onset (figures 2.6.4.2.1 and 2.6.4.2.2). The 100ms corresponds to the onset of the ERD irrespective of TMS conditions (see figure 2.6.4.2.1) and is consistent with previous reports (Pfurtscheller, et al., 1994). The offset of the visual noise and subsequent presentation of the questions from 400ms restricted quantification of the ERD to before this period. Rather than average power change (Pfurtscheller & Lopes da Silva, 1999) or area under the curve (Nogawa, Katayama, Tabata, Ohshio, & Kawahara, 1976), the rate of change of synchronisation was chosen as it was more representative of the desynchronisation aspect of the ERD and it should theoretically be less susceptible to outlying data points and potential confounds such as prior (Romei, et al., 2010; van Dijk, Schoffelen, Oostenveld, & Jensen, 2008) or subsequent synchronisations (Pfurtscheller, 1992). The gradient was obtained by a linear regression applied to the data collapsed across the frequency pairs in the alpha range. The channel most clearly (greatest negative gradient) demonstrating an ERD was selected to represent the ERD for each data set, as, owing to the localisation of the ERD (See SAM analysis below) many channels would not have expressed the desynchronisation.

Tomographic images of the localisation of the ERD were produced using Synthetic Aperture Magnetometry (SAM, Robinson and Vrba 1990). Although synthetic channels themselves were not used in the final statistical analysis, this allowed us to localise the oscillatory power changes in terms of group level anatomy (figure 2.6.4.3.3). This analysis made use of the same parameters as in the sensor space analysis of the ERD (baseline from -500ms to stimuli onset, 6-12Hz band width, 100-400ms period of interest) but applied to data sets that combined all post-TBS blocks. Co-registration of this data to subjects' anatomical MRI images was achieved with MRIViewer software (Krish Singh) in conjunction with digital photographs that depicted the position of the localisation coils (one above the nose bridge and two over the tragi). The resulting outputs were pseudo T maps (values below 0.4 were removed from the image), which were averaged across subjects and overlaid on a FreeSurfer (Harvard MA) inflated brain mesh using mri3dx (Krish Singh) (figure 2.6.4.3.3).
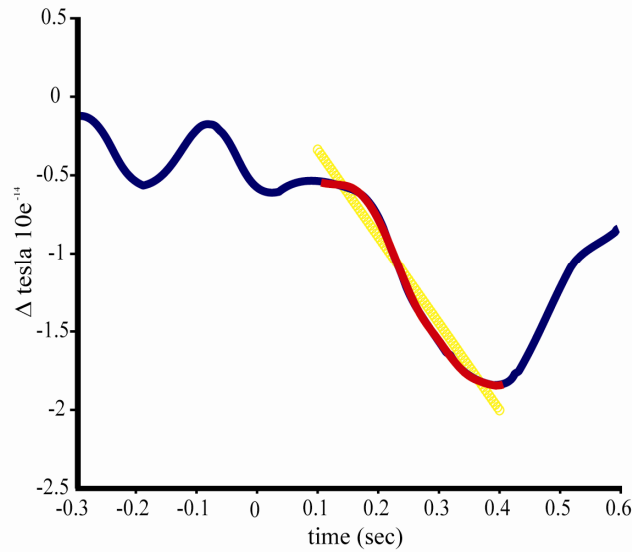
Figure 2.6.4.2.2. Example of linear fit applied to individual time frequency analysis of the ERD in each data set. The fit (yellow) was applied to the data (red) between 0.1 and 0.4 seconds post stimulus averaged across the 6-12Hz alpha frequency. The gradient of this fit was the dependent variable and the channel selected was the channel expressing steepest gradient in the occipital / parietal cluster.

2.6.4.3 Results

The ERD, being a drop in synchronisation was quantified by the rate of change of phase independent oscillatory power over fixed period relative to stimuli presentation. Its presentation was consistent with previous reports (Pfurtscheller, et al., 1994) in terms of frequency and temporal range (figure 2.6.4.2.1). The gradient of the ERD was found to increase significantly (become more negative) following the application of cTBS relative to controls ($F_{(1,16)}$=6.828, p=0.019, figure 2.6.4.31-3), indicating facilitation of the alpha band response. These analyses were conducted over the occipital / parietal cluster of sensors and additionally localisation of the ERD to the targeted occipital lobe is illustrated in figure 2.6.4.3.3. No time dependent effects were observed (time: $F_{(3,48)}$ =0.863, p=0.45, site × time: $F_{(3,48)}$ = 0.399, p=0.678). One outlier was excluded from this analysis.
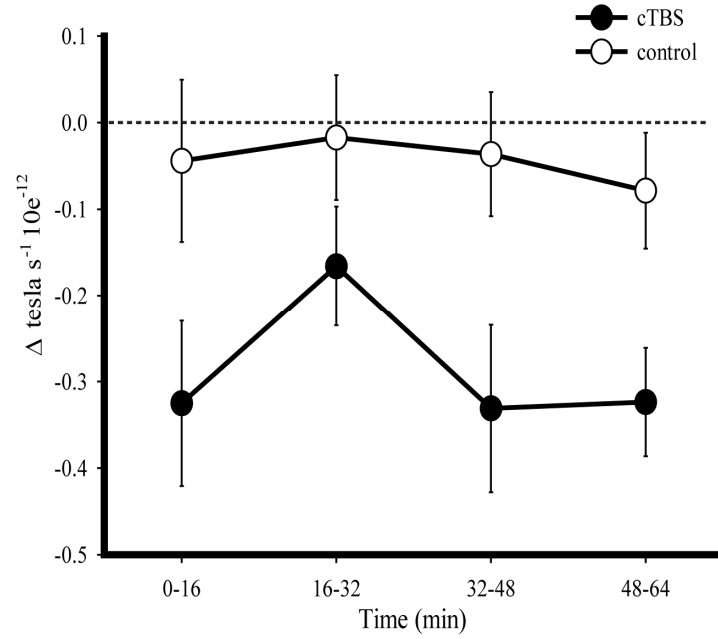
Figure 2.6.4.3.1. Change in gradient of the ERD. The more negative the gradient the steeper the onset of the shift from synchronisation to desynchronisation.
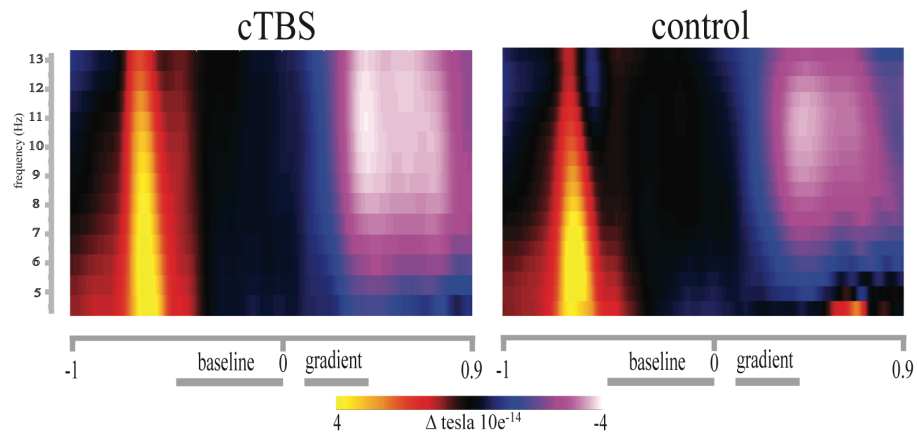
Figure 2.6.4.3.2. Group level time frequency plots collapsed over all post TMS blocks in control and cTBS conditions, covering the alpha frequencies.
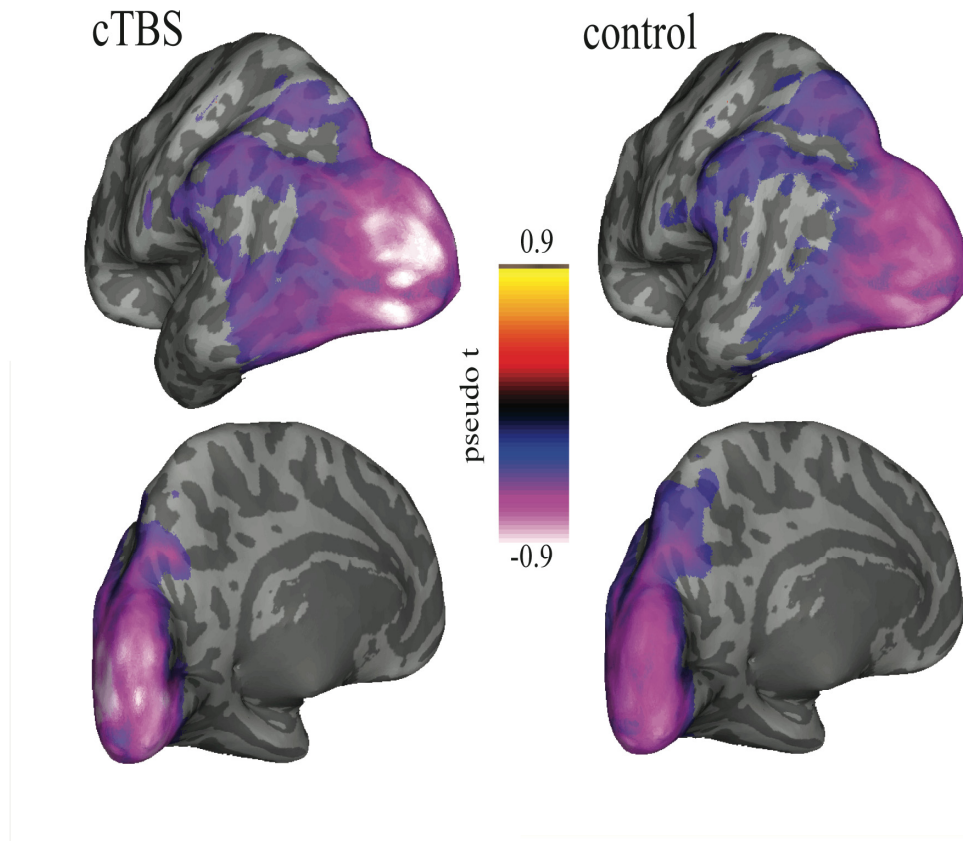
Figure 2.6.4.3.3. Synthetic aperture analysis illustrating localisation of the ERD to the occipital lobe (left hemifield depicted). Applied to data sets following cTBS and control stimulation.

The drop in synchronisation in the occipital lobe (the ERD) was more rapid following the application of the cTBS, indicating potentiation of the active relief from suppression. This mirrors the gating-by-inhibition hypothesis.

2.7 General Off-line Discussion

Experiment 2A showed that subjects' reported awareness of stimuli became more frequent following the application of cTBS, whereas their forced choice 'unseen' ability remained unaffected. This behavioural effect was replicated in Experiment 2D. Experiment 2B revealed an elevation of phosphene threshold following the application of cTBS, replicating Franca, et al., (2006). Experiment 2C showed a corresponding cTBS-induced increase in the inhibitory neurotransmitter GABA, consistent with the findings of Stagg et al., (2009). Experiment 2D additionally demonstrated that the onset of ERD in the alpha band became steeper following the application of the cTBS.

The occipital lobe has long been known to play a prominent role in visual consciousness. This study refines our understanding of that role by emphasising the mechanism of gating in the formation of conscious percepts. As with blindsight, the effect can be said to be relatively specific to conscious processing in that the contrast state of 'unseen' discrimination was above chance but unaffected by the intervention. Determination of whether or not a subject consciously perceives stimuli is likely to involve selection, which in turn involves discarding representations that do not enter the frame. The suggestion is that it was this negative aspect of selection which was facilitated by the application of the cTBS.

Selective gating and active suppression are associated with oscillatory changes in the alpha band (Foxe & Snyder, 2011; Jensen & Mazaheri, 2010). Decreases in alpha power are observed when the requirement for processing is raised; for example, alpha oscillations decrease when the eyes are open relative to closed (Berger, 1929). Alpha oscillations have also been implicated in the framing of individual percepts, where subjects are more likely to group percepts together if they are presented during the negative trough of alpha oscillations than they are during the peak (Varela, et al., 1981). Furthermore, the presence of ERDs in the alpha band has been specifically associated with the presence of explicitly conscious processing, whereas event related synchronisations have been observed following both conscious and unconscious processing (Pfurtscheller, 1992). The governance of perceptibility by alpha rhythms and their dissolution in the form of the ERD therefore seemed to exemplify the kind of gating mechanism that might be responsible for the observed increase in

conscious detection following the application of cTBS. Indeed the ERD dropped more rapidly following the application of the cTBS, indicating a facilitation of the gating response.

The balance of activity at the biochemical level is played out by the balance between excitatory and inhibitory neurotransmitters. The hypothesis put forward is that this balance is tipped towards inhibition by the application of cTBS, increasing the contrast or gain (Katzner, Busse, & Carandini, 2011) between representations. A similar interpretation has been offered following the demonstration that microiontophoretically (micro injection and measurement using ionic currents) applied GABA suppressed spontaneous discharges, but did not affect evoked response to auditory stimuli in rats (Ebert & Ostwald, 1995).The increase in observed GABA concentration following cTBS was therefore consistent with the gating-by-inhibition hypothesis, suggesting that an important factor in determining whether or not a representation is realised as conscious is the active suppression involving GABA.

Of the three rival hypotheses we explored, none received support in the data. Clearly arousal, as measured by pupil diameter, did not appear to correlate with the application of the cTBS (Kindler et al., 2012), although arousal or rather lack of arousal did influence the data.

 If the late evoked MEG responses are seen as representative of conscious signals, then increased peak ERF as a result of cTBS can be seen as consistent with the alternative hypothesis that cTBS increased excitatory content-bearing responses. Furthermore, if the claim were to be that general neuronal activity was suppressed (irrespective of function) owing to the application of cTBS, then we might have expected the ERP to have been reduced. No clear evidence of such changes was apparent.

A general suppression of neural activity caused by the cTBS is not sufficient in itself to explain the observed increase in conscious detection. Rather, the proposal is that the mechanisms which gate consciousness involve suppression / inhibition under normal conditions, and it was these selection mechanisms that were specifically facilitated by occipital cTBS. This may be why the application of cTBS resulted in the decreased detection of phosphenes and increased detection of normal external stimuli. Detection of phosphenes differs fundamentally from detection of external stimuli in that phosphenes are the result of direct pervasive stimulation of neurons in visual areas (Thielscher, et al., 2010), whereas conscious detection of external stimuli is the result of refined teleological mechanisms. Therefore, the increase in phosphene threshold is consistent with there being a general reduction in excitability and it being this inhibitory element - as part of the mechanism that

produces conscious percepts - which benefits from the application of cTBS. The increase in phosphene threshold also replicates the work of (Franca, et al., 2006) and detracts from the alternative hypothesis that cTBS might increase excitation and/or noise (stochastic resonance).

Although the effect of cTBS increasing conscious detection was replicated, two points of caution are worth emphasising. First, the stimuli used here were presented among luminance noise and the effect we claim cTBS to have can be interpreted in terms of suppression of superfluous noise. It is therefore possible that the increase in conscious detection may only be realised during tasks that embed task-relevant stimuli within noise, with the facilitation arising through suppression of that noise. Second, the TMS intensity applied in these experiments was relatively low. This raises the question: would higher intensities cause suppression (Abrahamyan, et al., 2011). A correlation analysis involving motor threshold designed to probe this question is described in appendix A1.2. Essentially this failed to provide conclusive evidence with respect to this question. It therefore remains a potential avenue for future research.

One of the more common assumptions in the design and interpretation of TMS experiments is that TMS operates as some form of 'virtual lesion' (Hilgetag, Theoret, & Pascual-Leone, 2001; Pascual-Leone, Bartres-Faz, & Keenan, 1999; Pascual-Leone, et al., 2000; Walsh & Cowey, 2000). The finding that this form of repetitive TMS facilitates conscious processing in contrast to the blockage of activity caused by actual lesions (as in classic blindsight, where consciousness is lost) indicates that the virtual lesion interpretation would be an oversimplification if applied here, and therefore might not be an appropriate default assumption of repetitive TMSs mode of operation (Hilgetag, et al., 2001).

The facilitation of conscious detection following cTBS is by no means the first demonstration of improvements in cognitive capacities following repetitive TMS (Cazzoli et al., 2012; Cazzoli, Wurtz, Muri, Hess, & Nyffeler, 2009; Funke & Benali, 2011; Kindler, et al., 2012; Marzi et al., 1998; Tegenthoff et al., 2005; Thompson, Mansouri, Koski, & Hess, 2008; Waterston & Pack, 2010; Yamanaka, Yamagata, Tomioka, Kawasaki, & Mimura, 2010). Some of these investigations are less relevant to the interpretation of current experiments than are others. This is because the frequency at which some experiments applied TMS are thought to be excitatory protocols, leading to increased activation of corresponding cortical representations (Funke & Benali, 2011; Tegenthoff, et al., 2005;

Thompson, et al., 2008; Yamanaka, et al., 2010), unlike cTBS (Franca, et al., 2006; Huang, et al., 2005). For example, Tegenthoff et al used an excitatory 5Hz TMS protocol to improve tactile discrimination when applied to somatosensory finger areas, and showed that this effect correlated with increased blood-oxygen-level-dependent activation in the corresponding region as measured by functional magnetic resonance imaging. Other studies where improvements have been observed can be attributed to the artefactual effects of the TMS such as auditory or inter-sensory facilitation (Marzi, et al., 1998).

Of greater relevance here is the study by Waterson et al., (2010) who reported that cTBS boosted visual sensitivity when applied at similar intensity to the current study. However, the findings of that study are questionable. In particular, they report in the text of the paper that 6 out of 7 subjects showed an improvement following the application of cTBS to primary visual areas. However, in the figure they present, drawn from the same data, the effect relative to controls is only apparent in 3 out of 7 subjects and it appears to be these 3 subjects that drove the group-level mean difference in performance. Although there are statistical concerns with this study, the authors' conclusions are consistent with the data collected over the current experiments. In particular, they describe cTBS as being effective in improving 'coarse' perceptual judgements (using large angular displacement, low contrast gratings) and not 'fine' judgments (small angular displacement, high contrast gratings). It is notable that the detection task employed here more closely resembled the coarse task, consistent with the findings of Waterson et al., (2010).

The most promising areas in which improvements in performance following cTBS might be applied is within the clinical context; for instance, cTBS has been shown to relieve symptoms of unilateral spatial neglect (Cazzoli, et al., 2012; Cazzoli, et al., 2009). It has been proposed that the cause of spatial neglect is a lack of inter-hemispheric inhibition, leading to hyperactivity with the contralesional hemisphere and the maintenance of attention to the corresponding contralesional visual field (Cazzoli, et al., 2009; Koch et al., 2008; Reuter-Lorenz, Kinsbourne, & Moscovitch, 1990). The interpretation offered by Cazzoli et al, is that cTBS increased the mechanism of inhibition leading to a reduction in the symptoms of neglect (Cazzoli, et al., 2012; Cazzoli, et al., 2009). This is consistent with the interpretation of the improvements based on facilitation of inhibitory gating mechanisms offered here. A similar interpretation (potentiation of intercortical inhibition mechanisms) could explain the finding that other forms of rTMS can help alleviate the symptoms of amblyopia (see Thompson, et al., 2008). The increase in conscious detection here is also consistent with the

recent demonstration that a similar cTBS protocol can speed verbal responses for aphasic patients following stroke (Kindler, et al., 2012). In conjunction with the data here, the further speculative suggestion can be made that the observed changes may reflect part of a homeostatic response orientated towards the return of optimal conditions following the presumably high levels of enforced activity that occurred during the application of the cTBS.

In summary, this study found cTBS to increase conscious detection of stimuli, contrary to the expected profile of blindsight. The confirmation of cTBS being neuronally suppressive and the demonstration of increases in measures associated with active regulation, suggests that mechanisms which gate, select and bring into relief consciousness are those that were influenced by occipital cTBS.

2.8 Limitations

2.8.1 Behavioural Task Limitations

This section will set out the principal limitations encountered. Resulting modifications in task design are discussed in section 3.2.1.

There were two interrelated problems encountered in these experiments. The first is the extent to which reportedly 'unseen' abilities can be regarded as *un*conscious. This in turn has two facets – one conceptual and one practical. The conceptual problem of what behavioural pattern should be interpreted as unconscious (and conscious) is discussed in section 4.3. The practical aspect is the possibility of subjects being aware – consciously – of a portion of the arrow (figure 2.8.1.1), and are therefore provided with directional information, yet, this representation is insufficient to elicit a categorical 'yes' response to the 'arrow?' question. It is therefore possible that the TMS may cause such partial ablation, resulting in the demonstration of above chance reportedly 'unseen' ability, yet the processing upon which these demonstrations are based would be conscious.

Figure 2.8.1.1 Illustration of the possibility that TMS could cause a partial scotoma which might prevent acknowledgment of conscious awareness, yet facilitate task performance on the basis of partial conscious awareness.

The second problem concerns the application of SDT when relatively few false alarms are produced. Subjects' response criteria to the 'yes/no' question was very conservative (BrC mean = 0.067 ± 0.087SD, where Br = 0 is 0 false alarms and 1 is responding 'yes' on all trials). That is, subjects were unlikely to report being aware of stimuli that had not been presented. This feature of 'yes/no' tasks has previously been described (Macmillan & Creelman, 1990) and is to be expected given the meta-cognitive aspect of the question. Meta, here, refers to the requirement for a secondary level of report and/or measurement referring to a primary level representation. The secondary response is unlikely to be positive in the absence of a primary representation (see section 4.4.1.1 for further details). The resulting willingness of subjects to report 'No/Unseen' relates to the difficulty of defining these responses as categorically unconscious (1.3.2.3). The problem this poses for the application of classic parametric SDT is that small changes such as single false alarms will have a disproportionately large effect on the data. Single trial differences in the context of low overall rates will correspond to the tail ends (most non-linear) of the normal distributions applied. Although the application of non-parametric SDT reduces this problem, the capacity of any form of SDT to extract bias from sensitivity measures and to track changes independently, is limited if no false alarms are produced.

Fortunately, the impact of these problems upon the conclusions of the off-line TMS experiments in this chapter is arguably limited. The central findings of these experiments all concern changes upon the measure of conscious detection, not 'unseen' discrimination. The relation between the measure of conscious detection and conscious awareness is only made *more* tenable through subjects' adoption of conservative criteria, as any token reported awareness is more likely to be veridical. Furthermore, the contrast between conscious awareness and perception lacking awareness is maintained even if the perception lacking awareness is not described as *un*conscious. This is because what is required for the contrast, revealing of conscious processing, to function as intended, is the relative difference between when the subject reports awareness and reports 'no' awareness, not the absolute *absence* of awareness when reporting 'no'. In this way, consciousness can still be said to be suppressed in the case of the partial scotoma.

Although these problems do not impact greatly upon the interpretation of the experiments described in this Chapter, they are important for the demonstrations of TMS-induced blindsight in Chapter 3. Task adaptations applied to lessen the impact of these problems are discussed in section 3.2.1.

2.9 Summary of off-line TMS study

These experiments have developed from the initial counterintuitive finding that suppressing cortical excitability with TMS can increase the conscious detection of stimuli. This facilitation was replicated and was found to be coincident with potentiation of the ERD and the elevation of occipital GABA concentration. Together these findings suggest that inhibitory gating is an important element in the formation of visual consciousness.

Chapter 3. On-line experiments

Using event related TMS to track the temporal dynamics and anatomy of information flow in visual perception.

3.1 Introduction

3.1.1 Overview

This series of experiments comprised two broad sections: a between-subjects study (Experiment 3A, section 3.3) and a single subject case study (section 3.4-6). The single subject study was conducted over three individual experiments (3B, 3C, 3D). The experimental tasks and aims were essentially the same across the two studies, so this chapter is prefaced with a general introduction followed by separate descriptions of the individual experiments.

The over-arching aim of this chapter was to look at the timing and pathways of activity that support conscious vision and visual perception lacking in conscious awareness. The main novel findings are that input from the superior colliculus is not necessary for the residual abilities of TMS-induced blindsight, and that these abilities appear to depend upon later recurrent activity, rather than the early feedforward sweep of activity as has been previously suggested (Lamme, 2001). In contrast to reportedly 'unseen' abilities, the conscious detection of stimuli appears to be supported by input from the superior colliculus and/or magnocellular routes of the lateral geniculate nucleus, during very early and very late stages of occipital processing.

3.1.2 The timing of visual activity

Individual TMS pulses applied to the occipital lobe, at times relative to visual stimuli, can track the passage of information through the visual system in a temporally precise manner. Amassian and colleagues were the first to use TMS in this mode, applying stimulation at 20ms intervals between 0 and 200ms during a letter recognition task (Amassian, et al., 1989). They showed that TMS caused a suppression or dip in performance accuracy when it was applied between 80-100ms after stimulus onset. They speculated that there is an early unconscious stage of processing and a second consciousness-related stage, corresponding to the 80-100ms dip, which is more susceptible to the effects of TMS (Amassian, et al., 1989). This strategy, of applying TMS at times relative to stimuli in order to reveal the timing of activity, became a theme running through a great deal of subsequent work and inspired this series of experiments.

The drop in detection as a result of applying TMS at ~ 100ms, first reported by Amassian et al, is probably the most widely replicated effect to have emerged from the non-motor TMS literature (e.g. Corthout, Uttl, Juan, Hallett, & Cowey, 2000; Corthout, Uttl, Walsh, Hallett, & Cowey, 1999; Corthout, Uttl, Ziemann, Cowey, & Hallett, 1999; Kammer, Puls, Strasburger, Hill, & Wichmann, 2005; Koivisto, Railo, & Salminen-Vaparanta, 2011; Overgaard, Nielsen, & Fuglsang-Frederiksen, 2004; Sack, van der Mark, Schuhmann, Schwarzbach, & Goebel, 2009; Thielscher, et al., 2010). It has also been suggested that there are other periods over which TMS is effective in reducing perception. For example, papers produced by Erik Corthout and colleagues suggest that there are several additional periods during which TMS is effective before the ~100ms period (Corthout, et al., 2002; Corthout, Hallett, & Cowey, 2003; Corthout, et al., 2000; Corthout, Uttl, Walsh, et al., 1999; Corthout, Uttl, Ziemann, et al., 1999). Others have suggested that there are periods where TMS is effective, in addition to ~100ms, at a much later stage, beyond 200ms (Camprodon, Zohary, Brodbeck, & Pascual-Leone, 2010; Chambers, Allen, Maizey, & Williams, 2012; Heinen, Jolij, & Lamme, 2005). Recently, it has even been suggested that the period of TMS efficacy at ~100ms can be divided into sub-components via the demonstration of additional dips in performance over the 60 to 120ms range (de Graaf, Herring, & Sack, 2011).

The central concern regarding the demonstrations of periods of TMS efficacy in addition to the ~100ms dip is consistency in replication. For example, work produced by the same researchers has appeared to provide evidence both for and against the existence of additional periods of TMS efficacy upon visual processing (cf. Sack, et al 2009., with de Graaf, et al., 2012; and Koivisto, Mantyla, & Silvanto, 2010; with Koivisto, et al., 2011). The early period of TMS efficacy (20-60ms post stimulus Corthout, Uttl, Walsh, et al., 1999) has also been resistant to replication with many studies capable of demonstrating efficacy of TMS applied at these times (e.g. Amassian, et al., 1989; Koivisto, Henriksson, Revonsuo, & Railo, 2012) but few having achieved such replication (Koivisto, et al., 2010). The efficacy of TMS when applied at times after the ~100ms period does seem slightly more consistent than the very early effects, although relatively few TMS experiments have targeted activity beyond 200ms. These experiments are explored in greater depth in section 3.5.3.

One reason for inconsistencies in the reporting of multiple periods of TMS efficacy may be statistical. In order to demonstrate two separable periods of TMS efficacy, the orthodox strategy is often to show two periods which express significant drops in performance (often performed with t-tests) with an intervening period where TMS is ineffective (e.g. Corthout, Uttl, Ziemann, et al., 1999). However, this may not be a valid approach because the demonstration of multiple periods rests on the demonstration of the intervening period when TMS is ineffective and the application of orthodox Fisher-based statistics does not address this null hypothesis (Dienes, 2008; Nieuwenhuis, Forstmann, & Wagenmakers, 2011); the absence of a significant effect of TMS applied over the intervening period does not mean there is no effect. Indeed, it is often the case with demonstrations of multiple periods that the statistics referring to the intervening period are simply absent from the published report (e.g. Juan & Walsh, 2003).

Another, and possibly more influential, reason for inconsistencies of effects may be differences in the behavioural tasks employed. Intuitively, different behavioural tasks put different demands upon early cortical areas in order to form correct responses. In particular, long exposures to stimuli and more complex tasks may require later activity, whereas speeded and reflexive responses are likely to depend upon earlier activity (cf. Ro, et al., 2004 with Chambers, et al., 2012).

The demarcation between early and late activity is most widely recognised as corresponding to the neuronal processes of feedforward followed by recurrent activity, as

described by Victor Lamme. He describes feedforward activity as corresponding to pre-attentive or unconscious vision and the later recurrent phase as conferring conscious awareness (Lamme, 2001; Lamme & Roelfsema, 2000). This differentiation maps on to the distinction between 'unseen' discrimination vs. conscious detection, in the current experiments, and the dissociation of blindsight type phenomena generally. Indeed, Lamme cites in support of his theory the disruption to recurrent processing in blindsight, the resulting lack of conscious awareness, the preservation of pathways that might still provide an early feedforward sweep of activity, and the preservation of unconscious abilities (Lamme, 2001).

TMS-induced blindsight studies tend to adopt the feedforward followed by recurrent – conscious processing model in their interpretation of data (Boyer, et al., 2005; Jolij & Lamme, 2005 ; Ro, et al., 2004). That is, the disruption caused by applying TMS from approximately 100ms post stimulus onset disrupts recurrent processing and therefore leads to a suppression of conscious awareness. One of the main aims of the current experiments has been to uncover whether there are one or multiple periods of TMS efficacy in response to this arrow based task, and to ascertain whether the early and late distinctions correspond to the distinctions between consciousness and perception lacking in consciousness.

3.1.3 The pathways of visual activity

According to Lamme (2001), the retinotectal pathway via the superior colliculus (SC) is the main route responsible for the feedforward sweep of activity and residual perceptual abilities of blindsight (Lamme, 2001). However, whether or not the residual abilities of blindsight are supported by this pathway, as opposed to other pathways such as preserved projections from the lateral geniculate nucleus (LGN), is not clearly resolved and this debate forms a central question in the blindsight literature (Cowey, 2010 see figure 3.1.3.1). The current series of experiments addressed this question by demonstrating TMS-induced blindsight and using stimuli types which do or do not activate the SC (Sumner, et al., 2002). I will briefly discuss the evidence in support of this stimulus intervention isolating collicular tracts before addressing the evidence in relation to the role of the SC in blindsight.
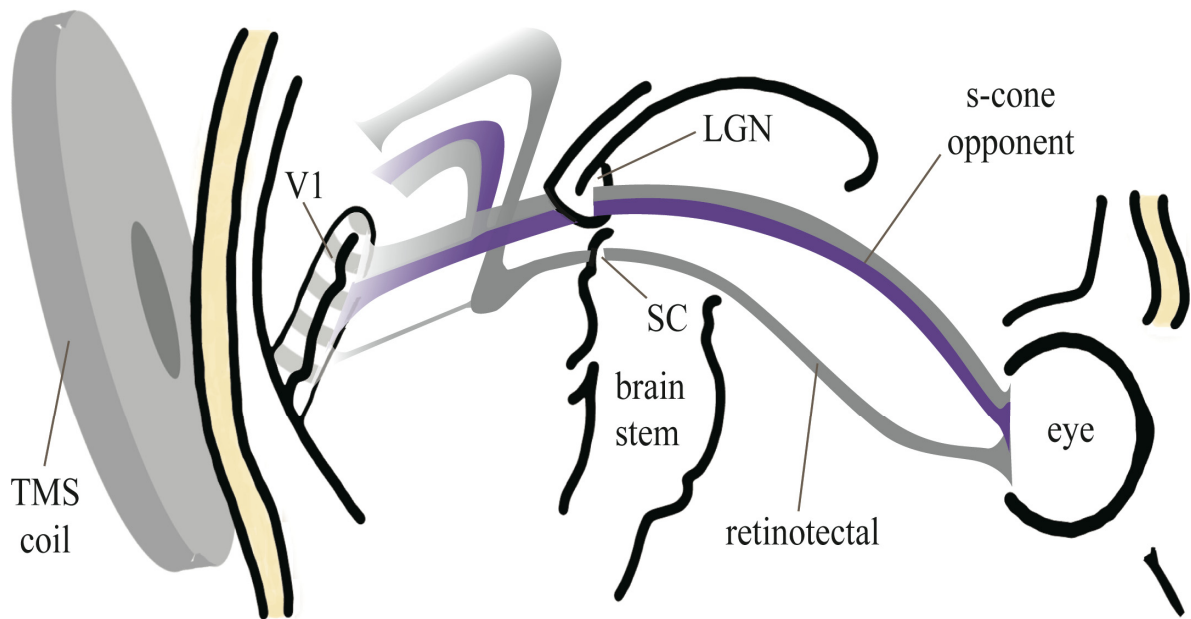
Figure 3.1.3.1. Illustration of the experimental configuration with TMS functionally targeted at visual cortices (see Thielscher, et al., 2010)(V1). The main geniculate pathway, via the lateral geniculate nucleus (LGN), is represented in grey / purple (purple depicting the s-cone opponent projection via the koniocellular layers of the LGN). These pathways project mainly to V1, but are also thought to contribute some inputs directly to extrastriate areas, represented by the junction between the LGN and V1 (e.g. Fries, 1981; Sincich, Park, Wohlgemuth, & Horton, 2004). The collicular / retinotectal pathway, also suggested as supportive of the residual abilities of blindsight, is illustrated as projecting via the superior colliculus (SC), here represented in grey below the geniculate pathway. From the SC, this route makes onward projections to pulvinar and several cortical areas including extrastriate, parietal and frontal regions (Fries, 1984).

The critical stimulus intervention in these studies is based on the principle that, of the three classes of colour receptors in the retina, short wavelength responsive cones (s-cones) are thought to be the only cells that do not activate the tectum of the SC in a colour opponent manner (Sumner, et al., 2002). In addition to bypassing the retinotectal (collicular) pathways, s-cones are also thought not to project via the magnocellular layers of the LGN, and possibly parvocellular layers (Dacey, 2000; Mollon, 1989). These non-s-cone retinotectal and magnocellular pathways are therefore collectively denoted henceforth as 'R/M'. This anatomical peculiarity has previously been exploited to demonstrate the roles of these pathways. The use of s-cone specific stimuli impoverishes R/M input, and differences in

responses, relative to comparable luminance stimuli, can be attributed to a functional role of the isolated pathways (e.g. Sumner, et al., 2002; Sumner, Nachev, Vora, Husain, & Kennard, 2004).

Evidence supporting the absence of s-cone input to the retinotectal pathway has emerged largely from primate electrophysiological work carried out in the 1970s. For instance, de Monasterio (1978) probed the responsiveness and connectivity of retinal ganglion cells by displaying monochromatic colours and antidromic stimulation of various areas of the monkey visual system. Although he found cells responsive to all wavelengths of visible light and cells activated by stimulation of the SC, he found no cell type that was responsive to both blue (short wave) light and to stimulation of the SC. Similarly, Schiller and Malpeli (1977) found a lack of colour opponency in ganglion cells projecting to the SC (and suggested the segregation of a blue 'on' opponent pathway to the LGN). Consistently, Marrocco and Li (1977) showed no colour opponent processing when probing the activity of the SC using microelectrodes, and argued that previous work (Humphrey, 1968; Wolin, Massopust, & Meder, 1966) which appeared to demonstrate colour processing in the SC could be attributed to luminance differences between coloured stimuli that did not appear to be adequately controlled for in the earlier experiments, and to which the SC is sensitive. Lesions of the SC in monkeys have also been shown to impair responses to some visual tasks, but not those that involve s-cone opponent (blue-yellow) discriminations (Butter, 1974).

The magnocellular layers of the LGN provide an additional pathway whose function can be addressed using s-cone stimuli. A significant proportion of magnocellular projections bypass primary visual areas, directly innovating extrastriate regions such as the middle temporal cortex (MT, Chatterjee & Callaway, 2002). As with the SC, it is not that magnocellular projections lack s-cone input but that they lack colour opponency in their response profiles, which enables their functional isolation (Derrington, Krauskopf, & Lennie, 1984). The presence of s-cone input was demonstrated in about 10% of magnocellular LGN neurons sampled by Chatterjee and Callaway (2002), but these responses summated in the presence of other colours rather than showing colour opponency. Since these pathways respond to changes in luminance, the use of s-cone stimuli to isolate the R/M pathways mandates the inclusion of luminance noise in order to mask the stimuli in terms of any luminance change. Instead of a change in luminance the s-cone stimuli are presented as an increment of s-cone colour saturation, rendering the R/M pathways blind to these stimuli.

Note that here 's-cone stimuli' is used to denote the transition between gray and purple in the context of luminance noise, rather than the presence of the purple arrow per se.

These s-cone stimuli provide an avenue of enquiry into one of the central debates in the blindsight literature. Namely, which pathways support the residual vision when primary visual areas are disrupted? Collicular pathways fulfilling this role were first suggested by Weiskrantz (1986). A series of psychophysical (Leh, Mullen, et al., 2006) and imaging (Leh, Johansen-Berg, & Ptito, 2006; Leh, Ptito, Schonwiesner, Chakravarty, & Mullen, 2010; Tamietto et al., 2010) experiments claimed to show that this was indeed the case via the use of s-cone stimuli. They demonstrated that collicular projections were preserved when geniculate routes were damaged in patients with blindsight (Leh, Johansen-Berg, et al., 2006). The superior colliculus responded to visual stimulation (Leh, et al., 2010; Tamietto, et al., 2010) and, most importantly, blindsight type effects were suppressed when s-cone stimuli were used (Leh, Mullen, et al., 2006; Tamietto, et al., 2010); that is, when R/M input including SC input was diminished, so too were residual unconscious abilities. It has also been suggested that the SC plays a critical role in the preserved perceptual responses of TMS-induced blindsight (Ro, et al., 2004). Ro et al demonstrated the preservation of saccadic responses during a TMS-induced suppression of awareness, whereas a manual response task was affected by the application of the TMS (Ro, et al., 2004). Because Ro et al., link saccadic response to role of the SC and manual responses to activity provided by other routes, they concluded that input from the SC is necessary for the residual abilities of TMS induced blindsight.

In contrast to these demonstrations of s-cone and collicular dependency of blindsight, it has been demonstrated that reversible blockade of the LGN is coupled to the reversible elimination of blindsight responses (Schmid, et al., 2010). This may be the clearest demonstration that the residual abilities of blindsight can depend on the LGN rather than the SC. This is also consistent with previous demonstrations that 'spared islands' of striate cortex (in the geniculostriate stream) respond to stimuli presented to the blind field of blindsighted patients (Fendrich, Wessinger, & Gazzaniga, 1992). Recently, it has also been demonstrated that in the context of a TMS-induced suppression of conscious detection, subjects' forced choice discrimination ability showed no preference for red vs. blue stimuli (Railo, Salminen-Vaparanta, Henriksson, Revonsuo, & Koivisto, 2012). If this ability were dependent upon SC input then we might expect a relative impairment in the selection of blue stimuli. The results of Railo et al., therefore provide weak evidence in support of the LGN providing the input for

the residual abilities of TMS-induced blindsight. However, because Railo et al., made no attempt to mask the stimuli in terms of a luminance change, the conclusions that can be drawn are limited.

These conflicting and equally compelling demonstrations of the neural basis for the preserved abilities of blindsight indicate that the debate over the pathways is far from resolved. In the current experiments, applying this s-cone intervention therefore allowed assessment of the role of collicular tracks in the preserved abilities of TMS-induced blindsight.

3.1.4 Synopsis of experimental aims for the on-line TMS experiments

The aims of this series of experiments were first to establish the time course of causally relevant information flow through visual cortical areas in the support of visual consciousness and reportedly 'unseen' perception. Second, the intention was to investigate the contributions made to such processing from sub-cortical pathways, particularly those via the SC and LGN. We therefore sought to investigate the correspondence between early feedforward and unconscious processing, as well as the dependency of conscious processing upon recurrent activity. If this correspondence holds, and the effects are consistent with previous reports, then the disruption of recurrent processing (~100ms) should lead to TMS-induced blindsight. Within the context of TMS-induced blindsight the dependency of the preserved abilities upon input from the SC can be obtained via the use of s-cone stimuli, where – if the collicular hypothesis is correct – 'unseen' abilities should be diminished relative to those in response to luminance stimuli.

3.2 Task adaptations for on-line TMS study

3.2.1 Modifications

The difficulties with the task encountered in the off-line experiments approximate to the possibility of there being some residual awareness that might facilitate discrimination ability when stimuli are reportedly 'unseen', and low incidence of false alarms (highlighted in section 2.8). In an attempt to alleviate some of these difficulties two modifications of the task were implemented.

The first and most important task modification was the introduction of a further question: 'Did you see something?' This was added to limit the impact of the problem, common in blindsight studies, that there is the possibility that some residual awareness of the target stimulus could support 'unseen' discrimination performance, even when the level of awareness is insufficient to elicit a categorical 'yes' response (Campion, et al., 1983). The additional question was explained to subjects as corresponding to lower levels of awareness. They were instructed to respond positively to this question if they saw something that may have been an arrow and/or were unsure (see appendix A4 participant instructions). Alternatively, if the subjects thought the trial to be equivalent to a stimulus-absent condition then they were to respond 'No' to both 'Did you see the arrow?' and 'Did you see something?' In this way, any accessible awareness that the subject holds of the stimulus on a given trial should lead to the exclusion of that trial from the category of 'unseen' discriminations. This experimental manipulation not only reduces the concern that residual conscious awareness could facilitate 'unseen' discrimination judgements; it also allows the subject to register degrees of awareness against which the efficacy of interventions can be compared.

Another reason for the inclusion of the additional 'something?' question was that some subjects also reported a degree of uncertainty during some trials of the off-line TMS experiments. The addition of the 'something?' question is therefore partly intended as a response to this phenomenological lead (Gallagher, 2003; Lutz, 2007). In addition, type II blindsight abilities have been associated with awareness that 'something' occurred in the

visual field without subjects being aware of what that 'something' was (Cowey, 2010; Weiskrantz, 1996). The presence of some vague feeling that something may have been presented in contrast to clear explicit detection of arrows may also be seen as corresponding to Ned Block's differentiation between phenomenal and access consciousness (Block, 2007), which is discussed in section 4.6. The 'something' question was therefore also intended to reflect these potential aspects of blindsight. The inspiration for the addition of this question should, in addition to the suggestions made by the subjects, also be partly attributed to the Perceptual Awareness Scale of Ramsoy and Overgaard (2004), which measures four levels of reported awareness (as opposed to the three included here), but which does not make use of SDT in its application.

| Luminance Left Arrow | S-cone Right Arrow | Noise Stim Absent | Luminance Non-arrow |

Figure 3.2.1.1. Examples of stimuli used, illustrating arrows for luminance and s-cone conditions, a stimulus-absent frame and a non-arrow luminance frame.

The second task modification was the addition of a non-arrow condition to the stimulus-absent condition. These were bars, equal in luminance (or s-cone colour in the s-cone condition) to that of the arrow, but containing no directional information, see figure 3.2.1.1. These control stimuli were presented on half the stimulus-absent trials (25% of all trials). This manipulation increased the number of false alarm responses in arrow detection irrespective of TMS conditions, which has the advantage of increasing the reliability of SDT analyses (Macmillan & Creelman, 1990). Subjects were instructed to respond positively to the 'something' question and negatively to the arrow question if they were aware of this non-arrow being present, as was reflected in the application of SDT (see below and appendix A4). The logic of using the something question in this way firstly ensured that the subjects were

performing the task correctly, i.e. not simply responding to the 'something?' in the same way as the arrow question, nor could they simply be responding 'yes' to 'something?' on all trials, as these response patterns would have prevented them from achieving threshold levels of detection. Secondly, if subjects' criteria were such that they were responding positively to a change in luminance in the range of the arrows' presentation (or the presence of colour in s-cone conditions) rather than when they actually saw an arrow, the presence of the non-arrow, classed as stimulus-absent, would again prevent them from obtaining threshold levels of detection.

The imbalance of stimulus-absent conditions and the now two classes of non-targets further necessitates the use of non-parametric versions of SDT which can cope with such imbalances, thus confirming the use of Pr as the preferred measure of sensitivity and Br as the preferred measure of bias (Corwin, 1994).

There was also an additional modification to the stimuli, which did not relate to matters concerning the off-line experiments. This was that the background colour on which the stimuli were presented was changed from black to the mean grey of the luminance noise stimuli (Smithson, 2003)(25 cd/m^2). This was the same background luminance as used in the s-cone calibration tasks (see 3.3.2).

The other notable difference between the off-line and on-line experiments (other than the mode of TMS) was in statistical approach adopted. Since some of the hypotheses of the on-line section required tests for both variance and invariance between experimental conditions, we employed a combination of orthodox significance tests and complementary Bayesian analyses (Dienes, 2008; Gallistel, 2009). Unlike conventional significance testing, Bayesian analyses can estimate the likelihood of the null hypothesis given data (Dienes, 2008; Rouder, Speckman, Sun, Morey, & Iverson, 2009).

3.2.2 Measures used

The addition of the 'something?' question and the non-arrow stimuli increased the number of SDT measures that could potentially be applied to the task. To minimise the

number of comparisons, the measures to which statistics are applied should be reduced to those that directly correspond to the logic of the task and address the hypotheses.

The primary measure of subjects' conscious awareness of the arrow applied the following responses to the SDT classes as follows:

Table 3.2.2.1. SDT classes for measures of conscious awareness (PrC) and bias in conscious detection (BrC). FA = False Alarms; CR = Correct Rejections. Hit Rate = Hits/(Hits + Miss); False Alarm Rate = FA/(FA+CR); Pr = Hit Rate – False Alarm Rate; Br = False Alarm Rate /(1-Pr) (Corwin, 1994).

	RESPONSE		
STIMULUS	*Something?*	*Arrow?*	*SDT Class*
Arrow	Yes	Yes	HIT
	No	No	MISS
Non-arrow	Yes	Yes	FA
	Yes	No	CR
	Yes	Yes	FA
Nothing	Yes	No	FA
	No	No	CR

The task aimed to quantify the subjects' conscious awareness of the arrow, so 'hits' could only be achieved by reporting awareness of the arrow (and something) in the arrow's presence. Conversely, a 'miss' could be achieved by denying awareness of the 'arrow' and 'something' in the presence of an arrow. 'False alarms' could be achieved by reporting that the arrow was present when it was not, and by reporting that 'something' occurred when no stimuli were present. This meant that even small shifts in bias as reflected by their use of the 'something' question might be reflected in the derived measures. Registering 'correct rejections' also made use of the 'something' question: reporting having not seen an arrow when none was presented was a 'correct rejection', as was responding 'no' to the arrow question and 'yes' to the 'something' question in the presence of a non-arrow. This profile was not classed as a 'hit' (which it was in terms of the non-arrow) because the principal task instruction referred to awareness of the arrow stimulus. There are of course many other ways in which responses could be classified to bring out different aspects of the task, but the above scheme was most closely aligned to the logic and phenomenology of the task. This

classification of SDT was also discussed with the subjects after their completion of the experiment or piloting, and took the descriptions they offered in reference to their experience of the task into account (Gallagher, 2003; Varela, 1996). The classification was established prior to the analysis of the data and used for the primary measure of conscious detection (PrC) and bias in detection (BrC).

The presence of the 'something' question meant that trials in which subjects reported not seeing the 'arrow' and not seeing 'something' could be assuredly categorised as lacking in conscious awareness of the stimuli. The use of these double 'unseen' discriminations to illustrate blindsight type effects means that the responses are more akin to type I blindsight than type II and hence of more utility in revealing the contrast of interest with overt conscious experience. The measure of 'unseen' perceptual ability was taken as proportion correct in terms of the left/right discrimination, when subjects reported both not having seen the 'arrow', nor 'something' (PcU).

These two measures (PrC and PcU) are the principal dependent variables of the following experiments, used to illustrate conscious detection and reportedly 'unseen' discrimination abilities. Fluctuations in secondary measure of response criteria (BrC) are also considered. There are also other measures derived from responses to the task that may shed light on secondary questions, the principal additional measure discussed being subjects' propensity to report awareness of 'something', which is discussed in Chapter 4.6 and appendix A5.

3.3 Between-subjects study, experiment 3A.

3.3.1 Between-subjects introduction

As noted in section 1.2.2, blindsight is the intriguing phenomenon in which observers can successfully identify visual stimuli for which they deny awareness (Weiskrantz, 1986). This condition typically arises from damage to the primary visual cortex (Sanders, et al., 1974) and can also be mimicked through reversible interference using TMS (e.g. Boyer, et

al., 2005). By dissociating visual ability from awareness, blindsight offers a unique insight into the neural basis of consciousness and perception lacking in awareness (Dretske, 2000), yet a consensus view remains elusive. Two central problems remain unsolved. The first, long-standing, question concerns which neural pathways are crucial for vision lacking in awareness (Cowey, 2010). The second, more recent, question concerns the temporal dynamics of vision – whether residual ability in blindsight is supported by early feedforward processing, and whether the later recurrent phase of processing corresponds to conscious vision (Lamme, 2001).

A classic point of contention in blindsight has been whether the residual vision depends on the ancient retinotectal pathway between the eye and superior colliculus (SC), which remains the main visual pathway in reptiles and birds (Cowey, 2010). Information via this route can bypass the visual cortex rapidly to reach frontal and parietal brain regions (Fries, 1984). In primates, however, the collicular pathway has become weakened compared with the abundant pathways from retina to visual cortex via the lateral geniculate nucleus (LGN). Thus, an alternative to the collicular hypothesis is that blindsight is supported by spared fibres or areas within geniculate pathways (Schmid, et al., 2010). Both the collicular and geniculate hypotheses have received empirical support from a range of studies (see section 3.1.3), with the majority in favour of the collicular hypothesis (Leh, Johansen-Berg, et al., 2006; Leh, Mullen, et al., 2006; Leh, et al., 2010; Rafal, Smith, Krantz, Cohen, & Brennan, 1990; Ro, et al., 2004; Tamietto, et al., 2010), cf. (Radoeva, et al., 2008; Railo, et al., 2012; Schmid, et al., 2010; Wessinger, Fendrich, & Gazzaniga, 1997).

In this study we directly contrasted the collicular and geniculate accounts of TMS-induced blindsight by comparing responses to stimuli that do or do not activate the collicular pathway, following the approach introduced by Sumner et al., (2002). As shown in figure 3.1.3.1, the collicular pathway is thought not to receive colour opponent input from short-wave cones (s-cones) of the retina (de Monasterio, 1978; Schiller & Malpeli, 1977; Sumner, et al., 2002). This means that the use of s-cone specific stimuli, when masked by luminance noise, will prevent information from passing via the collicular route.

Recently, the question of how neural processing produces vision has focused on the temporal dynamics of information flow. According to a popular account (Lamme, 2001; Lamme & Roelfsema, 2000), the onset of a visual stimulus triggers a rapid feedforward volley that is sufficient for some degree of unconscious processing; this is then followed

(from ~100ms post-stimulus) by recurrent feedback that facilitates awareness. This theory is linked to the anatomical question of blindsight as it has implicated collicular tracts in the unconscious feedforward sweep.

To determine how these temporal dynamics of visual processing interact with neuroanatomy, not only were visual stimuli manipulated either to activate or to bypass the collicular pathway, but the onset time of occipital TMS was also varied. In the paradigm, active or sham (control) TMS was applied in 25Hz pulse pairs over the occipital cortex (see figure 3.1.3.1 and 3.3.1.1). To modulate the earliest phases of visual processing, TMS was applied at 0-40ms and 40-80ms after stimulus onset. Recurrent processing was targeted by applying TMS at 90-130ms and during a later period at 280-320ms. These four periods were tested to capture possible independent early and late phases (e.g. Corthout, et al., 2002). The target stimulus was an arrow embedded in luminance noise, presented on 50% of all trials (figure 3.3.1.1). On each trial, subjects were asked if they were aware of the arrow stimulus, whether they were aware of 'something' having been presented (see section 3.1.1) and in which direction it was pointing (left or right). Demonstration of blindsight would thus manifest as a TMS-induced impairment of conscious detection while leaving the ability to discriminate reportedly 'unseen' arrows above chance and relatively unaffected by TMS.

If feedforward and recurrent phases of processing map directly onto unconscious and conscious vision (Lamme, 2001) respectively, then we expected early TMS to impair 'unseen' abilities through the selective disruption of feedforward processes, and possibly also to reduce awareness through the transmission of impoverished feedforward information. Later disruption, on the other hand, was predicted selectively to impair feedback, producing the characteristic signature of blindsight in which conscious detection is specifically impaired without influencing 'unseen' abilities.

Both the collicular and geniculate hypotheses predict that conscious detection is more likely to be disrupted by occipital TMS during later periods (e.g. ≥100 ms) than during earlier periods. In addition, both hypotheses predict that corresponding 'unseen' discrimination of luminance stimuli should be spared – resulting in blindsight – because these stimuli are processed by all visual pathways. Where these competing accounts can be dissociated is through 'unseen' discrimination performance for *s-cone* stimuli. According to the collicular hypothesis, there should be no blindsight for s-cone stimuli: whenever TMS impairs conscious detection, 'unseen' discrimination should also suffer because the collicular route is

unable to provide residual information. In contrast, according to the geniculate hypothesis, blocking the collicular pathway should be inconsequential and thus s-cone stimuli should be equally likely to produce blindsight.
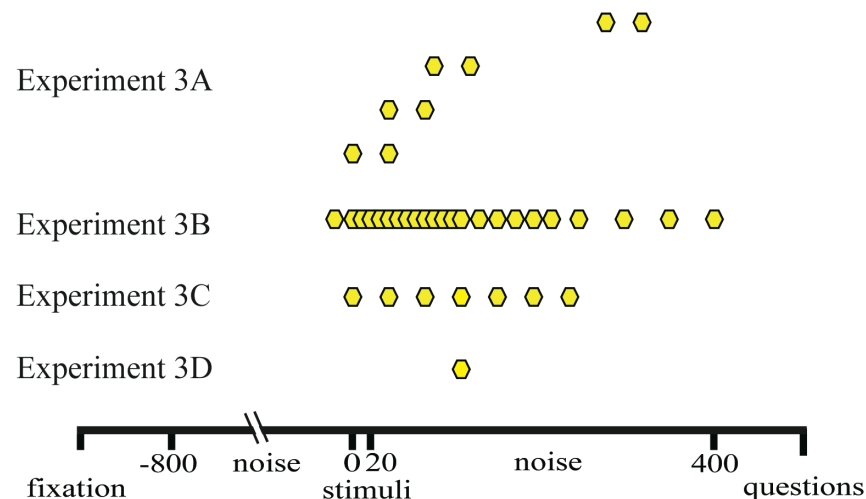
Figure 3.3.1.1. Illustration of the times at which TMS was delivered relative to stimulus presentation in the four experiments presented in this chapter.

3.3.2 Between-subjects materials and methods

3.3.2.1 Subjects

Sixteen neurologically healthy subjects with normal colour vision were recruited (7 male, aged 20 to 32, mean 24). All procedures were approved by the research ethics committee at the School of Psychology, Cardiff University.

3.3.2.2 Stimuli and task

Each trial commenced with a fixation cross (1.5 s) followed by luminance noise, which was produced by alternating the luminance of the bars at 50Hz (figure 3.3.1.1 and 3.2.1.1). The noise continued for 800ms before the onset of the target stimulus (20ms; or noise of equal duration on stimulus-absent trials). Following target offset, the noise continued for a further 380ms (figure 3.3.1.1). The three questions were then asked and responses collected on a standard computer keyboard. The questions and measures are described and discussed in section 3.2.

Two classes of target stimuli were employed (figure 3.2.1.1): *luminance* stimuli, which were defined by an increase in luminance, and *s-cone* stimuli, which were defined by an increase in colour saturation that stimulated s-cones, but not (or minimally) long and medium wave cones (Sumner, et al., 2002; Sumner, Nachev, Castor-Perry, Isenman, & Kennard, 2006; Sumner, et al., 2004). To mask the s-cone stimuli from R/M channels (Bompas & Sumner, 2009; Sumner, Nachev, et al., 2006), all stimuli were presented against a background of luminance noise (mean luminance of noise and background: 25 cd/m^2; range of noise: 10 cd/m^2). The noise occupied a 1.91^o square at fixation, while the arrows and non-arrows subtended 0.81 × 1.43^o. To aid stimulus comparability of stimulus classes, the noise also contained low levels of colour noise along the tritan axis (range: ±8% of mean s-cone saturation of stimuli). Arrow stimuli were present on 50% of all trials, while non-arrows and stimulus-absent trials comprised 25% of trials each. The use of one or other of the two possible non-arrows was randomised.

3.3.2.3 Procedure

Prior to the main experiment, each subject completed approximately four hours of preparatory testing. During this phase, stimulus levels were calibrated to ensure that all subjects performed at comparable levels of detection, subjects were screened for suitability for TMS, and the phosphene threshold was obtained (PT, Franca, et al., 2006). Subjects' individual tritan lines in colour space were also established, enabling s-cone stimuli to be

presented (Smithson, 2003, detailed below). The reason subject-specific tritan lines must be established is that there are individual differences and variance across retinal locations within subjects in the wavelength of light that specifically activate s-cones (Smithson, 2003). Thus, were a non-specific s-cone colour to be used (e.g. 0^0 in MacLeod-Boynton 1979 space) such a colour may have activated additional channels in some subjects.

Subjects then completed six experimental sessions of ~1 hour. Each session consisted of 4 blocks, one for each of the experimental conditions: s-cone stimuli or luminance stimuli combined with sham or active TMS, with the order randomised between sessions. Blocks consisted of 80 trials, including the 4 stimulus types (left arrow, right arrow, non-arrow and stimulus-absent) and 4 TMS onset times (0-40, 40-80, 90-130 and 280-320ms) in a randomised order.

The order of the response questions (Arrow? Something? Left/Right?) was different for each session, so that across the experiment each subject completed a fully counterbalanced set of question orders. Subjects were given ten practice trials at the start of each session to become familiar with the order. The sequence in which these question orders was presented was randomised between subjects.

3.3.2.4 Equipment and TMS procedure

Cortical stimulation was delivered with a Magstim high-power 90mm round coil in conjunction with a Magstim Rapid2 biphasic stimulator. Pulse delivery was controlled via a Cambridge Research Systems (CRS) Visage running Real Time Sequencer software on a Matlab platform, which also governed stimulus presentation on a gamma-corrected 21" Mitsubishi CRT monitor (100Hz vertical refresh rate). Pairs of TMS pulses were applied at 95% of PT (group mean 54.4% of maximum stimulator output ±11.5% SD). The same method for determining PT was used as described for a single analysis block in Experiment 2C, with the exception that there was no time critical aspect to the calibration applied here. This TMS intensity was adopted because higher intensities in pilot experiments caused blinks, twitches and phosphenes that the subjects reported as visible during the task. In pilot experiments, lower intensities and single pulses of TMS were found not to produce reliable effects on performance. The coil was oriented with the handle pointing upward and side 'B'

facing the subject, so that the induced current passed initially in a left-to-right direction. Coil positioning was initially based on anatomical proximity to the calcarine sulcus, localised in individual structural MRI scans. Immediately prior to each active TMS block, the intensity was set to 120% of PT and the coil was moved so that it produced a phosphene that the subject reported as being 'reasonably clear' and 'at least in part, covering the centre of their visual field' with their eyes closed. This procedure usually required ~10 pulses. The coil position was then recorded using a Brainsight system (Rogue Research Inc.) and used for the subsequent block of trials. If the subject moved beyond a 5mm tolerance of the original position, then the block was paused and the coil repositioned to the recorded site. An approximation of this position was used in the sham condition but with the coil perpendicular to the scalp so that the rim pointed toward the head, with a 9mm plastic spacer inserted between the scalp and coil to replicate the contact artefact.

To exclude effects of TMS-induced blinks on performance, eye tracking was undertaken throughout the experiment using a CRS chin-rest mounted infrared eye tracker (250Hz). Trials were excluded on the basis of blinks identified by a shift in pupil position followed by a transitory loss of pupil signal, coincident with the stimulus presentation. Overall, 241 of a total 30,720 trials were excluded from the analysis (<0.8%).

3.3.2.5 Statistical analyses

Effects were assessed using repeated measure ANOVAs with TMS site (2 levels: active and sham), stimulus type (2 levels: luminance and s-cone) and time of TMS relative to stimulus onset (4 levels: 0-40,40-80,90-130 and 280-320ms) as factors, conducted separately on measures of conscious detection (PrC), 'unseen' discrimination (PcU) and bias in conscious detection (BrC). Analyses of simple main effects exploited the relevant Greenhouse-Geisser corrected error terms from the ANOVAs (Winer, Brown, & Michels, 1991), and the sham condition as a baseline. 'Unseen' discrimination performance was assessed relative to chance using single-sample t-tests. All t-tests and analyses of simple main effects applied the Holm–Bonferroni method of correction for multiple comparisons (Holm, 1979).

Complementary Bayesian analyses were based upon models representing the critical hypotheses. The Bayes factor (B) represents the strength of support for the alternative hypothesis (H1) relative to the null (Dienes, 2008), with B>1 indicating evidence in favour of H1 and B<1 indicating evidence in favour of the null. In practice, values of B greater than 3 or smaller than 1/3 are regarded as providing substantial evidence in favour of H1 or the null hypothesis, respectively (Dienes, 2011; Jeffreys, 1961). Bayesian tests were conducted separately on both positive and negative changes from the sham baseline in order to avoid specifying opposing changes within a single model, with the hypotheses represented by a uniform distribution (Dienes, 2008). Since the critical measures were calculated relative to a baseline, zero was selected as the starting point for the distributions. The maximum reasonable shift on both measures was 0.5 and so was used as the upper limit for the theoretical distributions. To assess TMS-dependent effects, the analysis was applied to the change in measures from sham at each level of TMS onset time and stimulus condition. To assess TMS independent effects of stimulus type on 'unseen' abilities we analysed data from the sham condition only. This collapsed data across TMS onset times, and utilised the s-cone condition as the baseline for comparison. This was to ensure that baseline differences in task performance could not explain effects in the presence of the active TMS condition.

To distinguish between the collicular and geniculate accounts we tested for differences in 'unseen' performance between stimulus types, during TMS-induced blindsight (as observed 90-130ms after stimulus onset). For this reason, the comparisons involving 'unseen' discrimination at this time were reported individually (s-cone and luminance relative to sham) and a further analysis was undertaken which directly compared the change in PcU from sham between luminance *vs.* s-cone stimuli. The vector for this Bayesian analysis was thus calculated as: $(\text{Luminance}_{occipital - sham}) - (\text{s-cone}_{occipital - sham})$. The hypothesis for this analysis was represented by a uniform distribution between 0 and 0.5, corresponding to a benefit bestowed by the additional input of luminance stimuli. Also, for these comparisons involving 'unseen' discrimination with TMS at 90-130ms, the complementary t-tests were reported, as were the B statistics representing effects in the opposite direction. Because each comparison involved 16 independent data points (subjects), appropriate standard error adjustments were applied, as recommended by Dienes (2008).

Tests for outlier rejection were applied at a subject level using Chauvenet's criterion (Taylor, 1997). No subjects were excluded using this criterion.

3.3.2.6 Task calibration

Detection thresholds were calculated using a psychophysical method of constants. Following a period of familiarisation with the task, subjects completed approximately ten blocks over a range of stimulus intensities, separately for s-cone and luminance stimuli. For luminance stimuli, the threshold was obtained by adjusting stimulus intensity, whereas for s-cone stimuli the colour saturation was varied along the s-cone axis / tritan line. Data were then regressed using a sigmoidal or linear function (depending on fit quality), solving for the intensity which produced a PrC of 0.5; this corresponded to the stimulus level at which subjects reported consciously seeing the stimulus on 50% of target-present trials if no false alarms were reported. Following the initial calibration, subjects then completed several blocks at the derived intensity and minor adjustments were made to the stimulus levels to compensate for any learning effects, thus maintaining PrC≅0.5. At the start of each experimental session, subjects completed at least one pre-block with sham TMS (40 trials), and similar adjustments were made as required (criteria >±0.1 PrC from 0.5). If subjects completed two sessions in one day, then the mean value for the previous sham block was taken as the starting point for subsequent recalibration. The mean luminance for luminance stimuli was 36.5 cd/m^2 (±3.5 SD). Subject-specific s-cone stimuli were produced following the methods described in Smithson et al (2003) and used in several previous experiments (Anderson, Husain, & Sumner, 2008; Bompas, Sterling, Rafal, & Sumner, 2008; Bompas & Sumner, 2008, 2009, 2011; Sumner, 2006; Sumner, et al., 2002; Sumner, Nachev, et al., 2006; Sumner, et al., 2004). This procedure yielded individual s-cone colours with the following coordinates in CIE 1931 space: Mean (±SD): 0.249 (0.009), 0.186 (0.022), 25.067 (0.809) and is described below.

3.3.2.7 S-cone Calibration

Calibration of s-cone stimuli within each subject followed the methodology outlined by Smithson et al., (2003). In brief, this procedure included a four-alternative forced-choice

detection task to ascertain the hue to which subjects were maximally insensitive during a period of 'transient tritanopia'(Mollon & Polden, 1975): their tritan line. Transient tritanopia is a brief (~2 seconds) period where the retina (and/or retinal ganglion cells) are rendered insensitive to short wave input, following a period of exposure to a yellow-field stimulus and subsequent darker field exposure. During this period, subjects should be selectively insensitive to s-cone input; hence, the elevation of threshold for detecting coloured stimuli should be maximal in the presence of stimuli that specifically activate s-cone channels.

To ensure that subjects were performing this task on the basis of s-cone sensitivity rather than changes in luminance, the luminance of each colour stimulus was adjusted in a separate 'minimum motion task' (Sumner, 2000), designed to equate coloured stimuli with the surrounding grey patches. During this task, bars of s-cone colour and luminance grey were presented vertically above one another, alternating every 50ms, and giving rise to the illusion of motion. Using an up-down staircase to adjust the luminance of the bars, this apparent discrepancy between the coloured and grey bars was minimised, resulting in luminance values that were then used in the production of subsequent s-cone stimuli. The calibration stimuli appeared within the same spatial location as the noise of the experimental stimuli, used the same background luminance, and were presented using the same equipment as the main experiment.

Chromaticity vectors for the initial calibration were set to a range of -15º to +15º, in 5º intervals of MacLeod-Boynton space, where 0º is the theoretical average tritan line for foveal viewing. The procedure commenced by estimating the point of equi-luminance for these seven vectors (-15º, -10º, -5º, 0º, +5º, +10º, +15º standard MacLeod-Boynton 1979 space) using the minimum motion task, including 40 trials for each vector that were randomly interleaved in up-down staircases (step-size of 2cd/m^2). The resulting luminosities were then used to construct the seven stimuli for the estimation of transient tritanopia. This phase involved 25 trials per stimulus, which were again interleaved. The initial adapting exposure to the yellow-field lasted 25 seconds, with a 3 second exposure between trials. Immediately following adaptation, target stimuli were then presented for 250ms in one of four locations. The other target locations displayed a change in luminance within a 2 cd/m^2 range. The vector that produced the most elevated threshold was taken as an approximation of the tritan vector and used for a subsequent more accurate calibration in which three to five vectors were selected around the initial peak and the procedures repeated. If no peak was observed, or if a trend indicated that the peak might lie outside the -15º to +15º range, more eccentric

hue angles were tested. This iterative procedure was stopped when a peak was observed with lower threshold levels ±1° to either side. The final vector was taken as the subject's tritan line and the minimum motion task was repeated to ensure equi-luminance. If no reliable peak could be established, subjects were excluded from the subsequent experiments. Three subjects were excluded in addition to the 16 who participated in the full experiment.

3.3.2.8 Rationale for TMS Pulse Timing

Pilot experiments indicated that pairs of TMS pulses were more effective than single pulses in suppressing perception. In addition, pairs of pulses did not produce reportable phosphenes during the task (consistent with Rauschecker, Bestmann, Walsh, & Thilo, 2004). As noted, the timing of pulses was selected carefully to target early and late stages of visual processing. More specifically, TMS has been reported to disrupt the earliest stages of visual processing from 20ms after stimulus onset (Corthout, Uttl, Walsh, et al., 1999; Koivisto, et al., 2010) motivating the 0-40ms intervention. Although controversial, two temporally distinct suppressions or 'dips' in perception have also been described (Camprodon, et al., 2010; Corthout, et al., 2002; Corthout, Uttl, Walsh, et al., 1999; Corthout, Uttl, Ziemann, et al., 1999; Heinen, et al., 2005; Juan & Walsh, 2003; Koivisto, et al., 2010; Stevens, McGraw, Ledgeway, & Schluppeck, 2009). In order to demonstrate independence of these two periods, an intervening, unaffected period needs to be demonstrated, which motivated in part the inclusion of the 40-80ms intervention. In addition, it has been suggested (Camprodon, et al., 2010; Sack, et al., 2009) that the early feedforward period occurs later than described in some previous studies (Corthout, Uttl, Walsh, et al., 1999; Koivisto, et al., 2010); if so, then the 40-80ms intervention would be expected to capture such later feedforward activity.

The intervention at 90-130ms corresponds to the most clearly replicated period of TMS-induced visual disruption, which peaks at ~100ms (e.g. Amassian, et al., 1989; Kammer, et al., 2005). Since the input from s-cone stimuli is thought to arrive in the occipital cortex slightly later (~20ms) than the input from luminance stimuli (Bompas & Sumner, 2008; Cottaris & De Valois, 1998; Maunsell et al., 1999), and since our aim was to achieve comparable suppression in detection across stimuli types, the mean pulse time for this epoch (110ms) was slightly later than 100ms. Activity within the occipital cortex at ~300ms and

beyond has been widely linked to higher cognitive processes and re-entrant processing (e.g. Chambers, et al., 2012; Liddell, et al., 2004; Sergent, et al., 2005). Relatively few studies have applied occipital TMS as late as 300ms, although the published exceptions demonstrate suppressive effects during similar late epochs (Camprodon, et al., 2010; Chambers, et al., 2012; Heinen, et al., 2005). The 280-320ms intervention was therefore chosen to target this very late period.

3.3.3 Between-Subjects Results

As an overview, conscious detection was suppressed by TMS during the middle and later time periods, whereas 'unseen' performance remained above chance and statistically unaffected, indicating characteristic TMS-induced blindsight. The key result was that s-cone and luminance stimuli produced the same pattern of results during the principal blindsight epoch (~100ms), indicating that the residual abilities of TMS-induced blindsight did not rely on the collicular pathway. Beyond this, it was also demonstrated that the TMS appeared to facilitate rather than impair conscious detection at the earliest time point (0-40ms), for luminance stimuli only. Finally, TMS impaired conscious detection at the latest time point (280-320ms), also only occurred for luminance stimuli.
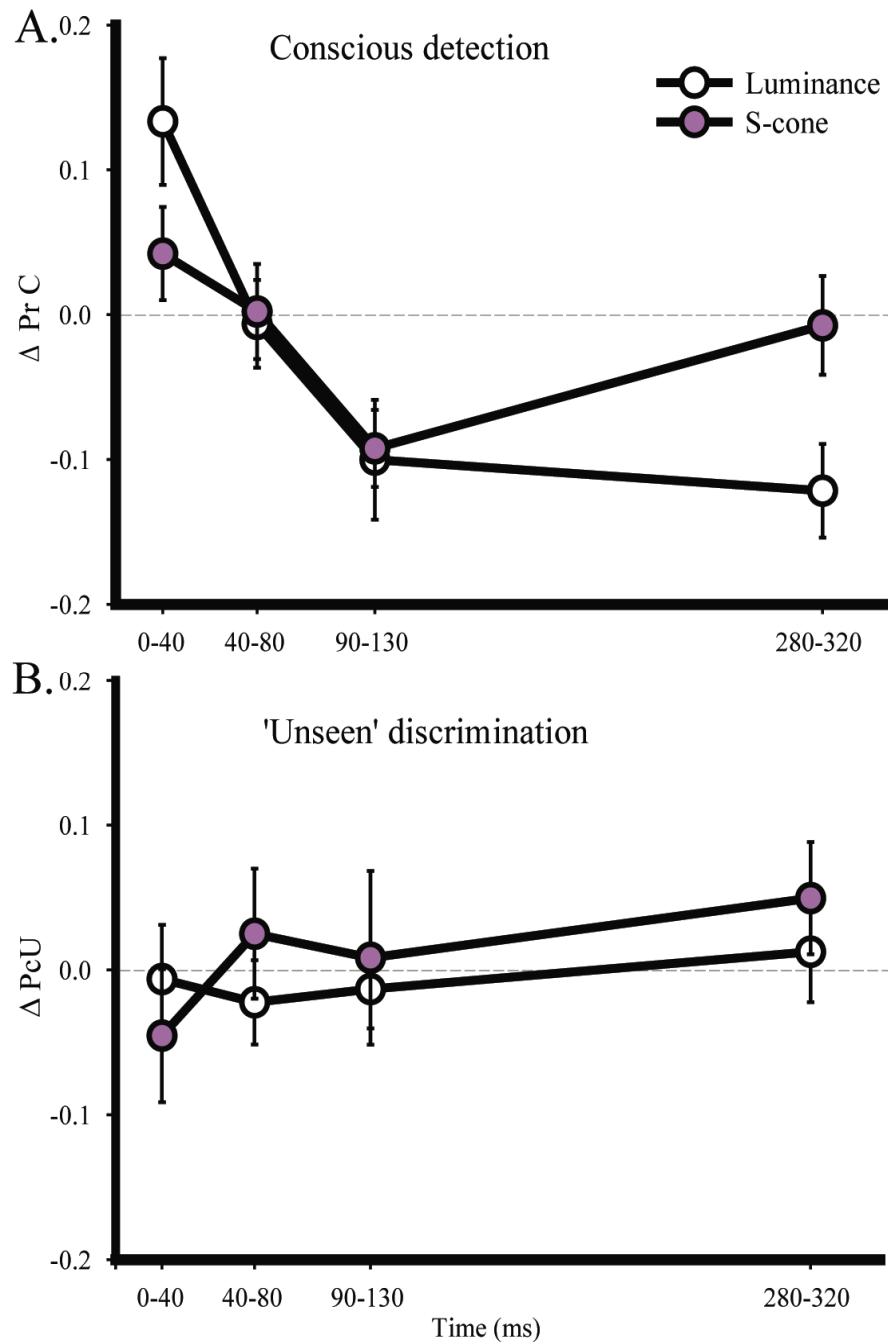
Figure 3.3.3.1. Sham-normalised effects of occipital TMS on (A) conscious detection (PrC) and (B) 'unseen' discrimination (Proportion correct 'Unseen' PcU), plotted according to stimulus type (s-cone or luminance) and TMS onset time. PrC is calculated as a non-parametric equivalent of d' (Corwin, 1994). Error bars are the within-subjects standard error (Loftus & Masson, 1994). Occipital TMS selectively enhanced (0-40ms) or impaired (280-320ms) PrC for luminance stimuli relative to s-cone stimuli, while impairing PrC for both stimulus types at 90-130ms. No significant effect of TMS was observed for PrC at 40-80ms, or for PcU at any TMS onset times.

3.5.3.1 Conscious detection

If conscious processing is supported by recurrent activity then TMS during middle and later time periods (≥100ms) should reduce conscious detection performance. Although no specific predictions were cast in terms of s-cone dependency, such effects are clearly a possibility, given the significant contribution of R/M pathways to cortical input (Nealey & Maunsell, 1994). Indeed, the modulation of awareness by TMS was both time- and stimulus-dependent: a significant three-way interaction ($F_{(3,45)}$=3.3, p=0.028; see figure 3.3.3.1.A) was revealed between factors of TMS condition (occipital, sham), TMS onset time (0-40, 40-80, 90-130, 280-320) and stimulus type (luminance, s-cone).

For luminance stimuli, occipital stimulation significantly *facilitated* conscious detection relative to sham when delivered early after stimulus onset (0-40ms; $t_{(15)}$=5.59, p<0.001), while having no significant effect at a slightly later time (40-80ms; $t_{(15)}$=0.27, p=0.792), and then impairing detection at even later times of 90-130ms and 280-320ms($t_{(15)}$ =4.19, p<0.001; $t_{(15)}$=5.09, p<0.001). For s-cone stimuli, however, occipital TMS significantly reduced conscious detection only when applied at 90-130ms ($t_{(15)}$=4.00, p=0.002; all other t <1.7, p>0.098). Additional comparisons of the change in conscious detection relative to sham revealed that both the early potentiation (0-40ms) and later impairment (280-320ms) were significantly greater for luminance stimuli compared with s-cone stimuli ($t_{(15)}$=2.70, p=0.016. $t_{(15)}$=3.53, p=0.004 respectively). At 90-130ms, however, the impairment in conscious detection did not significantly dissociate between stimulus types ($t_{(15)}$=0.23, p=0.822).

A convergent analysis of TMS-induced changes in conscious detection was undertaken using Bayesian methods. Corresponding analyses of change in conscious detection at each level of TMS onset time and stimulus type, baselined to sham, confirmed an early potentiation (0-40ms) and late impairment (280-320ms) that was specific to luminance stimuli (B=8.96 and B=169.49, respectively) and not observed for s-cone stimuli at equivalent times (B=0.31 and B=0.11). At 90-130ms, occipital TMS also impaired conscious detection for s-cone stimuli (B=59.98) and less robustly for luminance stimuli (B=1.58). All other comparisons of conscious detection between occipital TMS and sham revealed

evidence in favour of the null (all B<0.33). This outcome aligns with the orthodox analysis, indicating that the early (0-40ms) and late (280-320ms) modulation of conscious detection was dependent on inputs via the R/M pathways, whereas the intermediate period of disruption (90-130ms) was pathway-independent.

3.3.3.2 'Unseen' discrimination

The key question in this study is whether above-chance 'unseen' discrimination is maintained in the context of the TMS-induced impairment of conscious detection – i.e. whether TMS-induced blindsight occurred – and whether these 'unseen' abilities are stimulus independent, as predicted by the geniculate hypothesis, or differentially impaired in the presence of s-cone *vs.* luminance stimuli, as predicted by the collicular hypothesis. Furthermore, a TMS-induced early suppression of unseen abilities, compared to later time periods would support the role of the feedforward sweep in unconscious processing.

Discrimination performance on 'unseen' trials was consistently above chance at all 16 combinations of TMS condition, TMS onset time, and stimulus type (mean proportion correct = 0.71, ± 0.15 SD, all $t_{(15)}$ >3.14, p<0.007 with Holm–Bonferroni correction), indicating that the measure tracked a form of perception. In contrast to the analysis of conscious detection, a three-way ANOVA of 'unseen' discrimination revealed no significant interaction of TMS condition × TMS onset time × stimulus type ($F_{(3,45)}$ =0.36, p=0.75; figure 3.3.3.1.B), and no significant main effects or lower-order interactions (all F<2.37, all p>0.142). Bayesian analyses of the difference in 'unseen' discrimination between occipital TMS and sham supported the null hypothesis of equivalence in all conditions (all B<0.35). Since this analysis included any potential early suppression of 'unseen' abilities, these results do not directly support a correspondence between early occipital processing and unconscious vision (Lamme, 2001, 2006b).

During the critical time-period where conscious detection was suppressed for both stimulus types (90-130ms), 'unseen' ability was unaffected by occipital TMS in the presence of both s-cone stimuli ($B_{occipital < sham}$=0.18, $B_{occipital > sham}$=0.15, $t_{(15)}$ =-0.14, p=0.89) and luminance stimuli ($B_{occipital < sham}$=0.07, $B_{occipital > sham}$=0.12, $t_{(15)}$ =0.39, p=0.70). Additional comparisons between sham-normalised discrimination performance (occipital – sham)

confirmed no reliable difference between stimulus types ($B_{luminance > s-cone}$=0.16, $B_{luminance < s-cone}$=0.24; $t_{(15)}$=0.40, p=0.77). Moreover, the strongest trend in this comparison was a tendency for *reduced* performance to luminance over s-cone stimuli (B=0.24), contrary to the pattern predicted by the collicular hypothesis. Taken together, the findings for 'unseen' discrimination ability are thus consistent with the characteristic signature of blindsight in which conscious detection is selectively impaired while 'unseen' discrimination performance is spared.

It is notable that PcU is calculated from 'unseen' trials only, whereas more trials contribute to the measure of conscious detection (PrC calculated from a total of 22,275 trials, whereas PcU was calculated from 5,525 trials). Thus, it could be argued that the analysis of PcU is hindered by a relative lack of statistical power. To test whether this was the case, an additional analysis was undertaken. The estimates of standard error based on these specific trial numbers (the number of data points convolved for the number of contributing trials) can be compared in order to test whether heterogeneity of variance could be responsible for the observation of significant effects on one measure and not another. Keppel (1982) recommends a 3:1 ratio of such variances (Fmax ratio) as the cut-off point, beyond which the variances should be considered heterogeneous and the conditions non-comparable. For the data, the Fmax ratio was calculated as 1.84:1 (0.032/0.017), thus satisfying this requirement.

If 'unseen' abilities were dependent upon input from the SC without disruption of the occipital cortex, then this might negate the demonstration of collicular independency, because the difference indicating a role for the collicular pathway would be present in the baseline. In order to test whether this was the case we contrasted s-cone and luminance stimuli in the sham condition the results of which did not reveal any such difference ($t_{(15)}$=1.51, p=0.15, B=0.39). There was a weak non-significant trend, observed in the sham condition, for higher 'unseen' discrimination performance in the presence of luminance *vs.* s-cone stimuli; this suggests that R/M pathways could contribute to 'unseen' abilities when there is no disruption to early visual areas. However, given that this potential difference occurred in the presence of clearly above-chance performance when of R/M inputs are blocked, general (non-blindsight) 'unseen' abilities are unlikely to be crucially dependent upon such input. Furthermore, any potential difference between 'unseen' performances across stimuli classes in the sham condition could (speculatively) be attributed to the observed criterion difference between stimuli types: lower performance on 'unseen' trials, in the presence of s-cone stimuli, may be expected if the criteria subjects used to identify 'unseen' trials are more stringent. That is, the

analysis of bias (3.3.3.3 below) demonstrated subjects to be more likely to report s-cone stimuli as 'seen' independently of what has been presented to them. This meant that when they reported s-cone stimuli as 'unseen' they may have been doing so on the basis of less information than was the case for luminance stimuli. This possibility therefore highlights the necessity of baselining active TMS conditions to control conditions, such as sham, because doing so cancels out the bias difference between stimuli types.

3.3.3.3 Response Bias

It is possible that the pattern of behaviour that characterises blindsight could potentially be attributed to changes in response bias together with there being a disruption of 'normal' vision (Campion, et al., 1983; Nisbett, 1977). To alleviate these concerns we applied signal detection theory, which provides a measure of conscious detection that is essentially independent of response criteria (Pr), and a measure of response bias (Br) that is independent of sensitivity (Corwin, 1994). The analysis of Br revealed no higher level interactions, but there were main effects: subjects were more prone to respond positively in the presence of occipital over sham TMS (site effect: $F_{(1,15)}$ =5.235, p=0.037; replicating Rahnev, Maniscalco, Luber, Lau, & Lisanby, 2011), s-cone over luminance stimuli (stimuli type effect: $F_{(1,15)}$ =16.642, p=0.001) and at earlier TMS onset times (time effect $F_{(3,45)}$ =14.931, p<0.001, figure 3.3.3.3). The reduced propensity to respond 'yes' when TMS was applied at later times, independently of coil orientation (occipital or sham) is consistent with the effects of auditory distraction. Likewise, the stimulus-dependent effects may reflect differences in subjects' attribution of the stimulus noise across the stimuli conditions. The posing of experimental questions can affect subject's experience of the task, independently of stimuli conditions. This is the observer paradox (Labov, 1972). The application of active TMS is equally capable of producing such a reactive (Ericsson, 2003) effect, potentially independently of direct neuronal effects. This therefore could explain the increased propensity to respond positively under the active condition. These contingent bias effects are relatively uninformative and would have been conflated with effects on veridical conscious perception were it not for the application of SDT. These effects are considered in greater depth in section 4.3.
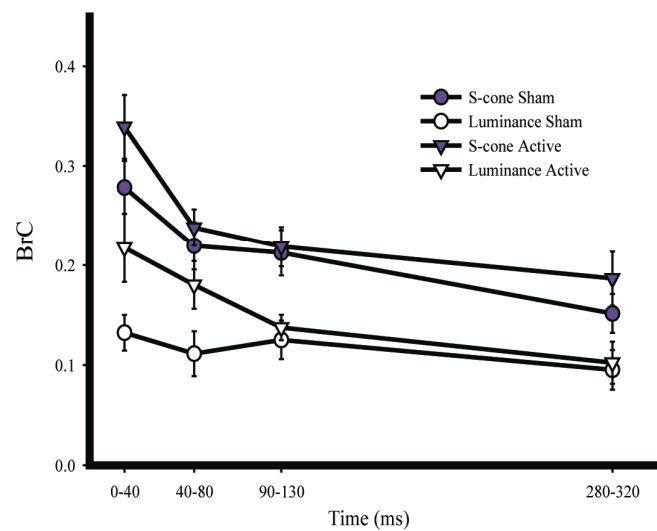
Figure 3.3.3.3. Bias in conscious detection (BrC) for all TMS and stimulus conditions. Higher values indicate greater propensity to respond 'yes'. Error bars are ±1 SEM (Loftus & Masson, 1994).

3.3.4 Between-Subjects Discussion

The results of Experiment 3A suggest that the residual abilities of TMS-induced blindsight shown here are predominantly geniculate (Schmid, et al., 2010) rather than collicular (e.g. Ro, et al., 2004) in origin. Blindsight was demonstrated most clearly by the suppression of reported awareness at 90-130ms for both stimulus types, in the presence of above chance 'unseen' discrimination. This corresponds to the classically reported epoch of visual suppression by occipital TMS (Amassian, et al., 1989). The central question posed was whether or not the concurrent 'unseen' abilities were dependent upon stimulus type, i.e. whether or not the residual abilities were maintained when input via the SC was blocked. The preservation of such abilities in the presence of s-cone stimuli indicates that the residual abilities in this case of TMS-induced blindsight are supported by chromatic geniculate pathways rather than the collicular route. Moreover, it seems unlikely that the absence of stimulus-specific effects on the measure of 'unseen' discrimination was due to the stimuli not isolating the R/M pathways or the TMS being ineffective, because both TMS effects and

stimulus specificity were reliably demonstrated upon the contrasting measure of conscious detection.

How can the previous evidence in favour of the collicular hypothesis be reconciled with the current evidence in support of the geniculate hypothesis? Much of the evidence consistent with collicular mediation does not in fact logically exclude a geniculate role, but in patients where the geniculate pathways are surgically cut the evidence for collicular mediation is particularly strong (Leh, Johansen-Berg, et al., 2006; Leh, et al., 2010; Tamietto, et al., 2010). Our results are not inconsistent with a role for the collicular pathway in certain kinds of residual ability in some contexts, particularly those following permanent brain injuries where plasticity may alter the functional contribution of different pathways (Huxlin, 2008; Mittmann & Eysel, 2001). Yet, our findings do negate the claim of collicular necessity (Ro, et al., 2004). Rather than exclusively supporting the geniculate hypothesis and disconfirming the collicular hypothesis, the position taken is that the data disconfirms the core idea implicit in the long debate between these theories – that because blindsight has been categorised as a single phenomenon we should expect a single pathway to support it. Different pathways are likely to support unconscious abilities that depend on different types of information (Danckert & Rossetti, 2005).

Although the results provide no evidence for collicular dependency of unconscious abilities, R/M pathways did make a time-specific contribution to conscious vision. Interestingly, the earliest effect we observed (0-40ms) was a facilitation of conscious detection. Early TMS has been previously reported to cause either suppression (Corthout, Uttl, Walsh, et al., 1999; Corthout, Uttl, Ziemann, et al., 1999) or enhancement (Abrahamyan, et al., 2011; Schwarzkopf, et al., 2011) in perception. Such enhancement effects have been attributed to the addition of neuronal noise, producing stochastic resonance and driving cortical activity across a threshold for detection (see Schwarzkopf, et al., 2011). Earlier potentiation and later suppression of functionality can be further understood if the role of the early period is seen as preparatory (Marr, 1982), potentially in the service of later conscious processing. Because the state of activation leading up to the arrival of information is crucial in determining whether or not that information is perceived (Mathewson, Gratton, Fabiani, Beck, & Ro, 2009; Romei, et al., 2010), the application of early TMS may thus provide a pedestal of activity, priming the region for subsequent processing. This might not be the case for later applications because increasing activity without interfering with information is only possible when TMS is applied *prior* to the arrival of that information.

Here this rapid facilitation was only realised in the presence of luminance stimuli, indicating that this early period was dependent upon fast (Bompas & Sumner, 2008; Cottaris & De Valois, 1998; Maunsell, et al., 1999; McKeefry, Parry, & Murray, 2003) input provided by R/M pathways.

The observed disruption of conscious detection at later TMS onset-times is consistent with existing proposals (Lamme, 2001) that delayed re-entrant processing supports conscious vision. The latest of these effects (280-320ms) occurred only in the presence of luminance stimuli, indicating that information conveyed via R/M pathways supports this later period. What then provides this late R/M input? Our findings do not provide a definitive answer for this question. However, it seems unlikely to be dependent upon feedforward or recurrent activity that has first passed through the affected early visual cortex, because s-cone inputs are thought to be integrated with other signals within these regions (Cottaris & De Valois, 1998; Sincich & Horton, 2005). It is more plausible that R/M information is first passed through unaffected, more rostral, areas – possibly via the aforementioned extrastriate pathways – and is then fed back to interact with feedforward processing. In this way, the later stage of conscious processing may have been susceptible to later TMS because it is supported by R/M inputs that initially enter frontoparietal networks (Corbetta, Patel, & Shulman, 2008) before later feeding back into the disrupted occipital cortex.

The differential directions of early and late effects upon conscious detection, together with the Bayesian analysis indicating an intervening (40-80ms) absence of effects, suggests that there are multiple periods during which TMS can affect perception. However, because the first period is expressed by an increase in the measure, the existence of more than a single 'dip' or period over which TMS directly interferes with responses to arrows, is not supported by this data. This issue is discussed further in section 3.5.3.

In conclusion, the residual visual abilities of TMS-induced blindsight shown here originate primarily from spared geniculate projections. These findings thus stand in contrast to the popular explanation of blindsight that relies on the collicular pathway (Ro, et al., 2004), instead adding weight to evidence that spared regions in the geniculate pathway can support unconscious vision following occipital disruption (Schmid, et al., 2010). The results also indicate that conscious vision depends on activity within the early occipital cortex from ~100ms after stimulus onset, and that R/M pathways make a contribution to awareness during early preparatory (0-40 ms) and later feedback (280-320 ms) stages of visual processing.

3.4 Single-subject introduction

Experiments 3B, 3C and 3D were developed to address very similar questions to those posed in Experiment 3A, but with higher temporal resolution than is feasible with multiple subjects. In terms of their development over the course of this research, it is worth noting that these experiments were originally designed as between-subject studies and were initially piloted on three subjects. However, it soon became apparent that applying this quantity of testing to multiple subjects was impractical, mainly owing to subject availability and their willingness to participate in experiments with such a prolonged exposure to TMS. The experiments were therefore adjusted to become single subject case studies where the subject underwent approximately 50 hours of testing, conducted over 1 to 3 hour sessions, occurring roughly weekly. Data were collected concurrently for the between-subject and single-subject studies.

The single subject study comprised three experiments. The first, Experiment 3B, aimed to establish a high-resolution time course for conscious and reportedly 'unseen' perceptual processing. This involved the application of single pulses (hence described as the single-subject single-pulses experiment) of TMS at 24 times relative to stimulus onset (see figure 3.3.1.1), designed to cover the range of previously reported effects from very early effects ~ 0ms (e.g. Corthout, et al., 2003) to late recurrent effects ~350-400ms (e.g. Chambers, et al., 2012), with the exception of the effects of TMS when applied 70-50ms before stimuli onset that have been attributed to blinks (Corthout, Uttl, Walsh, et al., 1999). The primary question was whether the data would conform to a 'single dip' in conscious detection as has been reported (e.g. Amassian, et al., 1989), or whether it would be better described by two (e.g. Lamme & Roelfsema, 2000) or more (e.g. Corthout et al 1999 -2002) periods of suppression, as has also been reported. This experiment also served to provide a Blindsight Inducing Pulse (BIP) which was exploited in subsequent sections. Once a reliable method for demonstrating TMS-induced blindsight had been established (the BIP) it was then possible to apply additional manipulations to study the role of different phases of activity and pathways in relation to TMS-induced blindsight.

The second of the single-subject experiments (3C) applied the BIP and additional TMS pulses before and after the BIP (described as the double-pulses experiment). This

permitted analysis of different phases of activity in the provision of the residual abilities of TMS-induced blindsight. Specifically, it has been proposed that the residual abilities of blindsight correspond to an early feedforward sweep of activity prior to the conscious recurrent phase (Lamme, 2001), presumably being disrupted by the BIP. If so, the expectation would be that application of additional pulses *prior* to the BIP to cause a greater suppression of 'unseen' abilities than comparable pulses applied *after* the BIP (Lamme, 2006b).

The third section was an s-cone experiment (3D), where performance was compared between stimuli types in the context of the BIP. This employed the same stimuli manipulation as in the between subjects study (3A), isolating the SC through the use of stimuli which do not activate the SC in a colour opponent manner. If the residual abilities of TMS-induced blindsight depend on input from the collicular pathway, then we would expect there to be a reduction or elimination of 'unseen' abilities in the presence of s-cone stimuli relative to luminance stimuli.

3.5 Single-subject single-pulse, experiment 3B.

3.5.1. Single-pulse methods

3.5.1.1 General methods

Methods are as in the between subjects study unless otherwise stated.

The subject was male, aged 28 - 29 at the time of the experiment, with normal vision and no known neurological abnormalities. He provided informed consent to participation in this experiment, which was approved by the ethics committee of the School of Psychology, Cardiff University.

Prior to experimental sessions the subject's conscious detection performance was calibrated to Pr = 0.6 in the presence of luminance stimuli. Ability to discriminate arrows above chance when reported 'unseen' was assessed through the concatenation of later practice and calibration blocks and application of χ^2 tests – in order to ensure the possibility of demonstrating blindsight type effects. Structural MRI scans and motor threshold (based on the methods set out by Rossini et al., 1994) were collected on separate sessions. The tritan-line and susceptibility to phosphenes were also assessed prior to experimental sessions.

The subject completed 12 experimental sessions, each involving 5 blocks of active TMS and one of sham, the order of which was counter balanced across sessions. Each session lasted approximately 2 hours and each block consisted of 100 trials. TMS was applied at 24 time conditions relative to stimulus onset (-20, 0, 10 ,20 ,30 , 40, 50, 60, 70, 80, 90, 100, 110, 120, 130, 140, 160, 180, 200, 220, 250, 300, 400 ms) and one condition where TMS was not applied, yielding a total of 25 TMS conditions. As mentioned, there were four stimulus conditions (arrow left, arrow right, non-arrow and stimulus-absent). Each block thus contained a full set of stimulus and TMS conditions, the order of which was randomised. The order in which the three questions were presented was counterbalanced between sessions. At the start of each session the subject was given a block of 40 practice trials in order to re-familiarise himself with the task and the new question order. The practice blocks and sham blocks were used to re-calibrate over the course of the experiments. That is, if the subject's detection behaviour was outside a tolerance of $\pm$ 0.1 of the Pr value of 0.6, the stimuli were adjusted accordingly, so that detection performance might be maintained at around 0.6 in the subsequent sessions. This meant that the luminance of the target ranged from 53 to 59 cd/m^2 with a mean of 56.5 ($\pm$1.9 SD).

The TMS intensity at which pulses were applied during the experiment was set at 120% of PT (70% of maximum stimulator output, the subject's PT being 58% of output) for the pulses applied during the experiment. This intensity did not appear to cause facial twitches when applied to occipital areas and rarely caused blinks in this subject. Eye tracking flagged 43 trials for exclusion from a possible 6000 active TMS trials, on the basis of blinks concurrent with visual stimuli exposure.

3.5.1.2 Statistical methods

Analyses were applied to single SDT or proportion correct data for each time point, resulting in a variable representing the time course of the effects of TMS on the measure of conscious detection (PrC) and 'unseen' discrimination (PcU). The main comparisons of interest were between active and sham TMS for both PrC and PcU measures. Comparisons were made by analysis of where data points sat in relation to the variance of the corresponding state of contrast: equivalent to z-tests for each time point. Therefore, if TMS effects upon conscious detection were to be assessed relative to sham, then z-scores were produced from PrC for each time point with active TMS, relative to variance in task performance across all points with sham TMS. These z-scores were converted to p values, which were assessed for significance relative to a Holm-Bonferroni adjusted alpha level (with 25 time points the most stringent p value required to reject the null hypothesis was < 0.002). Points that pass these criteria are denoted as significantly different (S) in figure 3.5.2.1.

Demonstration of 'unseen' discrimination ability (PcU) was provided by assessment of the subjects' forced choice proportion correct performance, coincident with reporting not seeing the 'arrow' nor 'something', relative to a binomial distribution. As with PrC, Holm-Bonferroni correction for the 25 time point comparisons was also applied. Note that this differs from the single sample t-test methodology applied in the between-subjects study because the absence of multiple mean values at each data point mandates the use of an alternative binomial test. Bayesian analyses were not applied to this data set because of the absence of variance information at each data point / condition.

Gaussian models were fitted to the data, using the curve fitting toolbox of MATLAB (Mathworks Inc) and Sigmaplot (SPSS Inc). These models described one or two distinct epochs as single and double peaks. Quality of fit was used to assess models on the basis of least means squared regression values.

The BIP was determined as being the minima of the lowest drop in conscious detection modelled. This value was rounded up to the closest CRT refresh (100Hz) in order to avoid the presentation artefact caused by the electromagnetic pulse of TMS when applied mid-frame. Another reason for the BIP being rounded *up* was because the input provided by s-cone stimuli is relatively slow (Bompas & Sumner, 2008; Cottaris & De Valois, 1998;

Maunsell, et al., 1999). Hence in order that the BIP that was comparable in terms of its effect across stimuli types, a slightly later timing for the BIP (120ms), was used than would have been chosen for purely luminance stimuli (110ms, see below).

3.5.2 Single-pulse results

3.5.2.1 Conscious detection

Active TMS, relative to sham, reduced conscious detection when applied at 10 times between 50 and 160ms after stimulus onset. The chance that these were drawn from the same population as the sham results is less than 0.2% (2.8782 standard deviations; figure 3.5.2.1).
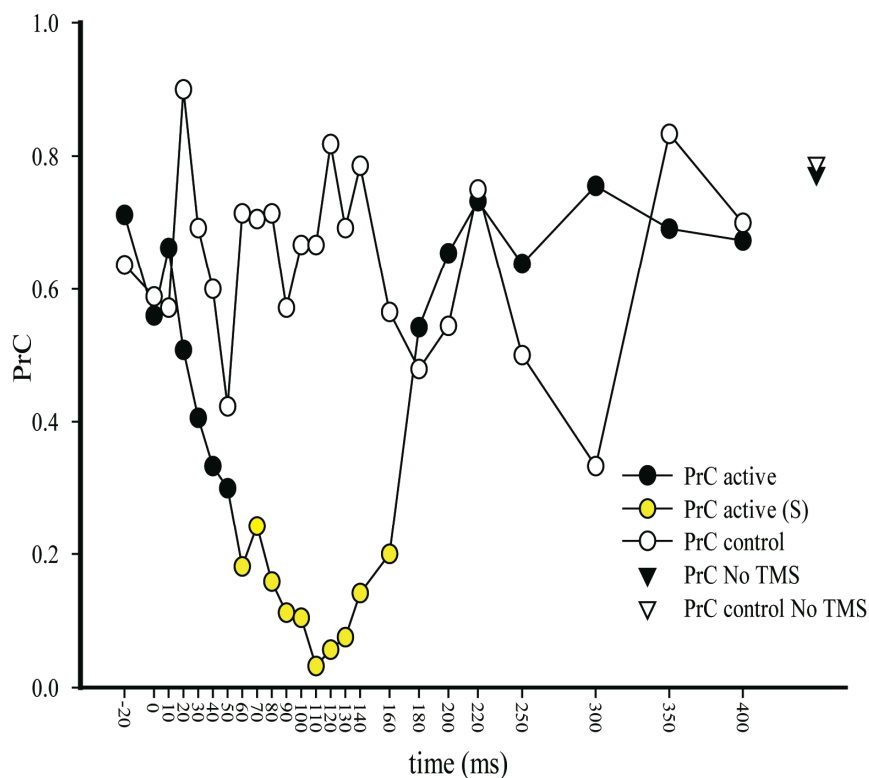
Figure 3.5.2.1. Conscious detection performance (PrC) in Experiment 3B over the range of times where TMS was applied relative to stimulus onset, under both active and sham (control) TMS conditions. Yellow points indicate times which significantly differ from the mean of the sham detection measure (S).

The simplest best fit to the conscious detection data is a Gaussian model expressing a 'dip' (Adjusted R^2 = 0.928, $F_{(3,20)}$ =99.314, p<0.001, see figure 3.5.2.3). The model is specified in equation 3.1. The fit for a double Gaussian model was not as good (model specified in equation 3.2, adjusted R^2 = 0.875)

$$Y = Y_O + ae^{\left[-0.5\left(\frac{X-X_O}{b}\right)^2\right]}$$

eq 3.1

Coefficients a =-0.70, b= 55.05, Xo=102.5 and Yo = 0.74

$$Y = Y_O + a1e^{\left[-0.5\left(\frac{X-X_{O1}}{b1}\right)^2\right]} + a2e^{\left[-0.5\left(\frac{X-X_{O2}}{b2}\right)^2\right]}$$

eq 3.2

Coefficients Yo =0.74 a1 =0.87, b1= 140, Xo1=306.4, a2 =0.81, b2=71.1 and Xo2 = -32.48
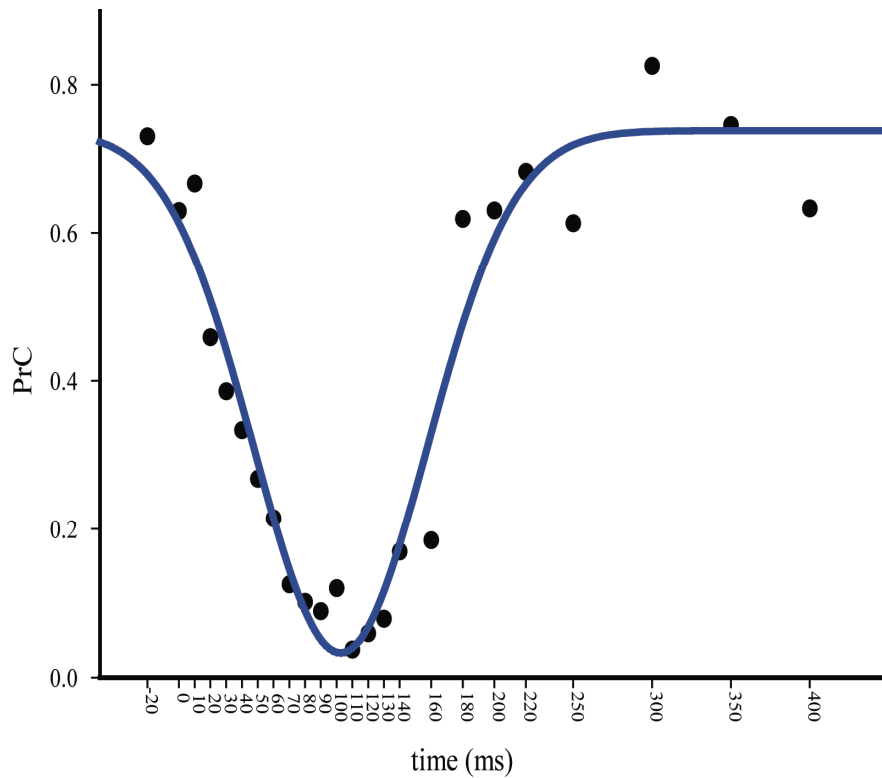
Figure 3.5.2.3. Conscious detection measure over time fitted with a Gaussian model (blue line, eq 3.1). Adjusted $R^2 = 0.9276$, $F_{(3,20)} = 99.3136$, $p<0.001$.

3.5.2.2 'Unseen' discrimination
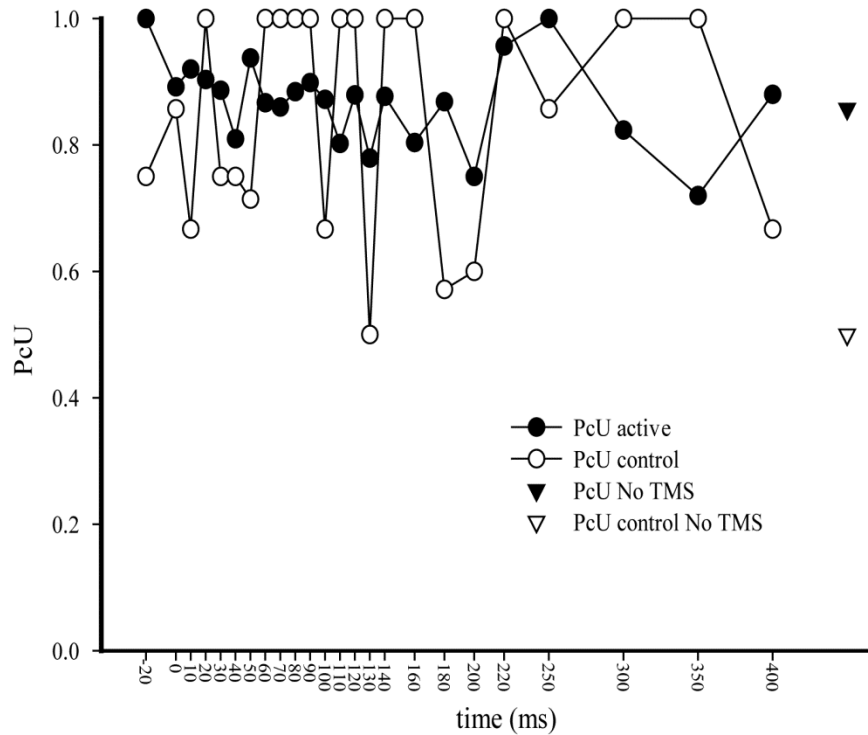
Figure 3.5.2.2. Proportion correct in 'unseen' discrimination (PcU) in Experiment 3B over the range of times where TMS was applied relative to stimulus onset, under both active and sham (control) TMS conditions. No data points were shown to be significantly affected by the active TMS.

The subject displayed above chance 'unseen' discrimination ability at all time points (2-tailed binomial $t_{(>16)}>2.4$, $p<0.024$), all of which survived correction for multiple comparisons. The same analysis as was applied to PrC, when applied to PcU 'unseen' discrimination, revealed no data points that were affected by the TMS according to these criteria (figure 3.5.2.2). Curve fitting using the Gaussian models failed to produce reliable fits to the 'unseen' discrimination data. However, a linear fit indicated a trend indicating lower

'unseen' abilities when TMS was applied at later times (the gradient of a linear fit being -0.25 PcU units per second, adjusted $R^2 = 0.118$, F of regression $_{(1,22)} = 4.08$, p = 0.056). Although only marginally significant the trend is worthy of note as it is consistent with and relevant to the double-pulse data discussed in the next section (Experiment 3C, section 3.6).

3.5.2.3 Response bias

Analysis of bias was impractical in this experiment as the subject only made 15 false alarms (in contrast to 1678 correct rejections) over the 25 time conditions. Attributing effects on the basis of the distribution of these 15 trials was considered invalid.

The Fmax ratio based on contributing trial numbers which compared the variance of PrC to PcU across times was 1.09. However, it can be argued that the application of Fmax here is not valid, as the variance of PrC is disproportionately increased (reducing the ratio) owing to PrC's susceptibility to the TMS. The alternative comparison of variance, at each data point, is not possible.

3.5.3 Single-pulse discussion

This section demonstrates that single pulse TMS can cause an unambiguous suppression of conscious detection of stimuli (see figures 3.5.2.1). As with the between-subjects study, Experiment 3A, this is consistent with the classically reported occipital suppression at ~100ms after stimulus onset (Amassian, et al., 1989). Concurrent above chance 'unseen' discrimination ability was maintained, indicating specificity of the suppression to the measure of conscious detection and, thus, a blindsight type effect.

This suppression is best described as a 'single dip', as all the time points which statistically illustrate this suppression are grouped over a single period of TMS interference and the fit provided by a single Gaussian model was extremely high. Given this consistency in the data pattern, it seems reasonable to suggest that if there were two or more distinct and separable periods of TMS efficacy, then this should have been apparent within this data set.

Although the conclusions that can be drawn from single subject data are limited, there is a degree of inconsistency in the literature over the existence of two or more discrete periods of TMS interference, lending weight to this interpretation. What follows is a brief synopsis of experiments that have previously claimed to show two or more distinct periods of TMS efficacy when applied to visual cortical areas.

Juan and Walsh (2003) appear to show that TMS applied from 0 to 120ms was effective in reducing perceptual abilities (p 0.01-0.05) as was applying pulses at 200 and 240ms (p=0.053) but they suggest that TMS applied at 140 and 180ms was ineffective. However, these results are difficult to interpret mainly because of the absence of statistical information that could have supported the claim of two distinct periods. No corrections for the multiple comparisons were applied (of which there were at least six capable of informing the experimental hypothesis). The 200-240ms intervention did not show a clear suppression, and, more importantly, no information was presented regarding any changes due to the 140-180ms intervention, upon which demonstration of two independent periods rests.

As mentioned there has been a series of experiments, carried out by Corthout and colleagues, in which at least four independent periods during which TMS applied to visual areas was reported as causing disruption in a letter discrimination task. Some of these periods are not relevant here and others seem to be inconsistent across subjects to the extent that their existence is questionable. The first of these periods of sensitivity occurred 70 to 50 ms prior to the onset of the stimulus (hence irrelevant to this study) and was attributed to TMS induced blinks, as this effect was independent of coil position (Corthout, Uttl, Walsh, et al., 1999; Corthout, Uttl, Ziemann, et al., 1999). The second and third 'dips' are reported as having occurred at approximately the same time as the stimuli (-10 to 0ms) and at 20ms post stimuli onset. The earliest of these was proposed in the 2003 paper (Corthout, et al., 2003). However none of the three participants showed a dip around -10ms to 0ms that was clearly independent of either the earliest dip, prior to stimuli onset or the later 20ms dip. The existence of the 20ms dip appears to be slightly more consistent as it was reported in four separate papers (Corthout, et al., 2002, 2003; Corthout, et al., 2000; Corthout, Uttl, Walsh, et al., 1999), however it is difficult to assess the extent to which data and participants overlapped between papers. The statistical methodology for testing the existence of this 'dip' only appeared in one of these papers (Corthout, et al., 2002) which applied a t-test to the difference between the 20ms dip relative to a 50ms period which appeared to be unaffected. Why precisely these

periods were chosen for entry into the analysis is unclear, as are the reasons for which the authors suggest that the 20ms period is independent of the period of TMS efficacy prior to the stimulus onset. Furthermore, across these papers the demonstration of the 20ms dip and two independent periods of TMS efficacy was expressed by the data pattern in, at best, three out of four subjects (Corthout, Uttl, Walsh, et al., 1999). In all of these experiments the classic ~100ms disruption was reported, highlighting the fragility of the demonstration of secondary early periods of TMS efficacy in causing suppressions.

Several studies using moving stimuli have demonstrated two distinct periods of TMS efficacy, which are difficult to reconcile with the results presented here (Pascual-Leone & Walsh, 2001; Silvanto, Cowey, Lavie, & Walsh, 2005; Silvanto, Lavie, & Walsh, 2005; Stevens, et al., 2009). For example, Koivisto et al (2010) reported two distinct periods in which awareness of moving stimuli was disrupted by TMS applied to early visual areas at 20ms and at 60ms as well as showing a intervening period at 40ms where TMS applied to V1/V2 was ineffective, but was effective when applied to motion sensitive area V5/MT – indicating recurrent feedback[7]. Although there are statistical concerns (regarding the absence of reporting effect sizes corresponding to the intervening period), the reproducibility of the 'double dip' when tasks involving moving stimuli are used does support the existence of separable periods.

In contrast to studies which have claimed to demonstrate two distinct periods of TMS disruption, more studies have shown a single period during which TMS applied to early visual areas is effective in causing disruptions replicating the classic study of Amassian et al (1989), (e.g. Kamitani & Shimojo, 1999; Kammer, 2007; Kammer, et al., 2005; Koivisto, et al., 2012; Koivisto, et al., 2011; Overgaard, et al., 2004; Sack, et al., 2009; Thielscher, et al., 2010). Most relevant to the data described here are those studies that have used discrimination and/or detection of arrows in their demonstration of perceptual abilities and have shown a single period of TMS efficacy (Koivisto, et al., 2011; Sack, et al., 2009).

---

[7] Note: the discrepancy between these times and those reported elsewhere is likely to be due to the use of the difference between stimuli offset and the TMS pulse being used to quantify the temporal distribution of TMS pulses, as opposed to the more conventional stimuli onset asynchrony.

The question this brief discussion of previous experiments was aimed at addressing was whether or not there exists more than one period over which perception is susceptible to TMS applied to the occipital lobe. The data here, together with the majority of other on-line experiments, points to a single period of causal importance in perception, where simple detection tasks are employed. However, the performance of other tasks may well involve two or more periods of processing in occipital cortical areas (e.g. Chambers, et al., 2012; Heinen, et al., 2005), the clearest examples of which are probably the demonstrations of independent periods in motion detection tasks, where an initial appraisal of the state of primary areas followed by a secondary recurrent monitoring phase might aid task performance (Pascual-Leone & Walsh, 2001; Silvanto, Lavie, et al., 2005; Stevens, et al., 2009). The overall picture, then, is that there likely exists a feedforward phase of processing which is then followed by a recurrent phase. However, whether or not TMS is capable of isolating them, and indeed, whether or not the two phases are temporally separate and do not overlap (Heinen, et al., 2005), is likely to depend upon the task involved. It would seem that simple detection and discrimination tasks do not allow for reliable independency to be demonstrated.

Perception lacking consciousness was demonstrated by the above-chance performance on 'unseen' trials. However, no effects of TMS on PcU were observed at any time period. This null effect suggests that the activity upon which the residual perception depends persists on the basis of lower levels of activity than are affected by occipital TMS. There was, however, a trend toward disruption of 'unseen' discrimination as a result of later TMS as opposed to earlier TMS. This is potentially weak evidence against the prediction made by Lamme (2006b) that TMS should affect unconscious abilities at earlier times and is addressed more thoroughly in the next section (Experiment 2C).

The minima of the suppression of conscious detection (at 110ms with a PrC value of 0.0375) is well within 1 standard deviation of zero (no detection at all). The absence of almost any sign of conscious detection, whilst coincident with consistently above chance 'unseen' discrimination ability, can be interpreted as one of the most conservative demonstrations of blindsight type abilities in this series of experiments. This time point was taken to establish the BIP but rounded up to the nearest stimuli refresh of the CRT display (120ms).

3.6 Single-subject double-pulse, experiment 3C.

The single pulse experiment (3B) established the blindsight inducing pulse (BIP), at 120ms. This pulse was present in all trials of the double pulse experiment, which allowed us to investigate at the roles of different phases of activity within the context of a blindsight type effect. Specifically, we sought to address the question of whether the residual, reportedly 'unseen' abilities of TMS-induced blindsight were dependent upon activity prior to the phase of conscious processing disrupted by the BIP. As noted, 'unseen' abilities of TMS induced blindsight have been linked to the early feedforward sweep prior to the recurrent conscious phase (Lamme, 2001). Interfering with these periods involved adding pulses in addition to the BIP, the prediction being that if the early period is responsible for unconscious / 'unseen' processing then the application of pulses *prior* to the BIP should suppress 'unseen' discrimination to a greater extent than pulses applied *after* the BIP, (Lamme, 2006b).

3.6.1 Double-pulse methods

Methods were identical to the single pulse experiment (3B) unless otherwise stated. The pulses in addition to the BIP were at 40ms intervals ranging from 0 to 240ms (that is, pairs of pulses at [0,120], [40,120], [80,120], [120,160], [120,200], [120,240] ms, see figure 3.3.1.1). This provided six TMS conditions of pulse pairs that were equivalent in terms of inter-pulse interval and balanced across the period prior to the BIP (early) and following the BIP (late). The reported effects of early TMS in particular appear to be highly variable (cf. Experiment 3A vs. Corthout et al 1999-2002); for this reason, a range of timings was covered.

Blocks consisted of 120 trials, yielding five repetitions of each stimulus and TMS condition with each block. The subject completed eight blocks of active TMS and eight

blocks of sham TMS on each session. Six sessions were completed on different days, separated by at least 24 hours, which produced six data points (one per session) for each measure and TMS conditions upon which analysis was based. The order of questions was counterbalanced across the six sessions. TMS intensity was again 120% of subjects PT for both pulses on every trial. Eye tracking resulted in the removal of 12 out of a possible 5760 trials.

Primary analysis of this data involved ANOVAs applied to each measure with TMS site (sham vs. active) and time (additional pulses applied before vs. after the BIP) as factors, using the individual sessions as data points. This meant there were 18 data points before and after the BIP. If time-dependent effects of the TMS were observed, then these were further probed with 2-tailed t-tests; comparing active to sham TMS effects, at each of the *six* time points, correcting for these six comparisons with the Holm-Bonferroni method.

'Unseen' discrimination ability relative to chance performance was assessed in two ways: single sample t-tests were applied to the data at each of the twelve time points (6 active and 6 sham), contrasting PcU performance relative to chance. This mirrored the analysis applied to the between-subjects study (3A). Also, a binomial test was applied to the correct/incorrect 'unseen' choices. This analysis collapsed the discriminations made across sessions and matched the analysis applied to the BIP task (3B).

In addition, the temporal spacing of pulses was analysed using ANOVAs with site (active vs. sham) and the three pulse spacings as factors. Pulse spacing here refers specifically to the difference between additional pulses and the BIP, i.e. the spacing of the 40,120ms intervention is 80ms and its comparable pair is 120, 200ms. This temporal spacing analysis did not directly relate to the experimental hypothesis, as it addressed whether or not pulses closer together were more effective than those further apart, which may be of interest in itself. However, entry of this temporal spacing factor (pulse spacing) as a covariate into the ANOVAs which compared early vs. late effects allowed us to eliminate variance due to the temporal spacing of pulses, the effects of which are reported separately as an ANCOVA (3.6.2.4).

The Bayesian analyses were designed to mirror those of the conventional analysis above and followed the methods of the between-subject experiment (3A) in application. This meant that Bayesian analysis was applied to this data in two ways to each of the primary measures. First, the early effects were compared to the later effects to substantiate any

difference between these two periods. A single vector representing this comparison was derived using change from sham data and subtracting post-BIP data averaged across the three late times from averaged pre-BIP data, the variance of which was taken across sessions (data points). Thus the vector was calculated as $(Early_{occipital - sham}) - (Late_{occipital - sham})$. Therefore, on this vector a suppression specific to additional early stimulation would result in positive values and would be represented by a positive prior. Effects in the opposite direction (i.e. later TMS being more effective than earlier) were also tested $((Late_{occipital - sham}) - (Early_{occipital - sham}))$. Second, the effects of the TMS were analysed at each of the six time points. This probed the efficacy of the TMS and how it evolved over time. The dependent variable was the change from sham data, and the prior represented a uniform suppression of perceptual abilities. Priors involving PcU ranged from 0 to 0.5 and those involving PrC ranged from 0 to 0.6 because threshold levels resulted in those being the maximum reasonable range.

Outlier rejection was again based on Chauvenet's criterion applied to each data condition, using the variance across the sessions. In the event, no data was shown to be out-lying in this experiment.

3.6.2 Double-pulse results

3.6.2.1 Conscious detection

Conscious detection was suppressed by the application of the double pulses (site effect: $F_{(1,68)} = 144.301$, $p<0.001$), and although there appeared to be a time dependent effect where later pulses resulted in lower detection (time effect: $F_{(1,68)} = 4.184$, $p=0.045$), this did not interact with the type of TMS and so cannot be attributed to the magnetic / cortical effects of TMS (site × time interaction: $F_{(1,68)} = 1.291$, $p=0.260$).

Bayesian comparison between early and late effects was inconclusive as the Bayes factors representing effects being expressed flowing later TMS was weak (B= 0.8495). The Bayes factor modelling an earlier effect of TMS, relative to late efficacy, favoured the null (B= 0.043), indicating an absence of a difference between the early and later periods on the

PrC measure. Bayesian priors which modelled a drop in PrC at each of the time points closely represented the data (B ranging from 1396 to $2.712e^{11}$, see table 3.6.2.2).
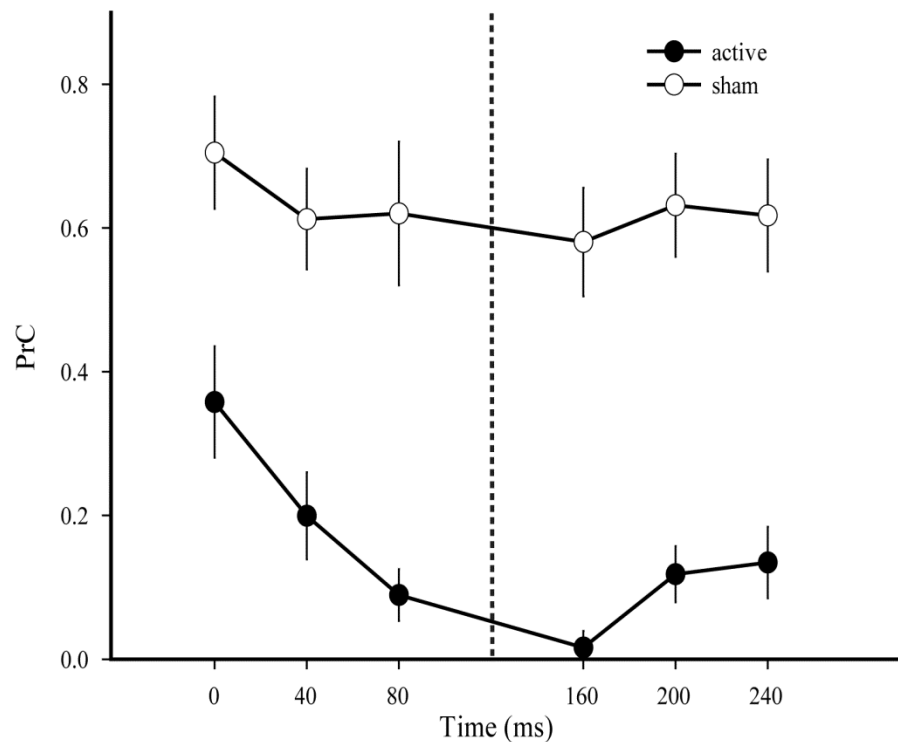
Figure 3.6.2.1. Conscious detection (PrC) as affected by double pulse interventions. The vertical dotted line indicates the time at which the BIP was applied relative to onset of the visual stimulus. Other timings represent the application of pulses in addition to the BIP. Error bars are standard error across session.

3.6.2.2 'Unseen' discrimination

For 'unseen' discrimination the effect of the active TMS was time dependent (site × time interaction: $F_{(1,68)}=8.681$, p=0.004, see figure 3.6.2.2). 'Unseen' discrimination was also shown to be suppressed by the TMS (site effect: $F_{(1,68)}= 19.462$, p<0.001). Independent of TMS site effects, no effect of when the pulse were applied (early vs. late) were observed (time effect $F_{(1,68)}=0.014$, p =0.907).

The greatest contributions to this effect were made by the two suppressions in PcU as a result of TMS applied at [120,160] ($t_{(5)}$=4.607, p= 0.006) and [120,200] ($t_{(5)}$=4.963, p=0.004). T-tests which compared active to sham conditions applied to other times were non-significant (t<1.828, p>0.127). The Bayesian analysis reflected this; first, comparing early to later effects indicated strong support for the effect being greatest as a result of late TMS (B = $7.4 \times 10e^{5}$), support for the effect being expressed earlier strongly supported the null (B=0.007). When applied to the individual time points, the analysis favoured the null of there being no effect of the TMS during the early period of intervention (B 0.07 to 0.35). TMS applied following the BIP, however, was effective in suppressing PcU, particularly during the 160,120ms (B=25.48) and 200,120ms (B=46.12) interventions. For the 120,240 intervention the effect was inconclusive (B=0.58, $t_{(5)}$=1.828, p=0.127). These analyses are summarised in table 3.6.2.2.

'Unseen' performance was consistently above chance as assessed by single sampled t-tests applied to all 12 data points (6 active and 6 sham)($t_{(5)}$ >4.63, p<0.006) or binomial statistics (all p <0.001).
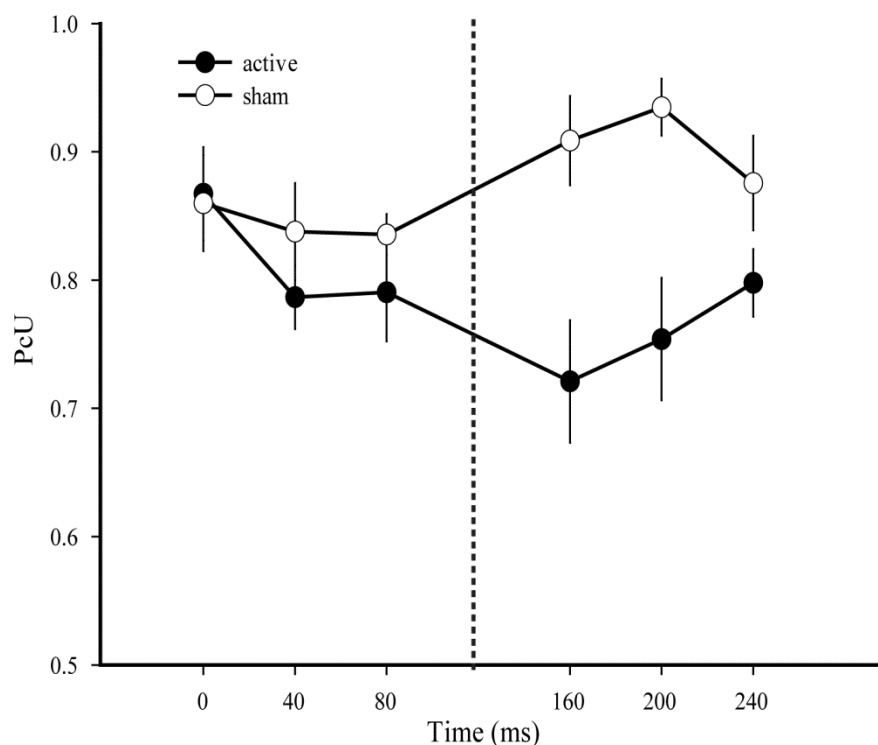
Figure 3.6.2.2. 'Unseen' discrimination as affected by double pulse interventions in Experiment 3C. This figure conforms to the same structure as figure 3.6.2.1.

Table 3.6.2.2. Bayes factors representing the effects of the double pulse interventions. Applied to individual times the Bayes factors represent the strength of support for the hypothesis that the TMS caused suppression in measures. When applied to the early vs. late analysis the Bayes factor represents the strength of support for TMS being effective when applied earlier vs. later, and vice versa.

Bayes factor (B)

Time TMS							
0,120	40,120	80,120	160,120	200,120	240,120	Early>Late	Late>Early
PrC 1394.626	$2.73 \times 10e^8$	2211.119	$4.68 \times 10e^8$	$3.22 \times 10e^8$	$2.71 \times 10e^{11}$	0.0434	0.8495
PcU 0.07151	0.347722	0.317055	25.48766	46.12056	0.582341	0.0071	$7.38 \times 10e^5$

The Fmax ratio of variance in PcU/PrC was 1.45 based on average contributing trial numbers.

3.6.2.3 Response bias

The subject responded positively less often with active TMS (site effect: $F_{(1,68)}$ =12.056, p=0.001). This effect did not interact with whether the additional pulse was applied before or after the BIP (site × time interaction: $F_{(1,68)}$=0.579, p=0.449). Neither was there a main effect of pulse timing (time effect: $F_{(1,68)}$=0.352, p=0.555). The analysis of the temporal spacing of pulses indicates that the temporal proximity of pulses affected response criteria (TMS site × pulse spacing interaction $F_{(2,66)}$ =4.001, p=0.023). This is potentially suggestive of a distracting effect in which temporally close TMS pulses distract more than temporally distant ones. However, given the high degree of variability in the sham condition (see figure 3.6.2.3.) this result should be treated with caution. Here 127 false alarms were made in contrast to 3308 correct rejections.
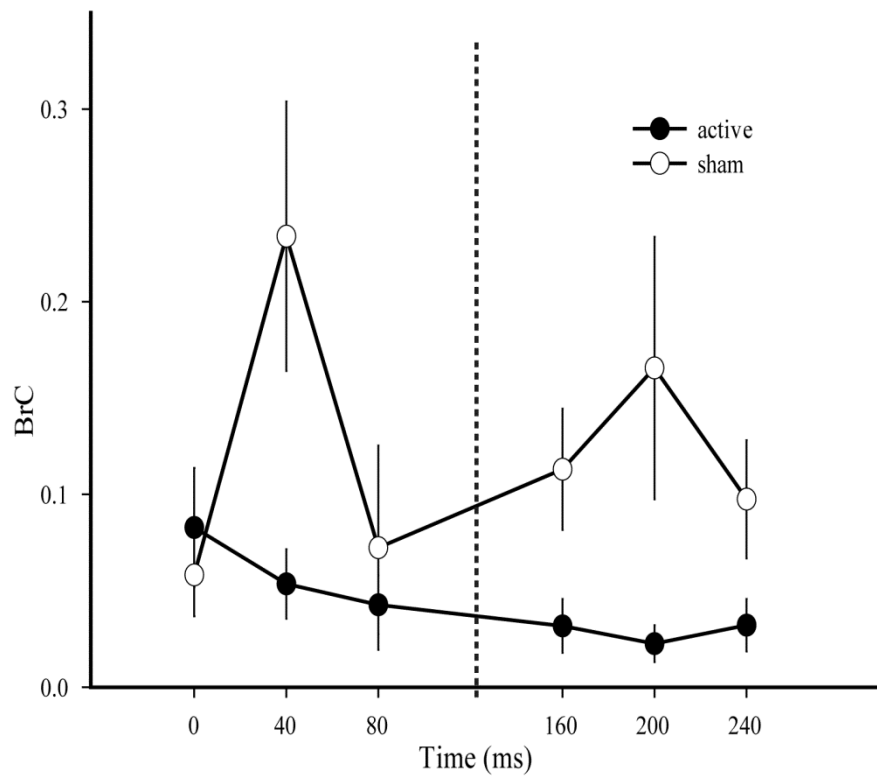
Figure 3.6.2.3. Bias in conscious detection as affected by double pulse interventions. This figure conforms to the same structure as 3.6.2.1.

3.6.2.4 ANCOVA

The spacing of pulses indicated that pulses placed closer together were more effective in suppressing conscious detection ($F_{(1,67)}$=7.55, p=0.008). A similar but non-significant trend was observed for 'unseen' discrimination ($F_{(1,67)}$=2.178, p=0.145). The addition of this as a covariance factor into the main site × time ANOVAs, resulted in an ANCOVA, which did not reveal any main effects or interactions not observed in the original analysis.

3.6.3 Double-pulse discussion

Again, the application of occipital TMS produced a blindsight type effect. Conscious detection was suppressed to a greater extent than 'unseen' discrimination. The 'unseen' discrimination measure, on the other hand, showed time-dependent effects which were not expressed upon the measure of conscious detection. The specificity of different effects to both measures can be considered a double dissociation and is therefore the clearest demonstration in this series of experiments that the measures used are sensitive to independent properties. The demonstration that the PcU measure is sensitive to the interventions supports the conclusions that can be drawn from other experimental sections, as it suggests that the absence of effects, for example, in reference to the between subjects s-cone intervention (Experiment 3A), were not a result of insensitivity of the measure, but rather arose due to the measure being unaffected by that particular intervention.

The time-dependent effect of the TMS upon 'unseen' discrimination has implications for the proposed correspondence between unconscious processing and the early feedforward sweep (Lamme, 2001; Lamme & Roelfsema, 2000). In particular, the observed effects had the opposite temporal signature to that which would have been expected; 'unseen' discrimination was reduced when TMS was applied later, during the more recurrent phase, rather than as a result of early TMS. Indeed, no clear effect of the TMS was observed when TMS was applied early and was only realised when TMS was added after the BIP. This study can therefore be seen as providing evidence against the idea that the basis of 'unconscious' vision lies in the early feedforward sweep of information through the visual system, as it is more likely to depend on activity occurring after 120ms.

The element of Lamme's theory that relates conscious awareness to later recurrent processing is well supported by evidence. However, there is relatively little direct evidence in Lamme's theory to support the early-unconscious claim. In terms of evidence for the theory, he cites the preservation of fast collicular routes (collicular input being early to arrive) in classic blindsight (Lamme, 2001) - evidence which the s-cone interventions applied here contradict. In addition, he discusses electrophysiological evidence, which shows that the earliest evoked activations (prior to later recurrent / conscious processing - in reference to

which the results here agree) carry information to sensory areas such as orientation selectivity in vision (Lamme & Roelfsema, 2000; Super, Spekreijse, & Lamme, 2001).

The current results might be seen as ruling against Lamme's theory (2001). However, rather than ruling against the claim of a correspondence between early processing and unconscious processing, it may be possible to reconcile this theory with the current data in two ways: by considering the definition of unconscious, and by considering what might contribute toward the residual abilities of TMS-induced blindsight. It should be noted that what follows is a tentative attempt at reconciliation between the data and the theory.

What is it that is referred to in Lamme's theories when he uses the term 'unconscious' and does this match the forced choice 'unseen' abilities demonstrated here? It is possible to read Lamme's use of 'unconscious' as referring to something more akin to 'subpersonal'. (Sub-personal is a level of description first expounded by Dennett (1969), which is a way of describing processes at the purely physical level, such as neuronal processes, in contrast to personal level descriptions which refer to the experience of the subject.) In which case, early to arrive activity might therefore be sub-personal and carry information such as orientation selectivity. Such activity might be at such a basic level that the TMS was incapable of affecting it at the levels applied here (e.g. columnar segmentation). Therefore on the 'sub-personal' reading of Lamme's theory the TMS applied here did not affect the feedforward sweep and the theory is maintained.

It is also possible that what is measured by the forced choice discrimination of arrows when reportedly 'unseen' is neither unconscious nor subpersonal. It is possible that what it is (the perception) that is referred to by the PcU measure actually depends on or is facilitated by the reactive (Ericsson, 2003) processes caused by the posing of the question. That is, when subjects are asked in which direction the arrow was pointing, they consciously reflected upon this question and although the activations caused by the stimuli were insufficient to elicit a categorical 'seen' response, it was the reflection itself that supported the discrimination decision, to the extent that it was above chance over a number of trials. The later conscious reflection, then, was the target of the perturbation caused by later TMS and the reason for which reportedly 'unseen' abilities were diminished at these later times. In which case, this experiment simply may not confront Lamme's theory. This argument is expanded upon in Chapter 4.

It is important to note that this entire discussion may reflect an over-interpretation of the data, and implies that we have not targeted purely 'unconscious' processing (or phenomenal consciousness in Block's sense (Block, 1995), discussed in Chapter 4.6). Even so, what does appear to have been tracked by this measure is perception which is *less* conscious than when the subject acknowledges veridical awareness (Cohen & Dennett, 2011). This difference is all that is required to make the contrast which reveals our target of conscious awareness, upon which the majority of the conclusions drawn here have been based.

3.7 Single-subject s-cone intervention, experiment 3D.

The aim of this experiment was to probe the pathways that contribute to processing in the context of TMS-induced blindsight. The principal question was whether or not the preserved abilities critically depend upon input via the superior colliculus. In doing so it made use of the same s-cone intervention as the between-subjects study (3A) and applied the BIP to reveal blindsight type responses.

3.7.1 S-cone intervention methods

Unless otherwise stated, procedures and equipment are the same as in previously described experiments. Prior to commencing these experiments the subject's tritan line and point of equiluminance were ascertained using the method described in Experiment 3A. The subject's tritan line was found to be at an angle of 9° in MacLoed-Boynton (1979) space with an equiluminance point of 24.7 cd/m^2. The colour used in CIE1931 coordinates was 0.256, 0.230, 17.62. Eye tracking resulted in the elimination of 26 trials from a possible 1920.

Single pulses of TMS were applied at the BIP (120ms) in either active or sham coil orientation. Stimuli were either s-cone or luminance. This yielded four conditions that were acquired in separate blocks over the course of a single session, the order of which was randomised. The subject completed 6 sessions, each involving the use of a different question order of presentation. Each block consisted of 80 trials and resulted in a single data point per block entered into the subsequent analysis.

The primary analyses applied ANOVAs to each measure with TMS site (active vs. sham) and stimuli type (s-cone vs. luminance) as factors. Analysis of above-chance performance on 'unseen' trials was performed both with a single sample t-test (matching the between subject study, 3A) and binomial tests of unseen abilities (matching the double pulses section, 3B), including Holm-Bonferonni correction for multiple comparisons.

The Bayesian analysis again made use of the variance across sessions. It was first applied to the change from sham data for each stimulus type and measure separately. Then, using the change from sham data and taking away the s-cone data from the luminance data, a vector was produced that, if positive, would represent the benefit bestowed by the additional input provided by luminance stimuli from RM pathways. The Fmax ratio between the PrC and PcU measures was 0.98.

3.7.2 S-cone intervention results

3.7.2.1 Conscious detection

Conscious detection was suppressed by the application of TMS (TMS site effect: $F_{(1,20)}$=129.1, p<0.001), but this effect did not depend on the type of stimuli used (site × stimuli interaction: $F_{(1,20)}$=0.572, p=0.458), nor was there a discernible difference between detection of stimuli types independent of TMS effects (stimuli effect: $F_{(1,20)}$=2.553, p=0.126), see figure 3.7.2.1 . BrC was not analysed for this section due, again, to the low numbers of false alarms (19 false alarms across all conditions, compared to 578 correct rejections).
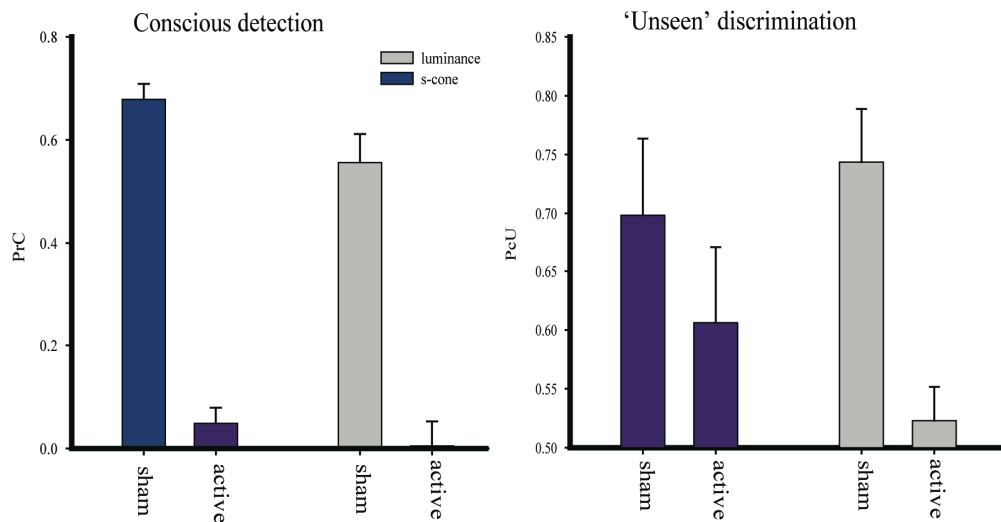
Figure 3.7.2.1. Conscious detection and 'unseen' discrimination in the presence of s-cone and luminance stimuli in active (BIP at 120ms) and sham conditions. Error bars are standard error.

3.7.2.2 'Unseen' discrimination

Changes in 'unseen' discrimination mirrored those found for conscious detection. PcU was suppressed by the TMS (TMS site effect $F_{(1,20)}$=10.26, p=0.004) and showed no dependency upon stimulus type (site × stimuli interaction: $F_{(1,20)}$=1.724, p=0.204, and stimulus effect: $F_{(1,20)}$=0 .150, p=0 .702).

Binomial tests indicated that ability to discriminate above chance when stimuli were reportedly 'unseen' was only clearly preserved for luminance sham condition (p<0.001), and was completely absent for the active luminance condition (p=0.689). Conscious detection being around zero for luminance stimuli (see figure 3.7.2.1) implies that all perception was eliminated at this data point. For s-cone stimuli in the active and sham conditions, above chance performance was only demonstrable relative to an uncorrected alpha level of 0.05 (sham p= 0.029, active p= 0.048). These values do not withstand correction for multiple comparisons. The same pattern of above-chance responses was also observed using the single sample t-test method with above-chance performance for luminance stimuli in the sham condition ($t_{(5)}$= 5.385, p=0.003) but not others (s-cone sham $t_{(5)}$=3.044, p=0.029, s-cone

active $t_{(5)} =2.152$, p=0 .084, luminance active $t_{(5)} =0.804$, p=0.458, again relative to alpha of 0.0167).

3.7.2.3 Bayesian analysis

The Bayesian analysis indicated that TMS suppressed conscious detection for both stimulus types (s-cone B=$4.26 \times 10e^{16}$, luminance B=$5.19 \times 10e^5$). 'Unseen' discrimination was also suppressed for luminance stimuli (B=$1.62 \times 10e^5$) but not in the presence of s-cone stimuli (B=0.63). The direct comparison of stimulus types, using of change from sham data, revealed no conclusive support for more pronounced suppression in the presence of s-cone stimuli compared with luminance stimuli (For PrC B=0.44, PcU B=0.14). In fact, the opposite effect, of 'unseen' ability being suppressed to a greater extent when the additional input provided by the luminance stimuli is present, was more strongly supported by the data (B=0.93).

3.7.3 S-cone intervention discussion

The conclusions that can be drawn from this experiment are limited. The absence of above-chance performance for 'unseen' trials in three out of four conditions means that it cannot be claimed that these data represent a form of blindsight.

A point worth highlighting is that 'unseen' performance dropped to a greater extent in the presence of luminance stimuli than it did for s-cone stimuli; an effect in the opposite direction to that which would have lent weight to the principle of collicular support for 'unseen' abilities (Schmid, et al., 2010). Alternatively, it could be argued that the absence of 'unseen' above chance performance in the s-cone sham condition could be taken as evidence that all 'unseen' abilities in this subject (irrespective of occipital disruption) operate via the superior colliculus (e.g. Fuchs & Ansorge, 2012). However, this seems improbable, as the reason for which 'unseen' performance in this condition was not considered above chance was the correction for multiple comparisons, and performance here was at 70% correct.

Additionally in the between-subjects study (3A) this subject's performance in the s-cone sham condition was at 68%. Therefore insufficient statistical power appears to be the primary reason for failing to demonstrate above-chance performance, and therefore failure to reveal TMS-induced blindsight.

3.8 General on-line conclusions

The experiments have demonstrated dissociations where conscious detection was suppressed by TMS applied at ~100ms while 'unseen' discrimination was preserved; corresponding to previous reports of TMS-induced blindsight (e.g. Boyer, et al., 2005). There are two aspects of these demonstrations that can be seen as contrary to prevailing views as to the neural basis of blindsight. First, the residual unconscious abilities of blindsight do not depend upon input via the superior colliculus. Second, neither do they appear to depend upon the early feedforward sweep of activity; rather, later interference suggests a later causal basis.

If, however, the unconscious abilities associated with blindsight are more nuanced than an early-collicular-unconscious relationship, then the data presented here might not contradict previous descriptions of the neural basis of blindsight. This might be the case if different tasks that fall under the banner of the unconscious abilities of blindsight call upon different resources. Some demonstrations might show collicular dependency (Leh, Mullen, et al., 2006) while others may not (here and Schmid, et al., 2010). Likewise, some early information that facilitates 'unseen' abilities might be at such a basic level of partition (e.g. of columnar segregation) or may directly pass to unaffected regions, that it is resistant to the early TMS intervention. Late conscious reflection upon the percept, on the other hand, which could mediate the residual abilities, may have been what it is that is susceptible to late TMS.

These issues of task dependency and, in particular, the problem that the reportedly 'unseen' abilities could be mediated by a conscious reaction to the task, suggests that a more indirect task may be preferable in demonstrations of blindsight type effects. This is because tasks involving, e.g. attentional cueing effects (see for review Mulckhuyse & Theeuwes, 2010; Shipp, 2004) might be thought of as less susceptible to this problem of reactivity (the posing of the question itself affecting the processing under analysis, see Ericsson, 2003;

Gagne & Smith, 1962, and section 4.3). Such attentional cueing tasks, especially those involving peripheral cues, are also more likely to illustrate independence of anatomical function (Ansorge, 2003; Leh, Mullen, et al., 2006; Mulckhuyse & Theeuwes, 2010), (but also see Bompas, et al., 2008). Indeed, comparisons of cued measures to overt forced choice paradigms are promising avenues for research.

However, what is gained by the use of indirect measures might not outweigh what is lost through the rejection of overt forced choice measures. What is lost is the proximity of the contrast used to reveal conscious processing. That is, for the contrast to be maximally informative, it is required that the difference between the two states / tasks be as minimal as possible in all respects, other than the difference of interest (*ceteris paribus* ,Chalmers, 1996). If a contrast state were to be chosen that is known to operate via a different set of pathways to consciousness under all conditions then the informative nature of the contrast is weakened. For example, we demonstrated R/M – potentially collicular – dependent effects upon the measure of conscious detection. If we had chosen a task to demonstrate the preserved abilities that are supported by the SC irrespective of interventions (as with Leh, Johansen-Berg, et al., 2006; Tamietto, et al., 2010), such as unconscious attentional cueing task (Mulckhuyse & Theeuwes, 2010; Shipp, 2004), then it may not have been possible to show that these pathways also play a role in specifically conscious processing. This would be because both 'unseen' abilities and conscious detection would presumably demonstrate s-cone dependent effects. In such a circumstance it would only be possible to claim that perception in general is supported by collicular and/or magnocellular input.

There is therefore no contradiction between the data presented here and previous work other than in prior claims of necessity (e.g. Ro, et al., 2004). Rather this work furthers our understanding by demonstrating that conscious processing, in particular, benefits from R/M input during early and late phases of occipital activity. Moreover, consciousness is more susceptible to direct interference of occipital cortical areas than are 'unseen' abilities, which can persist on the basis of presumably lower levels of input provided by a variety of routes that are likely to be task-dependent.

Chapter 4.

General discussion

4.1 Overview of experimental findings

These experiments sought to probe the neuronal processes involved in visual consciousness. The strategy developed for doing so involved deriving a measure that referred to subjects' conscious experience of arrow stimuli and to compare variation in this measure with a measure of perception where consciousness was lacking. The variations were the result of the direct causal intervention of TMS. Similar paradigms have previously been used to demonstrate the phenomena of blindsight (Boyer, et al., 2005; Lau & Passingham, 2006; Sanders, et al., 1974; Weiskrantz, 1986).

The 'off-line' experiments applied repetitive TMS (cTBS), which caused suppression in cortical excitability (Franca, et al., 2006; Huang, et al., 2005) while, increasing subjects' conscious detection of stimuli (Experiment 2A). This effect on conscious detection was replicated (Experiment 2D) and a number of other experiments were developed to probe this effect. Experiment 2B confirmed that the regions affected by the cTBS became less susceptible to direct stimulation. This was further consistent with the demonstration cTBS increased concentration of the inhibitory neurotransmitter GABA, as measured by MRS in Experiment 2C. The replication of the behavioural experiment was conducted in conjunction with MEG. Late evoked responses to stimuli, as measured by MEG, can be seen as representative of conscious signals (Koivisto & Revonsuo, 2007; Sergent, et al., 2005). There was no clear evidence that occipital cTBS altered the passage of such signals. Instead effects were observed in the oscillatory alpha band. The alpha band is thought to play an important role in the gating and segmentation of percepts (Jensen & Mazaheri, 2010; Pfurtscheller, et al., 1983; Varela, et al., 1981). In particular, the alpha ERD can be seen as representative of a gating mechanism, where a shift occurs from a period in which the power / inhibitory

influence of the alpha oscillation is high to a desynchronised, disinhibited state following the presentation of the stimuli. The rate of onset of this ERD became steeper following the application of the cTBS. Collectively, these experiments suggested that the facilitation of conscious detection observed may have been the result of a potentiation of the mechanisms involved in the formation of conscious percepts through gating of representations. The interpretation offered is that discrete conscious percepts are brought into relief by inhibitory processes, involving GABA and the ERD, and it is this function that benefited from the application of occipital cTBS.

The 'on-line' experiments addressed a long running debate in the blindsight literature as to whether conscious vs. unconscious processes map functionally onto anatomically distinct pathways via the LGN and SC (Cowey, 2010; Lamme, 2001, 2006b). Contrary to this correspondence, Experiment 3A showed that the residual, reportedly 'unseen' abilities of TMS-induced blindsight, can persist when informative input via the SC was removed. This experiment also confirmed previous demonstrations (e.g. Amassian, et al., 1989) that the critical period for TMS efficacy, when applied to visual areas, operates from ~100ms after stimulus onset. It has also been previously suggested that the level of activity in early visual cortical areas when the first feed-forward sweep of input arrives can determine whether or not stimuli are perceived (Romei, et al., 2010; van Dijk, et al., 2008). Experiment 3A extends our understanding of what this period may involve by showing that TMS applied around the same time as the stimuli (0-40ms) might add additional activity upon which early input, from the faster R/M pathways (Bompas & Sumner, 2008; Cottaris & De Valois, 1998; Maunsell, et al., 1999; McKeefry, et al., 2003) can be boosted (Schwarzkopf, et al., 2011). It was also shown that the latest (280-320ms) period of recurrent processing targeted may benefit from input from these pathways, which may first pass information to more frontal regions, before re-entry.

The single-subject single-pulse study (Experiment 3B) revealed that conscious detection was supported by a single disenabled period of processing in occipital cortex, consistent with previous demonstrations (Amassian, et al., 1989; Koivisto, et al., 2011; Sack, et al., 2009). This was perhaps the clearest demonstration of TMS-induced blindsight throughout these experiments. The findings of Experiment 3C appear contrary to previous suggestions that the early feedforward phase is the basis of 'unconscious' processing and that the later recurrent period corresponds to the conscious phase of processing (Lamme, 2001; Lamme & Roelfsema, 2000). Although a correspondence between later activity and

conscious processing was consistent with the data, TMS applied at later times suppressed reportedly 'unseen' abilities to a greater extent than did its earlier applications, indicating a later recurrent basis for the activity that supports 'unseen' abilities. It is, however, possible to reconcile this data with previous models (Lamme, 2001) if the measures of 'unseen' discrimination do not correspond to the intention behind Lamme's use of 'unconscious' and/or the 'unseen' measure is in some way mediated by conscious processing. Such discrepancies in the interpretation of these measures and the conceptual problems encountered in experiments directed at consciousness are the focus of this chapter, particularly sections (4.3-4.4). Before addressing the conceptual problems, some of the practical limitations of these experiments are first considered (4.2). Sections 4.5 to 4.8 discuss three of the influential positions in the literature that relate to these conceptual concerns, interpretation of the data and status of subjective reports. This is followed by a brief overview of some potential directions for future research (4.9).

4.2 Practical concerns with the experiments described

If the experiments were to be repeated or extended, there are two important changes to procedures that should be made.

1.   Following each subject's participation in the experiments, informal discussions were conducted. This involved discussing with subjects how they approached they approached the experiments and explored any adverse reactions to TMS they may have experienced (see Maizey, In Press). In retrospect, these informal discussions should have been formalised and recorded beyond the recording of adverse reactions. This is because the discussions influenced the derivation of the PrC and BrC measures in the on-line experiments (Chapter 3) in accordance with the neurophenomenological position adopted (see section 1.2.1 and 4.8). That is, the way in which response patterns were allocated to signal detection classes was influenced by these discussions. Responding 'Yes' to the 'something?' question and 'No' to the 'arrow?' question, in the presence of a non-arrow, can logically be assigned as either a 'hit' (in terms of the non-arrow stimulus) or a 'correct rejection' (in terms of the arrow stimulus). The majority of subjects reported that this response profile should not count as

seeing the target (a hit), as detection of the arrow was their primary concern. Therefore, this response pattern was classed as a 'correct rejection' (see table 3.2.2.1). Also, subjects acknowledged that responding positively to the 'something?' question, when nothing had been presented, was incorrect, hence classed as a 'false alarm', although it could also logically have been classed as a 'correct rejection' of the arrow's presence. These allocations were established prior to the analysis of the data and followed the logic of the task instructions (appendix A4). Therefore, although these influential discussions should have been recorded as a matter of scientific rigour, the way in which the discussions influenced the analysis should not detract from the experimental findings.

2.   There are statistical limitations to these experiments which may have been alleviated through increased trial or participant numbers. These limitations are most apparent in three aspects of the experiments. First, it may be thought desirable to combine the two main measures involved in the demonstration of TMS-induced blindsight (PrC and PcU) in a single analysis or to derive a combined measure from the two, representing the extent to which reportedly 'unseen' discrimination outpaces conscious detection[8]. The derivation of such a measure or the application of a combined analysis, would add an additional factor to the analysis, increasing the likelihood of a type II error (Smith, Levine, Lachlan, & Fediuk, 2002). This is especially the case here, where the highest order interaction would be most sensitive to a cross-over interaction (double dissociation); whereas the prediction made on the basis of the definition of TMS-induced blindsight is a single dissociation (statistically conforming to a reduction in PrC and absence of effect on PcU, section 1.3.3). The absence of a higher order interaction would prevent analysis of the lower order effects (as probed by the analyses implemented here), which may more accurately describe the hypothesis. The cost of the additional factor of 'measure' may be overcome through additional power in the analysis (Smith, et al., 2002), but because these experiments were designed only to be capable of illustrating effects on measures independently, such an analysis was not implemented. It should also be noted that none of the previous studies which have clearly

---

[8] A combined measure could be derived by taking one measure (e.g. PrC) away from the other (e.g PcU). Alternatively, the two measures could also be treated as factors where the interaction terms involving 'measure' might describe dissociations between the measures. If so, then any interaction term using a difference score would be mathematically identical to the higher level interaction term involving 'measure'. These two options are therefore considered as equivalent.

claimed to demonstrate TMS-induced blindsight have exploited a factorial / combined measure (Boyer, et al., 2005; Christensen, et al., 2008; Jolij & Lamme, 2005; Ro, et al., 2004). In the demonstration that cTBS reduced metacognitive confidence when applied to the prefrontal areas, Rounis et al (2010) derive a combined measure of meta-d' – d' (see appendix A6 for further details). However, p-values they reported were halved, because their hypothesis was unidirectional (Rounis, et al., 2010). This is questionable, as it seems likely that they would have reported effects with the opposite directionality had they arisen, as described in the off-line experiments (Lombardi & Hurlbert, 2009). The derivation of combined measures or a factorial approach across measures is not required to demonstrate blindsight, as an acceptable statistical definition (section 1.3.3) is adequately served by analysis of the two measures independently. However, derivations of such measures are certainly a promising avenue for future research.

The second concern related to statistical power pertains to the replication in the off-line experiments only (section 2.6, Experiment 2D). The replication employed two control conditions (iTBS and sham). The inclusion of all three TMS conditions in the corresponding ANOVA would again have led to increased type II error rates (Smith. et al., 2002). This reduced capacity of such an analysis to demonstrate effects of interest would be exacerbated if a single 3-level ANOVA were to test the differences between all three TMS conditions. This is because the hypothesis of interest is described by a pairwise difference between cTBS and the other conditions. Because differences between the three TBS conditions may be of interest, they are described in appendix A3. It is, however, worth emphasising that analyses described in appendix A3 are multiple comparisons and therefore should be treated with caution.

Finally, the capacity of any form of SDT to dissociate sensitivity from response criteria is severely limited if no 'false alarms' are produced. This is discussed further in section 4.4.1. Across the between subjects experiments there were 188 out of 686 cells (data points) containing no false alarms. Increasing the number of trials may have reduced this problem.

4.3 Problems encountered in directing experiments at consciousness

Many psychologists have been sceptical about questions related to conscious experience (e.g. Skinner, 1965). This failure to confront consciousness stems from differences concerning the epistemic status of subjective reports, which is the focus of this section. However, a contributing factor in the avoidance of consciousness as a topic of enquiry may be the misconception within psychology that there are broader philosophical reasons which directly block its investigation. While the metaphysical nature and epistemic strategy for the investigation of consciousness are by no means settled, there has to some extent been a consensus reached within modern philosophy of mind, to the point where it is possible to say that regardless of the status of consciousness, its empirical investigation can be informative and should be of use. Disparate points of view as to what constitutes an investigation of consciousness, and what consciousness is, can be to some extent reconciled in the agreement that consciousness is a legitimate target for investigation (e.g. compare Lutz, (2003) and Dennett, (2003)). There is the 'hard problem' of consciousness which suggests there may be something further to explain than purely objective science is capable of (Chalmers, 1996), discussed in section 4.7. But even Dave Chalmers, who established the notion of the 'hard problem', advocates the application of empirical methods to consciousness as a fruitful endeavour.

The central criticism of phenomenologically led approaches, and a reason why behaviourism was adopted by many psychologists, can be to some extent reduced to one of response bias. That is, any response to an introspective probe may not reflect the state to which the probe was targeted, but instead correspond to the response criteria of the subject, to which such probes are often responsive (Nisbett, 1977; Wilson & Nisbett, 1978). This questions the status of subjective reports as being capable of referring to conscious experiences of the subject, which is fundamentally problematic for the research presented here.

The central dilemma is that it is possible that there are differences between being aware, knowing you are and reporting that it is so. If there are such differences, then any

instance of a report may not correspond to its referent in the experience, because the criteria for awareness may be met, yet the subject may have no access to this awareness upon reflection and thus incapable of forming a report that reflects the original conscious percept (Nisbett, 1977). The construction of a report being additional to, and dependent upon, different substrates, than the conscious experience at which the experiments were directed, limits the conclusions that can be drawn. This is because the report and experience are intertwined in the responses provided by the subject, and there appears to be a limit to obtaining measures that refer to the experience of the subject independently of the report process. This problem may cut even deeper than the conflation between report process and the conscious experience of the subject. There is also the possibility that the requirement for a behavioural response may itself affect the experience, again preventing any measure based on subjective reports as referring directly to an instance of conscious (or even unconscious) experience (Gagne & Smith, 1962), this is also known as the 'observer paradox', (Labov, 1972). This problem of the report aspect of behavioural responses being additional to the experience itself is not trivial as it prevents these experiments from directly referring to their target of internal subjective states of the subject. The problem of there being additional elements related to report has been described as the problem of 'reactivity' (Ericsson, 2003), which is the term I will use to refer to it in the following discussion of some of the strategies that can and have been deployed to reduce the impact of this issue on this research.

4.4 Attempts at solutions

4.4.1 Signal detection theory and response allocations

The motivation for the development of the two principal measures, applied here, was primarily to obtain measures of subjects' conscious experience of stimuli and their perceptual ability which lacked conscious awareness; thus revealing consciousness through the contrast between them (Baars, 2005; Dretske, 2000). The secondary objective was to derive measures that were robust to the influence of potential confounds. This section will discuss how these measures were developed with respect to the problems encountered. The measure of

conscious detection is first discussed (4.4.1.1), followed by the 'unseen' discrimination measure (4.4.1.2). Finally, I discuss the issues of interpretation with respect to measures used to quantify blindsight more generally (4.4.1.3).

4.4.1.1 The conscious detection measures

Reactivity limits the extent to which measures can be attributed to their intended referent in conscious experience and it provides a source of potential confounds. As the previous section suggests, if any intervention (e.g. a disruption caused by TMS) is effective at the level of the report process, rather than upon the intended target of conscious awareness, then the effect in terms of the subject's response is the same – a negative report is produced, hence conflating the report and the conscious experience. The application of signal detection theory (SDT) was intended to help alleviate this problem by delineating measures of sensitivity and response criteria. By taking into account the state of the stimuli it is possible to reduce the conflation between the two processes. If TMS is effective in reducing the subject's conscious experience of the stimuli then the TMS should cause negative responses when the stimuli are presented. If, on the other hand, it is the report process that was targeted by the intervention then we might also expect a negative report, but we have no reason to think that such an occurrence of negative report should depend upon the stimuli's presence or absence (N.B. the same structure of argument applies to the production of positive reports). Hence, response criteria / bias as the propensity to acknowledge awareness independently of what was presented (Macmillan & Creelman, 1990) has been linked to changes at the level of reports, while sensitivity-based measures have been used to reflect subjects' conscious experience of the stimuli.

The principal experimental confounds in the on-line TMS experiments (Chapter 3) were the auditory and tactile artefacts accompanying TMS discharge, independent of direct neuronal effects. Additionally, there were differences in the extent to which the noise of the stimuli may have been attributed by the subjects across the two stimuli classes (S-cone or luminance stimuli[9]). In the off-line experiments (Chapter 2) the effect of fatigue and

---

[9] Because the two stimuli classes were not independently calibrated in terms of response criteria, the noise of the stimuli may have appeared to more closely resemble the s-cone target than it did the

expectation[10] may have been the primary causes of uninformative changes in the data. That many of these factors were also present in the control conditions provides the first line of defence in drawing secure conclusions. However, subtraction from baseline may not fully alleviate the influence of these factors, particularly when the magnitude of the confound differs between experimental conditions. This was the case for the stimuli class difference in the on-line experiments and the potential expectation difference in the off-line experiments. Both of these can be considered aspects of the problem of reactivity and may be reduced through the application of SDT. If a subjects' expectations are such that they expect to see stimuli less often (or more often) following the application of active TMS relative to control TMS, then they should respond 'no' (or 'yes') more often under the active TMS condition irrespective of the stimuli's status, and hence express the expectation confound as a fluctuation in the measure of bias rather than sensitivity. This may explain the initial observation (section 2.2.3) and subsequent trend (section 2.6.1) for an increased propensity to respond positively in the off-line experiments. Likewise, if subjects are more likely to attribute the stimulus noise to the presence of s-cone arrows than they are to the presence of luminance arrows, then this difference also should be expressed upon the measure of bias, not sensitivity, as was observed (see figure 3.3.3.3).

The segregation of bias and sensitivity in reported conscious awareness, additionally aided the experiments as the variability between subjects in terms of their criteria was removed from the primary measure of interest (Macmillan & Creelman, 1990). Subjects may approach the task with differing levels at which they acknowledge awareness of the stimuli, which in turn could result in differences in their susceptibility to the interventions; this would then limit the extent to which the experiments are able to demonstrate significant differences of interest. Again, such differences, across subjects, constitute differences in criteria and

---

luminance target. This led to increased 'false alarm' rates for s-cone stimuli and an imbalance of criteria across stimuli conditions, see section 3.3.3. Furthermore, owing to the insensitivity of BrC to fluctuations in the stimuli noise (see below), apparent during piloting, it is questionable whether it would have been possible to calibrate response bias across stimuli classes.

[10] Subjects were almost certainly aware of the difference between control/sham and active TMS in the off-line experiments, and therefore may have expected the TMS to have an effect upon their performance, resulting in a general adjustment of response criteria. The same difference (between active and sham TMS) does not apply to the same extent in the on-line experiments because the effects of interest were also time-dependent, resting upon differences between TMS epochs as small as 40ms. It is highly unlikely that subjects were capable of distinguishing and subsequently reacting to, such differences.

should be expressed upon the measure of bias, from which the primary measure of interest (sensitivity, PrC) is effectively independent.

The application of SDT may reduce this problem of reactivity but does not eliminate it, for three reasons. First, because of the causal aspect to the problem of reactivity – that the requirement for report could actively change the experience in question – the problem is to some extent unavoidable when SDT measures are acquired (hence its description as a paradox; Labov, 1972). Second, the elimination is not complete because the measure of sensitivity may not cover the full extent of conscious experience; that is, there are conscious experiences of the task whose occurrence will result in fluctuations in the measure of bias, but not sensitivity. For instance, misrepresentation of the noise of the stimuli grouped into arrow-like percepts would lead to an increase in bias, not sensitivity, and can easily be understood as a form of conscious awareness. It has even been suggested that a key to understanding what it is to be an intentional agent is the ability to misrepresent information (see Dretske, 1986). This means that the connection made between conscious awareness and the measure of sensitivity is limited to veridical awareness of actual stimuli. Therefore differences in the measures of bias may express report level and/or confound differences (as described above), but the measures of bias may also fluctuate in accordance with certain non-veridical aspects of the subject's *conscious* experience of the task, such as those occurring during misrepresentation. Consequently, the separation of changes in conscious awareness and changes in reactive report might not perfectly match the division between sensitivity and bias measures. Even so, acknowledgement of the division and practical segregation of the two facets seems preferable to the alternative (Weiskrantz, 2001), being a measure of reported awareness which derives from subjects' reports irrespective of the stimuli status. Such a measure would conflate the veridical awareness with response criteria and, therefore, may be more susceptible to experimental confounds such as the expectations of the subjects. Third, as acknowledged in section 1.3.2.2 and 4.2, false alarm rates were low, often approaching floor levels. The capacity of any form of SDT to dissociate response criteria from sensitivity is severely hampered by such floor effects.

All applications of SDT applied here have been non-parametric, owing to the violation of the central assumptions of classic parametric SDT. These assumptions are: "(1) The signal and noise distributions are both normal, and (2) the signal and noise distributions have the same standard deviation" (pp140 Stanislaw & Todorov, 1999). Although it is often unclear as to what exactly it is that is referred to by these distributions, I take the noise

distribution to refer to the information provided on stimulus-absent trials, and the signal distribution to be representative of stimulus-present information. Due to imbalances in the number of stimulus-absent and stimulus-present conditions in both experimental behavioural sections (see 2.6.1.1 methods and 3.2.1 methods) the distributions of the two may not be equal (Stanislaw & Todorov, 1999; Swets, 1986). Furthermore, false alarms were relatively rare, limiting the applicability of classic SDT. This can be interpreted as the result of the subjects' criteria for this task being situated toward the tail end of the noise distribution (Macmillan & Creelman, 1990; Stanislaw & Todorov, 1999). In development of the stimuli, an attempt was made to increase the rate of false alarms by increasing the range of luminance noise; however, this was found to have little effect in piloting experiments. Although the range of the luminance noise for the stimuli used in the final experiments was relatively high, the false alarm rates were still relatively low (the range of noise displayed here was 10-15 cd/m^2, in a previous comparable study the range was 3.6 cd/m^2; Sumner, et al., 2002). By contrast, hit rates appeared to be highly sensitive to fluctuations in the luminance of the target, as differences in hit rates over a range of target intensities drove the psychometric functions used to calibrate task performance. The insensitivity of false alarm rates to changes in stimulus noise, relative to the sensitivity of hit rates, further indicates that the distributions of the noise and the signal may not have been comparable.

Subjects typically adopt conservative criteria in forced choice yes/no tasks (Macmillan & Creelman, 1990), even over a range of noise presentations. Why should this be? It may be a characteristic of higher cognitive processes, such as consciousness, that agents are unlikely positively to report seeing a stimulus when no stimulus is presented. If part of what it is to be conscious is to have a higher order process that *monitors* the content of first order representations, then such a mechanism should act to minimise the incidence of false alarms. This is the 'higher order thought' position discussed in section 4.5 below and follows the proposition that "a conscious mental state is a compound of two things: the mental state, which itself is not conscious and one's being transitively conscious of it" (pp 738 Rosenthal, 2002). Such a mechanism would compare the information received with a higher order representation and produce a positive response only if a match is achieved, with the default response in the absence of a match being negative. The structure of this monitoring/meta-cognitive relationship is therefore such that it limits positive responses when information is limited. It has also been suggested that minimising false alarms is precisely the type of function consciousness has evolved to perform (Millikan, 1984).

The non-parametric versions of SDT adopted here make no such assumptions in relation to noise and signal distributions and therefore do not suffer from the problems caused by low false alarm rates which hinder classic SDT to the same extent. For the Pr and Br measures used, Pr simply reflects the extent to which subjects correctly identified the presence of the stimuli, whereas Br reflects the subjects' propensity to identify the presence of the stimuli irrespective of what was actually presented (Corwin, 1994, see section 1.3).

4.4.1.2 Reportedly 'Unseen' discrimination

In addition to the Pr and Br measures, the other measure of interest in these experiments was subjects' forced choice discrimination ability when they reported not having seen the stimuli (PcU). Both PrC and PcU are measures of sensitivity and operate over the same range. This ensures that the difference between the two principal measures is as minimal as possible (section 1.2-3), with the key difference being that the PcU measure applies when the subject is not reportedly consciously aware of the stimulus presented. Section 1.3.2 discussed the reasons for using 'unseen' trials only in the derivation of this measure (ensuring that the perceptual ability under consideration was lacking in consciousness and was not directly influenced by variation in reported awareness, during which discrimination ability was at ceiling level). In the on-line TMS experiments, which developed from the off-line experiments, the PcU measures were derived from trials when the subject reported 'unseen' to both 'arrow?' and 'something?' questions, allowing more confidence to be placed in the concurrent processing lacking consciousness. Even so, this did not alleviate the problem of reactivity in reference to PcU: while the content of experience with respect to the exposure to stimuli may lack consciousness, the report process is unlikely also to lack consciousness. This conscious reflection upon the trial, required to produce the report, could mediate residual visual abilities. As described, this could be the cause of the late effects in Experiment 3C (the double-pulse single-subject case study, section 3.6).

However, before endorsing such an interpretation, it is worth highlighting some characteristics of the negative report. First, the problematic awareness under consideration and the cause of the problem of reactivity with respect to PcU, is awareness of the question. Awareness of the *arrow* is still lacking in trials that contribute to the PcU measure. Second,

awareness of the question is contingent upon a number of factors (e.g. subjects not becoming habituated to the questioning through practice, see section 4.8.2). In contrast, the presence of conscious awareness when the arrow was presented and the subject responding positively to 'Did you see the arrow?' is far less open to question. Awareness may or may not be present on trials which contribute to PcU, but it is almost certainly present in trials which contribute positively to PrC. Therefore, the difference between the two measures exemplifies a difference specific to conscious awareness of the stimuli, and this difference is sufficient to make the contrast revealing of consciousness.

4.4.1.3 Blindsight and criteria

One of the major criticisms of blindsight as a distinct neurological phenomenon is the possibility that it reflects an extreme response criterion where subjects do not 'acknowledge' awareness and blindsight is therefore not fundamentally distinct from degraded normal vision (Campion, et al., 1983)[11]. Such an argument can comprehensively describe blindsight without referring to there being an absence of conscious awareness. As possibly the chief proponent of the phenomenon of blindsight, Weiskrantz has argued that the degraded normal vision interpretation is not viable (Weiskrantz, 1986, 2001, 2009). He argues that the discrepancy between subjects' reported experience of stimuli and their forced choice performance represents a double dissociation, as both have been perturbed independently, indicating that blindsight is a deficit largely specific to consciousness (Kentridge, et al., 1999; Weiskrantz, 1986 chapter 16 and 19). The single dissociation of blindsight is simply that conscious awareness is affected by the lesion while the residual abilities remain. The double dissociation has been demonstrated by effects on the residual abilities independent of the subjects' conscious knowledge: for example, reportedly 'unseen' responses can be modulated by the presence of cues (Kentridge, et al., 1999). Furthermore, recent work has indicated that

---

[11] Informal conversations with blindsight researchers have also raised an additional related criticism of classic blindsight: some subjects' know that their participation in, and payment for potential future experiments rests upon their not acknowledging any awareness of stimuli presented to their blind field. This could result in their maintenance of extreme response criteria independent of experience. Fortunately, this was not a consideration here as subjects were not paid according to experimental conditions.

a well-studied blindsight subject (DB) has greater sensitivity to stimuli, within a two alternative forced choice task, in his impaired classically blindsighted field, than in his reportedly sighted field (Trevethan, Sahraie, & Weiskrantz, 2007). This makes it difficult to see how degraded normal vision could support the residual vision of blindsight. The application of SDT here was therefore partly intended to circumvent this criticism via the derivation of measures that could evince blindsight (PrC and PcU) and operate independently of subjects' criteria (BrC).

Nevertheless, it is also possible to formulate an explanation of blindsight based on extreme response criteria that does not detract from its being a distinct phenomenon, and which also results in an absence of – or inability to access – consciousness. This is Ko and Lau's SDT based explanation of blindsight (Ko & Lau, 2012). They regard blindsight as a failure of frontal areas to update criteria according to the information available to them, owing to the lesions of the primary visual cortex. The blockage of reciprocal processing, leads to the failure to modulate extreme response criteria by frontal regions that are responsible for the initiation of the report. This highlights an important feature of the discussion of response criteria, bias and signal-to-noise ratio in general: that there are many different meanings to the applications of these terms and they may not all correspond to one another, as discussed further in appendix A2. The Ko and Lau model is a 'subpersonal' (Dennett, 1969)[12] mechanistic description, where failure to update criteria leads to an absence of higher order consciousness, hence does not detract from blindsight as illustrating a specific lack of consciousness. The meaning of criteria / bias in the criticism of blindsight (Campion, et al., 1983), in contrast, applies at the 'personal' level (Dennett, 1969) of the subject where changes in the data may *not* reflect the presence / absence of consciousness, but *instead* fluctuate according to changes in the report process and strategy. Because the uses of criteria in these two interpretations of blindsight differ, they may not contradict or even inform one another.

However, because the SDT explanation of Ko and Lau and the gating-by-inhibition hypothesis put forward in Chapter 2 each apply at the 'sub-personal' level they can be

---

[12] The personal / sub-personal distinction was introduced by Dennett (1969) which pertains to different levels of explanation. The personal level applies at the level of the subject as a whole, to their sensations and reported responses. The sub-personal level is the mechanistic physical level describing the brain and nervous system. Dennett emphasises that descriptions offered at each level may not be directly applicable or comprehensible when applied to the other level.

interpreted as mutually consistent. In the off-line TMS experiments described here, the signal provided to frontal regions, as quantified by the ERF, was not significantly influenced by cTBS, whereas the gating mechanism that potentially reflected the criteria (expressed by the ERD and changes in GABA) was facilitated in the presence of enhanced conscious awareness, mirroring the mechanism described by Ko and Lau (2012).

The conflation between levels of description provides a reason why a deliberate attempt was made to avoid discussing the gating-by-inhibition hypothesis in terms of Signal-to-Noise Ratio (SNR), even though it is acceptable to interpret it as a suppression of neuronal noise, leading to increased SNR. The difficulty in framing the hypothesis in terms of SNR is that neuronal SNR cannot easily be extricated from the SNR of the imaging techniques themselves used to assess the neuronal processes (detailed in appendix A2).

Discussing gating-by-inhibition in terms of SNR also suggests a relationship with the SDT based behavioural measures, which straddles levels of explanation and is also potentially misleading. For example, increased gating could potentially be expected to result in a *decrease* in the measure of bias (BrC), following the application of cTBS. This would be because the mechanisms that govern entry into consciousness (and which set the criteria) become facilitated and the increased active inhibition might lead to a suppression of superfluous representations, thus reducing false alarms. Although superficially compelling this argument is flawed. This is because the incidence of false alarms (produced at the personal level) might *increase* following the suppression of noise at the subpersonal level. This would be because the neuronal representations of the stimulus noise may be facilitated relative to the representation of stimuli background, leading to increased misrepresentation of the stimulus noise as a target and increasing false alarms. Hence two predictions are made which are incompatible if an attempt is made to forge an explanation based on SNR that runs across levels of explanation.

The work of Ko and Lau highlights the point that differences in criteria may have implications for the mechanistic basis of changes in conscious awareness, and are therefore of interest. Yet, in part because of the potential for confusion over the definition of 'criteria', and more importantly because of the susceptibility of the bias measures to confounds (described above in section 4.4.1.1), fluctuations of bias in the current experiments are of secondary interest.

Another criticism of the experimental paradigm developed here, and blindsight more generally, is the use of a forced choice question procedure in relation to conscious awareness, particularly the two alternative "Yes/No?" detection question of the off-line experiments. The procedure itself could potentially be the source of any dissociation between the measures. This would be due to the bifurcation of consciousness enforced by the questions used not accurately describing the likely graduated nature of consciousness (Overgaard, Rote, Mouridsen, & Ramsoy, 2006). Within a more graduated model of awareness, there may be a non-categorical 'grey area' where there is sufficient consciousness for performance on 'unseen' trials to exceed chance levels, yet insufficient to elicit a categorical 'yes' response. This is consistent with the work of Overgaard et al (2008) who showed that a blindsighted subject (GR) admitted no awareness of the stimuli when asked '*Did you see it? Yes or No?*' but would nonetheless demonstrate some degree of awareness using a questioning procedure with a finer granularity of response options. This possibility motivated the application of the 'something?' question in the on-line experiments (Chapter 3). Subjects' performance was well above chance even when they reported having seen neither the 'arrow' nor 'something', indicating the existence of perception when consciousness was lacking. Furthermore, recent work using the perceptual awareness scale developed by Overgaard et al, has shown that although conscious confidence correlates with accuracy, above-chance performance can still be demonstrated when subjects use the lowest categories of reported awareness available to them (Sandberg, Timmermans, Overgaard, & Cleeremans, 2010).

Notwithstanding this defence of the current procedures, I feel this criticism of the forced choice method remains because there is clearly a limit to which simple behavioural descriptions, to which statistics are applicable, can capture the richness and variety of conscious experience (Overgaard, et al., 2006). Although cogent, this criticism does not prevent the 'unseen' category from referring to experiences that are relatively lacking in consciousness compared to when subjects acknowledge veridical awareness, which is the key relative difference upon which these experiments have been based (discussed further in section 4.8)

4.4.2 Other strategies

Clearly, there is a number of additional simple contingencies that were implemented which lessened the problematic possibility of reports not referring to consciousness. For example, lying could potentially be a source of confounds in this respect, but subjects are unlikely to lie if there is no reason to do so (Roepstorff & Jack, 2004). Similarly, a mismatch between experience and report might arise from complexities in the report process, which can be minimised by the use of relatively simple tasks (Ericsson, 2003; Gallagher, 2003; Hurlburt, 2004).

Another way to reduce the problem of reactivity is through the use of additional experimental conditions without the requirement for report (Gagne & Smith, 1962; Lau, 2008). The contrast of the two (with and without report) isolates the additional computation associated with the report (Lau, 2008). This manipulation is more easily applied in imaging studies than in TMS studies where the dependent variable is often based on some form of report. Because the MEG experiment involved a replication of the behavioural study and was time-restricted, this manipulation was not implemented. The capacity of MRS to demonstrate functional changes based on visual task differences is far from reliable at this point in the development of the technique (Puts & Edden, 2012), so this manipulation was not implemented in the MRS experiment either. Since a contrast between conditions requiring and not requiring a response was not implemented, it is not possible to extricate changes as a result of the experimental manipulation being at the level of the report rather than at the level of the conscious experience (or *unconscious* experience in the double-pulse TMS Experiment 3C, see section 3.6.7). However, because reports are required on all trials it is possible to reduce what may be one of the more difficult aspects of the problem of reactivity: that the observed changes are the result of applying a behavioural probe (Dehaene et al., 2001; Gagne & Smith, 1962). The probe is also present in the baseline conditions, thus controlling for general effects of questioning on performance (Rees, 2007).

Clearly whether or not one regards the problems raised above as cogent or straw men depends on the position one takes as to the nature of consciousness and reports – i.e. the ontology of mind that the investigator chooses to adopt. If one accepts that the reporting process of awareness differs from awareness itself, then there is a difficulty in the

interpretation of these experiments because it remains unclear as to which is being measured by the data. If on the other hand one's ontologies are such that what it is to be conscious is to have full access to content (as the introductory quotation suggests) then the problem dissolves, as there are few reasons to believe reported awareness and awareness are dissociable in any meaningful way. If so then the PrC measures directly refer to subjects' conscious awareness of the stimuli and PcU is lacking in consciousness. This is essentially the interpretation position adopted by higher order thought theorists (e.g. Baars, Rosenthal) and is discussed in the next section 4.5. Another, widely cited, way in which this issue has been cast has been in terms of a difference between phenomenal consciousness and access consciousness (Block, 1995, 2007). This is discussed in section 4.6. Following a brief discussion of the 'hard problem' (4.7), section 4.8 will then discuss the intermediate neurophenomenological approach I have chosen to adopt in these experiments.

4.5 Interpretations under Global Workspace and Higher Order Theories.

There are several versions of higher order thought (HOT) and global workspace (GW) theories of consciousness (Carruthers, 2007; Lau & Rosenthal, 2011; Rosenthal, 1993). This section will briefly summarise some of the principal treatments of these theories and their correspondence to the measures and results. Because the predictions made by both these theories (see Lau & Rosenthal, 2011 for review) appear to be equivalent in terms of the data presented here, they are described in a single section.

The basic premise, common and defining of HOT theories, is that what makes a state conscious is that there is a primary level representation (e.g. sensory information), call it X, and that there is a secondary level, contemporaneous, representation whose content is of being in X. Mental states are conscious by virtue of there being a higher order state which represents the primary state: "...a mental state's being conscious consists in one's being conscious of being in that state" (pp 209 Rosenthal, 1993).

Under HOT theories, being aware, knowing you are aware, and being able to report that it is so, are intertwined. Being aware and knowing one is aware, are to some extent, constitutive of what it is to be conscious. Although the report aspect does not have to be present, there should be no objections to reports when they occur in veridical cases (see 4.4.1.1 and 4.4.2) referring to higher-order / conscious processes, because the requirement for monitoring upon which the report depends is the same requirement as is made of processes for them to be conscious. The impact of the problem of reactivity on the interpretation of these experiments is therefore nullified because accessing percepts and forming reports are not considered additional to conscious awareness under HOT theories. Therefore, the measures of conscious detection employed here refer directly to conscious awareness, and the corresponding interpretation of reportedly 'unseen' abilities would be that they correspond to first order processing that is lacking in consciousness (Dehaene & Naccache, 2001).

Like HOT theories GW theories are based on operational definitions which originate from the premise that consciousness has limited capacity (Baars, 1988; Baars, 2007). Representations gain access to a GW which is likely to involve thalamo-cortical complexes (Baars, 2007). Such information is then available and can be broadcast throughout such complex networks (Baars, 2007). In this way consciousness can be seen as involving unconscious representations that are gated and selected for entry into the GW (Baars, 1988; Baars, 2007), which is consistent with the interpretation of the off-line TMS experiments reported here. This is particularly the case with respect to the ERD results (section 2.6.4), as Baars has linked similar alpha band changes to such a gating mechanism (Baars, 2007). In terms of the interpretation of the current experimental paradigm, as a whole, the PrC measure can be interpreted as tracking representations of task-relevant stimuli that gain access to the GW and become broadcastable / conscious. BrC can be seen as the criteria set for entry into the GW, and PcU as corresponding to the overflow in unconscious processing that fails to gain entry to the global network.

In contrast to PrC, the interpretation of 'unseen' abilities, particularly under HOT theories, is complicated by the problem of reactivity. 'Unseen' measures are based upon trials in which the subject reports 'No'. Do the questions applied (both 'left/right?' and 'yes/no?') and responses collected require second order representations (Sahraie, Hibbard, Trevethan, Ritchie, & Weiskrantz, 2010)? If they do, then it is difficult to use data collected when subjects report stimuli as 'unseen' to draw conclusions that equate this response type with an absence of consciousness. Indeed, it might be that the second order-conscious reflection on

the question that was the target of the disruption of 'unseen' abilities caused by the TMS applied at later times in Experiment 3C (section 3.6), i.e. first order processes may not have been tracked by the 'unseen' measure.

As with the interpretation offered in 3.6.3, the proposition that 'unseen' abilities are mediated by the conscious requirement for report may be an over-interpretation of that data set and is not consistent with the conclusions offered by proponents of HOT theories to similar data sets (Rosenthal, 2012). For the HOT theorist, the requirement for report can be seen as evoking an additional third order relationship which may be conscious (of the form of 'Did *I* see anything?'), but the second order relationship that is the target of the question is not conscious by virtue of there being no primary representation which is discernible to the second order process (Rosenthal, 2012).

An additional consideration in the interpretation of these experiments under HOT theories is that it is possible that the measures of criteria / bias (BrC) could refer to higher order processing (Dienes, In Press; Ko & Lau, 2012; Timmermans, Schilbach, Pasquali, & Cleeremans, 2012). This prompts the question: should BrC also be considered a measure of higher order / conscious thought, and thus be treated as a measure of central interest in these experiments? Within the context of these experiments the 'Did you consciously see the arrow?' question was targeted at secondary level *sensitivity* in its construction. So here it is perhaps more appropriate to relate the measure of sensitivity to the conscious processing which is the target of this investigation (Dienes, In Press). Bias / BrC can be seen as corresponding to third order relationships, which may be conscious, but whose content is more likely to involve the reactive process mentioned in the preceding paragraph. Additionally it has been demonstrated that the criteria subjects apply to a task can be manipulated even when the subject is unaware of the stimuli to which the criteria refer (Kentridge & Heywood, 2000). That is, criteria can change independently of the presence of conscious awareness. Such changes have been reported within the context of classic blindsight where there is a loss of awareness (Kentridge & Heywood, 2000). Therefore, because it appears that criteria can fluctuate in the absence of awareness, the PrC measure is perhaps the safer referent for second order – conscious – processing than are changes in BrC.

4.6 Interpretations under a distinction between phenomenal and access consciousness

This section will consider the distinction between phenomenal and access consciousness (Block, 1995, 2007). In addition to describing this distinction, the data produced over the on-line experiments in particular affords us the opportunity to test a prediction, made by Block, that informs the validity of the differentiation. One interpretation of the distinction (Block, 2007) suggests the derivation of an additional measure not previously described, the analysis of which can be found in appendix A5.

What it is that the measures described and responses to tasks refer to is always open to interpretation. The assumption made here has been that PrC corresponds to conscious awareness and PcU corresponds to perception that is distinctly lacking in conscious awareness. These assumptions were adopted because they were felt to be broadly consistent with many, if not the majority of contemporary models of mind (HOT,GW, Heterophenomenology (see Dennett, 1991; Dennett, 2003), Neurophenomenology (see section 4.7 and Varela, 1996; Varela, et al., 1991) and previous interpretations of similar data (Boyer, et al., 2005; Holt, 2003; Lamme, 2006b; Ro, et al., 2004; Weiskrantz, 1986, 1996)). Additionally, the measures themselves were adjusted according to reports collected from the subjects following their experience of the experiments, placing the epistemic referents of the measure upon the subjects themselves (Dennett, 2003; Varela, 1996).

One notable exception to this consistency is the theory put forward in a series of influential articles by Ned Block. The basic differentiation made by Block is between phenomenal and access consciousness. Phenomenal consciousness (p-consciousness) is commonly associated with the 'what it is like' (Nagel, 1974) or 'qualia' (Lewis, 1929) aspect of experience (pp 230 Block, 1995). Block clarifies this by defining p-consciousness as having content, as in the difference between seeing red and green (Block, 2005). Access consciousness (a-consciousness) on the other hand is that which the subject has full, reportable reflective access to – thus associated with the mechanisms of report (Block, 1995, 2005, 2007). In terms of evidence, Block cites blindsight as a case in which phenomenality has greater capacity or overflows conscious access, indicating their disjunction (Block, 1995).

Because blindsight subjects differentiate between contents of stimuli, they express p-consciousness, but because they are not reflectively aware of whether or not stimuli were presented, they lack a-consciousness (Block, 1995).

The a-consciousness side of Blocks' argument is relatively uncontroversial and can be seen as represented by the PrC measure, PrC varying in accordance with the extent to which subjects have reportable / accessible awareness of the stimuli. The on-line experiments (Chapter 3) therefore show that a-consciousness is primarily supported by recurrent late processing, as it was most susceptible to disruption when TMS was applied at later times. Additionally, a-consciousness may also receive support from R/M pathways practically during the very early and very late stages of visual processing. The off-line experiments (Chapter 2) could be interpreted as displaying a potentiation of a-consciousness, and it may be the mechanisms governing which representations become accessed and which benefit from the application of occipital cTBS. The more controversial side of Block's argument is the existence of, and the way in which, p-consciousness can be understood.

The problem Block's theory raises for this series of experiments and their interpretation is that what has been classed as lacking in consciousness might actually be a form of consciousness (p-consciousness). According to this interpretation, the experimental contrast between the two measures fails to reveal anything about consciousness *per se* but rather highlights a difference between different forms of consciousness. My aim in this section is not to show that Block's interpretation of the phenomenon of blindsight is incorrect (for arguments of that type see peer comments in Block, (1995, 2007) and Cohen & Dennett, (2011)). Rather, Block makes a prediction which some of the findings of these experiments may inform.

The prediction made by Block, which is tractable in these data sets, is that phenomenal consciousness will be accompanied by recurrent processing (pp 498 Block, 2007). The 'on-line' experiments disrupted what can be regarded as recurrent processing, therefore was concurrent phenomenal consciousness suppressed? In order to tackle this, the additional question must be posed: What responses made by the subjects should be associated with presence of phenomenal consciousness?

The attribution of phenomenally and access consciousness status to any of the measures is open to interpretation. In earlier writings of Block (circa 1995) p-consciousness is associated with 'unseen' discrimination ability (PcU) as Block cites blindsight as a case for

p-consciousness. If so, then the previous discussions in Chapters 2 and 3 can be framed in terms of the p- / a- consciousness disjunction (reading in p-conscious for reportedly 'unseen' abilities, PcU) and support for Block's theory can be found. Although the multiple subject study (Experiment 3A) did not support (or refute) the correspondence between recurrent processing and p-consciousness / PcU, Experiment 3B and 3C did. In particular, Experiment 3C showed that when TMS was applied over later rather than at earlier times, disrupting recurrent processing, suppression in PcU / p-consciousness was observed (see figure 3.6.2.2), mirroring Block's prediction (Block, 2007). This same pattern and support for Block's theory is also suggested by the trend for TMS to be disruptive of PcU when applied over later times in Experiment 3B (see figure 3.5.2.2).

In later writings (Block, 2007, 2011) Block appears subtly to shift his position by emphasising the way in which experiments such as those of Sperling (1960)[13] express phenomenal consciousness in the form of a vague feeling or gist-like quality. This is perhaps in response to criticisms made by many authors (see peer responses in both Brain and Behavioural Science articles (Block, 1995, 2007) and (Cohen & Dennett, 2011)) along the lines of what defines blindsight (which Block in earlier writings cites in support of p-conscious) is a lack of consciousness, not the presence of phenomenal consciousness. This has been interpreted as 'phenomenal consciousness lite' (pp 520 O'Regan and Myin in Block, 2007). One interpretation of this later position could be that this 'phenomenal consciousness lite' corresponds to a subject's responding positively to the 'something?' question, as opposed to their explicit veridical awareness of the 'arrow' (PrC) which might be thought of as representative of their conscious access.

The rate of positive responses to the 'something?' question at each data point was computed. This measure compiled the number of 'yes' responses to the 'something?' question only (i.e. when the subject did not acknowledge awareness of the arrow) relative to the number of opportunities for such a response. If responding positively to the 'something' question can be attributed to the presence of p-consciousness, which involves recurrent

---

[13] The Sperling experiments (Sperling, 1960) involved showing subjects arrays of up to 4×3 letters or numbers. When an attentional cue was provided the subject could report any sub-set of the content of the array – i.e. they had access. When no cue was provided they reported a gist-like impression of all the characters, yet were only able to report correctly a reduced sub-set of what had been presented. Block discusses this pattern of results as an illustration of p-consciousness outpacing a-consciousness (Block, 2007).

processing, then the resulting hypothesis would be that the rate of reporting of 'something' (in trials without a 'yes' to the 'arrow?' question) should be suppressed by TMS applied at later times, disrupting recurrent processing. There are of course many more potential ways in which the phenomenal / access difference can be understood in terms of allocation of responses to measures. However the use of rate of reporting 'something' seems to fit the descriptions offered Block and makes a minimal number of assumptions. The analysis of this measure is described in appendix A5; no clear evidence was obtained for a correspondence between subjects' acknowledging awareness of 'something' and relatively late occipital TMS.

The question posed in this section was whether or not p-consciousness was accompanied by recurrent processing, following Block's prediction (Block, 2007), if one were to adopt an allocation of response patterns to the p- / a- distinction offered in earlier writings (Block, 1995), then the data may be consistent with Block's differentiation. That is, the residual abilities of blindsight have been linked with p-consciousness (Block, 1995) and here, when recurrent processing was disrupted by later TMS so too were residual 'unseen' (p-conscious) abilities, matching Block's prediction.

If, on the other hand, the measures of 'something' represents p-consciousness, then the data shows no correspondence with the disruption of later recurrent processing (see appendix A5). Hence, this measure does not provide evidence in favour of Block's later interpretation. A further distinction offered by Block in later writings, is the acknowledgement of the existence of 'unconscious' processing in addition to the p- / a-distinction (Block, 2011), under this, unconscious processing might be seen as represented by the PcU measure. If so, the disruption of PcU's at later times (Experiment 3C) relative to the absence of later disruption of the 'something' measure can be seen as evidence against Block's theory. This is because the data here suggests that 'unseen' / unconscious abilities are more sensitive to (therefore more likely to be dependent upon) the disruption of recurrent processing than is the 'gist-like' (phenomenal) acknowledgement of seeing 'something'.

Perhaps the more widely acceptable interpretation of the data and measures here (that is consistent with both earlier and later writings of Block and Lamme's interpretation of the p-/a- distinction (see peer review responses pp 511 Block, 2007)), is that the PcU measure, rather than corresponding to a form of consciousness directly, is indirectly supported by access consciousness. That is, conscious reflection upon the *question* to which the subject has

conscious access mediates the reportedly unseen abilities. The TMS applied at later times being effective upon the conscious reflection on the question, which itself is likely to involve recurrent processing (section 3.6.3). Therefore, the PcU measure simply does not reflect p-consciousness. Furthermore, the option is available to Block and his proponents to argue that the reason for which the 'something' measure failed to demonstrate a correspondence with the later intervention was that it simply failed to capture the intention behind Block's description of p-consciousness.

Overall, if one chooses to adopt Block's differentiation, then the data do not refute that belief, particularly if the understanding is based on the early description of p-consciousness. The later differentiation appears to be less well supported by the data. However it should be noted that this discussion has largely focused on the single subject's data and the conclusions that can be drawn from case studies must always be tentative. Therefore further to elucidate Block's theory, a between-subjects replication of the single-subject double-pulse study (3C) would be required. The task might also be developed more closely to reflect Block's differentiation, possibly via the adaptation of a task similar to the Sperling experiments, but in conjunction with event related / on-line TMS. As it stands, it remains possible to interpret these experiments as exemplifying a difference between two forms of consciousness rather than a difference between consciousness and a lack thereof.

4.7 The hard problem

The 'hard problem' places limitations on the strength of conclusions that can be drawn from research such as that which has been presented here. No novel attempt will be made to surmount the 'hard problem' here. Rather a brief overview of the problem will be given. The next section will then point the reader towards literature that has claimed to neutralise the problem and has been used as the epistemic basis for this research.

The hard problem of consciousness is essentially that it is entirely possible that all the neuronal architecture that we describe as the physical basis of consciousness could occur in the absence of consciousness. This is the basis of Chalmers' 'zombie' argument: creatures with all the physical attributes of anyone but are devoid of consciousness (Chalmers, 1996).

The hard problem places limits on the extent to which physically orientated science can explain consciousness. It can be attributed to a derivation of Cartesian dualism, but, more specifically, is based on Kripke's demonstrations (1980): regardless of the description of the neural basis of pain the designation of the physical processes to pain will never be as rigid or necessary as, say, water = H_2O. This is because it is possible that there are creatures that experience pain but which share none of the physical attributes of the original designation.

A closely related element of the argument is the 'explanatory gap'(Levine, 1983), which is probably most clearly expressed in Nagel's demonstration that regardless of our state of scientific knowledge as to bats' echo location we might still never know 'what it is like to be a bat' (Nagel, 1974). The doubt raised by these arguments is that there is something further to explain in phenomenal first person experience than any explanation rooted in pure third person objective science can achieve (this notion of third person objective science is most clearly captured in Dennett's heterophenomenology; see Dennett, 1991; Dennett, 2003).

4.8 The relative position and neurophenomenology

The goal of understanding the relationship between consciousness and its physiological underpinnings is one of the greatest challenges in modern science (Miller, 2005) and philosophy (Metzinger, 1995).The strategy initially adopted and outlined in Chapter 1 seeks to explore the coupling between physiology and phenomenology, and is known as neurophenomenology (Varela, 1996; Varela, et al., 1991). This section starts by discussing one of the key contentions in consciousness research which places limitations on what can and cannot be said in this type of research: the use of 'necessity' in the 'hard problem' (Chalmers, 1996; Chalmers, 2000) and the search for the 'neural correlates of consciousness' (Koch 2004) (4.8.1). The neurophenomenological position is then clarified and discussed with respect to some of the other problems encountered (4.8.2). How neurophenomenology relates to the other two positions (HOT/GW and the differentiation between p- and a- consciousness) is discussed in the final section (4.8.3).

4.8.1 The use of necessity

A key move in the development of the hard problem is the rigidity or necessity of some relationships in science (e.g. water = H_2O) in contrast to others involving phenomenal consciousness (Chalmers, 1996; Kripke, 1980). My opinion, derived from the work of Varela, is that it is the call for necessity that is the cause of the intractability of the hard problem. Necessary relations are notoriously easy to refute, which is particularly the case in such a massively parallel system such as the brain.

The search for the 'Neural Correlates of Consciousness (NCC)'(Chalmers, 2000) is often framed in terms of identifying the 'necessary' and 'sufficient'[14] conditions under which consciousness exists (e.g. Koch, 2004). This use of necessity sets the bar for classification of NCC's too high, as claims are too easily refuted. For example, Ffytche and Zeki (Barbur, Watson, Frackowiak, & Zeki, 1993; Ffytche, Guy, & Zeki, 1996; Ffytche & Zeki, 2011; Zeki & Ffytche, 1998) claim that the primary visual cortex and its involvement in recurrent processing is not necessary for consciousness, as has been previously claimed (Lamme, 2001; Pascual-Leone & Walsh, 2001; Silvanto, Cattaneo, Battelli, & Pascual-Leone, 2008; Silvanto, Cowey, et al., 2005; Silvanto, Lavie, et al., 2005; Tong, 2003). They show that subjects with lesions to their primary visual cortex can express a markedly reduced form of visual awareness when they are presented with moving stimuli and the authors attribute the residual awareness to pathways which parallel the main striate projections, directly innovating areas such as MT (Riddoch syndrome e.g. Ffytche & Zeki, 2011). Whilst this reasoning does appear valid (negating previous claims of necessity), it does not follow that the primary visual cortex could not be the site at which many of the changes occur that are causally

---

[14] The other element to the search for the NCCs involves the requirement that the description covers all 'sufficient' conditions. As noted previously (1.2.1) this is not particularly useful in the development of criteria for determining what is involved in consciousness, unless it is qualified by the use of 'minimally sufficient' (Chalmers, 2000) as here, i.e. uncontroversial. This is because, for example, a circulatory system is required as a sufficient condition for the conscious system to operate, yet its function contributes little to our understanding of consciousness. Furthermore, it is entirely possible that what constitutes some conscious processes extends beyond the body boundary into the world (Clark & Chalmers, 1998) further reducing the utility of relying upon sufficiency as a criterion.

constitutive of consciousness[15]. This is especially relevant in cases where the form of consciousness under consideration is the everyday veridical visual experiences that were not displayed by the subjects in the Ffytche and Zeki experiments, but which were the target of the experiments developed here. The data presented here suggests a prominent role for early visual areas such as V1 in consciousness.

The on-line experiments also took advantage of the relative ease with which claims of necessity are refuted. Ro et al (2004) claimed the superior colliculus to be necessary for the residual abilities of TMS-induced blindsight and it was demonstrated here that such residual abilities can operate via projections in the geniculate stream. Importantly, no claim was made that the superior colliculus is *not* involved in the performance of some related tasks. For example, tasks involving saccadic eye movements seem likely candidates for the involvement of the superior colliculus (Ro, et al., 2004; Walker, Fitzgibbon, & Goldberg, 1995). By avoiding a reliance upon claims of necessity it is possible to develop a more nuanced understanding of brain processes and consciousness, where different pathways contribute different inputs, the content of which – and the contribution made by each – is dependent upon the content of the experience caused by the task requirements.

What then is the alternative to calling upon necessity? John Campbell has offered the interventionist approach, which discusses relations in terms of causally constitutive conditions and emphasises that the way in which these relations should be explored is via interventional techniques, such as TMS (see Campbell, 2007; Campbell, 2008). This has been one of the strategies adopted here (also see Chambers & Mattingley, 2005). Perhaps more pertinent to the current concerns regarding the hard problem and claims of necessity are the arguments and strategies developed by Varela et al, in terms of it being the relative difference between states that is of critical importance in the development of understanding of the nature of consciousness (Varela, 1996; Varela, et al., 1991). If our notion of water and H_2O arise in relation to one another and neither of them is fundamentally grounded beyond co-dependent observation, then we have no reason to think the relationship between water and H_2O is set under all possible conditions, and therefore all such relations are to some extent contingent (Varela, 1996; Varela, et al., 1991). This both implies that the use of necessity as a rigid fixed

---

[15] The same structure of argument applies to Crick and Koch's statement that the visual cortex is not necessary as the neural realiser of consciousness (Crick & Koch, 1995). See also Block (1996) for a similar argument in reference to Crick and Koch's claims.

designation is ill-advised, and that relations between brain processes and phenomenal experience is not as fundamentally intractable as the hard problem suggests. The dualist arguments of Kripke and Chalmers rest upon their bringing into question the necessity of the relationship between consciousness and physiology. If the emphasis is placed on this relationship, rather than the attempt to maintain claims of necessity, then the hard problem does not seem quite so intractable.

4.8.2 Neurophenomenology

Neurophenomenology has been presented as a methodological solution to the hard problem (Varela, 1996). Neurophenomenology is fundamentally an epistemic strategy which regards neither the level of subjective phenomenal conscious experience nor objective physical architecture as fundamentally grounded (Varela, et al., 1991). Rather they arise in relation to one another[16] and therefore the coupling between the two (1st and 3rd person perspectives) is where we should focus our research efforts. It is epistemic in that it accepts that the subject and the experimenter have different information available to them[17] and refocuses enquiry on the relations and coupling between the two.

Central to the neurophenomenological approach is the understanding that properties such as consciousness are relational, i.e. properties only exist in relation to other properties. Such a position lessens the impact of problems such as reactivity on the interpretation of experiments. Because the emphasis is placed upon the differences between states, the difference between the measures of acknowledged awareness and perception when the

---

[16] In terms of acknowledging the historical derivation of neurophenomenology, it owes much to the continental phenomenological traditions of Husserl and Merlo-Pontey, from which the notion of epoché (see section 1.3.2.3) is taken and an emphasis is placed on the physical embodiment of the subject. The methodology is also consistent with and indebted to elements of Buddhist scripture, in particular the notion of co-dependent arising and baselessness (Varela, et al., 1991), from which are derived the notion that neither the physical nor the phenomenal are fundamentally grounded (exist in isolation or are the *basis* for existence) and that the extent to which they do exist is in relation to one another (Varela, 1996; Varela, et al., 1991).

[17] To have an optimal understanding of the investigation we would have to be both experimenter and subject in the investigation (contra Dennett, 2003), in order to acquire both first person experiential knowledge and objective knowledge. It follows that I was a subject in the original TBS experiments. However, I was excluded from the on-line studies owing to my colour blindness.

subjects do not acknowledge awareness, is critical to the investigation. The fact that both measures may require some form of reflection in order to produce a response is secondary, reducing the problem of reactivity.

The high levels of practice / familiarity obtained by subjects, particularly during the calibration phase of the experiment, together with the simplicity of the task, further reduced the problem of reactivity. This aspect of the task follows suggestions made by proponents of neurophenomenology that the distinction between being aware and knowing you are (and consequentially being able to report that it is so) is reduced through practice and repetition as the report aspect becomes habituated (Dapraz, 1999; Gallagher, 2003; Lutz, 2003; Vermersch, 1999).

4.8.3 The middle level

The position of neurophenomenology in relation to HOT/GW theories and first order theories (such as Block's) has been described as 'intermediate'. Functionally intermediate theories describe the level at which consciousness arises as being at the level of grouping of elements into individual percepts (Jackendoff, 1987; Prinz, 2007). For example, in reading, low-level information such as edge detection may not be conscious (e.g. Marr's primary sketch, section 3.3.4, Marr, 1982), neither are high-level grammatical rules, which we have learned and have become habituated to. What we are conscious of is the intermediate level of processing of the words and their meaning (Jackendoff, 1987).

This intermediate position is neatly summarised by the introductory quotation, which is consistent with the view put forward by Varela et al, particularly in the *embodied mind* (Varela, et al., 1991) in the description of the 'middle way', and various other works describing neurophenomenology (Lutz, 2007; Varela, 1996; Varela, 1999).

The intermediate position expresses some of the attributes of both first order and higher order theories. Both first order theories (Block 2007, 1995) and intermediate theories consider the primary level at which consciousness arises to involve content and the other element to be access to that content. Both theories also accept there is an issue with the demonstration of phenomenal aspect of consciousness. For Block the difficulty lies in the demonstration of a form of consciousness which by definition can only be observed, in

isolation, in the presence of a negative report or absence of a positive report. Under neurophenomenology the difficulty in the demonstration of the phenomenal aspect is due to the epistemic structure of the investigation. Phenomenology tends to refer specifically to one side of the explanation – the contribution made by the subject. This hinders its interpretation from the third person perspective. Although the first order and intermediate positions are therefore related they are not completely comparable.

Intermediate theories such as neurophenomenology are probably more closely related to HOT theories (Baars, 2003; Lau & Rosenthal, 2011) than to first order theories. Varela et al., describe a theory of radical embodiment in which conscious phenomenology is embodied in the large-scale temporal dynamics of the brain (Thompson & Varela, 2001). More specifically, phase synchronisation of oscillations offers a mechanism for the large scale integration of information over multiple brain regions, thus providing a mechanistic and embodied framework for the implementation of properties such as awareness (Varela, Lachaux, Rodriguez, & Martinerie, 2001). These principles are similar to those of GW models in particular. Within this framework, alpha rhythms play an important role in the governance of perception via its segmentation. In one of the last articles Varela wrote, he suggested this is a promising avenue for research, noting that alpha "rhythms could provide the slower temporal framing for successive cognitive moments of synchronous assemblies" (pp 273 Varela, et al., 2001). This is consistent with the finding in Experiment 2D that the ERD in the alpha band was facilitated, coincident with the increase in conscious detection and with the interpretation of the ERD more generally as performing a gating function.

How this neurophenomenological strategy differs from HOT and GW theories is that HOT/GW theories are essentially third person descriptions only. Neurophenomenology, on the other hand, accredits the first person perspective with a domain of knowledge vital in the development of the enquiry. This places the first person and third person perspectives on equal footing in terms of the investigation as a whole (this directly contradicts the strategy outlined by Lamme (Lamme, 2006a) and Dennett (Dennett, 1991, 2003)).

First person subjective experience is part of the explanation offered by neurophenomenology: only the embodied subjects know precisely what is meant when any instance of report is made (Clark, 1997; Varela, 1996; Varela, et al., 1991). The strategy has been to couple the subjects' experience with a third person objective understanding, through the measures (Varela, 1996; Varela, et al., 1991). By acknowledging the epistemic structure

of enquiries oriented towards consciousness the possibility that there will always be something further to explain (4.6) may be rendered less problematic.

4.9 Future directions

This section is restricted to a list of potential developments from the findings and protocols specifically established here, rather than a general discussion of the future of this type of research.

4.9.1 Off-line experiments

- The intensity at which the cTBS is applied could determine under what conditions the cTBS protocol was inhibitory or facilitatory. Applying cTBS over a range of intensities may lead to a more highly differentiated understanding of its effects. This could aid its application in a clinical context (e.g. Kindler, et al., 2012) and allow for the use of cTBS in the improvement of other cognitive abilities.

- Application of the same TMS protocol over a range of tasks and questioning procedures might further elucidate the nature of the effect caused by the cTBS and its applicability to other domains. More specifically, is the increase in conscious detection only observable when the task in question involves the isolation of a target from a background noise? Is the increase in detection reflected by an increase in reflective confidence potentially measured by Overgaard's PAS scale or post-decision wagering (Sandberg, et al., 2010)?

- The ERD's involvement in the gating of awareness and the method of quantifying the ERD according to its gradient of onset are promising avenues for future research. The change in the ERD in relation to a psychometric function representing conscious detection across a range of stimuli intensities might be a relatively informative and simple initial experiment.

- Differences in the MEG data in terms of the contrast between reportedly 'seen' and 'unseen' yet correct trials should, theoretically, be capable of revealing processes involved in

consciousness specifically. This contrast was not implemented here because the division of the data between analysis blocks and TMS conditions meant that there was relatively low power (number of trials) for each condition ('seen' and 'unseen' correct). Additionally the power of each condition was determined by the behavioural changes in that block, which would have made it difficult to claim that any change in the MEG data was not driven by this power difference. Experiments specifically directed towards this difference between 'seen' and 'unseen' conditions may also reveal reactive processes.

- A common assumption in neuroimaging, particularly functional magnetic resonance imaging (fMRI), is that an increase in activity (as indicated by blood oxygen level dependent (BOLD) response) relates to increased processing (Logothetis & Pfeuffer, 2004). The claim here is that, following cTBS, active suppression is potentiated and conscious detection is facilitated. Might this therefore result in a decrease in basic BOLD signal in the affected region (Allen, Pasley, Duong, & Freeman, 2007), where increased processing is apparent through other measures? Could different analysis techniques (e.g. mass univariate vs. multivoxel pattern analysis, Kamitani & Tong, 2005) be compared within the context of such a change? Such a result might inform our understanding by attenuating the link between BOLD changes and functionality.

4.9.2 On-line experiments

- The finding that both very early and very late activity in visual areas, related to visual consciousness, appear to be supported by input from the retinotectal and/or magnocellular routes was made without an *a priori* hypothesis. Confirmation of this result through replication is therefore desirable.

- Finer granularity of pulse timings in the between subjects study would further elucidate the temporal dynamics of information-flow through the early visual system. Through the addition of a second TMS coil (e.g. Silvanto, Cowey, et al., 2005), potentially over frontal areas such as the dorsolateral prefrontal cortex (Lau & Passingham, 2006), the recurrent and frontal dependency of such processing might be further understood.

- A key demonstration of Experiment 3A was that the residual abilities of TMS induced blindsight did not require input from the retinotectal pathways. However, this leaves open the

question as to whether the retinotectal pathway contributes to the residual abilities required to perform other tasks that are illustrative of TMS-induced blindsight. The task here might be easily adapted to include a peripheral cue or saccadic response, both of which are more likely to involve input from the superior colliculus (Ansorge, 2003; Leh, Mullen, et al., 2006; Mulckhuyse & Theeuwes, 2010).

-The main finding of Experiment 3C, that 'unseen' discrimination abilities were susceptible to pairs of TMS pulses applied at later times, was an intriguing finding which requires confirmation in a larger cohort. However, it is possible that there may be difficulties in recruiting a large group of participants that can tolerate the high levels of TMS which may be required to produce such effects.

- It is possible to develop an automated task and calibration system, similar to that presented here, which could operate over the internet. This would allow us to probe the extent to which reportedly 'unseen' abilities are above chance across a large population and range of criteria differences. Such a potentially large data pool might also be able to probe the contribution made to perception when subjects report stimuli as 'unseen', made by different pathways as probed by chromaticity differences. Not only would such a design be capable of probing such questions, it would also provide an ideal opportunity for public engagement with consciousness research.

4.10 Closing remarks

This research has attempted to contribute to our understanding of consciousness. To some extent, this has been achieved. Results confirmed that visual awareness, specifically, is supported primarily by activity in early visual areas from ~100ms post stimulus onset. Additionally, activity prior to this period also appears to contribute to consciousness, possibly via the provision of a pedestal of activity originating from the fast retinotectal and/or magnocellular pathways. Furthermore, some input from these pathways may first pass to more frontal regions and then contribute to later re-entrant conscious processing. In contrast, perception lacking in consciousness may operate via geniculate pathways, as well as the previously demonstrated collicular projections (e.g. Leh, Mullen, et al., 2006). The increase

in specifically conscious detection following the application of the cTBS is also a novel finding. The cTBS protocol was also associated with a sharpening of the ERD gating response and an increase in the levels of the inhibitory neurotransmitter GABA. These results collectively suggest a prominent role of early visual areas in the formation of conscious percepts involving selection and gating by inhibition.

Appendix

A1. Cross correlation

A1.1 The relationships between measures

A1.1.1 Introduction

The multiple measures employed in Chapter 2 allowed us to probe the relationships between the processes measured. For example, do changes in the gradient of the ERD predict the increase in conscious detection? Further, can causal relationships between such variables, in terms of moderation and mediation, be established (Baron & Kenny, 1986)? The claims made in reference to the gating-by-inhibition explanation would clearly be greatly strengthened by the demonstration of such interrelations. This section offers a tentative exploration of these relationships.

The relationships between processes measured can be described by correlations, making use of the variability in effect size, across subjects. The previously demonstrated differences between TMS conditions contributed to the gating-by-inhibition explanation, and can also be described as correlations between TMS conditions. This involved treating active vs. control TMS as a dichotomous discrete variable, entry of which into a Pearson's linear correlation being equivalent to an unpaired two tailed t-test between conditions (denoted as the 'TMS' variable in figures). Partial analysis (Baron & Kenny, 1986) involves the removal of the variance of one variable from a correlation that describes the relationship between another pair, thereby assessing the role of the third variable in mediating the relationship between the original pair. For example, we have demonstrated that the increase in conscious detection and the ERD appear to result from the application of TMS. If this significant effect of TMS on PrC is eliminated by the removal of the variance observed in the ERD, but

removal of the variance in PrC from the analysis that demonstrated a change in the ERD had no effect, then the shift in conscious detection would appear to be dependent upon (and linearly mediated by) the alpha band ERD. That is, the effect of the TMS upon conscious detection would be explained through the changes in the ERD as a result of the TMS. Under such conditions the effect of the TMS upon the ERD and the subsequent partialling out of the variance of the change in conscious detection would be required to assess the causal direction of influence. This section implements such a partial analysis strategy and additionally describes the relationships between measures as affected by the TMS by depicting the correlation of effect sizes (change from control TMS condition) observed for each subject across the measures of interest.

The partial analyses and correlations of measures were restricted to data from those subjects who participated in multiple experimental sections. The dependent variable for each measure entered into the partial analysis comprised of a single vector, baselined to pre-TBS levels, collapsed across post-TBS blocks. The partial analysis was applied to the following pairs of measures which have contributed to the gating-by-inhibition explanation: conscious detection (PrC) with the gradient in the ERD (A1.1.2), PrC with the change in GABA concentration (A1.1.3) and the change in the ERD with the change in GABA concentration (A1.1.4). All of these partial correlations were studied in relation to the presence of active vs. control 'TMS' (the independent variable). A further analysis is also applied (A1.1.5) which takes the change from control data for all three dependent variables (PrC, ERD and GABA) and applies a partial analysis to these change from baseline (pre-TBS) measures. Because only three subjects participated in the phosphene threshold experiment who had participated in the other experimental sections the cross correlation involving this measure was not computed. Outliers in terms of the correlation were identified, using Cooks distance and excluded from the corresponding analysis (Cook, 1977). Because of the relatively low power of these analyses, they should be treated as exploratory and tentative.

A1.1.2 Analysis: PrC and ERD

Seventeen subjects' data was available for entry into the partial analysis, as one subject was shown to be outlying in terms of the change in the ERD. The partial analysis of these measures is summarised in figure A1.1.2.1. No correlation representative of the difference between TMS conditions was observed. Therefore, removal of variance in other measures cannot be interpreted. Furthermore, no relationship between the size of the changes in conscious detection and gradient of the ERD following the application of the cTBS were observed (see figure A1.1.2.2., $R^2 = 0.07$, p=0.30, R_S =0.196, p=0.45).
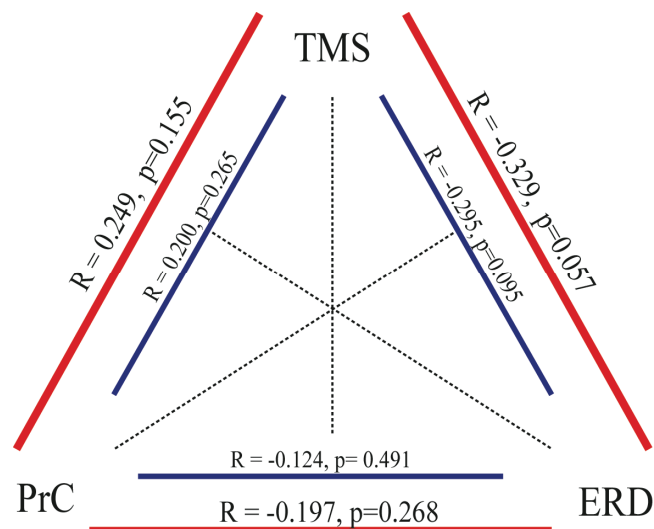
Figure A1.1.2.1 Summary of partial analyses applied to the measures of conscious detection, gradient of the ERD and the effect of the TMS. Red lines correspond to the bivariate Pearsons correlation between two measures and the blue lines correspond to the same correlation following the removal of the variance of the third variable from the partial analysis. Subsequent partial analyses conform to this structure of presentation.
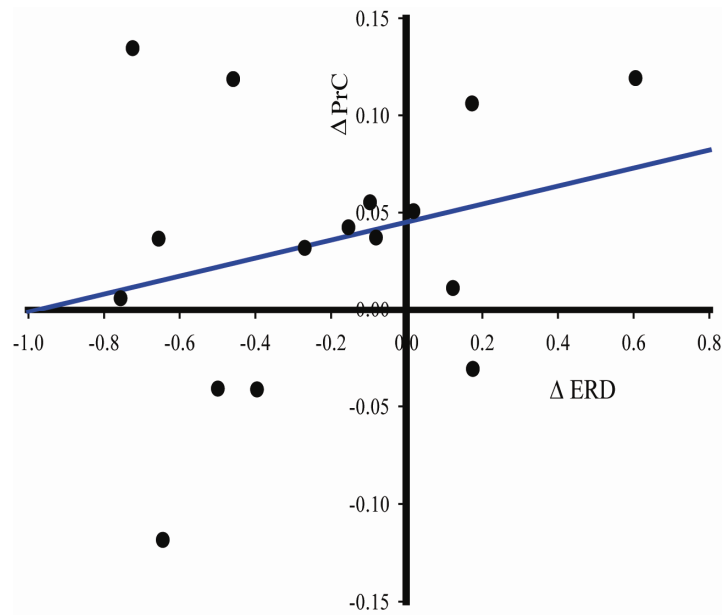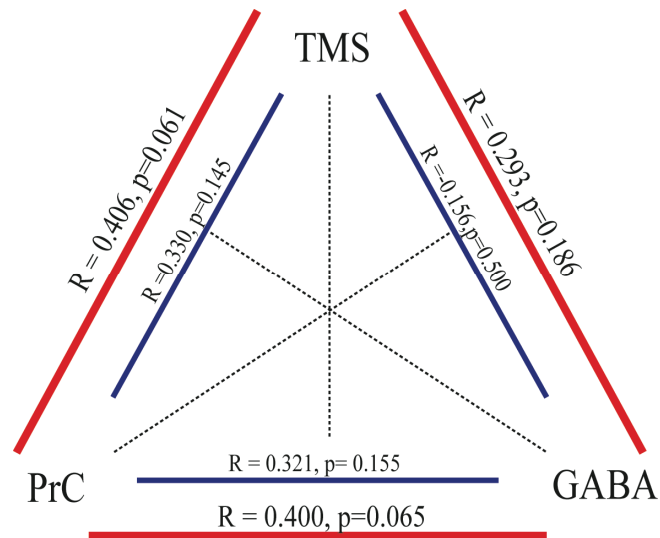
Figure A1.1.2.2 Linear correlation (blue line) between change in conscious detection (PrC, SDT derived units) and change in event related desynchronisation (ERD, Units tesla per second $\times 10^{-14}$). Data points are individual subject's data baselined to pre-TBS levels, then to control conditions, collapsed across post TBS blocks. Note: The more negative the change in the ERD the greater the shift or the steeper the gradient of the ERD.

A1.1.3 Analysis: PrC and GABA

12 subjects participated in the behavioural and MRS sections to these experiments, one of whom was excluded from the correlation analysis on the basis of being an outlier according to Cook's criteria. In order to match the data between the measures, only the first three post-TBS analysis blocks were collapsed across to form the vector representing the change in PrC (i.e. since there were three MRS acquisitions post-TBS). Three subjects participated in both original and replication behavioural experiments for which the relevant data was averaged across repetitions.

No significant primary relationships between the variables was revealed by the correlation analysis (see figure A1.1.3.1-2).

Figure A1.1.3.1 Summary of partial analyses of the relationship between the measures of conscious detection, GABA concentration, and application of the TMS.
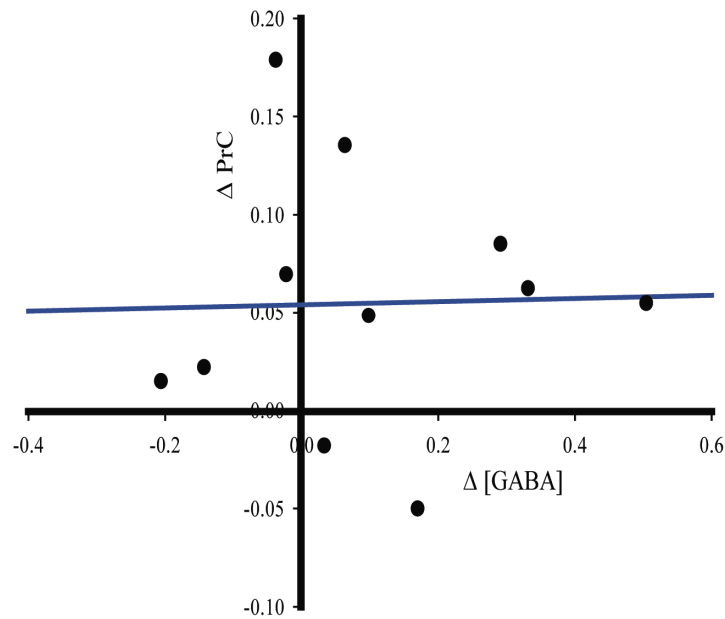
Figure A1.1.3.2 Linear correlation (blue line) between change in conscious detection (PrC, SDT derived units) and change in GABA concentration ([GABA], institutional units). No correlation was observed between the changes in these two measures following the TBS ($R^2 = 0.03$, p=0.94, $R_S = 0.091$, p=0.79).

A1.1.4 Analysis: ERD and GABA

Nine subjects participated in both the MEG and MRS experiment and so were entered into the partial correlation analyses. No significant correlations were observed.
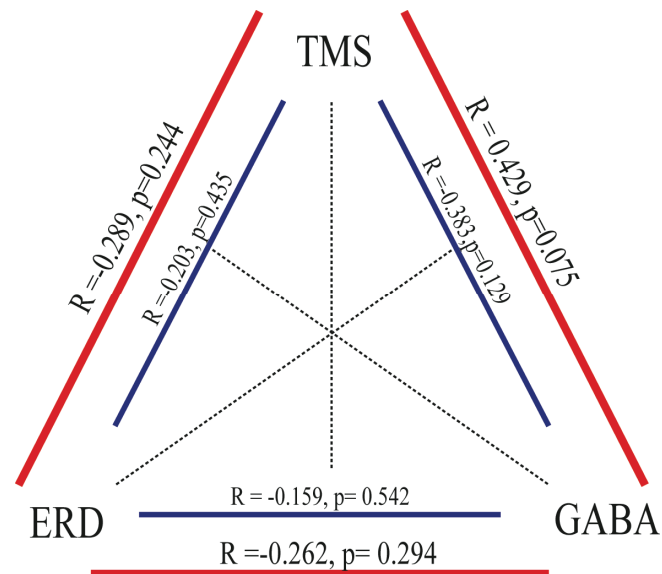
Figure A1.1.4.1 Summary of partial analyses applied to the relationships between the measures of the gradient of the ERD, GABA concentration, and application of the TMS.
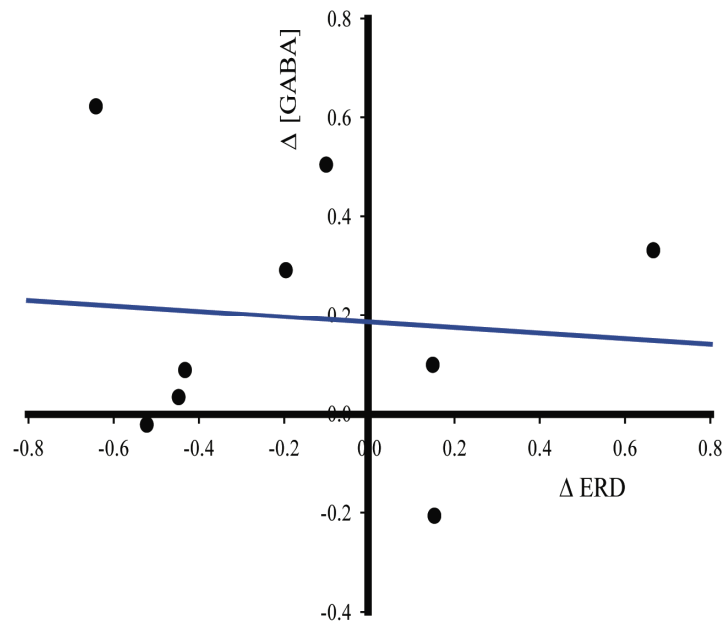
Figure 2.7.1.4.2. Linear correlation (blue line) between change in GABA concentration ([GABA], institutional units) and change in event related desynchronisation (ERD, Units tesla per second $\times 10^{-14}$) following application of cTBS (R^2 = -0.089, p=0.820, R_S=-0.050, p=0.898).

A1.1.5 Analyses of change in PrC, ERD and GABA

Nine subjects participated in all experimental sections so could be entered into this analysis. Each dependent variable here is a vector of the change in that measure, for each subject, from control levels. No significant correlations were demonstrated.
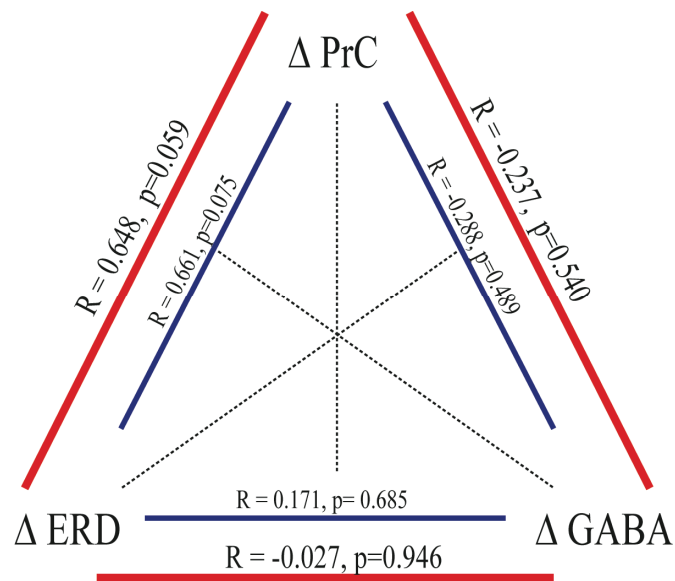
Figure A1.1.5. Summary of partial analyses applied to changes in the conscious detection PrC, ERD gradient and GABA concentration, from control levels. That is the between-subjects correlations between dependent variables and the effects of removing the variance of the third variable from these relationships.

A1.1.6 Discussion

Because no statistically demonstrable relationships were shown, no conclusions as to the interdependence between the processes targeted by the measures can be drawn here. This was not altogether unexpected for three reasons. First, there seems to be little theoretical reason to suppose that the relationship between the properties measured might be in any way linear: if other factors (such as individual differences in rates of metabolic turnover) intervene, between, for example, the change in the ERD and the change in GABA concentration, then we have no reason to think that a potentiation in one might predict a direct increase in the other.

Second, these analyses were lacking in sufficient statistical power. Primarily, this was limited by the availability of subjects to participate in experiments conducted over the course of roughly 2 ½ years. In an attempt to improve the power of these partial analyses we conducted the replication experiment on 18 subjects (the last 4 of whom were available for

both MRS and MEG acquisitions), whereas, a direct replication of the original experiment and strict enforcement of a stopping rule (see Dienes, 2008) might have meant the use of 16 subjects. The relatively low power of these analyses is likely to be the main cause of failure to demonstrate the TMS-related effects previously shown.

The final contribution made to this failure to demonstrate effects which had previously been shown was that the correlation method makes use of the variability between subjects. In the repeated measures analysis applied in previous experimental sections, this source of variance was effectively removed, improving the reliability of the repeated measures analysis. Therefore were this set of analyses to be perused further, a mixed-model approach should be adopted (e.g. West, Welch and Galecki, 2006).

A1.2. Correlation with Motor Threshold

A1.2.1 Introduction

The aim of this section is to address the question of whether the effect of cTBS on awareness is intensity-dependent, with relatively small levels of inhibition facilitating awareness while larger inhibitory levels are disruptive.

The mechanism proposed for the increase in conscious detection involves a facilitation of the mechanisms that select through inhibition. Because this mechanism is centred upon inhibition, it follows that were the levels of inhibition to have been increased further, beyond functionally optimal levels, then more pervasive suppression of representations, including conscious processing, may be expected. Consistent with such an effect are previous demonstrations of reversals of TMS effects, where lower levels of TMS can facilitate processing, whereas higher levels are shown to be suppressive (Abrahamyan, et al., 2011). Here, facilitation was displayed and the intensity at which the TMS was applied was relatively low, compared with previous demonstrations of suppressions; 42% of total output compared to 46% in the Franca et al., study (2006). This may be due to the method used here for obtaining the TMS intensity being based on motor threshold (via the observed resting contraction across hand muscles) which is one of the most sensitive threshold

measures and thus results in relatively low thresholds (Hanajima et al., 2007; Varnava, et al., 2011).

TBS applied to occipital regions has been described as being capable of reducing metacognative abilities, indicative of TMS-induced blindsight (Rahnev, et al., 2010). This has the opposite directionality to the effect described here. If lower intensity applications facilitate abilities whereas higher intensities are suppressive then the findings of Rahnev et al, and those presented here may be reconciled if the intensity they used was relatively high. The analysis of this section was designed to test this possibility. However, the description of the experiment by Rahnev et al (2010) has only appeared in abstract form, so it is very difficult directly to compare studies. The suggestion being made here is that the TMS intensity used here is likely to be less than those of the Rahnev et al, study or Rahnev et al, found the efficacy of the cTBS to be inconsistent over a range of intensities. The same principal investigator has published a more extensive account of the application of cTBS to the dorsolateral prefrontal cortex (DLPFC), which again reduces reports of meta-cognitive / conscious awareness of stimuli that dissociate from forced-choice ability (Rounis, et al., 2010). However the difference in target area, together with the absence of information about stimulator output levels again makes comparison difficult (Stokes, et al., 2005).

Subjects' motor thresholds[18] vary according to individual differences in cortical excitability. This limits the extent to which differences in the motor threshold can be expected to correlate with effect size differences in the dependent measures, if they are dependent on cortical excitability. However, because there appears to be no (Stewart, et al., 2001) or only weak (Deblieck, Thompson, Iacoboni, & Wu, 2008) correlation between motor and visual

---

[18] The reason TBS was applied as a fraction of motor threshold, here as opposed to phosphene threshold (which would be more likely to track occipital cortical excitability) was that phosphene thresholds are higher than motor thresholds. Mean phosphene threshold collected over experiment 2C = 71.1% stimulator output ±15.3 SD (based on all 12 subjects), mean motor threshold = 51.9% stimulator output , ±6.2 SD (based on 12 randomly chosen subjects from the behavioural experiment); this is consistent with (Stewart, Walsh, & Rothwell, 2001). Applying cTBS at this lower intensity reduced the risk of adverse reactions to the TBS and was essential to ensure that the stimulation fell within international safety guidelines (Maizey, In Press; Rossi, Hallett, Rossini, & Pascual-Leone, 2009; Wassermann, 1998). Furthermore, a high proportion of subjects simply cannot perceive phosphenes within the range of TMS intensities applicable. This also was the reason why only three subjects participated in both the phosphene experiment (2C) and the behavioural experiments (2 A and D). Additionally, phosphene thresholds are less consistent across subjects (Stewart, et al., 2001). For these reasons, motor threshold was used to determine the intensity of stimulation, which resulted in relatively low intensities at which the TBS was applied.

cortical thresholds for excitation and there are several additional sources of variation in the determination of motor threshold (e.g. subtle irregularities in angular coil placement (Kammer, Beck, Thielscher, Laubis-Herrmann, & Topka, 2001)), there is a degree of variability in the range of TMS intensities applied. This might reflect differences in the efficacy of the TBS when applied to the occipital lobe. Such variability might be sufficient to demonstrate correlations between the intensity of TMS applied and dependent measures, which might inform experimental hypotheses. With regard to the central question of this section: if the cTBS applied here increased conscious detection through a mild increase in inhibitory processes, then stronger potentiation of inhibition caused by higher intensity cTBS might lead to a reduction in conscious perception.

In order to probe this possibility two post-hoc analyses were applied which correlated the intensity at which TMS was applied with subjects' change in conscious detection. If relatively small increases in inhibition bestowed a benefit upon conscious detection, whereas larger increases in inhibition are suppressive, then a negative correlation between stimulator output (motor threshold) and the increase in conscious detection is predicted. The alternative method for describing this relationship was a median split, based on motor threshold. The hypothesis being that subjects in the group of lower motor thresholds would be more likely to express the increase in conscious detection than members of the higher TMS intensity group.

This method of correlation using motor threshold was not applied to the other dependent measures (ERD and GABA) for three reasons: i) such an analysis might constitute unnecessary multiple comparisons, ii) the difficulties with the correlations described in section A1.1.5 apply here, and, additionally iii) there is the limitation upon such a correlation owing to motor threshold's reflecting individuals' cortical excitability (mentioned above).

The analysis of the relationship between TMS intensity and the effect of the cTBS upon conscious detection made use of the data collected over both the original and replication results. Outlier rejection was performed according to Cook's distance (Cook, 1977) applied to the concatenated group and resulted in the elimination of 1 subject. With respect to subjects who participated in both behavioural experimental sections (4 subjects), the variables used applied an average across the two experiments. This resulted in the inclusion of 29 subjects in this analysis. Since cTBS was applied at 80% of motor threshold, the subjects' motor thresholds were correlated with a variable representing subjects' change in conscious detection. This variable comprised change from pre-TBS baseline in the PrC data, and was

subtracted from control data and averaged across post-TBS blocks. The correlation was a least-means-square linear correlation.

Additionally, a median-split analysis based on motor threshold was applied to the behavioural change in PrC. The median split divided the group into two (13 subjects) according to the subjects who received higher and lower levels of TMS. Because three subjects' motor threshold was at the median level (53% of total output), they could not be allocated to one or other of the groups so were excluded from the median split analysis.

A1.2.2 Results and discussion of correlation between change in PrC and motor threshold
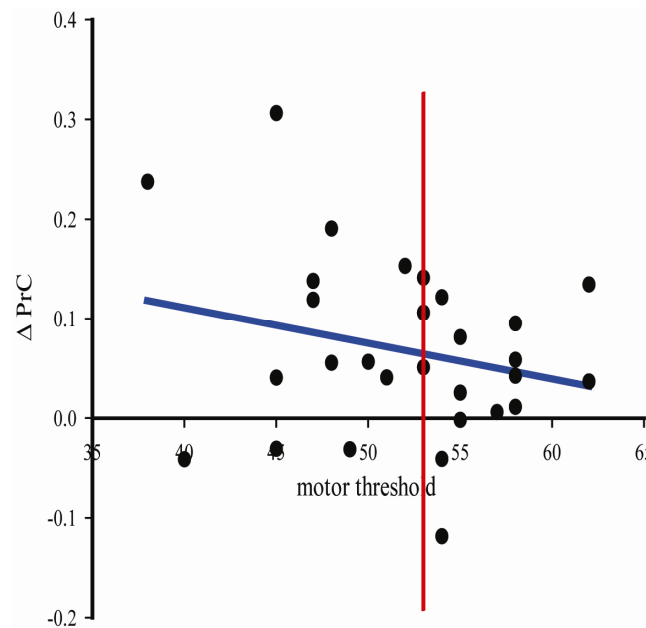
Figure A1.2.2. Correlation between motor threshold (the determinate of the intensity at which TBS was applied) and the change in conscious detection following the application of cTBS relative to controls. Blue line is the linear correlation where $R = -0.238$, Adjusted $R^2 = 0.021$, $F_{(1,28)} = 1.61$, $p = 0.216$. The red line indicates the position of the median split.

The correlation between the intensity at which the TMS was applied and the benefit in conscious detection was not significant ($R = -0.238$, Adjusted $R^2 = 0.021$, $F_{(1,28)} = 1.61$,

p=0.213), but it was in the negative direction, which suggested that increased conscious detection may be more likely to result from applications of cTBS when the intensity used is relatively low. Consistently, a significant increase in conscious detection was only observed in the group of subjects where a lower intensity was applied ($F_{(1,12)}$ = 10.08, p =0.008), whereas the group for whom the intensity applied was higher, no clearly significant effect of the cTBS was observed ($F_{(1,12)}$=3.35, p=0.092).

These analyses tested the hypothesis that lower levels of cTBS might facilitate reported awareness, whereas higher levels of TMS might lead to suppressions. The conclusions that can be drawn in relation to this question are limited. The correlation between the TMS intensity applied and the change in conscious detection was not significant and caution should always be applied when attempting to draw conclusions from p values around 0.1, as with the absence of effect in the median split analysis of the higher intensity group. However, the direction of the potential correlation suggests that the benefit bestowed by cTBS might be greater when the intensity used is relatively low. This would be consistent with the gating-by-inhibition hypothesis, where increased facilitation of gating might be brought about by relatively lower levels of increased inhibition; whereas higher levels are likely to lead to more pervasive suppression of vision.

The absence of a definitive answer to this question suggests that future experimentation may be profitable. Such an experiment could fruitfully include a range of TMS intensities to the same group of subjects using the behavioural paradigm in Experiment 2A.

A2. Additional MEG Analysis

Several analyses in addition to those described in the main experiment were undertaken. This section discusses them and the reasons for their exclusion from the main experimental discussion.

In the planning of the follow-on experiments (following section 2.3) we initially described the effect that we thought the TBS might have in terms of signal-to-noise ratio (SNR). That is, cTBS might increase signal-to-noise ratio via a suppression of neuronal noise leading to increased conscious detection (Waterston & Pack, 2010). On the grounds of simplicity, this description of effects at the neuronal level may be considered more appealing than the gating-by-inhibition description, although it is worth noting that the two explanations (SNR and gating-by-inhibition) do not necessarily contradict one another and are to some extent different ways of describing the same process.

The reason for reducing the emphasis placed upon an SNR-based explanation was the potential for confusion about which SNR is referred to by any particular measures. The problem is that there is an unavoidable conflation between neuronal SNR and SNR of the imaging technique itself. This means that when neuroimaging techniques are directed towards the investigation of the SNR explanation, the interpretation is potentially misleading. For example, gamma frequency (~40-80Hz) oscillations might be taken as representative of the carriers of signals (Bressler, 1990), therefore affording us the opportunity to probe the alternative hypothesis that it is increased signal (represented by a general increase in oscillatory power in the gamma band following cTBS) rather than suppression of noise that leads to the increased SNR. However, if neuronal noise is reduced, this too should lead to an increased ability to *detect* gamma frequency oscillations, irrespective of changes in the gamma band. The two competing hypotheses therefore make the same prediction. Likewise, the general levels of oscillatory power in the alpha band might be interpreted as a correlate of active suppression of noise. However, the observed increase in alpha power (figure A2.1) cannot dissociate any actual change in alpha power from increased ability to detect alpha oscillations. This difficulty in interpretation is the reason SNR was not emphasised as central to the explanation of the cTBS effect in question.
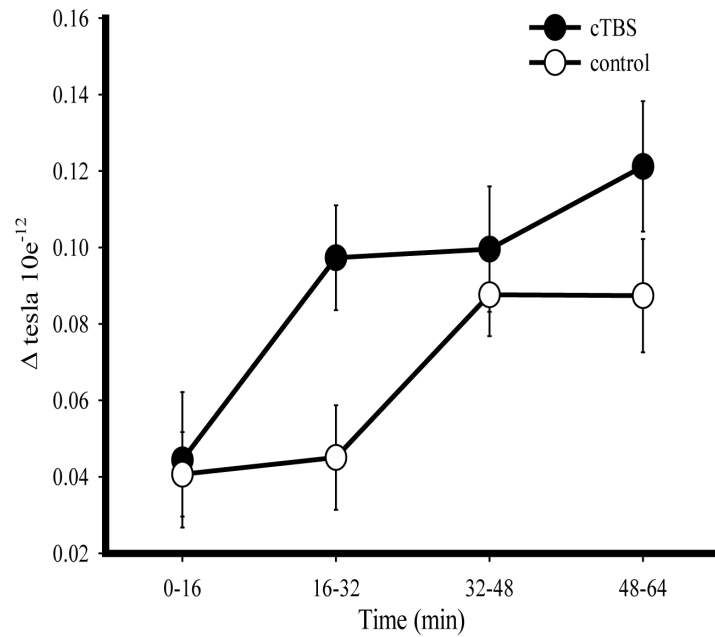
Figure A2.1. Alpha (6-12Hz) power changed over the course of Experiment 2D. Alpha power was not statistically dependent upon the application of the cTBS (site effect: $F_{(1,17)}$=1.940, p=0.182, site × time interaction: $F_{(3,51)}$=1.490, p=0.240). However it did increase over the course of the experiment (time effect: $F_{(3,51)}$=8.033, p=0.004), which is suggestive of increased fatigue (Steriade, McCormick, & Sejnowski, 1993).

The evoked and induced measures used did not suffer from this problem (of increased fidelity conflated with change in question) to the same extent as did the measures based on general oscillatory power (described above in the alpha and gamma bands). This is because the baseline for induced and evoked responses are taken within trial following the cTBS. Hence any increased fidelity as a result of reduced noise will be present prior to the stimuli presentation – hence removed from the data.

Induced gamma band responses to stimuli at around 100ms have previously been suggested as conveying conscious information (Crick, 1990; Fries, 2009). As such, they can be considered representative of conscious signals and therefore could have been incorporated into the experiment in a manner similar to that of the late evoked responses. Although a subset of subjects did display such gamma band responses to the stimuli, such responses were not by any means apparent across subjects and experimental blocks. Therefore the analysis of induced gamma in response to the stimuli was abandoned.

An additional way in which the SNR hypothesis was probed was via the derivation of a measure based on the variance at the sensor level, which might be representative of neuronal noise. This quantified levels of noise via the application of a simple root-mean-square analysis of variance to raw DC offset sensor data. Higher variance corresponding to increased noise. The SNR based explanation of the effect cTBS has would therefore predict a reduction in this measure. However, it displayed the same modulation as a measure of raw oscillatory power across frequency bands and – most importantly – it followed the general pattern of changes in the dominant alpha band (cf. figure A2.1 and A2.2). Given that oscillatory power as quantified via Hillbert transforms is based on such variance and the relative influence the alpha band has upon the sensors, this should have been expected. The result is that, again, it was not possible to dissociate competing hypotheses on the basis of this measure. Increases in the variance measure can be interpreted as increased noise and increased alpha suggests increased suppression of noise, yet the two correlate.
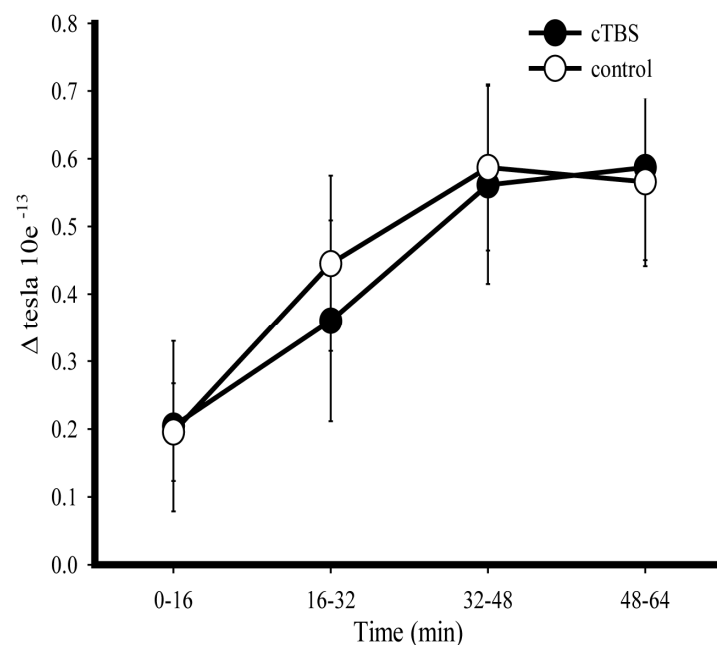
Figure A2.2. Root mean square analysis of variance applied to sensor level data. No significant effects of the TMS were observed (site: $F_{(1,17)}=0.012$, p=0.916, site × time interaction: $F_{(3,51)}=0.130$, p=0.913) but the variance did increase over the course of the experiment (time effect: $F_{(3,51)}=5.541$, p=0.009). Because of the correspondence between this and the alpha / oscillatory analysis, no conclusions can be drawn from these effects.

A3. Analysis of iTBS and sham control conditions in off-line experiment 2D

For completeness, this section reports the results for each of the three TMS conditions in Experiment 2D: cTBS, iTBS and sham. These analyses should not be interpreted as part of the main experiment as they incorporate multiple unnecessary comparisons.

F(df)	cTBS vs. iTBS	cTBS vs. sham	iTBS vs. sham
site (1,17)	6.13	2.86	0.24
time(3,51)	1.57	1.44	2.88
site × time (3,51)	1.67	0.93	1.24
p	cTBS vs. iTBS	cTBS vs. sham	iTBS vs. sham
site	0.02	0.11	0.63
time	0.21	0.25	0.06
site × time	0.19	0.42	0.31
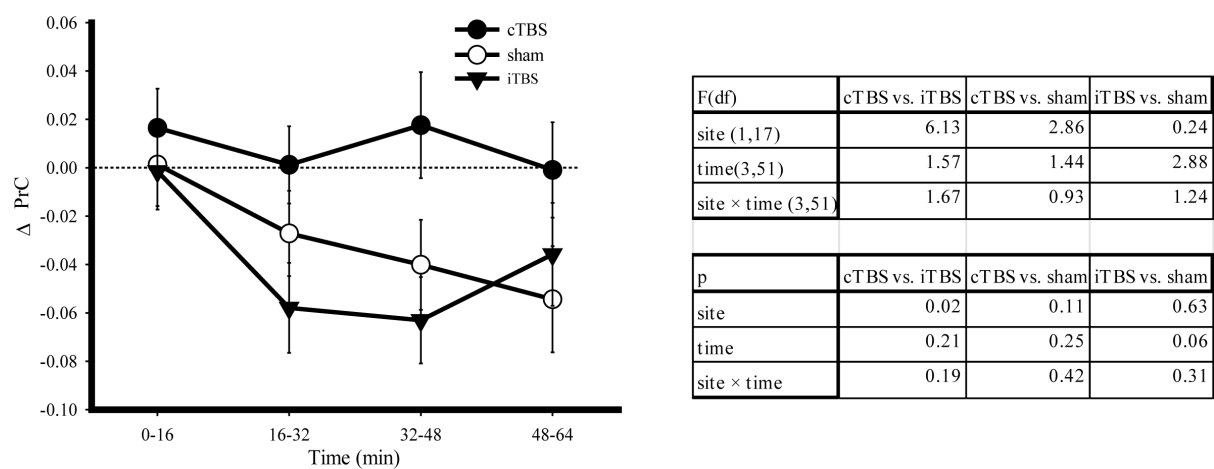
Figure and table A.3.1. Change in conscious detection (PrC) in the three TMS conditions (cTBS,iTBS and sham) of Experiment 2D, calculated relative to the session-specific pre-TBS baseline and plotted over the course of the experiment. Error bars are the within-subject standard error (Loftus & Masson, 1994; Masson & Loftus, 2003). Table provides F and p values from repeated measures ANOVAs applied in the same way as described for the experiments in Chapter 2, but contrasting all the possible combinations of TMS conditions. Subsequent plots and tables in this section all conform to this structure.
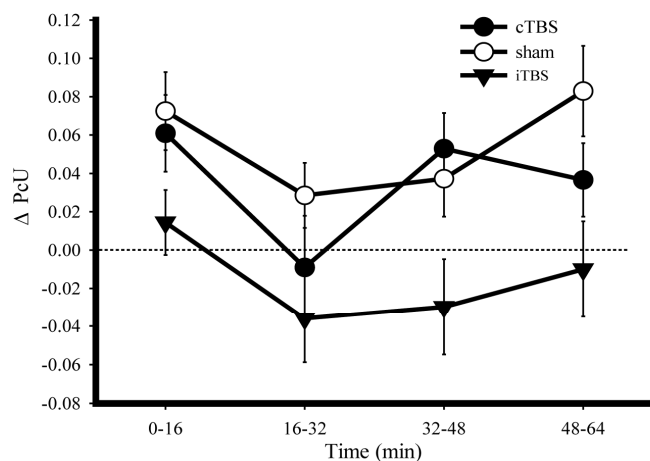
F(df)	cTBS vs. iTBS	cTBS vs. sham	iTBS vs. sham
site $(1,17)$	3.54	0.53	5.53
time $(3,51)$	3.54	6.43	4.65
site × time $(3,51)$	0.97	1.43	0.48

p	cTBS vs. iTBS	cTBS vs. sham	iTBS vs. sham
site	0.08	0.48	0.03
time	0.03	<0.01	0.01
site × time	0.41	0.25	0.66

Figure and table A.3.2. Change in 'Unseen' discrimination (PcU) form pre-TBS baseline in Experiment 2D.
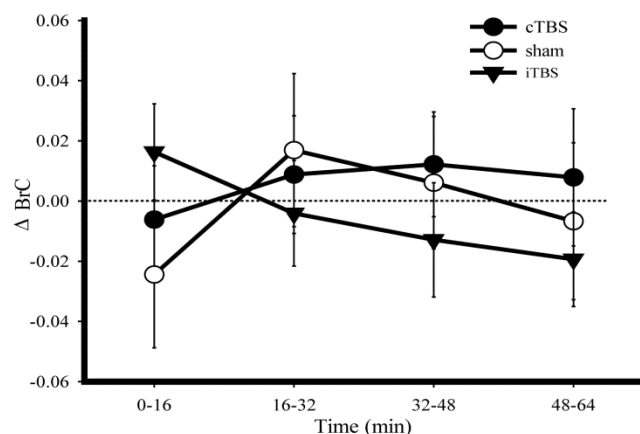
F(df)	cTBS vs. iTBS	cTBS vs. sham*	iTBS vs. sham*
site $(1,17)$	0.29	0.05	0.01
time $(3,51)$	0.26	1.63	0.82
site × time $(3,51)$	1.41	0.35	2.61

p	cTBS vs. iTBS	cTBS vs. sham*	iTBS vs. sham*
site	0.60	0.83	0.93
time	0.83	0.21	0.46
site × time	0.26	0.71	0.08

Figure and table A.3.3. Change in response criteria/Bias in conscious detection (BrC) form pre-TBS baseline in Experiment 2D. * denotes analysis in which a single outlier was rejected according to Chauvenet's criteria. Resultant degrees of freedom for corresponding site effects are 1,16 and for comparisons involving time it is 3,48. Graphs make use of data with this outlier excluded.
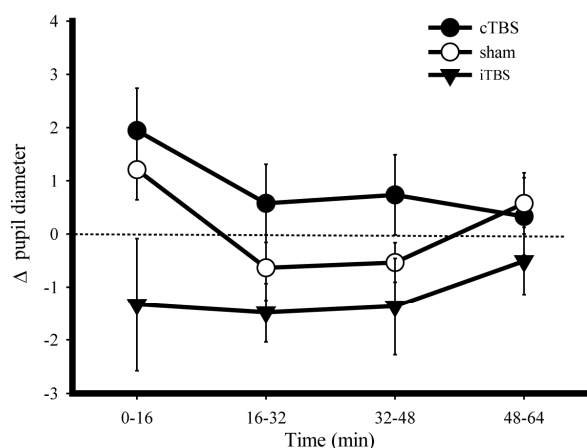
F(df)	cTBS vs. iTBS	cTBS vs. sham	iTBS vs. sham
site $(1,17)$	2.43	0.68	1.84
time $(3,51)$	1.16	7.92	3.02
site × time $(3,51)$	1.73	1.87	1.34

p	cTBS vs. iTBS	cTBS vs. sham	iTBS vs. sham
site	0.14	0.42	0.19
time	0.33	0.00	0.06
site × time	0.20	0.17	0.28

Figure and table A.3.4. Change in pupil diameter (arbitrary units) form pre-TBS baseline in Experiment 2D.

F(df)	cTBS vs. iTBS*	cTBS vs. sham*	iTBS vs. sham
site (1,17)	0.24	2.88	2.68
time(3,51)	0.32	1.20	0.44
site × time (3,51)	0.66	1.34	1.66
p	cTBS vs. iTBS*	cTBS vs. sham*	iTBS vs. sham
site	0.63	0.11	0.12
time	0.79	0.32	0.69
site × time	0.56	0.28	0.20
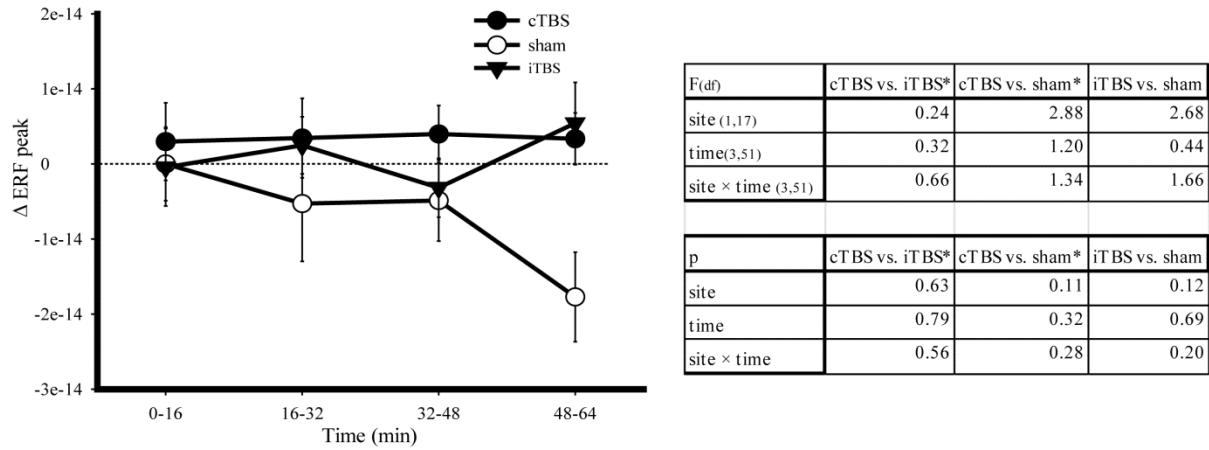
Figure and table A.3.5. Change in ERF peak amplitude form pre-TBS baseline in Experiment 2D. * denotes analysis in which a single outlier was rejected according to Chauvenet's criteria. Resultant degrees of freedom for corresponding site effects are 1,16 and for comparisons involving time it is 3,48. Graphs make use of data with this outlier excluded.

F(df)	cTBS vs. iTBS*	cTBS vs. sham*	iTBS vs. sham
site (1,17)	1.64	7.56	0.47
time(3,51)	0.61	1.57	0.15
site × time (3,51)	1.06	0.62	1.62
p	cTBS vs. iTBS*	cTBS vs. sham*	iTBS vs. sham
site	0.22	0.01	0.50
time	0.57	0.22	0.88
site × time	0.36	0.55	0.20

Figure and table A.3.6. Change in ERD gradient form pre-TBS baseline in Experiment 2D. * denotes analysis in which a single outlier was rejected according to Chauvenet's criteria. Resultant degrees of freedom for corresponding site effects are 1,16 and for comparisons involving time it is 3,48. Graphs make use of data with this outlier excluded.
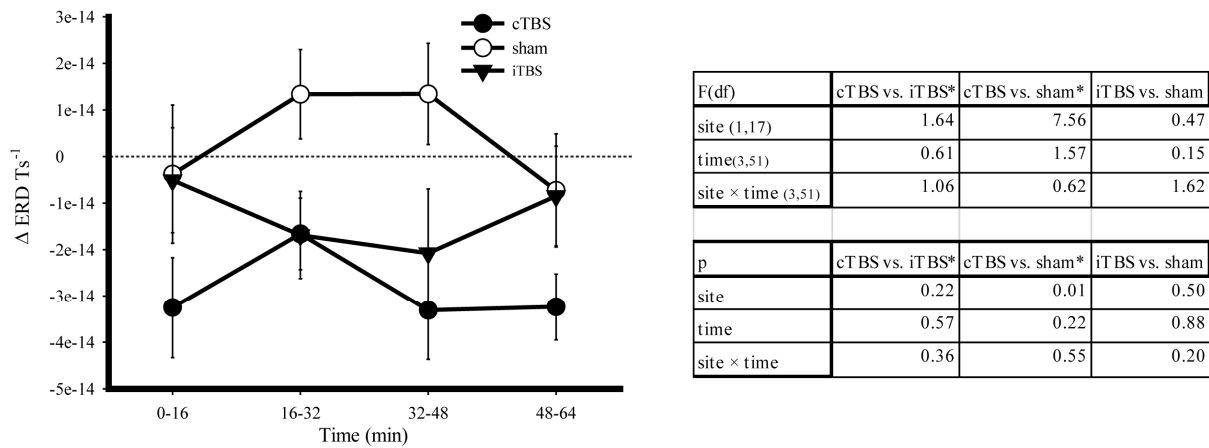
A4. Participant instructions

You will be shown arrows (figure. PI 1) that appear very briefly against a noisy background (figure. PI 2). Sometimes the arrow will not appear. Your task is to decide if you saw the arrow and which direction it was pointing in. You will be asked 3 questions about this:

Did you see the arrow?   Yes/No.
Which direction was it pointing in?   Left/Right.
Did you see something?   Yes/No.

It is important for the purpose of this study that you give your best guess for the direction discrimination question (Left/Right?), even when you don't see the arrow, i.e. do not simply press the same button when you think no stimulus has been presented.

Sometimes you will see something that might have been an arrow but you cannot be sure; this is why we ask the extra question: 'Did you see something?'. If you have the impression that you saw something which might have been the arrow, but cannot be sure, you should say 'Yes' to this question and 'No' to the 'Did you see the arrow?' question. You should say 'Yes' to the arrow question if you were consciously aware of the arrow. Also, we sometimes show non-arrows (figure PI 3). If you see this or think you might have, you should answer 'Yes' to the something question. If you saw nothing but the noise you should answer 'No' to both the arrow and the something questions.

The order of the questions will change from testing session to session. At the start of each session you will be given some practice trials to get you familiar with the order of the question and refresh your memory of the task.

So you will be asked 3 questions: 'Was the arrow pointing left or right?' which will appear on the screen as 'L/R', 'Did you see the arrow?' which will appear as 'Arrow Y/N' and 'Did you see something?' which will appear as 'something Y/N'. Responses will be taken on the keyboard number pad using the 1,3,4 and 6 keys (see figure PI 4). The buttons on the left (1 and 4) are for responding 'Left' and the ones on the right (2 and 6) are for

'Right'. The upper buttons (4 and 6) are for responding 'Yes' and the lower ones (1 and 3) are for responding 'No'. If at any point you enter the wrong response please tell the experimenter immediately so that she/he can make a note of it.

Sessions will involve four blocks of testing, each consisting of 80 trials. Two of these blocks will involve s-cone (purple) stimuli and two luminance (black and white) stimuli. To complete the experiment, you will complete 6 sessions over the course of 3 to 6 days, plus calibration days.

Calibration days will involve establishing thresholds for the task. This means we will vary the ease of the task so that we can get a stable level of performance, which we will then use in subsequent sessions. This will probably involve about 16 blocks of 80 trials and some of 40 trials. There will also be one of these shorter blocks at the start of each normal testing session to check that levels of performance are constant, and adjust them if necessary.

During the first days we will be assessing your susceptibility to phosphenes. Phosphenes are apparent flashes of light caused by the TMS. The experimenter will go through this with you in more detail during the first session. Also, during these calibration days we will also go through a series of tasks to calibrate the 'purple' colour stimuli. Again the experimenter will explain this.

If you have any questions, please ask.

Thanks for reading this and participating in our experiments.
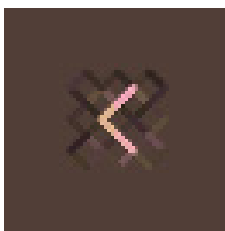

Figure PI 1.
The arrow that you will
be looking for.

Figure PI 2.
The noise in the absence
of an arrow.

Figure PI 3.
The non-arrow.
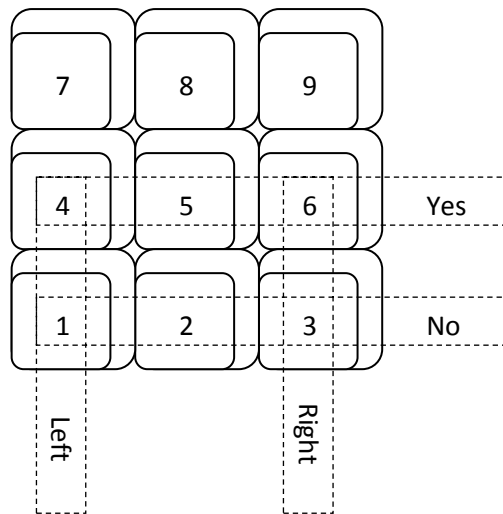
Figure PI 4.

The layout of the button responses on the
number pad of the keyboard.

A5. Analysis of the 'something' – p-conscious measure

The rate at which subjects reported awareness of 'something' is the measure derived to represent Blocks discussion of phenomenal (p-) consciousness (Block, 2007, 2011), discussed in section 4.6. Block makes the prediction that p–consciousness will be accompanied by recurrent processing (Block, 2007). This analysis tested the hypothesis that the application of TMS at relatively late times would lead to the disruption of recurrent processing and suppression in the frequency at which subjects reported 'something'. The three experiments capable of addressing this were those which applied TMS over a range of times relative to stimuli onset (Experiments 3 A,B,C). The analysis applied the same statistical methods as described in chapter 3 to the additional measure. The measure quantifies the incidence of subject's responding positively the 'something?' probe and negatively to the 'arrow?' question over the number of opportunities for such a response. The results are reported below, the discussion and interpretation of which can to be found in section 4.6.
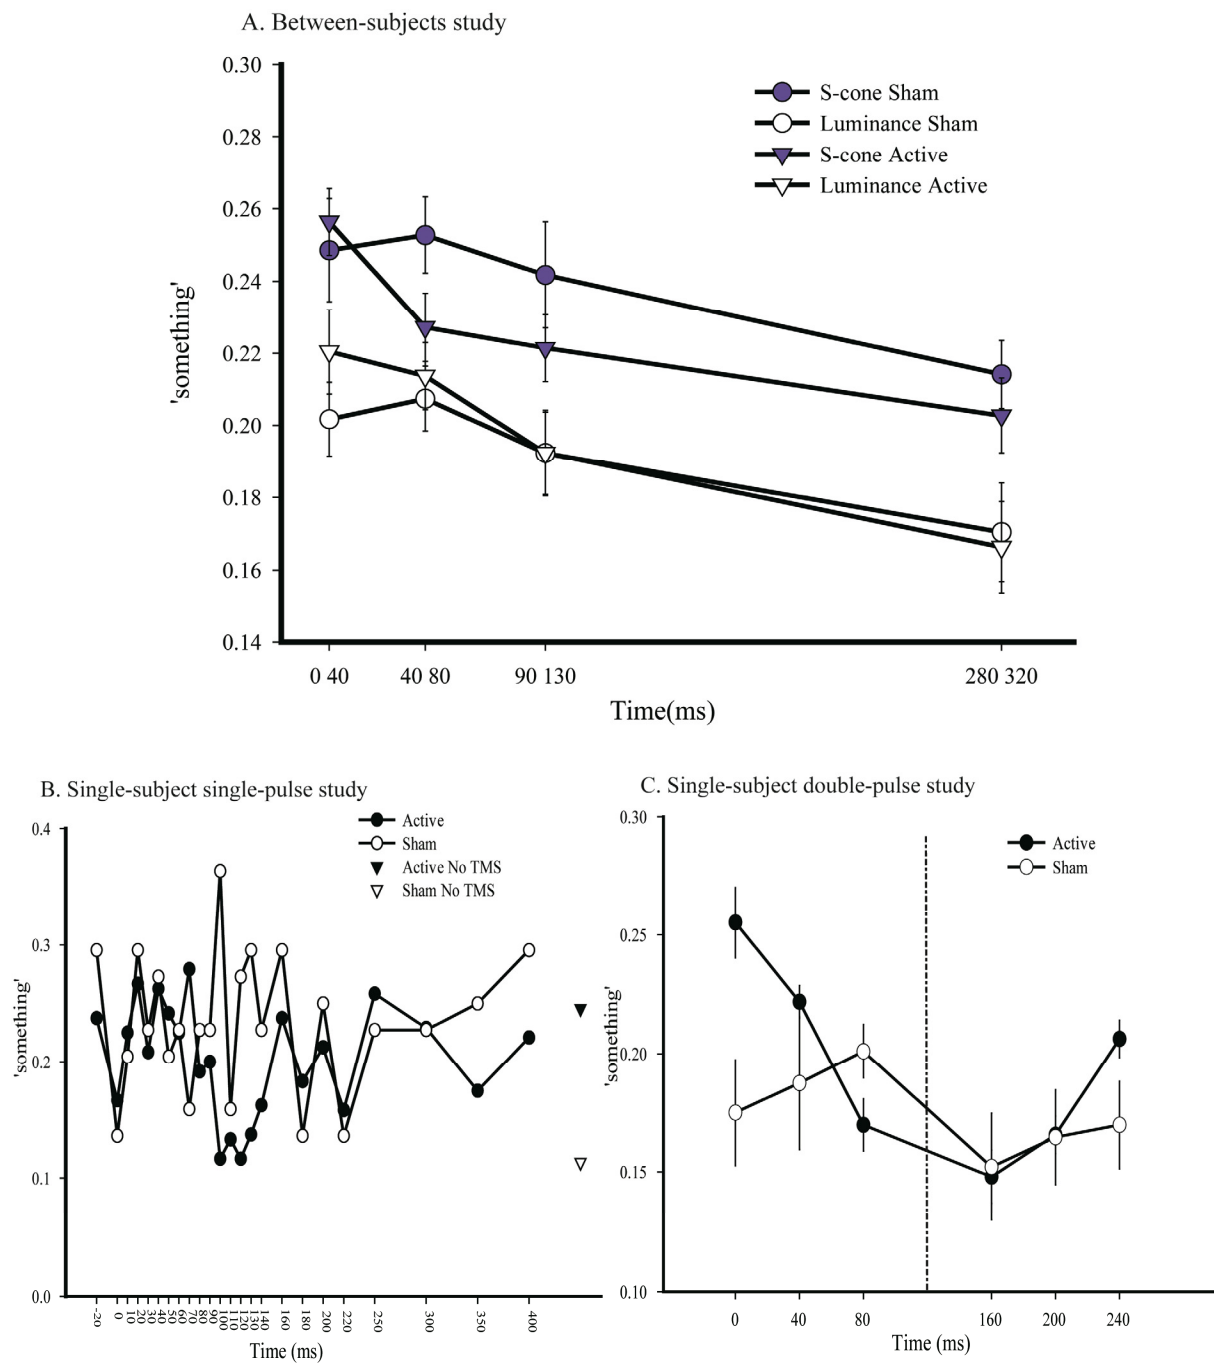
Figure A.5. Depiction of the main sections of the 'on-line' experiments with respect to the 'something' measure. The 'something' measure is the number of positive reports of 'something' only, over the number of trials at any given data point.

A. Rate of report of 'something' over the course of Experiment 3A. Error bars are the standard error across subjects according to (Loftus & Masson, 1994).

B. Rate of report of 'something' over the course of the single-subject single-pulse study (Experiment 3B).

194

C. Rate of report of 'something' over the course of the single-subject double-pulses study (Experiment 3C). Error bars are standard error across experimental sessions.

Experiment 3A, Between-subject study:

The critical effect in terms of relating p-consciousness (as measured by rate of reports of 'something') to recurrent processing would be a site (TMS vs. sham) effect that interacts with time, via a later suppression. See figure A5.1.A. No such effects or interactions were observed (site × time: $F_{(3,45)}$ =0.329, p= 0.762, site effect: $F_{(1,15)}$=0.260, p=0.617). There were however significant effects involving time ($F_{(3,45)}$ =9.383, p<0.001), again indicating the sensitivity of the measure and auditory distracting effects. Also, there was a significant effect of stimuli type (s-cone or luminance) ($F_{(1,15)}$=13.315, p=0.002), which simply indicates that subjects were more likely to respond positively to the 'something?' question in the presence of the s-cone stimuli. Since the s-cone and luminance stimuli were not independently calibrated in terms of propensity of the subjects to report 'something', this likely reflects a difference at the level of the stimuli rather than at the level of the pathways that the s-cone intervention can isolate. For this reason, together with the fact that no predictions were made in terms of s-cone dependency of the 'something' measure, the single subject s-cone experiment (3D) was not analysed in terms of this measure.

It is worth noting that there is a correspondence between the effects observed in the measure of 'something' and bias (BrC), which is due to the overlap between the attributions of responses over the two measures (both increase when subjects responded positively to the 'something?' question more often irrespective of stimuli presented). As these measures are applied to independent hypotheses, this form of auto-correlation is acceptable.

Experiment 3B, Single-subject, single-pulse study:

Based on the equivalent statistical tests applied to PrC and PcU measure in this experimental section 3.5.2 (which compared each data point in the active condition to that of the variance observed over the sham condition), no data points showed a significant effect of the TMS (Z < 1.79, p>0.037 relative to Bonferroni alpha of 0.002). See figure A5.1.B

Experiment 3C, Single-subject, double-pulse study:

The comparison of interest is the interaction between TMS site (active vs. sham) and time (early vs. late). The hypothesis supporting Block's theory is that there should be a later suppression of reports of 'something' if there is a correspondence between p-consciousness and recurrent processing. This analysis is of the same structure as applied to PrC and PcU measures, in section 3.6.2 and is depicted in figure A5.1.C.

The ANOVA indicated that there was no site × time interaction ($F_{(1,1)}$ = 0.63, p=0.43), no significant effect of the TMS ($F_{(1,1)}$ = 3.41, p = 0.07), but a significant effect of the time at which TMS was applied ($F_{(1,1)}$ = 10.46, p = 0.002). The non-significant trend for suppressed reports of 'something' in the presence of active TMS suggests that the TMS may have been effective in suppressing p-consciousness. However, the absence of a time-dependent interaction does not support the links between p-consciousness – awareness of 'something' – and any particular stage of processing, such as recurrence. The effect over time illustrates the sensitivity of the measure. Because this drop occurred over the later interventions and is expressed to a similar extent in the sham TMS condition as the active TMS condition, the effect is most easily attributable to an auditory distracting effect (e.g. Terao et al., 1997).

A6. An alternative response attribution

This section outlines an alternative way in which behavioural response could have been allocated to measures based on previously described experiments (e.g. Rounis, Maniscalco, Rothwell, Passingham, & Lau, 2010).

The principal dependent variables in these experiments have been conscious detection and 'unseen' discrimination. These two measures approximate to non-parametric versions of type-2 and type-1 measures in the SDT literature, respectively (Clarke, Birdsall, & Tanner, 1959; Galvin, Podd, Drga, & Whitmore, 2003; Pollack, 1959). Type-1 decisions corresponding to the information the subject can objectively display in forced choice tasks. Type-2 (also known as meta-d' (Maniscalco & Lau, 2012)) abilities correspond to the confidence that subjects place upon their primary discussion or their subjective judgment (Galvin, et al., 2003; Pollack, 1959). Independent differences between these two measures of task performance (type-1 and-2) have been used to demonstrate blindsight type phenomena (Dienes & Berry, 1997; Rounis, et al., 2010).

The match between type-2 / type-1 abilities and the detection / 'unseen' discrimination measures applied here is incomplete and was not pursued in these experiments for the following reason: the 'Did you consciously see the arrow?' question is a detection task which refers to presence / absence of subjects' conscious experience under presence / absence truth conditions. In contrast, the truth conditions under which type-2 probes are commonly tested is the discrimination difference (e.g. left/right) not the detection judgment (Galvin, et al., 2003; Maniscalco & Lau, 2012). Combining the two (detection and discrimination) in order to form a type-2 measure with these data sets is possible (see table A6.1), but it would result in the measure no longer tracking the conscious detection question posed, but rather a combination of responses to both questions, rendering interpretation difficult.

Two further, interrelated, practical concerns also limited the feasibility of applying type-1/2 style analysis to the data sets produced over these experiments: first, as described in section 1.3.2.3, the number of trials which the subject reported as being 'seen' yet the direction decision was incorrect was extremely low. As this response profile contributes to

produce meaningful type-2 measures (Galvin, et al., 2003; Maniscalco & Lau, 2012) this paucity of trials presents a significant practical limitation. Second, the construction of type 1 measures commonly use all discrimination judgments in their derivation (Galvin, et al., 2003; Maniscalco & Lau, 2012). As such they conflate 'seen' and 'unseen' trials. Because 'seen' performance is invariably at ceiling levels, changes in conscious awareness will be reflected in the type 1 measure (as discussed in section 1.3.2.3).

Table A6.1 Table of SDT classification for the task in order to produce a meta-d' / type-2 measure. Type-1 measure would consist of all forced choice discrimination (left/right) decisions.

	Hit		False alarm		Correct rejection		Miss	
Response	'seen'		'seen'		'unseen'		'unseen'	
	'left'	'right'	'left'	'right'	'left'	'right'	'left'	'right'
Stimulus	left	right	right	left	left	right	right	left

Bibliography

Abrahamyan, A., Clifford, C. W., Arabzadeh, E., & Harris, J. A. (2011). Improving visual sensitivity with subthreshold transcranial magnetic stimulation. *The Journal of neuroscience, 31*(9), 3290-3294.

Allen, E. A., Pasley, B. N., Duong, T., & Freeman, R. D. (2007). Transcranial magnetic stimulation elicits coupled neural and hemodynamic consequences. *Science, 317*(5846), 1918-1921.

Amassian, V. E., Cracco, R. Q., Maccabee, P. J., Cracco, J. B., Rudell, A., & Eberle, L. (1989). Suppression of visual perception by magnetic coil stimulation of human occipital cortex. *Electroencephalography and Clinical Neurophysiology, 74*(6), 458-462.

Anderson, E. J., Husain, M., & Sumner, P. (2008). Human intraparietal sulcus (IPS) and competition between exogenous and endogenous saccade plans. *NeuroImage, 40*(2), 838-851.

Ansorge, U. (2003). Asymmetric influences of temporally vs. nasally presented masked visual information: evidence for collicular contributions to nonconscious priming effects. *Brain and Cognition, 51*(3), 317-325.

Baars, B. J. (1988). *A cognitive theory of consciousness*. Cambridge: Cambridge University Press.

Baars, B. J. (2003). The global brainweb: An update on global workspace theory. *Science and Consciousness Review*

Baars, B. J. (2005). Global workspace theory of consciousness: toward a cognitive neuroscience of human experience. *Progress in brain research, 150*, 45-53.

Baars, B. J. (2007). The Global Workspace Theory of Consciousness. In M. S. Velmans, S. (Ed.), *the Blackwell companion to consciousness*. oxford: Blackwell Publishing Ltd.

Babiloni, C., Vecchio, F., Bultrini, A., Luca Romani, G., & Rossini, P. M. (2006). Pre- and poststimulus alpha rhythms are related to conscious visual perception: a high-resolution EEG study. *Cerebral cortex, 16*(12), 1690-1700.

Babiloni, C., Vecchio, F., Miriello, M., Romani, G. L., & Rossini, P. M. (2006). Visuo-spatial consciousness and parieto-occipital areas: a high-resolution EEG study. *Cerebral cortex, 16*(1), 37-46.

Barbur, J. L., Watson, J. D., Frackowiak, R. S., & Zeki, S. (1993). Conscious visual perception without V1. *Brain, 116 (Pt 6)*, 1293-1302.

Barker, A. T., Jalinous, R., & Freeston, I. L. (1985). Non-invasive magnetic stimulation of human motor cortex. *Lancet, 1*(8437), 1106-1107.

Baron, R. M., & Kenny, D. A. (1986). The moderator-mediator variable distinction in social psychological research: conceptual, strategic, and statistical considerations. *Journal of Personality and Social Psychology, 51*(6), 1173-1182.

Baulac, S., Huberfeld, G., Gourfinkel-An, I., Mitropoulou, G., Beranger, A., Prud'homme, J. F., . . . LeGuern, E. (2001). First genetic evidence of GABA(A) receptor dysfunction in epilepsy: a mutation in the gamma2-subunit gene. *Nature genetics, 28*(1), 46-48.

Berger, H. (1929). Über das Elektrenkephalogramm des Menschen. *European Archives of Psychiatry and Clinical Neuroscience, 87*(1), 527-570.

Block, N. (1995). On a Confusion About a Function of Consciousness. *Behavioral and Brain Sciences, 18*(2), 227-247.

Block, N. (1996). How not to find the neural correlate of consciousness *Royal Institute of Philosophy lectures, edited by Anthony O'Hear.*

Block, N. (2005). Two neural correlates of consciousness. *Trends in Cognitive Sciences, 9*(2), 46-52.

Block, N. (2007). Consciousness, accessibility, and the mesh between psychology and neuroscience. *Behavioral and Brain Sciences, 30*(5-6), 481-548.

Block, N. (2011). Perceptual consciousness overflows cognitive access. *Trends in Cognitive Sciences, 15*(12), 567-575.

Boksem, M. A., Meijman, T. F., & Lorist, M. M. (2005). Effects of mental fatigue on attention: an ERP study. *Brain research. Cognitive brain research, 25*(1), 107-116.

Bompas, A., Sterling, T., Rafal, R. D., & Sumner, P. (2008). Naso-temporal asymmetry for signals invisible to the retinotectal pathway. *Journal of neurophysiology, 100*(1), 412-421.

Bompas, A., & Sumner, P. (2008). Sensory sluggishness dissociates saccadic, manual, and perceptual responses: an S-cone study. *Journal of vision, 8*(8), 10 11-13.

Bompas, A., & Sumner, P. (2009). Oculomotor distraction by signals invisible to the retinotectal and magnocellular pathways. *Journal of neurophysiology, 102*(4), 2387-2395.

Bompas, A., & Sumner, P. (2011). Saccadic inhibition reveals the timing of automatic and voluntary signals in the human brain. *The Journal of neuroscience, 31*(35), 12501-12512.

Boyer, J. L., Harrison, S., & Ro, T. (2005). Unconscious processing of orientation and color without primary visual cortex. *Proceedings of the National Academy of Sciences of the United States of America, 102*(46), 16875-16879.

Bradley, M. M., Miccoli, L., Escrig, M. A., & Lang, P. J. (2008). The pupil as a measure of emotional arousal and autonomic activation. *Psychophysiology, 45*(4), 602-607.

Bradshaw, J. (1967). Pupil size as a measure of arousal during information processing. *Nature, 216*(5114), 515-516.

Bressler, S. L. (1990). The gamma wave: a cortical information carrier? *Trends in neurosciences, 13*(5), 161-162.

Butter, C. M. (1974). Effect of superior colliculus, striate, and prestriate lesions on visual sampling in rhesus monkeys. *Journal of Comparative & Physiological Psychology, 87*(5), 905-917.

Campbell, J. (2007). An Interventionist Approach to Causation in Psychology *Causal Learning: Psychology, Philosophy and Computation, (eds.)Alison Gopnik and Laura Schulz* (pp. 58-66). Oxford: Oxford University Press.

Campbell, J. (2008). Interventionism, Control Variables and Causation in the Qualitative World. *Philosophical Issues: Interdisciplinary Core Philosophy, 18*(1), 19.

Campion, J., Latto, R., & Smith, Y. M. (1983). Is Blindsight an Effect of Scattered-Light, Spared Cortex, and near-Threshold Vision. *Behavioral and Brain Sciences, 6*(3), 423-447.

Camprodon, J. A., Zohary, E., Brodbeck, V., & Pascual-Leone, A. (2010). Two phases of V1 activity for visual recognition of natural images. *Journal of Cognitive Neuroscience, 22*(6), 1262-1269.

Carruthers, P. (2007). Higher-Order Theories of Consciousness. In M. S. Velmans, S. (Ed.), *The Blackwell Companion to Consciouness*. Oxford: Blackwell Publishing Ltd.

Cazzoli, D., Muri, R. M., Schumacher, R., von Arx, S., Chaves, S., Gutbrod, K., . . . Nyffeler, T. (2012). Theta burst stimulation reduces disability during the activities of daily living in spatial neglect. *Brain*.

Cazzoli, D., Wurtz, P., Muri, R. M., Hess, C. W., & Nyffeler, T. (2009). Interhemispheric balance of overt attention: a theta burst stimulation study. *European Journal of Neuroscience, 29*(6), 1271-1276.

Chalmers, D. J. (1996). *The conscious mind : in search of a fundamental theory*. New York ; Oxford: Oxford University Press.

Chalmers, D. J. (2000). What is a Neural Correlate of Consciousness? In T. Metzinger (Ed.), *Neural Correlates of Consciousness: Empirical and Conceptual Questions* MIT press.

Chambers, C. D., Allen, C. P., Maizey, L., & Williams, M. A. (2012). Is delayed foveal feedback critical for extra-foveal perception? *Cortex*.

Chambers, C. D., & Mattingley, J. B. (2005). Neurodisruption of selective attention: insights and implications. *Trends in Cognitive Sciences, 9*(11), 542-550.

Chatterjee, S., & Callaway, E. M. (2002). S cone contributions to the magnocellular visual pathway in macaque monkey. *Neuron, 35*(6), 1135-1146.

Christensen, M. S., Kristiansen, L., Rowe, J. B., & Nielsen, J. B. (2008). Action-blindsight in healthy subjects after transcranial magnetic stimulation. *Proceedings of the National Academy of Sciences of the United States of America, 105*(4), 1353-1357.

Clark, A. (1997). *Being there : putting brain, body, and world together again*. Cambridge, Mass ; London: MIT Press.

Clark, A., & Chalmers, D. J. (1998). The extended mind. *Analysis, 58*(1), 7-19.

Clarke, F. R., Birdsall, T. G., & Tanner, W. P. (1959). 2 Types of Roc Curves and Definitions of Parameters. *Journal of the Acoustical Society of America, 31*(5), 629-630.

Cohen, M. A., & Dennett, D. C. (2011). Consciousness cannot be separated from function. *Trends in Cognitive Sciences, 15*(8), 358-364.

Cook, D. (1977). Detection of Influential Observation in Linear Regression. *Technometrics, 19*(1), 3.

Corbetta, M., Patel, G., & Shulman, G. L. (2008). The reorienting system of the human brain: from environment to theory of mind. *Neuron, 58*(3), 306-324.

Corthout, E., Hallett, M., & Cowey, A. (2002). Early visual cortical processing suggested by transcranial magnetic stimulation. *Neuroreport, 13*(9), 1163-1166.

Corthout, E., Hallett, M., & Cowey, A. (2003). Interference with vision by TMS over the occipital pole: a fourth period. *Neuroreport, 14*(4), 651-655.

Corthout, E., Uttl, B., Juan, C. H., Hallett, M., & Cowey, A. (2000). Suppression of vision by transcranial magnetic stimulation: a third mechanism. *Neuroreport, 11*(11), 2345-2349.

Corthout, E., Uttl, B., Walsh, V., Hallett, M., & Cowey, A. (1999). Timing of activity in early visual cortex as revealed by transcranial magnetic stimulation. *Neuroreport, 10*(12), 2631-2634.

Corthout, E., Uttl, B., Ziemann, U., Cowey, A., & Hallett, M. (1999). Two periods of processing in the (circum)striate visual cortex as revealed by transcranial magnetic stimulation. *Neuropsychologia, 37*(2), 137-145.

Corwin, J. (1994). On measuring discrimination and respones bias: unequal numbers of targets and distractors and two classes of distractors. *Neuropsychology, 8*(1), 7.

Cottaris, N. P., & De Valois, R. L. (1998). Temporal dynamics of chromatic tuning in macaque primary visual cortex. *Nature, 395*(6705), 896-900.

Cowey, A. (2010). The blindsight saga. *Experimental brain research, 200*(1), 3-24.

Crick, F., & Koch, C. (1995). Are we aware of neural activity in primary visual cortex? *Nature, 375*(6527), 121-123.

Crick, F. K., C. (1990). Towards a neurobiological theory of consciousness. *seminars in the neurosciences, 2*, 12.

d'Arsonval, J. A. (1896). Dispositifs pour la mésure des courants alternatifs de toutes fréquences. *Comptes Rendus de l Academie des Sciences. Biologies (Paris), 2*, 2.

Dacey, D. M. (2000). Parallel pathways for spectral coding in primate retina. *Annual review of neuroscience, 23*, 743-775.

Danckert, J., & Rossetti, Y. (2005). Blindsight in action: what can the different sub-types of blindsight tell us about the control of visually guided actions? *Neuroscience & Biobehavioral Reviews, 29*(7), 1035-1046.

Dapraz, N. (1999). The Phenomenological Reduction As Praxis. In F. J. Varela, J. Shear (Ed.), *The View From Within*: Imprint Academic.

de Graaf, T. A., Goebel, R., & Sack, A. T. (2012). Feedforward and quick recurrent processes in early visual cortex revealed by TMS? *NeuroImage, 61*(3), 651-659.

de Graaf, T. A., Herring, J., & Sack, A. T. (2011). A chronometric exploration of high-resolution 'sensitive TMS masking' effects on subjective and objective measures of vision. *Experimental brain research, 209*(1), 19-27.

de Monasterio, F. M. (1978). Properties of ganglion cells with atypical receptive-field organization in retina of macaques. *Journal of Neurophysiology, 41*(6), 1435-1449.

de No, R. L. (1933). Vestibulo-ocular reflex arc. *Archives of Neurology and Psychiatry, 30*(2), 245-291.

Deblieck, C., Thompson, B., Iacoboni, M., & Wu, A. D. (2008). Correlation between motor and phosphene thresholds: a transcranial magnetic stimulation study. *Human brain mapping, 29*(6), 662-670.

Dehaene, S., & Changeux, J. P. (2004). Neural Mechanisms for Access to Consciousness. In Gazzaniga (Ed.), *The cognitive neurosciences* (Vol. 3).

Dehaene, S., Changeux, J. P., Naccache, L., Sackur, J., & Sergent, C. (2006). Conscious, preconscious, and subliminal processing: a testable taxonomy. *Trends in Cognitive Sciences, 10*(5), 204-211.

Dehaene, S., & Naccache, L. (2001). Towards a cognitive neuroscience of consciousness: basic evidence and a workspace framework. *Cognition, 79*(1-2), 1-37.

Dehaene, S., Naccache, L., Cohen, L., Bihan, D. L., Mangin, J. F., Poline, J. B., & Riviere, D. (2001). Cerebral mechanisms of word masking and unconscious repetition priming. *Nature neuroscience, 4*(7), 752-758.

Del Cul, A., Baillet, S., & Dehaene, S. (2007). Brain dynamics underlying the nonlinear threshold for access to consciousness. *PLoS biology, 5*(10), e260.

Dennett, D. C. (1969). *Content and consciousness*. London: Routledge & Kegan Paul.

Dennett, D. C. (1991). *Consciousness explained*. Boston ; London: Little, Brown.

Dennett, D. C. (2003). Who's On First? Heterophenomenology Explained. *Journal of Consciousness Studies, 10*(9-10), 12.

Derrington, A. M., Krauskopf, J., & Lennie, P. (1984). Chromatic mechanisms in lateral geniculate nucleus of macaque. *The Journal of Physiology, 357*, 241-265.

Dienes, Z. (2008). *Understanding Psychology as a Science: An Introduction to Scientific and Statistical Inference*: Palgrave Macmillan.

Dienes, Z. (2011). Bayesian Versus Orthodox Statistics: Which Side Are You On? *Perspectives on Psychological Science, 6*(3), 274-290.

Dienes, Z., & Berry, D. (1997). Implicit learning: Below the subjective threshold. *Psychonomic Bulletin & Review, 4*(1), 3-23.

Dienes, Z. S., A.K. (In Press). The conscious and the unconscious. In G. T. Koob, R.E. Le Moal, M. (Ed.), Encyclopedia of Behavioral Neuroscience.

Dretske, F. (1986). Misrepresentation. In R. Bogdan (Ed.), *Belief: Form, Content, and Function* (pp. 17-36): Oxford University Press.

Dretske, F. (2000). Perception Without Awareness. In T. S. G. J. Hawthorne (Ed.), *Perceptual Experience*. Oxford: Oxford University Press.

Ebert, U., & Ostwald, J. (1995). GABA can improve acoustic contrast in the rat ventral cochlear nucleus. *Experimental brain research, 104*(2), 310-322.

Ericsson, K. A. (2003). Valid and Non-Reactive Verbalization of Thoughts During Performance of Tasks. *Journal of Consciousness Studies, 10*(9-10), 18.

Eriksen, C. W. (1960). Discrimination and learning without awareness: a methodological survey and evaluation. *Psychological Review, 67*, 279-300.

Fendrich, R., Wessinger, C. M., & Gazzaniga, M. S. (1992). Residual vision in a scotoma: implications for blindsight. *Science, 258*(5087), 1489-1491.

Ffytche, D. H., Guy, C. N., & Zeki, S. (1996). Motion specific responses from a blind hemifield. *Brain, 119 (Pt 6)*, 1971-1982.

Ffytche, D. H., & Zeki, S. (2011). The primary visual cortex, and feedback to it, are not necessary for conscious vision. *Brain, 134*(Pt 1), 247-257.

Foxe, J. J., & Snyder, A. C. (2011). The Role of Alpha-Band Brain Oscillations as a Sensory Suppression Mechanism during Selective Attention. *Frontiers in psychology, 2*, 154.

Franca, M., Koch, G., Mochizuki, H., Huang, Y. Z., & Rothwell, J. C. (2006). Effects of theta burst stimulation protocols on phosphene threshold. *Clinical neurophysiology, 117*(8), 1808-1813.

Fries, P. (2009). Neuronal gamma-band synchronization as a fundamental process in cortical computation. *Annual review of neuroscience, 32*, 209-224.

Fries, W. (1981). The projection from the lateral geniculate nucleus to the prestriate cortex of the macaque monkey. *Proceedings of the Royal Society of London. Series B, Containing papers of a Biological character. Royal Society, 213*(1190), 73-86.

Fries, W. (1984). Cortical projections to the superior colliculus in the macaque monkey: a retrograde study using horseradish peroxidase. *The Journal of comparative neurology, 230*(1), 55-76.

Fuchs, I., & Ansorge, U. (2012). Unconscious Cueing via the Superior Colliculi: Evidence from Searching for Onset and Color Targets. *Brain Sciences, 2*(1), 33-60.

Funke, K., & Benali, A. (2011). Modulation of cortical inhibition by rTMS - findings obtained from animal models. *The Journal of physiology, 589*(Pt 18), 4423-4435.

Gagne, R. M., & Smith, E. C., Jr. (1962). A study of the effects of verbalization on problem solving. *Journal of experimental psychology, 63*, 12-18.

Gallagher, S. (2003). Phenomenology and Experimental Design Toward a Phenomenologically Enlightened Experimental Science. *Journal of Consciousness Studies, 10*(9-10), 85.

Gallistel, C. R. (2009). The importance of proving the null. *Psychological Review, 116*(2), 439-453.

Galvin, S. J., Podd, J. V., Drga, V., & Whitmore, J. (2003). Type 2 tasks in the theory of signal detectability: Discrimination between correct and incorrect decisions. *Psychonomic Bulletin & Review, 10*(4), 843-876.

Green, D. M., & Swets, J. A. (1966). *Signal detection theory and psychophysics*. New York ; London: Wiley.

Hamalain, M. S., Hari, R. (2002). Magnetoencephalographic (MEG) Characterization of Dynamic Brain Activation. In A. W. Toga, Mazziotta, J.C. (Ed.), *Brain Mapping : the Methods* (2 ed.): Academic Press.

Hanajima, R., Wang, R., Nakatani-Enomoto, S., Hamada, M., Terao, Y., Furubayashi, T., . . . Ugawa, Y. (2007). Comparison of different methods for estimating motor threshold with transcranial magnetic stimulation. *Clinical neurophysiology, 118*(9), 2120-2122.

Hautus, M. J. (1995). Corrections for Extreme Proportions and Their Biasing Effects on Estimated Values of D'. *Behavior Research Methods Instruments & Computers, 27*(1), 46-51.

Heinen, K., Jolij, J., & Lamme, V. A. (2005). Figure-ground segregation requires two distinct periods of activity in V1: a transcranial magnetic stimulation study. *Neuroreport, 16*(13), 1483-1487.

Hilgetag, C. C., Theoret, H., & Pascual-Leone, A. (2001). Enhanced visual spatial attention ipsilateral to rTMS-induced 'virtual lesions' of human parietal cortex. *Nature neuroscience, 4*(9), 953-957.

Holm, S. (1979). A Simple Sequentially Rejective Multiple Test Procedure. *Scandinavian Journal of Statistics, 6*(2), 65-70.

Holt, J. (2003). *Blindsight and the nature of consciousness*: Broadview Press Ltd.

Huang, Y. Z., Edwards, M. J., Rounis, E., Bhatia, K. P., & Rothwell, J. C. (2005). Theta burst stimulation of the human motor cortex. *Neuron, 45*(2), 201-206.

Humphrey, N. K. (1968). Responses to visual stimuli of units in the superior colliculus of rats and monkeys. *Experimental Neurology, 20*(3), 312-340.

Hurlburt, R. T., and Heavey, C.T. (2004). To Beep or Not To Beep (Obtaining Accurate Reports About Awareness). *Journal of Consciousness Studies, 11*(7-8), 15.

Huxlin, K. R. (2008). Perceptual plasticity in damaged adult visual systems. *Vision research, 48*(20), 2154-2166.

Jackendoff, R. (1987). *Consciousness and the computational mind, British Psychological Society*. Cambridge, Mass ; London: MIT Press.

Jeffreys, E. T. (1961). *The theory of probability* (3 ed.): Oxford University Press.

Jensen, O., & Mazaheri, A. (2010). Shaping functional architecture by oscillatory alpha activity: gating by inhibition. *Frontiers in human neuroscience, 4*, 186.

Jolij, J., & Lamme, V. A. (2005). Repression of unconscious information by conscious processing: evidence from affective blindsight induced by transcranial magnetic stimulation. *Proceedings of the National Academy of Sciences of the United States of America, 102*(30), 10747-10751.

Juan, C. H., & Walsh, V. (2003). Feedback to V1: a reverse hierarchy in vision. *Experimental Brain Research, 150*(2), 259-263.

Kaernbach, C., Schroger, E., Jacobsen, T., & Roeber, U. (1999). Effects of consciousness on human brain waves following binocular rivalry. *Neuroreport, 10*(4), 713-716.

Kamitani, Y., & Shimojo, S. (1999). Manifestation of scotomas created by transcranial magnetic stimulation of human visual cortex. *Nature neuroscience, 2*(8), 767-771.

Kamitani, Y., & Tong, F. (2005). Decoding the visual and subjective contents of the human brain. *Nature neuroscience, 8*(5), 679-685.

Kammer, T. (2007). Masking visual stimuli by transcranial magnetic stimulation. *Psychological research, 71*(6), 659-666.

Kammer, T., Beck, S., Thielscher, A., Laubis-Herrmann, U., & Topka, H. (2001). Motor thresholds in humans: a transcranial magnetic stimulation study comparing different pulse waveforms, current directions and stimulator types. *Clinical neurophysiology, 112*(2), 250-258.

Kammer, T., Puls, K., Strasburger, H., Hill, N. J., & Wichmann, F. A. (2005). Transcranial magnetic stimulation in the visual system. I. The psychophysics of visual suppression. *Experimental brain research, 160*(1), 118-128.

Katzner, S., Busse, L., & Carandini, M. (2011). GABAA inhibition controls response gain in visual cortex. *The Journal of neuroscience, 31*(16), 5931-5941.

Kentridge, R. W., & Heywood, C. A. (2000). Metacognition and awareness. *Consciousness and cognition, 9*(2 Pt 1), 308-312; discussion 324-306.

Kentridge, R. W., Heywood, C. A., & Weiskrantz, L. (1997). Residual vision in multiple retinal locations within a scotoma: Implications for blindsight. *Journal of Cognitive Neuroscience, 9*(2), 191-202.

Kentridge, R. W., Heywood, C. A., & Weiskrantz, L. (1999). Effects of temporal cueing on residual visual discrimination in blindsight. *Neuropsychologia, 37*(4), 479-483.

Keppel, G. (1982). *Design and analysis : a researcher's handbook* (2nd ed.). Englewood Cliffs, N.J.: Prentice-Hall.

Kindler, J., Schumacher, R., Cazzoli, D., Gutbrod, K., Koenig, M., Nyffeler, T., . . . Muri, R. M. (2012). Theta Burst Stimulation Over the Right Broca's Homologue Induces Improvement of Naming in Aphasic Patients. *Stroke; a journal of cerebral circulation*.

Ko, Y., & Lau, H. (2012). A detection theoretic explanation of blindsight suggests a link between conscious perception and metacognition. *Philosophical transactions of the Royal Society of London. Series B, Biological sciences, 367*(1594), 1401-1411.

Koch, C. (2004). *The quest for consciousness : a neurobiological approach*. Englewood, CO: Roberts and Company Publishers.

Koch, G., Oliveri, M., Cheeran, B., Ruge, D., Lo Gerfo, E., Salerno, S., . . . Caltagirone, C. (2008). Hyperexcitability of parietal-motor functional connections in the intact left-hemisphere of patients with neglect. *Brain, 131*(Pt 12), 3147-3155.

Koivisto, M., Henriksson, L., Revonsuo, A., & Railo, H. (2012). Unconscious response priming by shape depends on geniculostriate visual projection. *The European journal of neuroscience, 35*(4), 623-633.

Koivisto, M., Mantyla, T., & Silvanto, J. (2010). The role of early visual cortex (V1/V2) in conscious and unconscious visual perception. *NeuroImage, 51*(2), 828-834.

Koivisto, M., Railo, H., & Salminen-Vaparanta, N. (2011). Transcranial magnetic stimulation of early visual cortex interferes with subjective visual awareness and objective forced-choice performance. *Consciousness and cognition, 20*(2), 288-298.

Koivisto, M., & Revonsuo, A. (2003). An ERP study of change detection, change blindness, and visual awareness. *Psychophysiology, 40*(3), 423-429.

Koivisto, M., & Revonsuo, A. (2007). Electrophysiological correlates of visual consciousness and selective attention. *Neuroreport, 18*(8), 753-756.

Kolb, F. C., & Braun, J. (1995). Blindsight in Normal Observers. *Nature, 377*(6547), 336-338.

Kostandov, E., & Arzumanov, Y. (1977). Averaged cortical evoked potentials to recognized and non-recognized verbal stimuli. *Acta neurobiologiae experimentalis, 37*(5), 311-324.

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.

Kripke, S. A. (1980). *Naming and necessity*. Cambridge, Mass.: Harvard University Press.

Labov, W. (1972). *Sociolinguistic patterns*. Philadelphia: University of Pennsylvania Press.

Lamme, V. A. (2001). Blindsight: the role of feedforward and feedback corticocortical connections. *Acta psychologica, 107*(1-3), 209-228.

Lamme, V. A. (2006a). Towards a true neural stance on consciousness. *Trends in Cognitive Sciences, 10*(11), 494-501.

Lamme, V. A. (2006b). Zap! Magnetic tricks on conscious and unconscious vision. *Trends in Cognitive Sciences, 10*(5), 193-195.

Lamme, V. A., & Roelfsema, P. R. (2000). The distinct modes of vision offered by feedforward and recurrent processing. *Trends in Neurosciences, 23*(11), 571-579.

Lamy, D., Salti, M., & Bar-Haim, Y. (2009). Neural correlates of subjective awareness and unconscious processing: an ERP study. *Journal of Cognitive Neuroscience, 21*(7), 1435-1446.

Lau, H. (2008). Are we studying consciousness yet? In L. Weiskrantz & M. Davies (Eds.), *Frontiers of Consciousness*: Oxford University Press.

Lau, H., & Rosenthal, D. (2011). Empirical support for higher-order theories of conscious awareness. *Trends in Cognitive Sciences, 15*(8), 365-373.

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.

Lee, C., Rohrer, W. H., & Sparks, D. L. (1988). Population coding of saccadic eye movements by neurons in the superior colliculus. *Nature, 332*(6162), 357-360.

Leh, S. E., Johansen-Berg, H., & Ptito, A. (2006). Unconscious vision: new insights into the neuronal correlate of blindsight using diffusion tractography. *Brain, 129*(Pt 7), 1822-1832.

Leh, S. E., Mullen, K. T., & Ptito, A. (2006). Absence of S-cone input in human blindsight following hemispherectomy. *European Journal of Neuroscience, 24*(10), 2954-2960.

Leh, S. E., Ptito, A., Schonwiesner, M., Chakravarty, M. M., & Mullen, K. T. (2010). Blindsight mediated by an S-cone-independent collicular pathway: an fMRI study in hemispherectomized subjects. *Journal of Cognitive Neuroscience, 22*(4), 670-682.

Levine, J. (1983). materialism and qualia: The explanatory gap. *Pacific PhilosophicalQuarterly, 64*, 8.

Lewis, C. I. (1929). *Mind and the world-order : outline of a theory of knowledge*. New York: Scribners.

Liddell, B. J., Williams, L. M., Rathjen, J., Shevrin, H., & Gordon, E. (2004). A temporal dissociation of subliminal versus supraliminal fear perception: an event-related potential study. *Journal of Cognitive Neuroscience, 16*(3), 479-486.

Loftus, G. R., & Masson, M. E. J. (1994). Using Confidence-Intervals in within-Subject Designs. *Psychonomic Bulletin & Review, 1*(4), 476-490.

Logothetis, N. K., & Pfeuffer, J. (2004). On the nature of the BOLD fMRI contrast mechanism. *Magnetic resonance imaging, 22*(10), 1517-1531.

Lombardi, C. M., & Hurlbert, S. H. (2009). Misprescription and misuse of one-tailed tests. *Austral Ecology, 34*(4), 447-468.

Luck, S. J. (2005). *An introduction to the event-related potential technique*. Cambridge, Ma. ; London: MIT Press.

Lutz, A. (2007). Neurophenomenology and the study of self-consciousness. *Consciousness and cognition, 16*(3), 765-767.

Lutz, A., Lachaux, J. P., Martinerie, J., & Varela, F. J. (2002). Guiding the study of brain dynamics by using first-person data: synchrony patterns correlate with ongoing conscious states during a simple visual task. *Proceedings of the National Academy of Sciences of the United States of America, 99*(3), 1586-1591.

Lutz, A. T., E. (2003). Neurophenomenology, Integrating subjective experience and brain dynamics in the neuroscience of consciousness. *Journal of Consciousness Studies, 10*(9-10), 21.

Macmillan, N. A., & Creelman, C. D. (1990). *Detection theory : a user's guide*. Cambridge: Cambridge University Press.

Macmillan, N. A., & Kaplan, H. L. (1985). Detection theory analysis of group data: estimating sensitivity from average hit and false-alarm rates. *Psychological Bulletin, 98*(1), 185-199.

Maizey, L., Allen, C.P.G., Dervinis, M., Verbruggen, F., Varnava, A., Kozlov, M., Adams, R.C., Stokes, M., Klemen, J., Bungert, A., Hounsell, C.A., Chambers, C.D. (In Press). Comparative incidence rates of mild adverse effects to transcranial magnetic stimulation. *Clinical Neurophysiology*.

Maniscalco, B., & Lau, H. (2012). A signal detection theoretic approach for estimating metacognitive sensitivity from confidence ratings. *Consciousness and Cognition, 21*(1), 422-430.

Marcel, A. J. (1983). Conscious and unconscious perception: an approach to the relations between phenomenal experience and perceptual processes. *Cognitive psychology, 15*(2), 238-300.

Marr, D. (1982). *Vision : a computational investigation into the human representation and processing of visual information*. San Francisco: W.H. Freeman.

Marrocco, R. T., & Li, R. H. (1977). Monkey superior colliculus: properties of single cells and their afferent inputs. *Journal of Neurophysiology, 40*(4), 844-860.

Marzi, C. A., Girelli, M., Miniussi, C., Smania, N., & Maravita, A. (2000). Electrophysiological correlates of conscious vision: evidence from unilateral extinction. *Journal of Cognitive Neuroscience, 12*(5), 869-877.

Marzi, C. A., Miniussi, C., Maravita, A., Bertolasi, L., Zanette, G., Rothwell, J. C., & Sanes, J. N. (1998). Transcranial magnetic stimulation selectively impairs interhemispheric transfer of visuo-motor information in humans. *Experimental brain research, 118*(3), 435-438.

Masson, M. E. J., & Loftus, G. R. (2003). Using confidence intervals for graphically based data interpretation. *Canadian Journal of Experimental Psychology, 57*(3), 203-220.

Mathewson, K. E., Gratton, G., Fabiani, M., Beck, D. M., & Ro, T. (2009). To see or not to see: prestimulus alpha phase predicts visual awareness. *The Journal of neuroscience, 29*(9), 2725-2732.

Maunsell, J. H., Ghose, G. M., Assad, J. A., McAdams, C. J., Boudreau, C. E., & Noerager, B. D. (1999). Visual response latencies of magnocellular and parvocellular LGN neurons in macaque monkeys. *Visual neuroscience, 16*(1), 1-14.

McKeefry, D. J., Parry, N. R., & Murray, I. J. (2003). Simple reaction times in color space: the influence of chromaticity, contrast, and cone opponency. *Investigative Ophthalmology & Visual Science, 44*(5), 2267-2276.

Melloni, L., Molina, C., Pena, M., Torres, D., Singer, W., & Rodriguez, E. (2007). Synchronization of neural activity across cortical areas correlates with conscious perception. *The Journal of neuroscience, 27*(11), 2858-2865.

Mescher, M., Merkle, H., Kirsch, J., Garwood, M., & Gruetter, R. (1998). Simultaneous in vivo spectral editing and water suppression. *NMR in biomedicine, 11*(6), 266-272.

Metzinger, T. (Ed.). (1995). *Conscious experience*. Exeter: Imprint Academic.

Miller, G. (2005). What is the biological basis of consciousness. *Science, 309*(5731), 79-79.

Millikan, R. G. (1984). *Language, thought, and other biological categories : new foundations for realism*. Cambridge, Mass: MIT Press.

Mittmann, T., & Eysel, U. T. (2001). Increased synaptic plasticity in the surround of visual cortex lesions in rats. *Neuroreport, 12*(15), 3341-3347.

Mollon, J. D. (1989). "Tho' she kneel'd in that place where they grew..." The uses and origins of primate colour vision. *The Journal of experimental biology, 146*, 21-38.

Mollon, J. D., & Polden, P. G. (1975). Colour illusion and evidence for interaction between cone mechanisms. *Nature, 258*(5534), 421-422.

Morgan, M. J., Mason, A. J., & Solomon, J. A. (1997). Blindsight in normal subjects? *Nature, 385*(6615), 401-402.

Mulckhuyse, M., & Theeuwes, J. (2010). Unconscious attentional orienting to exogenous cues: A review of the literature. *Acta Psychologica (Amst), 134*(3), 299-309.

Nagel, T. (1974). What Is It Like to Be a Bat. *Philosophical Review, 83*(4), 435-450.

Nealey, T. A., & Maunsell, J. H. (1994). Magnocellular and parvocellular contributions to the responses of neurons in macaque striate cortex. *The Journal of neuroscience, 14*(4), 2069-2079.

Niedeggen, M., Wichmann, P., & Stoerig, P. (2001). Change blindness and time to consciousness. *The European journal of neuroscience, 14*(10), 1719-1726.

Niehaus, L., Guldin, B., & Meyer, B. (2001). Influence of transcranial magnetic stimulation on pupil size. *Journal of the neurological sciences, 182*(2), 123-128.

Nieuwenhuis, S., Forstmann, B. U., & Wagenmakers, E. J. (2011). Erroneous analyses of interactions in neuroscience: a problem of significance. *Nature neuroscience, 14*(9), 1105-1107.

Nisbett, R. E. W., T.C. (1977). Telling more than we can know: Verbal reports on mental processes. *Psychological Review, 84*(3), 28.

Nogawa, T., Katayama, K., Tabata, Y., Ohshio, T., & Kawahara, T. (1976). Changes in amplitude of the EEG induced by a photic stimulus. *Electroencephalography and clinical neurophysiology, 40*(1), 78-88.

Overgaard, M., Fehl, K., Mouridsen, K., Bergholt, B., & Cleeremans, A. (2008). Seeing without Seeing? Degraded Conscious Vision in a Blindsight Patient. *PloS one, 3*(8), e3028.

Overgaard, M., Nielsen, J. F., & Fuglsang-Frederiksen, A. (2004). A TMS study of the ventral projections from V1 with implications for the finding of neural correlates of consciousness. *Brain and Cognition, 54*(1), 58-64.

Overgaard, M., Rote, J., Mouridsen, K., & Ramsoy, T. Z. (2006). Is conscious perception gradual or dichotomous? A comparison of report methodologies during a visual task. *Consciousness and cognition, 15*(4), 700-708.

Overgaard, M. N., J.F. Fuglsang-Frederiksen, A. (2003). A TMS study of the ventral projections from V1 with implications fro the finding of the neural correlates of consciousness. *Brain and Cognition, 54*, 8.

Pascual-Leone, A., Bartres-Faz, D., & Keenan, J. P. (1999). Transcranial magnetic stimulation: studying the brain-behaviour relationship by induction of 'virtual lesions'. *Philosophical transactions of the Royal Society of London. Series B, Biological sciences, 354*(1387), 1229-1238.

Pascual-Leone, A., & Walsh, V. (2001). Fast backprojections from the motion to the primary visual area necessary for visual awareness. *Science, 292*(5516), 510-512.

Pascual-Leone, A., Walsh, V., & Rothwell, J. (2000). Transcranial magnetic stimulation in cognitive neuroscience--virtual lesion, chronometry, and functional connectivity. *Current opinion in neurobiology, 10*(2), 232-237.

Pfurtscheller, G. (1992). Event-related synchronization (ERS): an electrophysiological correlate of cortical areas at rest. *Electroencephalography and clinical neurophysiology, 83*(1), 62-69.

Pfurtscheller, G., & Aranibar, A. (1979). Evaluation of event-related desynchronization (ERD) preceding and following voluntary self-paced movement. *Electroencephalography and Clinical Neurophysiology, 46*(2), 138-146.

Pfurtscheller, G., & Lopes da Silva, F. H. (1999). Event-related EEG/MEG synchronization and desynchronization: basic principles. *Clinical neurophysiology, 110*(11), 1842-1857.

Pfurtscheller, G., Neuper, C., & Mohl, W. (1994). Event-related desynchronization (ERD) during visual processing. *International journal of psychophysiology : official journal of the International Organization of Psychophysiology, 16*(2-3), 147-153.

Pfurtscheller, G., Schwarz, G., Pfurtscheller, B., & List, W. (1983). Quantification of spindles in comatose patients. *Electroencephalography and Clinical Neurophysiology, 56*(1), 114-116.

Pollack, I. (1959). On Indices of Signal and Response Discriminability. *Journal of the Acoustical Society of America, 31*(7), 1031-1031.

Pollack, I., & Norman, D. A. (1964). A Non-Parametric Analysis of Recognition Experiments. *Psychonomic Science, 1*(5), 125-126.

Polonsky, A., Blake, R., Braun, T., & Heeger, D. J. (2000). Neuronal activity in human primary visual cortex correlates with perception during binocular rivalry. *Nature Neuroscience, 3*(11), 1153-1159.

Prinz, J. (2007). The Intermediate Level Theory of Consciousness *The Blackwell Companion to Consciousness* (pp. 247-260): Blackwell Publishing.

Puts, N. A., & Edden, R. A. (2012). In vivo magnetic resonance spectroscopy of GABA: a methodological review. *Progress in nuclear magnetic resonance spectroscopy, 60*, 29-41.

Quiroga, R. Q., Mukamel, R., Isham, E. A., Malach, R., & Fried, I. (2008). Human single-neuron responses at the threshold of conscious recognition. *Proceedings of the National Academy of Sciences of the United States of America, 105*(9), 3599-3604.

Radoeva, P. D., Prasad, S., Brainard, D. H., & Aguirre, G. K. (2008). Neural activity within area V1 reflects unconscious visual performance in a case of blindsight. *Journal of Cognitive Neuroscience, 20*(11), 1927-1939.

Rafal, R., Smith, J., Krantz, J., Cohen, A., & Brennan, C. (1990). Extrageniculate vision in hemianopic humans: saccade inhibition by signals in the blind field. *Science, 250*(4977), 118-121.

Rahnev, D., Bahdo, L., Munneke, M., de Lange, F., & Lau, H. (2010). Theta-burst transcranial magnetic stimulation to V1 impairs subjective confidence ratings and metacognition. *Journal of Vision, 10*(7), 901.

Rahnev, D. A., Maniscalco, B., Luber, B., Lau, H., & Lisanby, S. H. (2011). Direct injection of noise to the visual cortex decreases accuracy but increases decision confidence. *Journal of neurophysiology*, *107* (6),1556-1563.

Railo, H., Salminen-Vaparanta, N., Henriksson, L., Revonsuo, A., & Koivisto, M. (2012). Unconscious and conscious processing of color rely on activity in early visual cortex: a TMS study. *Journal of Cognitive Neuroscience, 24*(4), 819-829.

Ramsoy. T.Z., a. O. M. (2004). Introspection and subliminal perception. *Phenomenology and the Cognitive Sciences, 3*, 23.

Rauschecker, A. M., Bestmann, S., Walsh, V., & Thilo, K. V. (2004). Phosphene threshold as a function of contrast of external visual stimuli. *Experimental brain research, 157*(1), 124-127.

Rees, G. F., C. (2007). Methodologies for identifying the neural correlates of consciousness. In M. S. Velmans, S. (Ed.), *the Blackwell Companion to Consciousness*. Oxford, UK: Blackwell.

Reuter-Lorenz, P. A., Kinsbourne, M., & Moscovitch, M. (1990). Hemispheric control of spatial attention. *Brain and Cognition, 12*(2), 240-266.

Ristic, J., & Kingstone, A. (2006). Attention to arrows: Pointing to a new direction. *Quarterly Journal of Experimental Psychology, 59*(11), 1921-1930.

Ro, T., Shelton, D., Lee, O. L., & Chang, E. (2004). Extrageniculate mediation of unconscious vision in transcranial magnetic stimulation-induced blindsight. *Proceedings of the National Academy of Sciences of the United States of America, 101*(26), 9933-9935.

Robichaud, L., & Stelmach, L. B. (2003). Inducing blindsight in normal observers. *Psychonomic Bulletin & Review, 10*(1), 206-209.

Roepstorff, A., & Jack, A. I. (2004). Trust or interaction? Editorial introduction. *Journal of Consciousness Studies, 11*(7-8).

Romei, V., Gross, J., & Thut, G. (2010). On the role of prestimulus alpha rhythms over occipito-parietal areas in visual input regulation: correlation or causation? *The Journal of neuroscience, 30*(25), 8692-8697.

Rorden, C., & Brett, M. (2000). Stereotaxic display of brain lesions. *Behavioural neurology, 12*(4), 191-200.

Rosenthal, D. (1993). Thinking that One Thinks. In M. H. Davies, G.W. (Ed.), *Consciousness, Readings in Mind and Language*. Oxford: Blackwell.

Rosenthal, D. (2012). Higher-order awareness, misrepresentation and function. *Philosophical transactions of the Royal Society of London. Series B, Biological sciences, 367*(1594), 1424-1438.

Rosenthal, D. M. (2002). A Theory of Consciousness. In N. F. Block, O. and Guzeldere, G. (Ed.), *The Nature of Consciouness, Philosphical debates.*: MIT press.

Rossi, S., Hallett, M., Rossini, P. M., & Pascual-Leone, A. (2009). Safety, ethical considerations, and application guidelines for the use of transcranial magnetic stimulation in clinical practice and research. *Clinical neurophysiology : official journal of the International Federation of Clinical Neurophysiology, 120*(12), 2008-2039.

Rossini, P. M., Barker, A. T., Berardelli, A., Caramia, M. D., Caruso, G., Cracco, R. Q., . . . et al. (1994). Non-invasive electrical and magnetic stimulation of the brain, spinal cord and roots: basic principles and procedures for routine clinical application. Report of an IFCN committee. *Electroencephalography and clinical neurophysiology, 91*(2), 79-92.

Rouder, J. N., Speckman, P. L., Sun, D., Morey, R. D., & Iverson, G. (2009). Bayesian t tests for accepting and rejecting the null hypothesis. *Psychonomic bulletin & review, 16*(2), 225-237.

Rounis, E., Maniscalco, B., Rothwell, J. C., Passingham, R. E., & Lau, H. (2010). Theta-burst transcranial magnetic stimulation to the prefrontal cortex impairs metacognitive visual awareness. *Cognitive Neuroscience, 1*(3), 165-175.

Sack, A. T., van der Mark, S., Schuhmann, T., Schwarzbach, J., & Goebel, R. (2009). Symbolic action priming relies on intact neural transmission along the retino-geniculo-striate pathway. *NeuroImage, 44*(1), 284-293.

Sahraie, A., Hibbard, P. B., Trevethan, C. T., Ritchie, K. L., & Weiskrantz, L. (2010). Consciousness of the first order in blindsight. *Proceedings of the National Academy of Sciences of the United States of America, 107*(49), 21217-21222.

Sandberg, K., Timmermans, B., Overgaard, M., & Cleeremans, A. (2010). Measuring consciousness: is one measure better than the other? *Consciousness and cognition, 19*(4), 1069-1078.

Sanders, M. D., Warrington, E. K., Marshall, J., & Wieskrantz, L. (1974). "Blindsight": Vision in a field defect. *Lancet, 1*(7860), 707-708.

Schiller, P. H., & Malpeli, J. G. (1977). Properties and tectal projections of monkey retinal ganglion cells. *Journal of Neurophysiology, 40*(2), 428-445.

Schmid, M. C., Mrowka, S. W., Turchi, J., Saunders, R. C., Wilke, M., Peters, A. J., . . . Leopold, D. A. (2010). Blindsight depends on the lateral geniculate nucleus. *Nature, 466*(7304), 373-377.

Schubert, R., Blankenburg, F., Lemm, S., Villringer, A., & Curio, G. (2006). Now you feel it--now you don't: ERP correlates of somatosensory awareness. *Psychophysiology, 43*(1), 31-40.

Schwarzkopf, D. S., Silvanto, J., & Rees, G. (2011). Stochastic resonance effects reveal the neural mechanisms of transcranial magnetic stimulation. *The Journal of neuroscience, 31*(9), 3143-3147.

Sergent, C., Baillet, S., & Dehaene, S. (2005). Timing of the brain events underlying access to consciousness during the attentional blink. *Nature neuroscience, 8*(10), 1391-1400.

Shipp, S. (2004). The brain circuitry of attention. *Trends in Cognitive Sciences, 8*(5), 223-230.

Sidis, B. (1898). The Psychology of Suggestion. *Science, 8*(188), 162-163.

Silvanto, J., Cattaneo, Z., Battelli, L., & Pascual-Leone, A. (2008). Baseline cortical excitability determines whether TMS disrupts or facilitates behavior. *Journal of neurophysiology, 99*(5), 2725-2730.

Silvanto, J., Cowey, A., Lavie, N., & Walsh, V. (2005). Striate cortex (V1) activity gates awareness of motion. *Nature Neuroscience, 8*(2), 143-144.

Silvanto, J., Lavie, N., & Walsh, V. (2005). Double dissociation of V1 and V5/MT activity in visual awareness. *Cerebral Cortex, 15*(11), 1736-1741.

Sincich, L. C., & Horton, J. C. (2005). The circuitry of V1 and V2: integration of color, form, and motion. *Annual review of neuroscience, 28*, 303-326.

Sincich, L. C., Park, K. F., Wohlgemuth, M. J., & Horton, J. C. (2004). Bypassing V1: a direct geniculate input to area MT. *Nature neuroscience, 7*(10), 1123-1128.

Singh, K. D. (2006). Magnetoencephalography. In C. R. Senior, T. Gazzaniga, M.S. (Ed.), *Method in Mind*: MIT press.

Skinner, B. F. (1965). *Science and Human Behavior*: pp. x. 461. Free Press: New York; Collier-Macmillan: London.

Smith, R. A., Levine, T. R., Lachlan, K. A., & Fediuk, T. A. (2002). The High Cost of Complexity in Experimental Design and Data Analysis: Type I and Type II Error Rates in Multiway ANOVA. *Human Communication Research, 28*(4), 515-530.

Smithson, H., Sumner, P., & Mollon, J. D. (2003). How to find a tritan line. In J. D. Mollon (Ed.), *Normal and Defective Colour Vision* Oxford University Press.

Snodgrass, J. G., & Corwin, J. (1988). Pragmatics of measuring recognition memory: applications to dementia and amnesia. *Journal of experimental psychology. General, 117*(1), 34-50.

Sperling, G. (1960). The Information Available in Brief Visual Presentations. *Psychological Monographs, 74*(11), 1-29.

Stagg, C. J., Wylezinska, M., Matthews, P. M., Johansen-Berg, H., Jezzard, P., Rothwell, J. C., & Bestmann, S. (2009). Neurochemical effects of theta burst stimulation as assessed by magnetic resonance spectroscopy. *Journal of Neurophysiology, 101*(6), 2872-2877.

Stanislaw, H., & Todorov, N. (1999). Calculation of signal detection theory measures. *Behavior research methods, instruments, & computers, 31*(1), 137-149.

Steriade, M., McCormick, D. A., & Sejnowski, T. J. (1993). Thalamocortical oscillations in the sleeping and aroused brain. *Science, 262*(5134), 679-685.

Stevens, L. K., McGraw, P. V., Ledgeway, T., & Schluppeck, D. (2009). Temporal characteristics of global motion processing revealed by transcranial magnetic stimulation. *The European journal of neuroscience, 30*(12), 2415-2426.

Stewart, L. M., Walsh, V., & Rothwell, J. C. (2001). Motor and phosphene thresholds: a transcranial magnetic stimulation correlation study. *Neuropsychologia, 39*(4), 415-419.

Stokes, M. G., Chambers, C. D., Gould, I. C., Henderson, T. R., Janko, N. E., Allen, N. B., & Mattingley, J. B. (2005). Simple metric for scaling motor threshold based on scalp-cortex distance: application to studies using transcranial magnetic stimulation. *Journal of neurophysiology, 94*(6), 4520-4527.

Sumner, P. (2000). *The salience of colour: Studies in visual ecology and psychophysics.* Ph.D., University of Cambridge., Cambridge UK.

Sumner, P. (2006). Inhibition versus attentional momentum in cortical and collicular mechanisms of IOR. *Cognitive neuropsychology, 23*(7), 1035-1048.

Sumner, P., Adamjee, T., & Mollon, J. D. (2002). Signals invisible to the collicular and magnocellular pathways can capture visual attention. *Current Biology, 12*(15), 1312-1316.

Sumner, P., Nachev, P., Castor-Perry, S., Isenman, H., & Kennard, C. (2006). Which visual pathways cause fixation-related inhibition? *Journal of neurophysiology, 95*(3), 1527-1536.

Sumner, P., Nachev, P., Vora, N., Husain, M., & Kennard, C. (2004). Distinct cortical and collicular mechanisms of inhibition of return revealed with S cone stimuli. *Currunt Biology, 14*(24), 2259-2263.

Sumner, P., Tsai, P. C., Yu, K., & Nachev, P. (2006). Attentional modulation of sensorimotor processes in the absence of perceptual awareness. *Proceedings of the National Academy of Sciences of the United States of America, 103*(27), 10520-10525.

Super, H., Spekreijse, H., & Lamme, V. A. (2001). Two distinct modes of sensory processing observed in monkey primary visual cortex (V1). *Nature neuroscience, 4*(3), 304-310.

Swets, J. A. (1986). Form of Empirical Rocs in Discrimination and Diagnostic Tasks - Implications for Theory and Measurement of Performance. *Psychological Bulletin, 99*(2), 181-198.

Tamietto, M., Cauda, F., Corazzini, L. L., Savazzi, S., Marzi, C. A., Goebel, R., . . . de Gelder, B. (2010). Collicular vision guides nonconscious behavior. *Journal of Cognitive Neuroscience, 22*(5), 888-902.

Tanner, W. P., Jr., & Swets, J. A. (1954). A decision-making theory of visual detection. *Psychol Rev, 61*(6), 401-409.

Taylor, J. R. (1997). *An introduction to error analysis : the study of uncertainties in physical measurements* (2nd ed.). Sausalito, California: University Science Books.

Tegenthoff, M., Ragert, P., Pleger, B., Schwenkreis, P., Forster, A. F., Nicolas, V., & Dinse, H. R. (2005). Improvement of tactile discrimination performance and enlargement of cortical somatosensory maps after 5 Hz rTMS. *PLoS biology, 3*(11), e362.

Terao, Y., Ugawa, Y., Suzuki, M., Sakai, K., Hanajima, R., Gemba-Shimizu, K., & Kanazawa, I. (1997). Shortening of simple reaction time by peripheral electrical and submotor-threshold magnetic cortical stimulation. *Experimental brain research, 115*(3), 541-545.

Thielscher, A., Reichenbach, A., Ugurbil, K., & Uludag, K. (2010). The cortical site of visual suppression by transcranial magnetic stimulation. *Cerebral cortex, 20*(2), 328-338.

Thompson, B., Mansouri, B., Koski, L., & Hess, R. F. (2008). Brain plasticity in the adult: modulation of function in amblyopia with rTMS. *Current biology 18*(14), 1067-1071.

Thompson, E., & Varela, F. J. (2001). Radical embodiment: neural dynamics and consciousness. *Trends in Cognitive Sciences, 5*(10), 418-425.

Timmermans, B., Schilbach, L., Pasquali, A., & Cleeremans, A. (2012). Higher order thoughts in action: consciousness as an unconscious re-description process. *Philosophical transactions of the Royal Society of London. Series B, Biological sciences, 367*(1594), 1412-1423.

Tong, F. (2003). Primary visual cortex and visual awareness. *Nature reviews. Neuroscience, 4*(3), 219-229.

Trevethan, C. T., Sahraie, A., & Weiskrantz, L. (2007). Can blindsight be superior to 'sighted-sight'? *Cognition, 103*(3), 491-501.

van Dijk, H., Schoffelen, J. M., Oostenveld, R., & Jensen, O. (2008). Prestimulus oscillatory activity in the alpha band predicts visual discrimination ability. *The Journal of neuroscience : the official journal of the Society for Neuroscience, 28*(8), 1816-1823.

Varela, F. J. (1996). Neurophenomenology: a methodological remedy for the hard problem. *Journal of Conciousness Studies, 3*(4), 20.

Varela, F. J., Lachaux, J. P., Rodriguez, E., & Martinerie, J. (2001). The brainweb: phase synchronization and large-scale integration. *Nature reviews. Neuroscience, 2*(4), 229-239.

Varela, F. J., Thompson, E., & Rosch, E. (1991). *The Embodied mind : cognitive science and human experience*: The MIT Press.

Varela, F. J., Toro, A., John, E. R., & Schwartz, E. L. (1981). Perceptual framing and cortical alpha rhythm. *Neuropsychologia, 19*(5), 675-686.

Varela, F. J. a. S., J. (1999). The View From Within. In J. A. Gouen (Ed.).

Varnava, A., Stokes, M. G., & Chambers, C. D. (2011). Reliability of the 'observation of movement' method for determining motor threshold using transcranial magnetic stimulation. *Journal of neuroscience methods, 201*(2), 327-332.

Vermersch, P. (1999). Introspection as practice. In F. J. Varela, J. Shear (Ed.), *The View From Within*: Imprint Academic.

Vrba, J., & Robinson, S. E. (2001). Signal processing in magnetoencephalography. *Methods, 25*(2), 249-271.

Vuilleumier, P., Sagiv, N., Hazeltine, E., Poldrack, R. A., Swick, D., Rafal, R. D., & Gabrieli, J. D. (2001). Neural fate of seen and unseen faces in visuospatial neglect: a combined event-related functional MRI and event-related potential study. *Proceedings of the National Academy of Sciences of the United States of America, 98*(6), 3495-3500.

Wagner, T., Rushmore, J., Eden, U., & Valero-Cabre, A. (2009). Biophysical foundations underlying TMS: setting the stage for an effective use of neurostimulation in the cognitive neurosciences. *Cortex, 45*(9), 1025-1034.

Walker, M. F., Fitzgibbon, E. J., & Goldberg, M. E. (1995). Neurons in the monkey superior colliculus predict the visual result of impending saccadic eye movements. *Journal of neurophysiology, 73*(5), 1988-2003.

Walsh, V., & Cowey, A. (2000). Transcranial magnetic stimulation and cognitive neuroscience. *Nature reviews. Neuroscience, 1*(1), 73-79.

Wassermann, E. M. (1998). Risk and safety of repetitive transcranial magnetic stimulation: report and suggested guidelines from the International Workshop on the Safety of Repetitive Transcranial Magnetic Stimulation, June 5-7, 1996. *Electroencephalography and Clinical Neurophysiology, 108*(1), 1-16.

Waterston, M. L., & Pack, C. C. (2010). Improved discrimination of visual stimuli following repetitive transcranial magnetic stimulation. *PloS one, 5*(4), e10354.

Weiskrantz, L. (1986). *Blindsight: a case study and implications* (Vol. 12): Oxford university press.

Weiskrantz, L. (1996). Blindsight revisited. *Current opinion in neurobiology, 6*(2), 215-220.

Weiskrantz, L. (1997). *Consciousness lost and found : a neuropsychological exploration*. Oxford: Oxford University Press.

Weiskrantz, L. (2001). Blindsight - putting beta on the back burner. In B. D. H. De Gelder, E.H.F. Heywood, C.A. (Ed.), *Out of Mind - varieties of unconscious processes*. Oxford: Oxford University Press.

Weiskrantz, L. (2009). Is blindsight just degraded normal vision? *Experimental brain research, 192*(3), 413-416.

Wessinger, C. M., Fendrich, R., & Gazzaniga, M. S. (1997). Islands of residual vision in hemianopic patients. *Journal of Cognitive Neuroscience, 9*(2), 203-221.

Wilson, T. D. C., & Nisbett, R. E. (1978). Accuracy of Verbal Reports About Effects of Stimuli on Evaluations and Behavior. *Social Psychology, 41*(2), 118-131.

Winer, B. J., Brown, D. R., & Michels, K. M. (1991). *Statistical principles in experimental design* (3rd ed.). New York: McGraw-Hill.

Wolin, L. R., Massopust, L. C., Jr., & Meder, J. (1966). Differential color responses from the superior colliculi of squirrel monkeys. *Vision Research, 6*(12), 637-644.

Wright, D. B., Horry, R., & Skagerberg, E. M. (2009). Functions for traditional and multilevel approaches to signal detection theory. *Behavior research methods, 41*(2), 257-267.

Yamanaka, K., Yamagata, B., Tomioka, H., Kawasaki, S., & Mimura, M. (2010). Transcranial magnetic stimulation of the parietal cortex facilitates spatial working memory: near-infrared spectroscopy study. *Cerebral cortex, 20*(5), 1037-1045.

Yordanova, J., Kolev, V., & Polich, J. (2001). P300 and alpha event-related desynchronization (ERD). *Psychophysiology, 38*(1), 143-152.

Zeki, S., & Ffytche, D. H. (1998). The Riddoch syndrome: insights into the neurobiology of conscious vision. *Brain, 121 (Pt 1)*, 25-45.