

Intentional Binding and the Sense of Agency: A review

James W. Moore^{1,2} & Sukhvinder S. Obhi^{3,4}

1. Department of Psychology, Goldsmiths, University of London, London, UK
2. Brain Mapping Unit, Department of Psychiatry, University of Cambridge, Cambridge, UK
3. Centre for Cognitive Neuroscience & Department of Psychology, Wilfrid Laurier University, Waterloo, Ontario, Canada
4. Institute of Cognitive Neuroscience, University College London, London, U.K.

Correspondence: sobhi@wlu.ca

PrePrint (uncorrected proof) – Consciousness & Cognition (2012)

Abstract

It is nearly 10 years since Patrick Haggard and colleagues first reported the ‘intentional binding’ effect (Haggard et al, 2002). The intentional binding effect refers to the subjective compression of the temporal interval between a voluntary action and its external sensory consequence. Since the first report, considerable interest has been generated and a fascinating array of studies has accumulated. Much of the interest in intentional binding comes from the promise to shed light on human agency. In this review we survey studies on intentional binding, focusing, in particular, on the link between intentional binding and the sense of agency (the experience of controlling action to influence events in the environment). We suggest that, whilst it is yet to be fully explicated, the link between intentional binding and the sense of agency is compelling. We conclude by considering outstanding questions and future directions for research on intentional binding.

1. Introduction

Humans are agents. That is, they have the capacity to bring about change in the external world through their own goal-directed behaviour. Often, humans also have a corresponding conscious experience of this capacity, which is referred to as the ‘sense of agency’. A key challenge for scientific investigations of the sense of agency is the discovery and use of appropriate measures. In 2002, Haggard and colleagues introduced a novel measure of the sense of agency based on an intriguing relationship between voluntary action and subjective time (Haggard, Clark, & Kalogeras, 2002; Haggard, Aschersleben, et al., 2002). This so-called ‘intentional binding’ measure has generated considerable interest and has been used in a number of experiments on the sense of agency. However, as a measure of the sense of agency it is not without its detractors (e.g. Buehner & Humphreys, 2009). We therefore feel that the time is right for a review of experiments using intentional binding to study the sense of agency. We first provide a brief historical background to the use of subjective time in experimental psychology. We begin by introducing Benjamin Libet’s seminal work and the ‘clock methodology’ he used (which forms the basis of the original intentional binding paradigm). Having set the scene we then review experiments using the intentional binding paradigm. We conclude this review by considering a) the validity of intentional binding as a measure of the sense of agency, and b) future directions for research.

1.1. The use of subjective time in Experimental Psychology: Libet’s work on volition

Measures based on the subjective experience of time have a long history in Experimental Psychology. In the 1880s Wilhelm Wundt developed his complication-clock apparatus to explore the time course of attention (Figure 1). Participants had to orient to a clock (or pendulum) when presented with a certain stimulus (such as an auditory click) and report the onset of that stimulus by noting the position of the clock hand (or pendulum) when the stimulus occurred. Wundt observed systematic differences in the perceived onset of the auditory stimulus: People either perceived the auditory event earlier or later relative to the position of the clock hand (or pendulum). Moreover, this difference was attributed to whether participants were attending to the clock hand or the auditory stimulus. Wundt’s chronometric methodology thus provides an invaluable tool for comparing subjective and objective stimulus onset timings.

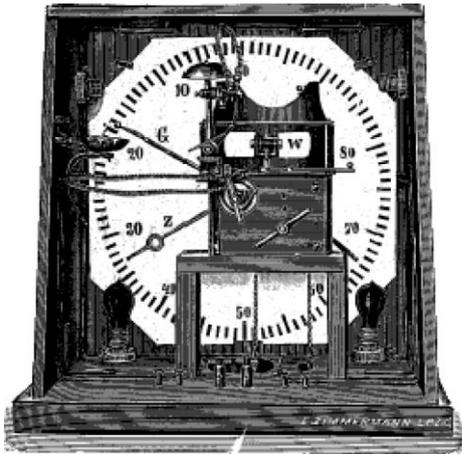


Figure 1. The complication-clock apparatus developed by Wundt.

In the 1980s Benjamin Libet adopted Wundt's classic methodology to explore human volition. In Libet et al's (1983) seminal study, participants sat in front of a clock face marked at regular intervals. During each trial a spot rotated around the clock-face at a speed of 1 revolution every 2.56s. Participants used the clock to judge the onset of certain events. In one condition participants flexed their wrist when they felt the urge, and judged the time they became aware of raising their hand. In a second condition participants again flexed their wrist when they felt the urge, and this time they judged the time they became aware of their conscious *intention* to raise their hand. In a third condition, a somatosensory stimulus was applied at an unpredictable time during the trial and participants judged the time that they felt this stimulus. To make these timing judgements participants reported the position of the spot on the clock face when they were perceived the event (intention, action or somatosensory stimulus). Libet and colleagues concurrently used electroencephalography (EEG) to record the readiness potential, a cascade of neural activity that reliably precedes the onset of voluntary movement. Libet found that, judgements of movements were slightly early compared to actual movement onset, whereas judgements of somatosensory stimuli were slightly delayed compared to their actual onset. Famously, although intentions were perceived as occurring prior to movements, they lagged behind the onset of the readiness potential. This suggests that the intention to act may arise after the brain has initiated an action. The implications of this work have been extensively discussed elsewhere. What is most relevant to the current paper is the 'clock methodology' that Libet used as it formed the basis of the original work on intentional binding, which is discussed in the following section.

1.2. Original experiment and some basic conditions for intentional binding

In the first journal article to report the intentional binding effect, Haggard, Clark & Kalogeras (2002) used the Libet clock method (see earlier description of this approach) to study the perceived time of actions and their consequent effects (see also Haggard, Aschersleben, Gehrke, & Prinz, 2002). There were four critical conditions in their experiment. In baseline conditions, participants either made voluntary actions or listened for the occurrence of an auditory tone (in the absence of action) while they watched a rotating clock hand on a computer screen. They were asked to report the position of the clock hand when they moved or when the tone occurred. In operant conditions, participants made a voluntary key press on every trial, but this time it was followed 250 ms later by an auditory tone. In specific blocks, participants were asked to judge either the time of their action or the time of the tone. The trial structure of a typical operant condition is shown in figure 2. The key comparison was the perceived times of actions and tones in baseline conditions with the perceived time of actions and tones in operant conditions. The authors found that, in operant conditions, the perceived time of their actions was later than in baseline conditions and the perceived time of the tone was earlier than in baseline conditions. Hence, the derived interval between the action and tone in operant conditions was compressed compared to the derived interval between these two events in baseline conditions. Critically, in an identical set of conditions involving involuntary movements induced via transcranial magnetic stimulation over the primary motor cortex, the binding effect was reversed such that the interval between action and effect actually increased in ‘operant’ conditions compared to baseline conditions. Figure 3 shows the classic pattern of intentional binding found in this seminal experiment.

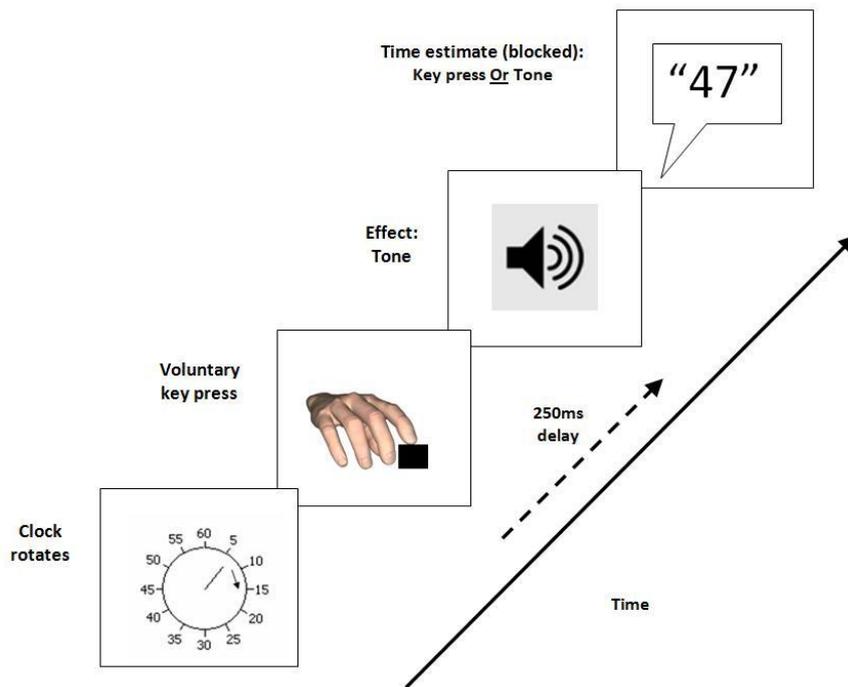


Figure 2. Typical trial structure in an *operant* condition (following Haggard, Clark, & Kalogeras, 2002). Participants pressed the key at a time of their choosing which produced a tone after a delay of 250ms. Participants judged where the clock hand was when they pressed the key or when they heard the tone, in separate blocks of trials

In a second experiment the authors varied the delay and the predictability of when the tone occurred after voluntary action. The results from this study indicated that temporal contiguity and predictability are important determinants for binding. On the basis of their findings the authors speculated that a specific cognitive function of the central nervous system is to bind together critical sensorimotor events that surround voluntary action, and that this function may be crucial for the normal experience of agency. There are two things to highlight about this early hypothesis. The first is the assumed link with sense of agency. The aim of this review is to scrutinise this assumption, and, while the evidence reviewed provides compelling support for this link, it should be noted that alternative proposals have been advanced. For example, intentional binding may reflect the ‘unity of perception’ (Yarrow et al., 2001) or be more generally linked to causality rather than agency (Stetson, Montague and Eagleman, 2006). The second thing to highlight is that this early hypothesis suggested that binding may help support inferences of agency. However, as we shall see in this review, more recent evidence suggests that inferences of agency can also support binding (e.g. Moore & Haggard, 2008). This serves to emphasise a recurring theme in this review; namely, that although there

is compelling evidence supporting a link between intentional binding and sense of agency, the exact nature of that relationship is yet to be fully understood

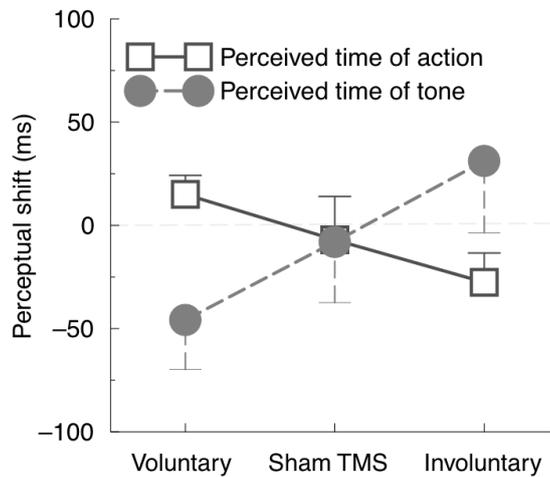


Figure 3. The classic pattern of intentional binding described by Haggard, Clark and Kalogeras (2002). The perception of onset of voluntary actions was shifted later in time and the perception of tone onset was shifted earlier. This was not the case for involuntary actions induced by TMS, which showed the opposite effects. Original figure used with permission.

Since the seminal journal article on intentional binding, many studies have been performed using the Libet clock method. However, another, more direct method has also been employed, yielding confirmatory results. Specifically, citing some of the methodological criticisms surrounding the Libet clock method, several authors have used a direct interval estimation procedure in which participants are instructed to simply report the perceived interval (in milliseconds) between an action and effect (see Figure 4). These studies, which we describe in more detail later, have shown that mean interval estimates are lower (indicating a shorter perceived interval) in voluntary vs. Passive conditions. This temporal illusion is just one example of perceptual distortion surrounding voluntary action. Another well-known example is the spatial compression effect in which targets flashed around the onset of a saccade are systematically mislocalised toward the saccade landing position (interested readers are referred to Ross, Morrone, & Burr, 1997 for more on this fascinating perceptual effect). In the context of this review, we shall refer to the direct estimation of temporal intervals as the (direct) interval estimation approach.

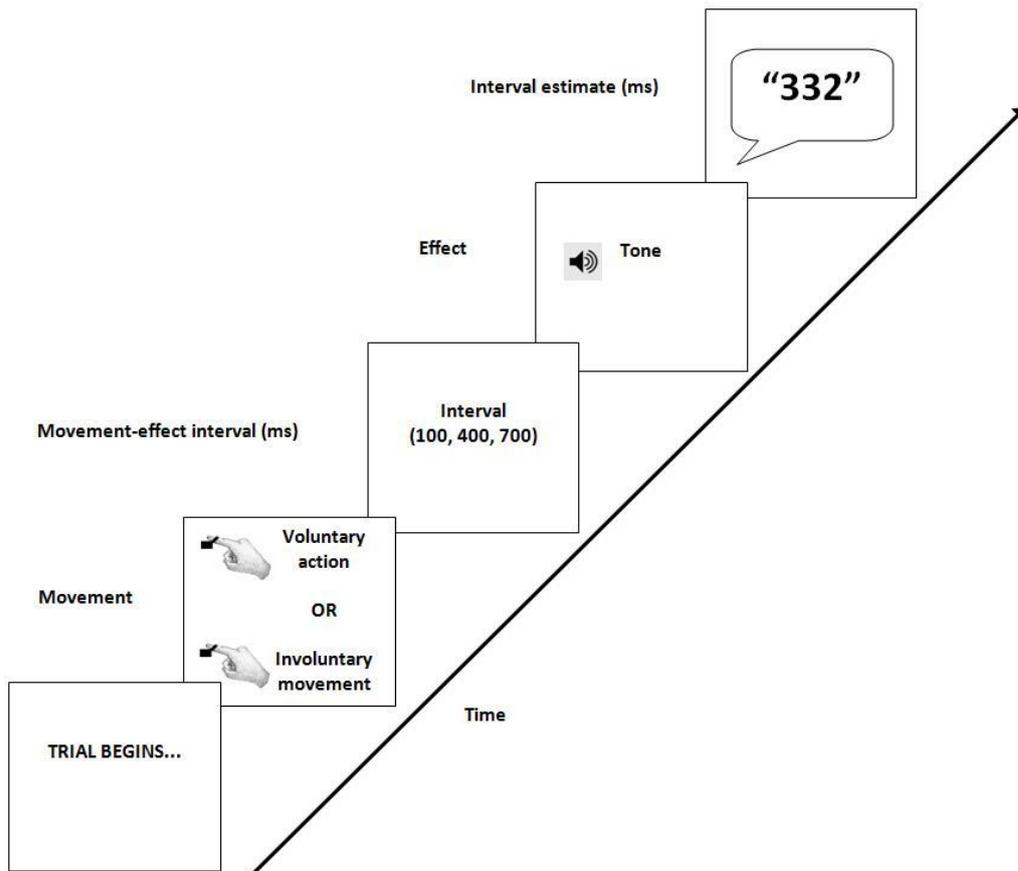


Figure 4. Schematic showing typical trial structure for the interval estimation procedure. Participants judge the duration of the movement-effect interval in milliseconds. This is an alternative to the Libet clock procedure (see Figure 2).

The original demonstration of intentional binding and the suggestion of a link to agency sparked considerable interest from researchers interested in the phenomenology of human action. The remaining sections of our review highlight some of the most significant work that followed the original journal article. As will be evident, intentional binding and its potential relationship with the sense of agency is far from fully understood, although with each additional experiment, a slightly clearer view of the significance of intentional binding in human action contexts and the factors that modulate intentional binding, is emerging.

2. Prediction and retrospective inference

The purported link between intentional binding and the sense of agency is perhaps the most tantalising aspect of the original article. To appreciate why the link was suggested, consideration must be given to current ideas about how the sense of agency might arise. In this regard, it is possible to distinguish between two theoretical positions on the

neurocognitive origins of the sense of agency. On the one hand there is the ‘*predictive*’ position. On this view the sense of agency is generated by processes dedicated to the control of voluntary action. Optimal motor control and learning require predictions of both the future states of the motor system and the sensory consequences of movement (Wolpert & Ghahramani, 2000). These predictions are derived from *internal forward models*, of which there are two classes: *forward dynamic* and *forward sensory*. The *forward dynamic model* captures the dynamics of bodily movement. The *forward sensory model* captures the causal relation between movements and their sensory consequences, generating predictions of the likely sensory consequences of movement based on efference copy of motor commands. According to the so-called, ‘comparator model’ of the sense of agency, this plays a key role in the sense of agency (Blakemore, Wolpert, & Frith, 2002): the sense of agency is produced when there is a match between predicted and actual sensory consequences of movement.

On the other hand there are those who downplay the specific contribution of the motor system. Instead, it is suggested that the sense of agency is generated by a process of *retrospective inference*. On this view a general-purpose inferential mechanism uses sensory information to establish the causal origins of actions and their effects. One influential version of this position is Wegner’s ‘theory of apparent mental causation’ (Wegner & Wheatley, 1999; Wegner, 2002; Wegner, 2003), which takes an explicitly Humean approach. According to this theory, the experience of willing an action arises if a thought (intention) 1) occurs prior to action, 2) is consistent with the action, and 3) is the most plausible cause of the action.

This same debate, prediction vs. retrospective inference, concerns the intentional binding effect (Haggard & Clark, 2003). Intentional binding could be produced by dedicated motor control mechanisms predicting the sensory consequences of an action. Alternatively, intentional binding could be the result of retrospective inference triggered by the actual presence of the sensory consequences of movement. The apparent specificity of the binding effect to voluntary movement strongly implies the involvement of predictive motor processes. Support for this was provided by Haggard and Clark (2003), who investigated whether disrupting the intention to produce an action impacted on intentional binding. Participants made voluntary key presses whilst watching the Libet clock. In some blocks transcranial magnetic stimulation (TMS) was randomly applied over motor cortex. The aim was to disrupt the completion of some of these intentional actions by triggering an identical

involuntary movement. Some actions were therefore truly intentional (the action precisely matched the participant's intention) and some were not (the action did not precisely match the participant's intention). Haggard and Clark found that disrupting intentions significantly weakened intentional binding. These results are consistent with the predictive interpretation of the effect: they suggest that intentional binding is supported by processes involved in the preparation of action. The fact that the mere co-occurrence of actions and outcomes was not sufficient to produce binding (in the intention disruption trials) was interpreted as evidence against the contribution of retrospective inference.

The idea that intentional binding is supported by predictive processes involved in the preparation of action is supported by Tsakiris and Haggard's (2003) study. Previous experiments have shown that, when the predictions of a forward model are matched by sensory feedback the 'strength' of the feedback is attenuated, and sensory effects are therefore subjectively experienced as less intense (e.g. Blakemore, Frith, & Wolpert, 1999). Tsakiris & Haggard (2003) demonstrated that intentional binding occurred for actions that were followed by somatic effects, and also that somatic effects following action were experienced as less intense than those following passive movement. It should be noted that these effects were not correlated in this experiment rather they merely co-occurred in the same sample. Nevertheless, this is suggestive of a link between intentional binding and the operation of predictive forward models.

Further support for the predictive interpretation came from a study by (Engbert & Wohlschläger, 2007), which focussed on the action component of intentional binding. Using the Libet clock method, they assessed the impact of changing certain properties of the action *outcome* on the perceived time of action. A first experiment looked at the impact of whether or not an outcome was intended. Of more relevance to the current discussion on prediction, Experiments 2 and 3 assessed the impact of outcome probability on the perceived time of action. In Experiment 2, the probability of the outcome (a tone) was manipulated, so that in one condition that probability was higher (80%) than in the other (20%). They found that actions were perceived later (i.e. intentional binding for actions was stronger) in the higher outcome probability condition. Interestingly, this delay in the perceived time of action was observed on trials with *and* without a tone, suggesting that the presence of the tone was not necessary for intentional binding. This emphasises the role of prediction in intentional binding: predicting the outcome was sufficient to induce binding. In a final experiment they

showed that this predictive effect was specific to voluntary action. When the movements were *passively induced* the perceived time of these movements did not significantly differ in the high and low outcome probability conditions. Engbert and Wohlschläger's study provides convincing evidence in support of the role of prediction in intentional binding, showing that intentional binding is sensitive to whether or not an outcome is predicted. Moreover, the mechanisms providing these contributions to intentional binding seem to be specifically motoric.

Moore and Haggard (2008) also investigated the contribution of prediction and retrospective inference to the action component of intentional binding. Participants made voluntary key-presses while watching the Libet-clock. In one condition the outcome was predictable: the key press produced a tone on 75% of trials. In another condition, the outcome was unpredictable: the key press produced a tone on 50% of trials. Using this design, Moore and Haggard could isolate the contribution of predictive and retrospective inferential processes to action binding. The contribution of prediction was confirmed by an increase in action binding on 'action only' trials (where the key press *did not* produce the tone) in the 75% vs. 50% conditions. This increase in binding could only be due to the increase in outcome probability. The contribution of retrospective inference was confirmed by an increase in action binding on 'action + tone' vs. 'action only' trials in the 50% condition. In this condition the contribution of prediction was minimal because the outcome was unpredictable. Therefore, the increase in binding on 'action + tone' trials could only have been due to the presence of the tone. This suggests that the tone retrospectively triggered a shift in the perceived time of action. In a follow-up analysis, the authors showed that predictive action binding – binding in the absence of the outcome - was modulated by the *recency* of reinforcement history. Only when an action had produced an outcome in the immediately preceding trial was there a significant shift in the perceived time of action on 'action only' trials. Overall, this study not only emphasises the importance of outcome prediction for intentional binding, but also provides the first evidence that retrospective processes are involved. Moreover, the recency effect suggests that learning may be important.

The study by Moore and Haggard (2008) suggests that retrospective inference plays a role in intentional binding. In other studies of motor experience the contribution of this process has been shown using priming paradigms. In a seminal study, Wegner and Wheatley (1999) showed that simply priming participants with words consistent with an action outcome

increased the participants feeling of controlling the action, even though they had not actually made the movement. According to the authors, these primes engendered thoughts about action consequences, and therefore acted as surrogate intentions. The fact that these primes were able to modulate the participants' feeling of control over movements they had not made was taken as evidence that the experience of control arises from a process that *infers* causal relations between events (intentions, actions, and outcomes). According to the authors, "this interpretive process is fundamentally separate from the mechanistic process of real mental causation". This clearly downplays the contribution of the motor control system to the sense of agency.

Moore, Wegner and Haggard (2009) subsequently investigated whether priming effects could also be shown for intentional binding using the interval estimation procedure. There were two movement conditions: voluntary and passive. Prior to the movement in each condition participants were presented with a prime: a high or low pitch tone. The movement then produced, after a brief delay, a high or low pitch tone. In this way, the authors were able to manipulate prime-outcome consistency. They found that intentional binding was increased when primes were congruent with the outcome. Given the putative link between intentional binding and the sense of agency, this finding was consistent with previous priming studies showing that the sense of agency was augmented for congruent primes. Despite the significant main effect of prime congruence, it was also found that the effect of the prime was significantly reduced when the movement was voluntary. This suggests that, although retrospective inference does play a role in intentional binding (and the sense of agency more broadly), its role is most important in the absence of predictive motor processes that normally dominate.

The studies reviewed in this section show that *both* prediction and retrospective inference contribute to intentional binding. However, the study by Moore, Wegner and Haggard (2009) suggests that the dichotomy between prediction and retrospective inference may be too crude. In response to this, Moore, Wegner and Haggard (2009) suggested that the sense of agency is based on different agency cues (including but perhaps not limited to sensorimotor predictions and external action consequences) and that the relative influence of these cues is determined by their reliability. In this way the sense of agency may be based on a mechanism that optimally integrates a *variety* of agency cues. As will be discussed later, this model may help explain patterns of agency experience associated with certain diseases, such as schizophrenia.

3. Causality and intentional binding

A number of studies have considered the role of *causality* in intentional binding. Causal inferences are likely to be based on properties of the action-outcome relation such as contiguity and contingency. Contiguity refers to the temporal proximity of the action and the outcome. The closer together in time these events appear the stronger the inferred causal relation (Hume, 1888). The role of contiguity in intentional binding is unclear. Haggard, Clark and Kalogeras (2002) examined intentional binding at action-outcome delays of 250ms, 450ms, and 650ms, and found that intentional binding was weaker at longer delays. In contrast, Buehner and Humphreys (2009) and Humphreys and Buehner (2009) found that intentional binding was preserved at action-outcome delays of up to 4s.

A novel take on the link between contiguity and intentional binding has also been made by Stetson, Montague and Eagleman (2006). Stetson et al found that when participants are adapted to a fixed delay between an action (key press) and a sensory consequence (colour flash), occasional presentation of the colour flash at a shorter delay can lead to an illusory reversal of the action and its sensory consequence. They suggest that this effect occurs because once sensory events (such as a flash of colour) are interpreted as *consequences* of actions the brain re-calibrates timing judgements to fit the prior expectation that actions and sensory events are contiguous. When an outcome is presented at a shorter delay, this recalibration can lead to an illusory reversal of action and outcome. Although their study was not directly concerned with intentional binding, Stetson et al also suggest that this recalibration mechanism underpins the intentional binding effect – reflecting the fact that the brain is recalibrating the timing of actions and outcomes to fit the prior expectation that these events should be contiguous. Stetson et al also investigated whether this recalibration occurred for sequences of sensory events in the absence of action. The sensory events were a tap applied to the finger followed by a flash of colour. Although there was numerical evidence suggesting the presence of illusory reversals in this condition, this was not quite statistically significant. This emphasises the importance of voluntary action in generating these shifts in perceived time.

Further to its implications for the role of contiguity in intentional binding, Stetson et al's study also highlights the apparent importance of contingency learning – recalibration is

thought to take place once a sensory event is interpreted as a *consequence* of moving. Goal-directed behaviour is mediated by contingency learning (Balleine & Dickinson, 1998). During contingency learning, organisms acquire knowledge about the causal relation between actions and outcomes; the extent to which an outcome is *contingent* on an action. There is a non-contingent relation when the outcome is equally likely in the presence and absence of the action. There is a contingent relation when the probability of the outcome is increased in the presence of action. By computing contingency information, organisms are able to purposefully control their environment.

Moore, Lagnado, Deal and Haggard (2009) investigated the relevance of contingency learning for intentional binding, by measuring its impact on the action component of intentional binding using the Libet clock method. They used the same probability manipulation as Moore and Haggard (2008) described above, with 75% and 50% outcome probability conditions. The key methodological change was that in each condition participants were given the choice of whether or not to press the key on each trial. The probability of the outcome in the absence of the key press (when subjects chose to withhold action) was manipulated, allowing the authors to manipulate action-outcome contingency. It was found that increasing action-outcome contingency significantly increased the contribution of predictive and retrospective inferential processes to intentional binding. This suggests that the occurrence of intentional binding (and the processes subserving it) is sensitive to the causal relation between actions and outcomes. Therefore, whilst the presence of intentional movement seems to be necessary for intentional binding, depending on the context, it may not be sufficient.

Buehner and Humphreys (2009) also demonstrated the importance of contingency in intentional binding using a novel intentional binding paradigm. In a non-causal condition, participants were initially exposed to a succession of two tones (t1 and t2), thus learning that t2 was not contingent on action. In a causal condition participants had to press a key, which caused a tone (t2), thus learning that t2 was contingent on action. In a subsequent training phase, participants were presented with a succession of tones (t1 and t2) and had to synchronise a key press with t1. In the non-causal condition t2 was contingent on t1 whereas in the causal condition t2 was contingent on the action. In a final test phase participants were instructed to synchronise a key press with t1 (as before) and also synchronise a second key press with t2. To assess intentional binding they measured the time of these synchronised key

presses relative to the onset of t1 and t2. Using this novel behavioural measure they found that intentional binding was increased when participants were trained in the causal condition than when trained in the non-causal condition. This emphasises the importance of contingency for intentional binding: binding was only found when the subjects learnt that the tone was *contingent* on the key press. According to the authors, these data show that *causality* is the principle factor driving intentional binding, downplaying the role of voluntary action. However, given the fact that intentionality was not manipulated in this study, this claim is rather hard to justify.

Contrary to Beuhner and Humphreys we would suggest that the data reviewed so far imply that causality and intentionality are both necessary for binding, but neither factor is itself sufficient. This is supported by Cravo, Claessens, and Baldo's (2009) study in which they assessed the contribution of both voluntary action and causality to intentional binding. A Michottean launching paradigm was used. In a *collision* condition a first disk travelling from left to right collided with a second disk, launching it to the right. In a *non-collision* condition a first disk moved from the centre of the screen to the left and then a second moved from the centre of the screen to the right. In each condition the movement of the first disk was either controlled by the computer (passive) or by the participant (active). Cravo and colleagues measured participants' causal ratings and also intentional binding (using the interval estimation procedure). They found that causal ratings were higher in the collision condition for both active and passive movements. However, intentional binding was only observed in the collision condition and only when the movement was active. This study shows that neither causality nor intentionality are sufficient for intentional binding, rather *both* conditions seem to play a role. This finding is echoed by a more recent study by Cravo, Claessens, and Baldo (2011). They used a novel intentional binding measure based on perceived synchrony of an auditory consequence of action (a tone) and a visual reference (a flash). Intentional binding in this paradigm is revealed when participants make 'synchrony' judgements for visual references occurring *prior* to the auditory consequence of movement (indicating that they perceive the later auditory tone to be synchronous with the earlier visual flash). They investigated the contribution of voluntary action, predictability, and contiguity to intentional binding, and found that intentional binding was strongest when subjects made voluntary movements and the consequences of those movements were predictable and contiguous. The effects of contiguity and predictability were abolished in the absence of

voluntary movement. Again, this emphasises that both causality and intentionality are necessary conditions for binding (but neither are sufficient).

The data reviewed in this section emphasise the importance of learning in intentional binding. That is, acquired knowledge of action-outcome relations impacts on intentional binding. The importance of learning is also recognised in certain models of the sense of agency. For example, according to the so-called ‘comparator model’ (described in section 2.), the sense of agency is informed by sensorimotor predictions generated as part of the normal system of motor control (Blakemore, Wolpert, & Frith, 2002; Frith, Blakemore, & Wolpert, 2000). Given that accurate predictions are the hallmark of successful learning, this model, at least implicitly, recognises the contribution of learning.

Despite the apparent importance of learning in the sense of agency, few studies have considered the nature of this learning. This is something Moore, Dickinson, & Fletcher (2011) wished to address, focussing on the possible role of associative learning mechanisms. Central to theories of associative learning is the role of surprise or prediction error (Mackintosh, 1975; Pearce & Hall, 1980; Rescorla & Wagner, 1972). According to these theories, learning only occurs when the outcome is unexpected – if the outcome is already fully predicted no new learning should occur. Moore, Dickinson and Fletcher used a modified ‘blocking’ procedure in which action-outcome associations were trained under different conditions of surprise. In a pre-training phase, participants freely made left or right key presses. These key presses each caused a specific flash of colour on the computer screen. Having acquired these action-colour associations participants then completed a training phase. In this phase, left and right key presses each caused a compound effect consisting of a specific flash of colour and a specific tone. For one action the resultant flash of colour was the same as that learned during pre-training - this represented a condition of *lower* surprise (only the tone was novel). For the other action the flash of colour was novel - this represented a condition of *augmented* surprise (both the tone and flash of colour were novel). In a test phase, participants made left and right key presses and only the tones from the previous training session were presented. Intentional binding was measured for these action-tone associations. It was found that intentional binding was increased when the action-tone association had been trained under the condition of augmented surprise. This finding is consistent with theories of associative learning. They also showed that the magnitude of this effect was negatively correlated with participants’ scores on schizotypy scales: the smaller

the effect of surprise on binding the higher the participant tended to scores on the schizotypy scales. This is an intriguing finding given the putative link between patterns of associative learning and dopaminergic dysfunction in patients with schizophrenia (discussed in section 5.).

Results from studies like the one above suggest that causality is an important factor in intentional binding. However, contrary to some suggestions, these data do not show that causality is sufficient for intentional binding. Instead, it appears that intentionality *and* causality need to be present for the occurrence of intentional binding. We would conclude, therefore, that motoric contributions to intentional binding (and the sense of agency in general) are embedded in a wider cognitive system that takes into account the causal structure of the environment. The study by Moore, Dickinson and Fletcher suggests that the acquisition of this causal knowledge may proceed (at least in part) by way of relatively simple associative learning mechanisms.

4. Top-down effects on binding

The previous sections have highlighted the roles of predictive and retrospective processes in producing intentional binding and possibly the sense of agency. They have also considered the relationship between intentionality and causality in intentional binding. In particular some of these studies highlighted the role of learned associations that result in outcome expectations (i.e., predictions). Expectations can be considered as ‘top-down’ knowledge, gained through previous experience that guides experience of the self and other aspects of the environment. To address the more general role of top-down processes, several studies have specifically sought to investigate the influence of beliefs on intentional binding.

4.1. The influence of causal beliefs on binding

In one study, Desantis, Roussel and Waszak (2011) investigated the effects of causal belief about context on intentional binding. In their experiment they induced participants to believe that they were responsible for producing a tone with their action, or that their partner was responsible, or that the cause of the tone was ambiguous. The authors showed that a prior causal belief that “I” will be responsible for the tone increased the degree of binding shown compared to causal belief that the other person would be responsible for producing the tone.

In actual fact, the participant's action always caused the tone, only the belief was manipulated by way of a pre-movement cue indicating that they would be responsible, or that the other person would be responsible, or that either person could be responsible. As alluded to earlier, intentional binding seems to be linked to mechanisms involved with predicting the sensory consequences of self-produced actions. Given this assumption, the results of Desantis et al suggest that a *causal belief* can directly affect the relatively low-level sensori-motor processes that are thought to (at least partially) mediate binding, or can at least affect the output of a comparator that assesses actual sensory feedback in the light of predicted sensory feedback. This study builds upon what is known about intentional binding and its relationship to the sense of agency by confirming that there are at least three sources of information that contribute to binding; 1. High-level beliefs about the context, 2. The operation of internal models involved in prediction of action outcomes, 3. Information made available after action onset (post-dictive information). What remains to be investigated is exactly what types of high level belief and post-dictive information modulate binding and more importantly, exactly how binding emerges from the three sources of information outlined. Across the studies reviewed so far, the emerging picture points to a process that integrates available sources of information to compute agency and register motor experience.

4.2. Intentional binding and attention

In a natural extension of the original binding study Haggard and Cole (2007) investigated the role of attention in intentional binding and whether binding occurs for intervals between intentions and action effects, or whether the phenomenon is limited to intervals between actions and their effects. Recall that one view holds that (causal) intentions are real precursors to actions and that they arise prior to movement out of the neural processes involved in motor preparation (predictive account, also known as the pre-constructionist account). On the other hand, reconstructive accounts such as that proposed by Wegner and colleagues (Wegner, 2002; Wegner & Wheatley, 1999), suggest that the experience of causation (of actions and effects) is generated post-hoc, on the basis of a particular sequence of thoughts, actions and effects, within the constraints of priority, consistency and exclusivity (see earlier discussion). Specifically, the re-constructionist view predicts that post-movement effects should influence the conscious experience of intentions. To examine this, Haggard & Cole used the Libet method and asked participants to judge the onset of the intention, the action, or the sensory effect (a tone). In addition, to assess the effect of focused attention on

binding, in some conditions, participants were not told which event to judge until the end of the trial. This clearly prevents the participant from strategically attending to one of the three possible critical events (intention, action, effect). Finally, to investigate the potential role of reafferent signals, binding was also assessed in a de-afferented patient in movements of an effector with no functional re-afferent projections (finger) and an effector with intact re-afference (jaw). Hence, there were several questions posed in this study and several important results emerged.

First, temporal judgments of intentions were not affected by action outcomes, whereas as with previous demonstrations, the judgment of actions was affected by the presence of a sensory effect. Second, intentional binding seems not to depend on focussed attention to a particular event as, when participants were prevented from focussing attention in this way, they showed a very strong binding effect in which percepts of the three events converged on an almost single point in time. Lastly, the de-afferented patient demonstrated binding suggesting that efferent signals, but not re-afferent signals were sufficient to trigger intentional binding in this case.

4.3. Action inhibition and binding

In another study of top-down modulation of binding, Haggard, Poonian and Walsh (2009) applied the intentional binding measure (operationalized in this experiment as a perceived *shift of the tone* back toward the action) to scenarios in which an intentional action was followed by a tone, and to scenarios in which an intentional action was endogenously cancelled but nevertheless a tone either linked to the originally intended action, or not linked to the originally intended action, occurred. To be clear, voluntary inhibition of an action is a top-down process in which a “high level” decision has to be made to withhold the planned action. The main result from the study was a perceived shift of the tone back toward the action for intentional actions that were followed by tones (i.e., a replication of classic intentional binding), but a trend toward a reversal of this effect (i.e., a shift of the perceived time of the tone away from the action) for trials in which an intentional action was inhibited, but nevertheless a tone was delivered. The authors suggested that endogenous inhibition of action results in an updated prediction that the effect should not occur.

5. Intentional binding in patients

Disordered experiences of agency are characteristic of certain psychiatric and neurological disorders. Given the putative link between intentional binding and the sense of agency, intentional binding offers a useful paradigm for exploring the aberrant cognitive and neural processes underpinning these unusual experiences.

In the first intentional binding patient study, Haggard, Martin, Taylor-Clarke, Jeannerod and Franck (2003) compared the magnitude of intentional binding in patients with schizophrenia to a healthy control group. They found that intentional binding was significantly stronger in patients compared to controls. This finding was surprising because schizophrenia patients suffering from unusual experiences of agency typically report feeling *less* in control of their actions. This would have been expected to manifest itself in the form of significantly weaker intentional binding compared with controls. Although unexpected, this finding has been replicated in subsequent patient studies (Voss et al., 2010). This hyper-binding effect has also been induced in healthy controls following infusion of the drug ketamine (Moore, Turner, Corlett, et al., 2011). Ketamine is an anaesthetic agent which acts as an NDMA-receptor antagonist. At sub-anaesthetic levels it has been shown to produce a state that resembles schizophrenia in several key ways (Corlett, Honey, & Fletcher, 2007; Corlett, Honey, Krystal, & Fletcher, 2010). Ketamine is therefore a useful drug model of schizophrenic illness. Moore, Turner, Corlett, et al. (2011) showed that intentional binding was significantly stronger in participants on ketamine compared with their own performance on a placebo. This replicates the hyper-binding found in patients with schizophrenia. Although hyper-binding is unexpected given the typical self-reports from patients with schizophrenia, it appears to be a robust finding. It is also consistent with other data from patients with schizophrenia based on experiments using different agency paradigms: these also show a tendency for patients to hyper-associate their actions and outcomes and to over-attribute the consequences of their movements to themselves (Daprati et al., 1997; Franck, Farrer, et al., 2001). These findings highlight the complex relationship between putative measures of the sense of agency, such as intentional binding, and higher-level experiences of agency.

Although changes in the magnitude of intentional binding are informative, this paradigm also allows us to investigate specific neurocognitive impairments that may contribute to the aberrant sense of agency in schizophrenia. Using Moore and Haggard's (2008) probability paradigm described above, Voss et al (2010) found that the predictive contribution to

intentional binding was absent in patients with schizophrenia. They also found that the inferential contribution was stronger in patients. Intriguingly the magnitude of the predictive deficit in patients correlated with the severity of certain positive symptoms. The deficit in prediction observed on this task is consistent with Frith's model of delusions, the so-called 'comparator model'. This model emphasises the role of aberrant sensorimotor prediction in the genesis of delusions. Closer inspection of Voss et al's data sheds further light on the nature of this predictive deficit. Rather than it being due to a total absence of prediction on 'action only' trials in both outcome probability conditions, these data show that it was due to strong prediction on 'action only' trials in *both* outcome probability conditions. Crucially the strength of prediction was not increased in the 75% condition. What this suggests is that patients were generating imprecise predictions rather than no predictions. That is, the strength of outcome prediction was not related to the objective likelihood of the outcome and instead, predictions were equally strong in the high and low outcome probability conditions. Imprecise prediction has also been observed in a separate study using a different agency paradigm (Synofzik, Thier, Leube, Schlotterbeck, & Lindner, 2010). Imprecise sensorimotor prediction may explain the stronger retrospective contribution to intentional binding in patients: given unreliable internal sensorimotor predictions, the patients' experience of agency may be strongly influenced by the more reliable external cues to agency. This can be captured by the optimal cue integration model introduced by Moore, Wegner and Haggard (2009), where the influence of cues is determined by their reliability.

Intriguingly, the pattern of predictive deficits in intentional binding seems to change during the progression of schizophrenic illness. Using the aforementioned probability paradigm, (Hauser et al., 2011) explored possible deficits in prediction and inference in a group of patients in the prodromal phase of psychotic illness. These were patients who presented with psychotic-like symptoms, but who did not yet meet full diagnostic criteria. Unlike patients with full-blown schizophrenic illness, these prodromal patients showed augmented predictive contributions to intentional binding. The authors suggested that these findings, hyper-prediction in the psychotic prodrome and hypo-prediction in established schizophrenic illness, were consistent with recent models of delusion formation emphasising the role of aberrant error signalling mediated by glutamatergic and dopaminergic pathways in the brain (Corlett, Honey & Fletcher, 2007; Corlett et al., 2010; Fletcher & Frith, 2009). In the early stages of the disease, augmented error signalling, mediated by glutamatergic pathways, could strengthen action-outcome associations leading to stronger predictions. Subsequent

dysregulation of dopaminergic neurotransmission in the later stages of the disease could corrupt error signalling by introducing noise to that signal. This would lead to the formation of strong and *indiscriminate* action-outcome associations leading to strong and indiscriminate (or noisy) predictions.

The importance of dopaminergic neurotransmission for intentional binding was established in another patient study, this time in patients with Parkinson's disease.. Parkinson's disease is particularly interesting in the context of agency research because disturbances in willed behaviour are both a consequence of the disease and a common side-effect of dopaminergic drug treatment used to treat the symptoms of the disease. Moore et al. (2010) used a drug withdrawal paradigm to investigate intentional binding in patients both off and on their dopaminergic drug therapy. The magnitude of intentional binding in Parkinson's patients off medication was statistically identical to a group of aged-matched controls. However, when patients were on medication intentional binding was significantly stronger than in controls and also the patients' own performance off medication. This suggests that dopaminergic drug therapy exaggerated the sense of agency in these patients. This is a fascinating result as impulse control disorders are a relatively common side-effect of medication. The authors suggested that the exaggerated sense of agency may be an important factor in the development of impulse control disorders – if a patient feels their actions are highly effective they may be less likely to inhibit them in future. The modulation of intentional binding by dopaminergic medication also shows that dopaminergic pathways are relevant to intentional binding. Given the role of dopaminergic neurotransmission in associative learning, this finding is consistent with previous data described above showing that intentional binding is sensitive to a modified blocking procedure (Moore, Dickinson & Fletcher, 2011).

The studies reviewed in this section show that intentional binding is a useful paradigm for exploring aberrant experiences of agency in certain psychiatric and neurological disorders. Most of the work carried out so far has focussed on schizophrenic illness. This work has shown aberrant performance on intentional binding tasks in patients with schizophrenia and those in the psychotic prodrome. The ketamine study by Moore et al (2011) also suggests that this drug may be a useful model of aberrant experiences of agency associated with the disease, particularly its early stages. Finally, the work on Parkinson's disease patients by Moore et al (2010) demonstrates the potential utility of exploring intentional binding in other disorders where there are deficits in the performance or experience of willed action.

6. Brain basis of intentional binding

Few studies have explicitly considered the neural basis of intentional binding. The studies described previously detailing intentional binding in patients have implicated dopaminergic (Moore et al., 2010) and glutamatergic pathways (Moore et al., 2011). The role of these pathways is likely to be in regulating learning processes involved in intentional binding.

In terms of specific brain structures supporting intentional binding, the only study that has directly addressed this is Moore, Ruge, Wenke, Rothwell and Haggard (2010). They used continuous theta burst stimulation to explore the brain basis of intentional binding. Continuous theta burst stimulation is a repetitive transcranial magnetic stimulation protocol that is known to suppress neural activity in the stimulated region for up to 60 minutes (Huang, Edwards, Rounis, Bhatia, & Rothwell, 2005). Moore and colleagues investigated the contribution of two specific target sites: the pre-supplementary motor area (pre-SMA) and primary motor cortex (M1). The pre-SMA is involved in higher-order cognitive aspects of self-generated action (Picard & Strick, 2001) and with the conscious experience of intending to act (Fried et al., 1991). In this sense it is likely to support predictive contributions to intentional binding. On the other hand, M1 processes signals that are involved in actual motor execution, signals that the authors suggest are required to support inferences of agency. In this sense disrupting the processing of these signals would be expected to impact on inferential contributions to intentional binding. The effect of stimulation of these target regions was compared with stimulation of a control site, sensory leg area. It was found that only stimulation of pre-SMA led to a significant reduction in intentional binding. Stimulation of M1 marginally reduced intentional binding, but this effect was not significant. The authors therefore concluded that pre-SMA is likely to play a key role in intentional binding. Because continuous theta burst stimulation only significantly reduced the binding of the outcome to the action, the authors also suggest that the role of pre-SMA may be to pre-emptively link intentions to the sensory consequences of action. Although this study is informative, more work needs to be done in order to establish the neural basis of intentional binding. Furthermore, it is possible that, rather than MI or pre-SMA subtending a specific component of intentional binding (predictive or inferential), these areas form part of a wider network of structures that support intentional binding.

7. How does voluntary action modulate time perception?

It is clear from the studies reviewed so far that intentional binding is a robust phenomenon surrounding voluntary action. Whilst patient work has revealed intriguing patterns of intentional binding in schizophrenics and neural intervention approaches have highlighted the role of regions involved in motor preparation in intentional binding, the fundamental question of why exactly voluntary action modulates time perception remains. In one attempt to address this question, Wenke and Haggard (2009) performed an elegant study in which they proposed two broad classes of possible explanations for intentional binding.

First, the effect could be due to changes in the rate of an internal clock, such that intentional binding is simply a result of a slowing of internal clock rate. Within this broad class of explanation, there are further possibilities, one involving a constant change of clock rate over the whole action-effect interval, and the other involving a transient change in clock rate after the voluntary action. Second, the effect could be due to a recalibration of the perceived onset of sensory events such that, for example, tones are pre-dated to maintain perceptual constancy, causal relations, or even agency. These two broad classes of potential mechanisms both predict an intentional binding effect for voluntary compared to passive actions, but make different predictions regarding the temporal discrimination of two sensory events occurring within the interval (no difference in temporal discrimination between active and passive movements for recalibration mechanism; change in temporal discrimination between active and passive movements for clock rate mechanism). The authors exploited this by including a temporal discrimination of shocks task either early in the action-effect interval or late in the interval.

By examining patterns of intentional binding and temporal discrimination of shocks, the authors demonstrated that a recalibration of the time of sensory events does not explain intentional binding, and neither does a constant slowing of clock rate during the entire action-effect interval. Rather, they found support for the idea that clock rate slows transiently and very briefly after a voluntary action and then speeds up again to compensate. To be clear, this is borne out by the fact that they found increased temporal discrimination thresholds for active movements compared to passive movements only when discrimination stimuli were delivered early after voluntary action and not when the stimuli were delivered late in the action effect interval. These results suggest that intentional binding may be due to transient

slowing of an internal clock that then accelerates to return to its initial rate, and is not the result of a re-calibration of the timing of sensory events themselves.

Despite the elegance of this experimental approach, it must be borne in mind that changes of internal clock rate represent just one mechanism by which voluntary action *might* modulate time perception. Hence, this initial result should be interpreted cautiously. For example, Moore & Haggard (2008) demonstrated that intentional binding appears to occur due to recalibration in the absence of sensorimotor prediction. Furthermore, previous work has shown that the system is capable of recalibration in a very similar task environment. Specifically, Stetson et al (2006) showed that after adaptation to a fixed action-effect delay, effects that were delivered much earlier than the adapted latency, but still after the action, were perceived to occur prior to the action. This illusory perception provides evidence that motor-sensory recalibration is indeed possible. Bearing these results in mind, Wenke & Haggard's initial attempt at identifying explanatory mechanisms underlying intentional binding is a welcome first step, but further experiments are warranted to fully determine whether or not a recalibration process contributes to intentional binding. The key point that emerges is that intentional binding can be influenced (and presumably mediated) by a number of factors, and that the precise nature of intentional binding in any given situation is probably a function of the specific task demands and the prevailing environmental context.

8. Binding and explicit agency

As mentioned earlier, intentional binding may represent an implicit measure of sense of agency. However, several authors have highlighted the multi-level nature of the sense of agency by referring to explicit (i.e., reflective) agency and low level (i.e., pre-reflective) agency (David, Newen, & Vogeley, 2008; Synofzik, Vosgerau, & Newen, 2008; Obhi & Hall, 2011a;b). At present though, it is unclear whether intentional binding and explicit agency are mediated by common or separable mechanisms, or whether similar brain areas are involved in the underlying computations. Recently however, several authors have started to investigate the potential relationship between these forms of agency.

For example, in two experiments, Ebert and Wegner (2010) investigated the relationship between intentional binding and explicit self-reported feelings of authorship. Citing the failure of previous studies to examine *both* these forms of agency indicators, the authors

developed a relatively naturalistic paradigm that involved pulling and pushing a joystick in order to move on-screen objects toward or away from themselves. In their first experiment, participants saw pictures of everyday objects such as an apple, and then pushed or pulled the joystick. On some ‘pull’ trials the apple appeared to move towards the participant, and hence the motion of the object was congruent with their action, whereas in others there was an inconsistency in which the object moved away from the person when they pulled the joystick and toward the person when they pushed the joystick. In order to examine the effect of temporal delay on binding and perceived authorship, the authors also employed three different actual delays (100, 400 and 700 ms). Intentional binding was measured using the interval estimation method. Explicit agency was measured by asking participants to report the extent to which they felt their movement had caused the movement of the object, on a 7-point likert scale.

The results showed that estimated intervals were briefer for consistent compared to inconsistent trials and that perceived authorship was greater for consistent versus inconsistent trials. While the binding measure and the perceived authorship measure were both taken on every trial in experiment 1, these measures were taken in separate blocks in experiment 2, in order to reduce the likelihood of one measure ‘contaminating’ the other. Overall, the results from both experiments showed that action-effect consistency affects binding and perceived authorship. However, and importantly, the results also showed that perceived authorship effects were relatively strong at longer action-effect delays compared to binding measures, which were stronger for shorter delays. Interestingly, the majority of participants were also unaware that their interval estimates were different for consistent versus inconsistent action-effect pairings. On the basis of their results, the authors claimed that binding occurs in situations where perceived authorship is experienced and that these two forms of authorship indicator are mediated by at least partially dissociable mechanisms that operate over different time courses.

9. Intentional binding in social contexts

The majority of intentional binding experiments have been conducted in settings where individuals act alone. In this respect they have perhaps fallen prey to what Richardson has called the “experimental quarantine” approach (Richardson et al., 2008). In many everyday situations, individuals do not act in isolation but rather act in social contexts, in many cases

even making “joint actions” with others. Just two examples are factory workers working together to assemble products and members of special operations teams acting in concert to effectively neutralise a threat. Very little is known about the phenomenology of agency and the occurrence of intentional binding in joint action situations. A key question concerns how the sense of agency changes in social settings when more than one agent is involved in bringing about an environmental change. Although still an understudied issue, some authors have begun to investigate the sense of agency and intentional binding in actions occurring in social contexts.

In an attempt to further understand the determinants of intentional binding, Engbert, Wohlschläger, Thomas and Haggard (2007) contrasted actions that caused somatic effects to the actor with those that caused somatic effects to another agent. The authors were specifically interested in whether binding is more dependent on the production of actions or experience of the sensory consequences of those actions. To shed light on this issue they employed a 2x2 design (action production: self, other; somatic effect: self, other). In their first experiment they showed that there was more binding when participants moved than when the experimenter moved, and there was no influence of who experienced the effect on binding. That is, binding was similar regardless of whether the participant’s action produced a passive movement of their own finger or of the experimenter’s finger. In their second experiment, the authors included a baseline condition comprising movement of a rubber hand and a non-somatic effect applied to the rubber hand. Again, the authors found that binding was stronger for self-produced movements and was independent of where the somatic effect occurred. That is, for self produced movements, somatic effects delivered to the participant’s hand, the experimenter’s hand or the rubber hand produced similar levels of binding. In experiment 3, the authors ruled out the possibility that the difference in binding between self and other movement conditions was due to temporal predictability of the action. Lastly, in experiment 3 the authors also found that intentional binding was greatest when the effect was applied to the rubber hand, suggesting a special significance of actions that cause effects in the external world, as opposed to somatic effects to the body, or the body of another.

Engbert, Wohlschläger, and Haggard (2008) examined whether binding effects were similar or different for intentional actions, passive actions and observed actions of another individual, and whether binding is affected by the modality in which a sensory effect occurs. In their first experiment they assessed intentional binding when participants made active or passive

actions that were followed by auditory tones, or watched another agent making active or passive actions that were followed by auditory tones. In observed active movements of the other agent, the same motor that was used in the passive action conditions initiated the action, but participants were not aware of this. Results clearly showed that participants judged the interval between the movement and the consequent tone to be shorter in trials involving their own intentional actions compared to passive movements of their own hand. This was not the case for observed action conditions. In particular, intervals involving actions of the other agent were generally longer than those from conditions involving movement of the participants' body, and there was no difference in the intervals between active and passive movement conditions of the other agent. The authors also demonstrated this pattern of intentional binding effects for active actions that were followed by somatic (passive actions of their own body) or visual effects, hence confirming that the binding effect is generalisable across different sensory modalities, and seems to depend on intentional action.

The two studies outlined above pertain to situations involving two individuals but do not actually address intentional binding in situations where two individuals act together to produce an effect. In the first study to examine intentional binding in "joint action" contexts Strother, House and Obhi (2010) used the Libet method and set up a task in which two individuals sat side by side and performed key presses on either end of a space bar. Each participant rested their right index finger on force sensitive resistor at one end of a spacebar and they were requested to press the spacebar at a time of their own choosing. An additional instruction was that if the other person pressed the bar, their partner should simply let their finger move down with the key but exert no force themselves. Explicit agency judgments and intentional binding was assessed on each trial. Interestingly, intentional binding was similar and indistinguishable for the initiator and the passive mover, despite the fact that only the initiator claimed any explicit agency for the action-effect. The authors pointed out that intentional binding occurred for the passive mover despite the fact that no change in the output of the force sensitive resistor was seen. This shows that intentional binding can occur in the absence of privately held efferent information and also that intentional binding and the explicit sense of agency can be dissociated (c.f., Ebert & Wegner, 2010). The same results were found in a blocked condition in which one participant was told not to intend to make an action at all and in conditions where the initiator emerged dynamically based on who pressed the spacebar first. What separates this experimental paradigm from others is that it was obtained in a specific scenario involving two individuals acting together. Hence, it may be

that such cases of “joint action” represent a special class of action contexts, in which private efferent information is not the primary trigger for the occurrence of intentional binding.

In other recent experiments, Obhi & Hall (2011a) manipulated the social role of the partners in a similar joint action task so that one was a leader and the other was a follower. In this experiment the person who emerged as a follower was instructed to press the key as fast as possible in response to the leader’s key press. In other conditions the role (leader, follower) was assigned at the beginning of the block and remained constant. Again, intentional binding was found for both individuals, regardless of whether they felt an explicit sense of agency or not. To be clear, even when the follower explicitly claimed no responsibility for the key press and subsequent effect, they demonstrated the same degree of intentional binding as the initiator, who did experience explicit agency.

In a further recent study, Obhi & Hall (2011b) used a similar task to compare the effects of being paired with a human or with a computer. Specifically, in this experiment participants sat side by side with another person (a confederate) and were unable to see them (a curtain separated them). The genuine participant and the confederate were seated in front of laptop computers and were instructed to make silent index finger touches of the touchpad and report the time of the touch (and in some conditions a tone that followed) using the Libet method. After each movement-tone complex a coloured dot was provided as feedback about whose action had actually caused the tone. This feedback was false, as in reality the experimental participant’s action always caused the tone. Explicit reports of agency revealed that participants always aligned with the feedback, but as in previous experiments, the participants showed intentional binding even when they explicitly believed the other person had caused the tone. In a further variation, the participant was told that their partner was a computer that would simulate the actions of a person and that could cause tones in an identical fashion. In this case, the effects of belief about the partner had dramatic effects on intentional binding. Whilst the participant still always aligned with the feedback in their explicit judgments (saying that they were responsible for the tone when the feedback indicated this, and that the computer was responsible for the tone when the feedback indicated that), they never showed intentional binding. Thus, when paired with a computer, even when participants explicitly knew they caused an effect, they did not show intentional binding. This result opens up a new line of investigation into how the sense of agency emerges in human-machine interactions. This is a promising and very topical area for future research.

On the basis of their joint action experiments Obhi et al (2011a;b) suggested that, when paired with another human, a new “we” identity is automatically formed. In this situation both individuals experience pre-reflective agency (i.e., demonstrate intentional binding) for any action produced by either partner. However, at a higher reflective level of processing they still believe that only one person caused the tone. Critically, this automatic formation of a new agentic identity seems not to occur when the other “person” is a machine, and furthermore individual agency experience also seems to break down in this scenario. Further experiments seem warranted to investigate the sense of agency and intentional binding in joint human actions and joint human-machine actions.

10. Conclusions and future directions

As this review has highlighted, since the seminal paper in which intentional binding was reported in 2002, many studies investigating the factors that affect intentional binding have been published, with a robust pattern of effects. Furthermore, interest in intentional binding and the sense of agency more generally appears to be growing which provides hope of deeper insights in the not too distant future. Whilst considerable progress has been made in determining how intentional binding is affected by a host of manipulations, it is fair to say that a comprehensive understanding of the intentional binding effect and particularly its relationship to the sense of agency has not yet been fully achieved. What has been shown is that there is a reliable link between intentional binding and the presence of efferent information (but recall the apparent exception of joint action).

It has also been conclusively demonstrated that other sources of information such as actual sensory feedback, outcome expectations and causal beliefs all modulate the degree of binding, and that the effects of these variables seem to depend on the exact task-operator characteristics. This idea has been formalised in the optimal cue integration theory (Moore, Wegner, & Haggard, 2009), which, in our view, represents the most parsimonious and simple account of how intentional binding - and more broadly, the sense of agency - arises in humans. This theory is outlined in detail in a recent paper by Moore and Fletcher (in press). Briefly, it posits that a range of different agency cues contribute to the sense of agency and that the influence of these cues is dependent on their reliability. This is a *statistically* optimal solution based on the maximum likelihood estimation rule (Ernst & Banks, 2002; see also

Moore & Fletcher, in press, for how this relates to sense of agency). These agency cues include, but are not limited to, efferent motoric information; proprioceptive information; action consequences; and the social context of action (see Wegner & Sparrow, 2004, and Synofzik et al., 2008, for a more extensive discussion of different agency cues). This optimal cue integration approach is able to account for patterns of agency experience from a range of studies, including those that have used the intentional binding paradigm. Importantly, it can explain patterns of agency experience associated with certain neurological and psychiatric diseases. For example, the sense of agency in patients with schizophrenia appears to be characterised by noisy sensorimotor prediction (Voss et al., 2010; Synofzik et al., 2010). Moreover, external sensory cues to agency dominate sense of agency in these patients. Optimal cue integration can explain this: external sensory cues to agency are more reliable than noisy sensorimotor prediction and will therefore dominate sense of agency in patients with schizophrenia.

Out of all the factors that have been linked to intentional binding, it appears that the presence of efferent information is most central to the manifestation of intentional binding as, when efferent information is not present, such as in passive movement conditions, or passive observation of others, intentional binding is reduced or absent. The key role of efferent information and the similarly strong role of efference-based prediction in neurocognitive accounts of agency perhaps provide the greatest theoretical link between intentional binding and the sense of agency. However, as we have highlighted in this paper, there are exceptions to this. It is these exceptions which have, at least in part, motivated the development of cue integration models of the sense of agency. One such exception seems to be joint action contexts where private efferent information is not required for the expression of intentional binding. This result has been discussed with reference to a potential “we” identity that is formed whenever two individuals (even nominally) act together to produce some outcome. A topic for future investigation concerns why it should be that, in some cases, private efferent information is not necessary for intentional binding to occur.

In contrast to the apparently strong roles of efference, sensory feedback, causal beliefs, and intentionality in intentional binding, what is much less clear from the work done so far is exactly what the relationship between intentional binding and explicit sense of agency actually is. On the one hand, there have been demonstrations of intentional binding and the explicit sense of agency co-occurring. On the other hand, there have been reports of

dissociations between intentional binding and the explicit sense of agency. At least two questions naturally arise: first, how does one go about validating a purportedly implicit measure of agency? That is, what are the necessary conditions that must be met in order for us to be confident that intentional binding really does index a pre-reflective sense of agency? Second, what exactly is the relationship between implicit agency (perhaps as indexed by intentional binding) and explicit reflective agency? Any account must be able to explain why the two measures sometimes co-occur and sometimes dissociate. These are hard issues to resolve. The complex relationship between intentional binding and explicit agency is further highlighted by the patient data in which schizophrenic patients show strong intentional binding but (often) weak explicit agency experience. There is an elegant explanation (see section 5.) To account for the stronger intentional binding expression in patients, but again, and critically, reconciling this with patient subjective experience remains a challenge.

Another challenge for future research is to shed light on the specific neural mechanisms underlying intentional binding. No research has yet directly addressed this issue, although established theories concerning the mechanisms of binding more generally might be informative. These include ‘oscillation theories’ which, at the neuronal level, suggest that integration emerges from temporal synchronization of neural activities using neural oscillations as a binding mechanism (Engel & Singer, 2001). Future research should explore the relevance of these ideas to intentional binding.

Despite these issues it should be clear from the studies reviewed in this paper that, since it was first introduced almost a decade ago, intentional binding has stimulated an impressive array of research and has moved us towards a better understanding of human agency. The challenge for the next decade of research is to further our understanding of the causal mechanisms of binding, identify the brain networks involved in binding, and determine how binding, low-level (pre-reflective) agency and explicit (reflective) agency are related. These are difficult questions, but we are confident that, with the foundation that has been built over the last ten years, the next ten years of research on intentional binding and, more generally the sense of agency, will yield even deeper insights.

Acknowledgements

SSO holds funding from the Natural Sciences & Engineering Research Council of Canada that enable his work in this area.

References

- Balleine, B. W., & Dickinson, A. (1998). Goal-directed instrumental action: contingency and incentive learning and their cortical substrates. *Neuropharmacology*, *37*(4-5), 407–419.
- Blakemore, S., Wolpert, D., & Frith, C. (2002). Abnormalities in the awareness of action. *Trends in Cognitive Sciences*, *6*(6), 237-242.
- Blakemore, S.-J., Frith, Chris D., & Wolpert, Daniel M. (1999). Spatio-Temporal Prediction Modulates the Perception of Self-Produced Stimuli. *Journal of Cognitive Neuroscience*, *11*(5), 551-559. doi:10.1162/089892999563607
- Buehner, M. J., & Humphreys, G. R. (2009). Causal binding of actions to their effects. *Psychological Science*, *20*(10), 1221.
- Corlett, P R, Honey, G D, & Fletcher, P C. (2007a). From prediction error to psychosis: ketamine as a pharmacological model of delusions. *Journal of Psychopharmacology (Oxford, England)*, *21*(3), 238-252. doi:10.1177/0269881107077716
- Corlett, P R, Honey, G D, & Fletcher, P C. (2007b). From prediction error to psychosis: ketamine as a pharmacological model of delusions. *Journal of Psychopharmacology (Oxford, England)*, *21*(3), 238-52. doi:10.1177/0269881107077716
- Corlett, Philip R, Honey, Garry D, Krystal, J. H., & Fletcher, Paul C. (2010). Glutamatergic Model Psychoses: Prediction Error, Learning, and Inference. *Neuropsychopharmacology*. Retrieved from <http://dx.doi.org/10.1038/npp.2010.163>
- Cravo, A. M, Claessens, P. M.E, & Baldo, M. V.C. (2009). Voluntary action and causality in temporal binding. *Experimental brain research*, *199*(1), 95–99.
- Cravo, A. M., Claessens, P. M. E., & Baldo, M. V. C. (2011). The relation between action, predictability and temporal contiguity in temporal binding. *Acta Psychologica*, *136*(1), 157-166. doi:16/j.actpsy.2010.11.005
- Daprati, E., Franck, N., Georgieff, N., Proust, J., Pacherie, E., Dalery, J., & Jeannerod, M. (1997). Looking for the agent: an investigation into consciousness of action and self-consciousness in schizophrenic patients. *Cognition*, *65*(1), 71-86.
- David, N., Newen, A., & Vogeley, K. (2008). The “sense of agency” and its underlying cognitive and neural mechanisms. *Consciousness and Cognition*, *17*(2), 523-534. doi:10.1016/j.concog.2008.03.004
- Desantis, A., Roussel, C., & Waszak, F. (2011). On the influence of causal beliefs on the feeling of agency. *Consciousness and Cognition*.

- Ebert, J. P., & Wegner, D. M. (2010). Time warp: Authorship shapes the perceived timing of actions and events. *Consciousness and cognition*, *19*(1), 481–489.
- Engbert, K., Wohlschläger, A., & Haggard, P. (2008). Who is causing what? The sense of agency is relational and efferent-triggered. *Cognition*, *107*(2), 693–704.
- Engbert, K., Wohlschläger, A., Thomas, R., & Haggard, P. (2007). Agency, subjective time, and other minds. *Journal of Experimental Psychology: Human Perception and Performance*, *33*(6), 1261.
- Engbert, Kai, & Wohlschläger, A. (2007). Intentions and expectations in temporal binding. *Consciousness and Cognition*, *16*(2), 255-264. doi:10.1016/j.concog.2006.09.010
- Engel, A. K., & Singer, W. (2001). Temporal binding and the neural correlates of sensory awareness. *Trends in Cognitive Sciences*, *5*(1), 16-25.
- Ernst, M. O., & Banks, M. S. (2002). Humans integrate visual and haptic information in a statistically optimal fashion. *Nature*, *415*(6870), 429–433.
- Franck, N., Farrer, C., Georgieff, N., Marie-Cardine, M., Daléry, J., d' Amato, T., & Jeannerod, M. (2001). Defective recognition of one's own actions in patients with schizophrenia. *The American Journal of Psychiatry*, *158*(3), 454-9.
- Fried, I., Katz, A., McCarthy, G., Sass, K. J., Williamson, P., Spencer, S. S., & Spencer, D. D. (1991). Functional organization of human supplementary motor cortex studied by electrical stimulation. *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience*, *11*(11), 3656-66.
- Frith, C D, Blakemore, S. J., & Wolpert, D M. (2000). Abnormalities in the awareness and control of action. *Philosophical Transactions of the Royal Society of London. Series B, Biological Sciences*, *355*(1404), 1771-88. doi:10.1098/rstb.2000.0734
- Haggard, P., & Clark, S. (2003). Intentional action: conscious experience and neural prediction. *Consciousness and Cognition*, *12*(4), 695–707.
- Haggard, P., & Cole, J. (2007). Intention, attention and the temporal experience of action. *Consciousness and cognition*, *16*(2), 211–220.
- Haggard, P., Aschersleben, G., Gehrke, J., & Prinz, W. (2002). Action, binding, and awareness. *Common mechanisms in perception and action: Attention and Performance XIX*, 266.
- Haggard, P., Clark, S., & Kalogeras, J. (2002). Voluntary action and conscious awareness. *Nature Neuroscience*, *5*(4), 382–385.
- Haggard, P., Martin, F., Taylor-Clarke, M., Jeannerod, M., & Franck, N. (2003). Awareness of action in schizophrenia. *Neuroreport*, *14*(7), 1081.

- Haggard, P., Poonian, S., & Walsh, E. (2009). Representing the consequences of intentionally inhibited actions. *Brain research, 1286*, 106–113.
- Haggard, P, Clark, S., & Kalogeras, J. (2002). Voluntary action and conscious awareness. *Nature Neuroscience, 5*(4), 382-5. doi:10.1038/nn827
- Hauser, M., Moore, J. W, de Millas, W., Gallinat, J., Heinz, A., Haggard, P., & Voss, M. (2011). Sense of agency is altered in patients with a putative psychotic prodrome. *Schizophrenia research, 126*(1-3), 20–27.
- Huang, Y.-Z., Edwards, M. J., Rounis, E., Bhatia, Kailash P, & Rothwell, J. C. (2005). Theta burst stimulation of the human motor cortex. *Neuron, 45*(2), 201-6. doi:10.1016/j.neuron.2004.12.033
- Hume, D. (1888). A treatise of human nature. In L. A. Selby-Bigge (Ed.), *Hume's Treatise of human nature*. Oxford: Oxford University Press, Clarendon Press. (Original work published 1739).
- Humphreys, G. R., & Buehner, M. J. (2009). Magnitude estimation reveals temporal binding at super-second intervals. *Journal of Experimental Psychology: Human Perception and Performance, 35*(5), 1542.
- Libet, B., Gleason, C. A., Wright, E. W., & Pearl, D. K. (1983). Time of conscious intention to act in relation to onset of cerebral activity (readiness-potential). The unconscious initiation of a freely voluntary act. *Brain: A Journal of Neurology, 106* (Pt 3), 623-642.
- Mackintosh, N. J. (1975). A theory of attention: Variations in the associability of stimuli with reinforcement. *Psychological Review, 82*(4), 276.
- Moore, J. W., & Fletcher, P. C. (in press). Sense of agency in health and disease: A review of cue integration approaches. *Consciousness and Cognition*. doi:10.1016/j.concog.2011.08.010
- Moore, J. W, Lagnado, D., Deal, D. C., & Haggard, P. (2009). Feelings of control: Contingency determines experience of action. *Cognition, 110*(2), 279–283.
- Moore, J. W, Ruge, D., Wenke, D., Rothwell, J., & Haggard, P. (2010). Disrupting the experience of control in the human brain: pre-supplementary motor area contributes to the sense of agency. *Proceedings of the Royal Society B: Biological Sciences, 277*(1693), 2503.
- Moore, J. W, Schneider, S. A., Schwingschuh, P., Moretto, G., Bhatia, K. P, & Haggard, P. (2010). Dopaminergic medication boosts action-effect binding in Parkinson's disease. *Neuropsychologia, 48*(4), 1125–1132.

- Moore, J. W, Turner, D. C., Corlett, P. R, Arana, F. S., Morgan, H. L., Absalom, A. R., Adapa, R., et al. (2011). Ketamine administration in healthy volunteers reproduces aberrant agency experiences associated with schizophrenia. *Cognitive Neuropsychiatry*, 99999(1), 1–18.
- Moore, J. W, Dickinson, A., & Fletcher, P. C. (2011). Sense of agency, associative learning, and schizotypy. *Consciousness and Cognition*, 20(3), 792-800. doi:10.1016/j.concog.2011.01.002
- Moore, J. W, Wegner, D. M, & Haggard, P.. (2009). Modulating the sense of agency with external cues. *Consciousness and Cognition*, 18(4), 1056-1064. doi:10.1016/j.concog.2009.05.004
- Moore, J., & Haggard, P.. (2008). Awareness of action: Inference and prediction. *Consciousness and Cognition*, 17(1), 136-44. doi:10.1016/j.concog.2006.12.004
- Morrone, M. C., Ross, J., & Burr, D. C. (1997). Apparent position of visual targets during real and simulated saccadic eye movements. *Journal of Neuroscience*, 17, 7941–7953.
- Obhi, S. S., & Hall, P. (2011). Sense of agency and intentional binding in joint action. *Experimental Brain Research*, 211, 655-62
- Obhi, S.S., & Hall, P. (2011). Sense of Agency in Joint Action: Influence of Human and Computer Co-actors. *Experimental Brain Research*, 211, 663-70.
- Pearce, J. M., & Hall, G. (1980). A model for Pavlovian learning: variations in the effectiveness of conditioned but not of unconditioned stimuli. *Psychological Review*, 87(6), 532-552.
- Picard, N., & Strick, P. L. (2001). Imaging the premotor areas. *Current Opinion in Neurobiology*, 11(6), 663-672.
- Reisberg, D. *Cognition: Exploring the Science of the Mind*, Norton, New York
- Rescorla, R. A., & Wagner, A. H. (1972). A theory of Pavlovian conditioning: Variations in the effectiveness of reinforcement and nonreinforcement. *Classical conditioning II: Current research and theory*, 64–99.
- Richardson DC, Hoover MA, Ghane A (2008) Joint perception: gaze and the presence of others. In: Love BC, McRae K, Sloutsky VM (eds) *Proceedings of the 30th annual*

- conference of the cognitive science society*. Cognitive Science Society, Austin, pp 309–314
- Stetson, C., Cui, X., Montague, P. R., & Eagleman, D. M. (2006). Motor-sensory recalibration leads to an illusory reversal of action and sensation. *Neuron*, *51*(5), 651-9. doi:10.1016/j.neuron.2006.08.006
- Strother, L., House, K. A., & Obhi, S. S. (2010). Subjective agency and awareness of shared actions. *Consciousness and cognition*, *19*(1), 12–20.
- Synofzik, M., Thier, P., Leube, D. T., Schlotterbeck, P., & Lindner, A. (2010). Misattributions of agency in schizophrenia are based on imprecise predictions about the sensory consequences of one's actions. *Brain*, *133*(1), 262.
- Synofzik, M., Vosgerau, G., & Newen, A. (2008). Beyond the comparator model: A multifactorial two-step account of agency. *Consciousness and Cognition*, *17*(1), 219-239. doi:10.1016/j.concog.2007.03.010
- Tsakiris, M., & Haggard, P. (2003). Awareness of somatic events associated with a voluntary action. *Experimental Brain Research*, *149*(4), 439–446.
- Voss, M., Moore, J., Hauser, M., Gallinat, J., Heinz, A., & Haggard, P. (2010). Altered awareness of action in schizophrenia: a specific deficit in predicting action consequences. *Brain*, *133*(10), 3104.
- Wegner, D M, & Wheatley, T. (1999). Apparent mental causation. Sources of the experience of will. *The American Psychologist*, *54*(7), 480-92.
- Wegner, D. M. (2002). *The Illusion of Conscious Will*. MIT Press.
- Wegner, D. M. (2003). The mind's best trick: how we experience conscious will. *Trends in Cognitive Sciences*, *7*(2), 65-69.
- Wegner, D. M., & Sparrow, B. (2004). Authorship processing. In M. Gazzaniga (Ed.), *The Cognitive Neurosciences*: 3rd Edition. Cambridge, MA: MIT Press, pp. 1201-1209.
- Wenke, D., & Haggard, P. (2009). How voluntary actions modulate time perception. *Experimental brain research*, *196*(3), 311–318.
- Wolpert, D M, & Ghahramani, Z. (2000). Computational principles of movement neuroscience. *Nature Neuroscience*, *3 Suppl*, 1212-7. doi:10.1038/81497
- Yarrow, K., Haggard, P., Heal, R., Brown, P., & Rothwell, J. C. (2001). Illusory perceptions of space and time preserve cross-saccadic perceptual continuity. *Nature*, *414*(6861), 302-305. doi:10.1038/35104551

