



Learning to infer the time of our actions and decisions from their consequences



Helena Matute^{a,*}, Carmelo P. Cubillas^b, Pablo Garaizar^c

^a *Facultad de Psicología y Educación, Universidad de Deusto, Bilbao, Spain*

^b *Facultad de Ciencias de la Salud y de la Educación, Universidad a Distancia de Madrid, Spain*

^c *Facultad de Ingeniería, Universidad de Deusto, Bilbao, Spain*

ARTICLE INFO

Keywords:

Intentional binding
Causal binding
Conscious will
Learning
Conscious decision
Action

ABSTRACT

Research shows that people infer the time of their actions and decisions from their consequences. We asked how people know how much time to subtract from consequences in order to infer their actions and decisions. They could either subtract a fixed, default, time from consequences, or learn from experience how much time to subtract in each situation. In two experiments, participants' actions were followed by a tone, which was presented either immediately or after a delay. In Experiment 1, participants estimated the time of their actions; in Experiment 2, the time of their decisions to act. Both actions and decisions were judged to occur sooner or later as a function of whether consequences were immediate or delayed. Estimations tended to be shifted toward their consequences, but in some cases they were shifted away from them. Most importantly, in all cases participants learned progressively to adjust their estimations with experience.

1. Introduction

A famous experiment by Benjamin Libet and his colleagues (Libet, Gleason, Wright, & Pearl, 1983) used an oscilloscope clock with a dot rotating around a sphere and asked their experimental participants to perform a quick movement of their finger or wrist at any time they felt they wished to. This procedure was repeated for a number of trials, and at the end of some of those trials participants were asked to indicate on the clock the position of the dot at the exact moment of their “conscious awareness of wanting to move”. In addition, electroencephalographic (EEG) activity was recorded in order to assess the exact time at which the readiness potential (an electroencephalographic component that precedes the initiation of voluntary actions) occurred. Contrary to what intuition would suggest, the sequence for self-initiated actions that was recorded did not start with the participants' reporting their conscious will to perform the action. Instead, the readiness potential was recorded first in the temporal chain (around 550 ms before the action), with participants reporting their conscious will as occurring much later in time (around only 200 ms before the initiation of the action). This seems to suggest that voluntary actions are initiated unconsciously at a neuronal level, and that the conscious decision to act occurs much later in time (about 350 ms later). This experiment, along with several others that confirmed, extended, or discussed the initial findings (Banks & Isham, 2009; Guggisberg & Mottaz, 2013; Haggard & Eimer, 1999; Rigoni, Brass, Roger, Vidal, & Sartori, 2013; Rigoni, Brass, & Sartori, 2010; Soon, Brass, Heinze, & Haynes, 2008) have fueled some rather heated debates about consciousness and free will in psychology, law, and related disciplines (Alquist, Ainsworth, & Baumeister, 2013; Baumeister, Masicampo, & DeWall, 2009; Guggisberg & Mottaz, 2013; Moore, 2016; Rigoni, Kühn, Gaudino, Sartori, & Brass, 2012;

* Corresponding author at: Facultad de Psicología, Universidad de Deusto, Avda. Universidades 24, 48007 Bilbao, Spain.
E-mail address: matute@deusto.es (H. Matute).

Vohs & Schooler, 2008). Libet's clock procedure has also been used in many experiments studying not only how people become aware of their own decisions but also of their own actions (Banks & Isham, 2009; Caspar & Cleeremans, 2015; Haggard, Clark, & Kalogeras, 2002; Moore, Lagnado, Deal, & Haggard, 2009; Rigoni et al., 2010; Tobias-Webb et al., 2017). A standardized, open source version of Libet's clock has recently been published (Garaizar, Cubillas, & Matute, 2016), which should facilitate reproducibility of these findings.

In spite of the many experiments that have pursued this line of research, there are still many questions that remain unanswered. Perhaps the most obvious problem is that it is impossible to assess directly the moment at which conscious will actually occurs. In particular, experimental participants do not have direct access to the time of their conscious decisions (Banks & Isham, 2009; Fahle, Stemmler, & Spang, 2011; Guggisberg & Mottaz, 2013; Rigoni et al., 2010). Moreover, we do not even know whether the temporal order in which the events are recorded in those experiments is the temporal order in which they actually occur. It might well be that people are poor at estimating the time of their conscious decision, so that perhaps they simply misjudge this time when asked to report it. Indeed, the only information that Libet et al.'s (1983) experiment (or, to our knowledge, any other experiment) can reveal in relation to the time of conscious will is the subjective report of when participants estimate that it took place.

Therefore, one crucial question is how people estimate the time of their decision. One possibility is that they directly perceive the moment at which they decide to act, whilst another is that they infer retrospectively the time at which they decided to act. Preliminary evidence for this latter possibility exists in the literature. For example, Banks and Isham (2009) added auditory feedback (i.e. a tone) to the participants' action and manipulated the time intervals between action and feedback. If people perceived the time of their decisions directly, then the time at which the tone occurred after the action should be irrelevant. However, their participants reported that their conscious decision to act occurred sooner or later as a function of whether the tone occurred sooner or later after the action. In other words, participants appeared to retrospectively infer, rather than perceive, the time at which they decided to act. Banks and Isham concluded that the generation of the action was largely unconscious, with participants inferring the time of their decision from observable cues, particularly from the apparent time of their own actions (see also Lau, Rogers, & Passingham, 2007; Rigoni et al., 2010).

But assuming that the time of decision is inferred from the estimated time of action leads us to another crucial question: How do people estimate the moment of their actions? Again, they might directly perceive their own actions or they might infer the occurrence of their actions from their consequences. This has been a central question in Experimental Psychology since the nineteenth century. According to the Ideomotor Theory proposed by James (1890), people use the sensorimotor properties of the action in order to estimate when they performed the act. It is easy to extend this logic to assume that people use not only the sensorimotor properties of their action, but also any other consequences of their action, such as a tone — which an experimenter might introduce into the situation — to infer when they have acted. Interestingly, this line of research has shown that when the consequences of the action (e.g. a tone) are delayed, then people estimate that their action occurred later in time (Banks & Isham, 2011; Haering & Kiesel, 2014; Haering & Kiesel, 2015; Wegner, 2002). Moreover, a very similar finding has been consistently reported using the so-called “action-effect binding” paradigm, which shows that actions and their consequences tend to be perceived as closer to each other in time than they actually are (Haggard et al., 2002). According to Haggard and his colleagues, the intentionality of the action is the crucial factor in producing this effect, but it has also been suggested that intentionality is not always necessary (e.g., Isham, Banks, Ekstrom, & Stern, 2011), and that it might be the causal relationship between the action and the outcome, rather than its intentionality, that is critical (Bechlivanidis & Lagnado, 2013; Buehner, 2012; Buehner & Humphreys, 2009). In any case, the finding that the estimated time of action is affected by the time of its consequences and it is shifted toward them is a robust result that has been reported in many different experiments (Buehner, 2012; Cravo, Haddad, Claessens, & Baldo, 2013; Garaizar et al., 2016; Haering & Kiesel, 2015; Haggard et al., 2002; Moore & Haggard, 2008; Nolden, Haering, & Kiesel, 2012).

Thus, it seems that it is not only the case that people cannot directly perceive the time of their decisions, they also seem to have difficulties in estimating the time of their own actions. Indeed, the estimation of the time of decisions and the time of actions appears to vary in the same manner and they both seem to be strongly influenced by the time at which the consequences occur (see also Banks & Isham, 2011). The most basic example of such consequences are sensorimotor consequences (such as visual and haptic cues that signal, for example, that a button has been pressed), but additional consequences such as added auditory stimuli can also be signals that the action (and decision) has occurred.

In order to infer the time of their actions, people appear to discount some time from the time at which the consequences occur, and then, they possibly discount some extra time from the apparent time of action in order to infer the time of decisions. But then the question is how they know how much time they need to subtract in each situation. Do people simply subtract some fixed, default, time, which is identical in all cases, or do they learn from experience how much time they need to subtract in each particular case? In daily life, it seems reasonable to expect that people will constantly need to learn how much time to discount in each particular situation. People may easily learn that a light of a room will turn on immediately after they press the correct switch, but they should as readily learn to predict a long delay for tomatoes to grow after they have planted the seeds. Indeed, it has been shown that humans and other animals have expectations of when consequences should occur after the potential cause has occurred, and these expectations affect their learning and behavior (e.g., Arcediano, Escobar, & Miller, 2003; Buehner & May 2002; Buehner & May 2003; Buehner & May 2004; Matzel, Held, & Miller, 1988; Miller & Barnet, 1993; Savastano & Miller, 1998). Quite possibly, these expectations can be learned through experience and are continuously adjusted for each action-outcome pair and context. By the same reasoning, people will possibly learn to assume that in some cases their action occurred immediately before its consequences whilst in other cases they might learn that their actions (and decisions) took place a long time before they observed the consequences. The present study will test whether people learn to infer the time of their actions and decisions or if, in contrast, they just subtract a fixed time upon observing their consequences.

Although the majority of research in the area of temporal binding implicitly assumes that learning plays an important role, we are aware of very few investigations that explicitly examine such learning. Indeed, most experiments report the dependent variable (e.g., subjective judgments or estimations of when the action occurred) averaged over the total number of trials (Haggard et al., 2002; Moore et al., 2009), and so there is typically a lack of information on what happens on a trial-by-trial basis during the course of learning. Thus, the possibility exists that action-effect binding might take place during the early trials due to pre-experimental biases, after which people might progressively learn to correct their error and improve the accuracy of their estimated time of actions (and perhaps also their decisions). Alternatively, it is possible that estimations do not improve with experience and instead show a fixed binding effect throughout the experimental session. This might occur, for instance, if actions were always, by default, subjectively shifted toward their effects. This is what most reports on temporal binding seem to suggest, given that they typically provide only one (averaged) binding value for each group in each experimental session.

We are aware of only two studies looking at learning in the binding tradition. One of them was reported by Cravo et al. (2013). They used three different intervals (250, 300, 350 ms) between the action and the outcome (i.e. Action condition), and between an external event and the outcome (i.e., No Action condition). They observed, first, that the mean reported interval between the action and the outcome was shorter than between the external event and the outcome, an effect that was evident from the very early trials. According to the authors, this implies that action-outcome binding effects are due to people expecting short intervals between actions and outcomes from the very first trials, suggesting that pre-experimental biases must be important. In addition, they observed no differences between the early and later trials when considering the whole experimental session. Thus, they concluded that previous biases are probably more important than learning in producing the action-outcome binding effect. Importantly, however, they also noted that some learning was taking place within each of the consecutive and independent training blocks of trials in which they had subdivided the training phase. Thus, although their experiment represents a good starting point, there are a number of reasons why we are unable to extract any firm conclusions about learning from this study. First, the fact that learning seemed to take place within each block but no differences were observed between early and later trials throughout the session suggests that learning might have occurred but was eroded each time a new training block started. If this were the case, participants might not have benefited from the entire learning session and might have started anew each time a new block of trials was presented. Thus, it is possible that the learning phase might have just been too short to observe any appreciable differences. In addition, there were two different conditions and three different intervals in each block (with a total of 20 trials for each interval in each block), which might also have made the problem too difficult to be learned in the short time allocated to each block. A potential way to solve these problems might be to make the task easier and to provide a larger number of learning trials, that is, to ensure a more effective learning experience. This is one strategy that we will follow to examine the course of learning in the present research.

The other study we are aware of was reported by Moore et al. (2009). They tested whether contingency learning played a role in the binding effect. Unfortunately, however, they used a complex design which did not allow them to observe the course of learning. As they put it: “we averaged our estimates across several trials because of the high variability of human timing performance. Therefore, we could not measure the time-course of the learning process but we can infer that causal learning occurs based on our contingency effects” (p. 282–283). This study also suggests that reducing the difficulty of the task might possibly be the first step needed to study the learning curve. The present research will take such steps in order to test whether people progressively learn to estimate the time of their actions and decisions with experience.

2. Overview of the experiments

In order to facilitate the observation of a learning curve, we will use only one continuous training session (i.e. blocks of trials will constitute a continuous and incremental training session, rather than independent learning experiences). In addition, we will use only two different conditions during training (i.e. immediate feedback vs delayed feedback). A third condition (i.e. no-feedback) will be added as a separate baseline phase to assess judgments in the absence of feedback, but it will be presented only after the training phase has already been completed, so that baseline assessment could not interfere with the critical learning experience.

2.1. Ethics statement

The computer program informed participants that their involvement was voluntary and anonymous. We did not ask participants for any data that could compromise their privacy, nor did we use cookies or software in order to obtain such data. The stimuli and materials were harmless and emotionally neutral, the goal of the study was transparent, and the task involved no deception. The ethical review board of the University of Deusto examined and approved the procedure used in this research, and the two experiments were conducted in accordance with the approved guidelines.

2.2. Data selection criterion

Data from participants who did not press the key more than 25% of trials on either one of the three feedback conditions described above (i.e., immediate, delayed, and no-feedback) were discarded.

3. Experiment 1

The purpose of this experiment was twofold. First, we tested whether people estimate the time of their actions accurately or




Fig. 1. Screen capture showing the assessment of the dependent variable at the end of a training trial in Experiment 1.

whether their estimate is affected by the time of consequences. We expected that a delay in the consequences of the action should produce a delay in the estimated time of action, thereby reproducing previous reports (e.g., Banks & Isham, 2011; Garaizar et al., 2016; Haggard et al., 2002). Most importantly, this experiment aimed to extend previous findings by assessing whether the estimation of the time of action was subject to learning. That is, we tested whether people will then progressively learn to correct their estimation errors and improve their reported time of action with experience.

3.1. Method

3.1.1. Participants

The sample was comprised of 60 Psychology students from the University of Deusto who volunteered to take part in the experiment in exchange for academic credit. Data from four of the participants were discarded according to the data selection criterion described above.

3.1.2. Apparatus

Participants performed the task on personal computers in a large computer room. They were seated approximately 1 m apart from each other. We developed a Visual Basic computer program based on Garaizar et al.'s (2016) open source HTML version of Libet's clock. The computers run the program on Microsoft Windows 7. A screen capture of this version is shown in Fig. 1. Its accuracy and settings were validated using the identical methodology published by Garaizar et al. Auditory stimuli were presented via headphones connected through the audio output of the PCs.

3.1.3. Procedure and design

The procedure and design were identical to those described by Garaizar et al. (2016). On each trial, participants observed a clock face on the computer screen with a dot rotating around the sphere at a constant speed of 2560 ms per cycle. There were two cycles of the dot per trial. Participants were asked to look at the center of the clock during the first cycle within each trial, and to press the space bar anytime they wished to during the second cycle on each trial. Pressing the bar did not stop the rotating dot, but was followed by auditory feedback at different delays as described below.

We used a within-subject design with two types of trials during training. On half of the trials (i.e., immediate condition) a 1000 Hz 200 ms tone was programmed to occur immediately following the registration of the action, using Microsoft Windows Multimedia Timers,¹ which have a resolution of 1 ms as assessed by popular experimental software packages such as SuperLab (Abboud, Schultz, & Zeitlin, 2006) and E-Prime (Schneider, Eschman, & Zuccolotto, 2002). On the other half of trials (i.e., delayed condition), the same tone was programmed to occur with a 500 ms delay after the registration of the action. There were 40 trials for each condition, and they were presented in pseudorandom order. Immediately after each one of the 80 trials was completed, participants observed an empty clock face and were asked to indicate the location at which the dot was when they pressed the space bar during that trial (see Fig. 1). This was their subjective judgment of their time of action. Participants answered this question using the mouse, after which the empty clock face was presented again and the inter-trial interval (ITI) began. The ITI lasted between 1000 and

¹ [https://msdn.microsoft.com/en-us/library/ms704986\(v=vs.85\).aspx](https://msdn.microsoft.com/en-us/library/ms704986(v=vs.85).aspx).




Fig. 2. Mean estimated time of action for conditions Immediate, Delayed, and Baseline in Experiment 1. Zero on the Y axis reflects the actual time of action. Thus, a positive value means that the action was perceived after it took place; a negative value means that the action was perceived before it took place. Error bars represent the standard error of the mean.

3000 ms. At the beginning of the ITI, a text prompting participants to be prepared for the next trial was presented along with a complex 1000 ms tone (500 ms at 250 Hz followed by 500 ms at 440 Hz). Then, after a variable interval of 0–2000 ms the next trial began. Finally, after all 80 trials had been completed a screen instructed participants that on the following trials no auditory feedback would be presented. Participants were then presented with 20 additional trials without auditory feedback in order to assess their baseline judgments without the provision of feedback (i.e., baseline condition).

3.2. Results and discussion

For comparison purposes, we will first describe our results as they have typically been described in previous research. Thus, we will first report the subjective estimates for the time of action averaged across the training session. Mean estimated time of action is shown in Fig. 2 against the actual recorded time for action (which is represented as zero on the Y axis). A positive value means that the action was estimated on average after it already had taken place, whereas a negative value means that the action was estimated on average before it had taken place. In addition, it should be noted that baseline judgments are often subtracted from target judgments in the literature and the data are reported after this transformation has already been applied (e.g., Moore et al., 2009; Tobias-Webb et al., 2017). However, in order to facilitate potential comparisons, we decided to depict the experimental trials and the baseline trials separately (rather than subtract baseline data from target action-outcome trials). Thus, Fig. 2 shows the raw average estimates for the time of action (in ms from the actual time of action) for both the action-outcome trials and the baseline trials separately. Means (and standard errors of the mean) are -11.16 (4.97), 81.09 (17.98), and -8.45 (8.15) for the immediate, delayed and baseline conditions, respectively.

Visual inspection of Fig. 2 shows that the participants' subjective judgment of the time of action was significantly longer when feedback was delayed. A one-way ANOVA confirmed this impression, revealing a main effect of condition, $F(2, 110) = 27.010$, $p < 0.001$, $\eta_p^2 = 0.329$. Further pairwise comparisons showed that judgments were significantly delayed when feedback was delayed in comparison with when it was immediate, $t(55) = 5.404$, $p < 0.001$, $d_z = 0.722$, or with when feedback was not provided (i.e. baseline condition), $t(55) = 5.451$, $p < 0.001$, $d_z = 0.728$. No significant differences were observed between the immediate and the baseline conditions, $t(55) = 0.381$, $p = 0.705$, $d_z = 0.05$. The observed difference in the subjective estimation of the time of action between the delayed and the immediate trials is an example of temporal binding, that is, the subjective estimation of the time of action was displaced toward the outcome when it was delayed. Consistent with previous research, this suggests that people inferred the time of their actions as a function of when the consequences of those actions occurred.

The subjective estimation of the time of action in the delayed condition was also significantly longer than the time at which the action actually occurred (i.e. zero on the Y axis of Fig. 2), $t(55) = 4.508$, $p < 0.001$, $d_z = 0.602$. Interestingly, there was also a discrepancy between the estimated time of the action and the occurrence of the action in the immediate condition, $t(55) = 2.246$, $p = 0.029$, $d_z = 0.3$. In those cases, however, participants reported that their action occurred before, rather than after, the actual time of the action. We will use the term *antibinding* to refer to this effect that occurs in the immediate condition, since the estimated time of action was shifted away from the outcome. No significant differences were observed during baseline trials between the estimated time of action and actual time of action, $t(55) = 1.037$, $p = 0.304$, $d_z = 0.138$.

The most critical and newest data in this experiment are shown in Fig. 3. It shows the learning curve for the subjective judgments of the time of action throughout the experimental session, depicted in blocks of 10 trials. As can be observed in this figure, during the first block of trials participants estimated the timing of their action quite accurately in the immediate condition. However, in the delayed condition, their subjective judgment was displaced toward the tone, showing a clear binding effect. Importantly, Fig. 3 also




Fig. 3. Estimated time of action in blocks of 10 trials, for conditions Immediate, Delayed, and Baseline, in Experiment 1. Zero on the Y axis reflects the actual time of action. Thus, a positive value means that the action was perceived after it took place; a negative value means that the action was perceived before it took place. Error bars represent the standard error of the mean.

shows that after those early trials participants progressively learned to correct their estimations in the delayed condition, and they became more accurate as they gained experience, showing a clear learning effect. These observations were confirmed by a 2 (condition: immediate vs. delayed) \times 4 (block of trials) ANOVA which yielded a main effect of condition, $F(1, 55) = 28.984$, $p < 0.001$, $\eta_p^2 = 0.344$, a main effect of block of trials $F(3, 165) = 15.936$, $p < 0.001$, $\eta_p^2 = 0.225$, as well as an interaction between these two variables $F(3, 165) = 0.783$, $p < 0.001$, $\eta_p^2 = 0.127$.

Subsequent planned comparisons showed that in the delayed condition there was a significant difference between the first and the last block of trials, which suggests that learning occurred and participants improved the accuracy of their estimations through the training session, $t(55) = 5.054$, $p < 0.001$, $d_z = 0.675$. Also, and even though learning and improvement are evident in the delayed condition, the participants' estimations are significantly different from the actual time of action, even on the fourth block of trials, $t(55) = 6.511$, $p < 0.001$, $d_z = 0.87$, for the first block; $t(55) = 2.637$, $p = 0.011$, $d_z = 0.352$, for the fourth block. The learning curve in Fig. 3 suggests that the accuracy of judgments could still improve if a more extended learning experience were provided.

The immediate condition also shows the significant difference between the first and the fourth block which reflects the occurrence of learning, $t(55) = 2.409$, $p = 0.019$, $d_z = 0.321$. In this condition, however, the estimation was at first quite accurate and did not differ from zero (i.e. the actual time of action) on the first block of trials, $t(55) = 0.247$, $p = 0.785$, $d_z = 0.033$, but as learning proceeded, estimations began to shift away from the outcome (and thus from the action as well) so that the antibinding effect became evident, and by the fourth block participants came to believe that they had acted significantly earlier than they had, $t(55) = 3.035$, $p = 0.004$, $d_z = 0.405$. Thus, the observed antibinding effect seems to be a learning effect.

The estimations of the time of action during the baseline phase did not significantly differ between the first and the last block of trials, $t(55) = 0.294$, $p = 0.77$. Moreover, no significant differences were observed between the estimated time of action and the actual time of action during baseline, $t(55) = 0.896$, $p = 0.374$, $d_z = 0.119$, for the first baseline block of trials, $t(55) = 0.7$, $p = 0.487$, $d_z = 0.093$ for the last block. That is, the time of action appears to be accurately judged throughout the baseline phase.

4. Experiment 2

Experiment 1 showed that participants inferred the time of their actions from the consequences of those actions, so that when the consequences were delayed participants estimated that they had acted later than they had. Most importantly, participants learned progressively to improve their estimations of their action on the delayed trials, possibly by subtracting some time from the time at which the auditory feedback was presented. However, as they did so, they also began to subtract some time from the time of feedback in the immediate condition, so they also began to assume that their action occurred earlier than it did when feedback was immediate.

Our next purpose was to explore whether those effects would also occur if we asked participants to estimate the time of their conscious decision to act rather than the time of their actions. If the subjective time of action can be so easily misled by the time in which the consequences of the action occur, then the subjective time of decisions should also be affected. As has been shown previously, participants should believe that their decision took place sooner or later as a function of whether the auditory feedback for their actions occurred sooner or later (Banks & Isham, 2009). This implies that people cannot be certain of when they made a decision. Instead, they possibly infer their time of decision upon observing the consequences of their behavior. Investigating whether people infer a fixed time for decisions as a function of when the outcome occurs or whether, in contrast, this time is learned and modifiable through experience is the aim of Experiment 2.




Fig. 4. Mean estimated time of decision for conditions Immediate, Delayed, and Baseline in Experiment 2. Zero on the Y axis reflects the actual time of action. Thus, a positive value means that the decision was perceived after the action had already occurred; a negative value means that the decision was perceived before the action took place. Error bars represent the standard error of the mean.

4.1. Method

4.1.1. Participants

The sample was comprised of 56 Psychology students from the University of Deusto, who volunteered to take part in the experiment in exchange for academic credit. Data from seven of them were discarded according to the data selection criterion described for Experiment 1.

4.1.2. Design, apparatus, and procedure

All aspects of the experimental design and procedure are identical to those of Experiment 1 except that we asked participants to report the time of their decision to act rather than the time of their action.

4.2. Results and discussion

As in Experiment 1, we will first report the subjective judgments averaged throughout the experimental session. In the present experiment, these judgments reflect the estimated time of decision rather than action. Again, in order to facilitate potential comparisons, we decided to depict the raw data from the experimental trials and the baseline trials separately. Fig. 4 shows the raw average estimates for the time of decision in ms from the recorded time of action (which is shown as point zero on the Y axis) for both the action-outcome trials and the baseline trials. Means (and standard error of the mean) are -11.33 (6.98), 27.92 (9.25), and -24.06 (6.90) for the immediate, delayed and baseline conditions, respectively.

As can be observed in Fig. 4, when the action was immediately followed by feedback participants estimated that their decision had occurred slightly before their action, which seems to be a sensible inference. However, when feedback following the action was delayed, participants then moved their estimated time of decision toward the tone. This replicates previous findings (Banks & Isham, 2009) and can also be regarded as an extension of the standard binding effect, as it shows that the shift of the subjective estimation toward the outcome occurs not only for actions but also for decisions.

A one-way ANOVA confirmed those impressions, showing a main effect of condition on the reported time of decisions, $F(2, 96) = 20.398$, $p < 0.001$, $\eta_p^2 = 0.298$. Judgments of the time of decision were significantly delayed when feedback was delayed in comparison with when it was immediate, $t(48) = 4.912$, $p < 0.001$, $d_z = 0.701$, and when compared to the baseline condition, $t(48) = 5.716$, $p < 0.001$, $d_z = 0.816$. No significant differences were observed between the immediate and the baseline conditions, $t(48) = 1.529$, $p = 0.133$, $d_z = 0.218$. Thus, participants similarly judged the time of their own decisions when there was no feedback or when feedback was immediate, but their estimation of their own decisions was delayed when the consequences of their action were delayed.

It is interesting to note that the estimated time of decision in the delayed condition was also significantly delayed when compared with the time at which the action actually occurred (i.e. zero on the Y axis of Fig. 4), $t(48) = 3.017$, $p = 0.004$, $d_z = 0.431$. We refer to this finding as a *superbinding* effect, as participants came to believe that their decision occurred *after* they had already acted, rather than before they acted. Thus, the mean estimation of the time of decision was shifted toward the tone, and beyond the action. This suggests that when estimating the time of their decisions, participants were more strongly influenced by the observation of consequences, which probably signaled the apparent time of action, than by their actual decisions (or their actual actions).

There was also a significant difference between the estimated time of decision and the actual occurrence of the action on the baseline trials, $t(48) = 3.486$, $p = 0.001$, $d_z = 0.498$ (see Fig. 4). In those cases, participants tended to report that their decision




Fig. 5. Estimated time of decision in blocks of 10 trials, for conditions Immediate, Delayed, and Baseline, in Experiment 2. Zero on the Y axis reflects the actual time of action. Thus, a positive value means that the decision was perceived after the action had already occurred; a negative value means that the decision was perceived before the action took place. Error bars represent the standard error of the mean.

occurred before the action, which is a sensible inference. This temporal priority of the reported conscious decision over the actual action was not observed in the immediate condition, $t(48) = 1.623$, $p = 0.111$, $d_z = 0.231$, perhaps due to the fact that the delayed condition was presented during the same block of trials. (Nevertheless, recall that these are averaged data for the entire session; the analyses of the learning curve are described below and show a different and interesting pattern).

The most critical and newest results in this experiment are shown in Fig. 5. It shows the learning curves for the subjective timing of conscious decision throughout the experimental session, depicted in blocks of 10 trials. The learning curves are very similar to those observed in Experiment 1, which suggests that the course of learning is very similar for actions and decisions. Although the binding effect for decisions in this experiment appears to be less pronounced than the effect for actions found in Experiment 1, it should be noted that the estimated time of decision should be located, at least in principle, before the action. Thus, the fact that the binding effect is still present, that is, decisions are still shifted toward the outcome and beyond the action, suggests that the binding effect for decisions is actually strong. As already mentioned, we are referring to it as a superbinding effect. Particularly during the early trials, participants seem to infer that their decisions occurred near the tone, and after their actions had already occurred. Quite possibly, they misjudged the timing of their actions, as was evident in Experiment 1, and therefore they also misjudged the location of their decisions. Then, as training proceeds, the participants appear to progressively learn that their decisions must have occurred earlier than they thought, so they gradually reduce the superbinding effect.

These observations were confirmed by a 2 (condition: immediate vs. delayed) \times 4 (block of trials) ANOVA which yielded a main effect of condition, $F(1, 48) = 19.221$, $p < 0.001$, $\eta_p^2 = 0.286$, of block of trials $F(3, 144) = 3.471$, $p = 0.018$, $\eta_p^2 = 0.067$, and no interaction between these variables, $F(3, 141) = 1.358$, $p = 0.258$, $\eta_p^2 = 0.28$. Judgments about when conscious decisions occurred were significantly influenced by whether the tone followed the action immediately or after a delay, as shown by the significant difference observed between the immediate and delayed conditions throughout the session, $t(48) = 3.207$, $p = 0.002$, $d_z = 0.458$ for the first block of trials, and $t(48) = 4.421$, $p < 0.001$, $d_z = 0.631$ for the last block. This represents a clear binding effect for the estimated time of decision, which in principle should not be affected by the consequences of actions but was significantly shifted toward them when they were delayed. The results also show a progressive learning effect in both the immediate and the delayed conditions. The analysis of the delayed condition shows that there is a significant difference between the first and the last block of trials, which indicates that learning occurred and that participants changed (and presumably improved) the accuracy of their estimated time of decision throughout the session, $t(48) = 3.774$, $p < 0.001$, $d_z = 0.539$. The immediate condition also shows a significant difference between the first and the fourth blocks of training, which reflects the occurrence of learning as well, $t(48) = 2.276$, $p = 0.027$, $d_z = 0.325$.

Interestingly, in the delayed condition decisions are judged to occur much later than the actual actions during the first block of trials, $t(48) = 4.089$, $p < 0.001$, $d_z = 0.584$, but this superbinding effect is then gradually corrected and by the fourth block of trials decisions are no longer judged to occur after their actions, $t(48) = 0.186$, $p = 0.853$, $d_z = 0.026$. Nevertheless, by the end of training decisions are not yet judged as preceding the action. It is quite possible that the estimation of the action is still suffering from some degree of binding by the end of the training session, which suggests that, like in Experiment 1, the accuracy of judgments could still improve if a more extended learning experience were provided. In the immediate condition, however, the estimated time of decisions did not differ from the actual time of action in the first block of trials, $t(48) = 0.625$, $p = 0.535$, $d_z = 0.089$, but by the fourth block participants had learned to locate their decisions significantly before their actions, $t(48) = 5.27$, $p < 0.001$, $d_z = 0.752$. This suggests that, as should be expected, the immediate condition was easier for them to learn.

The estimations of the time of decision during the baseline trials did not significantly differ between the first and the last block of trials, $t(48) = 0.752$, $p = 0.456$, $d_z = 0.107$, and thus no learning was observed during baseline. Estimated time of decision was significantly lower than zero (see Fig. 5) from the very first block of trials, $t(48) = 3.381$, $p = 0.001$, $d_z = 0.483$, and remained so

through the last block, $t(48) = 3.097$, $p = 0.003$, $d_z = 0.442$. That is, participants inferred that their decisions occurred before their actions throughout the baseline phase.

5. General discussion

The two experiments reported here suggest that participants infer the time of their actions and decisions from their consequences. By manipulating the time of those consequences we were able to change both the time at which they believed they had acted (Experiment 1) and the time they believed they had decided to act (Experiment 2). The results of Experiment 1 extend previous reports in the action-outcome binding literature (Garaizar et al., 2016; Haering & Kiesel, 2012; Haering & Kiesel, 2015; Haggard et al., 2002; Moore, 2016); and the results of Experiment 2 extend previous reports in the conscious will literature (Banks & Isham, 2009; Libet et al., 1983; Wegner, 2002). Taken together, the two experiments suggest that the subjective estimates of the time of action and decision are governed by similar processes. That is, both the subjective estimates of the time of actions and decisions were affected by their consequences and tended to be shifted toward them. Most importantly, we observed that estimations of the time of action and decision were subject to learning and were progressively adjusted as learning proceeded. We believe this is the first time that a learning curve has been reported to show how people estimate the time of their actions and decisions. Below we elaborate on these findings.

5.1. Learning, antibinding, and superbinding

The two experiments showed that inferring the timing of decisions and actions is subject to learning, since the participants were continuously adjusting their timing inferences. We believe this is the first time that a clear learning curve has been shown for the estimation of actions and decisions. Action-outcome binding effects found in previous studies were replicated in the current experiments for both actions and decisions, and not only when judgments were averaged across the experimental session (i.e. as they are typically reported; see Figs. 2 and 4), but were also observed during the early training trials (see Figs. 3 and 5). Importantly, however, those binding effects were progressively corrected as learning proceeded (see Figs. 3 and 5). In some cases the binding effects vanished for both actions and decisions as learning proceeded whilst in other cases this learning and adjusting process yielded what we have called an *antibinding* effect: The reported time of action tended to be progressively shifted away from the outcome, rather than toward it, and in the immediate condition the action ended up being judged as occurring significantly earlier than the actual action. Interestingly, as participants learned to improve their estimations in the delayed feedback condition, they also overcompensated (erroneously) their subjective inferences in the immediate condition. Thus, this antibinding effect is also the result of learning. Seemingly, participants learned that some time interval needed to be subtracted from the apparent time of action (which was strongly influenced by the time of feedback), and so they subtracted this time not only in the delayed condition but also in the immediate condition, thereby producing the antibinding effect on the immediate trials. Thus, this antibinding effect has possibly been favored by the fact that various intervals were presented in the same block, so that what was learned for one of them implied an error when applied to the other one.

It could be argued that the antibinding effect might simply be reflecting that the computer program is registering the time in which the action is completed whereas the participants might be reporting the time at which they initiated the action. It is true that the time at which an action is initiated should differ slightly from the time at which it is completed and registered, but even if participants were using the initiation of the action as the assessment point, we believe that this could not explain the observed effects. During the early trials the participants' estimate of their time of action in the immediate condition does not differ from the actual time of action recorded by the computer. It is only with time and experience that participants start to locate their actions away from the outcome (and from the actual action), and this happens only in the immediate condition. In the delayed condition the estimated time of action is shifted toward the outcome. If antibinding were simply due to participants reporting the time of the initiation of their action, then the learning curves should be identical (and flat) in the immediate and the delayed conditions. We believe the interpretation in terms of learning is more plausible.

Indeed, effects similar to these ones have also been reported in cases in which, instead of being trained on several delays simultaneously, participants are trained with one delay, so that they learn to expect a certain action-outcome interval, and then they are tested with different delays in a subsequent test phase. When tested with shorter delays, participants feel that the outcome occurs before their action, and sometimes they even feel that they have not caused the outcome, because it occurred too early for this to be the case (Haering & Kiesel, 2012; Haering & Kiesel, 2015; Heron, Hanson, & Whitaker, 2009; Stetson, Cui, Montague, & Eagleman, 2006). Moreover, in some cases participants may even distort their perception of temporal order in order to make it consistent with their causal beliefs. If the outcome occurs slightly before the action, participants will not realize that it does. They may even reorder the sequence of events so that they come to believe that the events occurred in the "correct" order, that is, action before effect (Bechlivanidis & Lagnado, 2013; Bechlivanidis & Lagnado, 2016; Desantis, Roussel, & Waszak, 2011; Heron et al., 2009).

One potential problem of any experiment using Libet's clock and related procedures is that participants could preplan their actions. Then the delayed feedback might be taken as indication of their inaccuracy to press the spacebar exactly at the planned time during the early trials, which might be the reason why they typically believe they acted later than they did. Note, however, that even though preplanning is certainly possible, participants would still be taking the tone as a signal of when the action (and decision) was performed, so we believe that the possibility of preplanning does not undermine our main proposal that the time of actions and decisions is inferred from consequences. Moreover, the role of learning is still evident in the learning curve and shows that participants learn to gradually correct their estimation error (regardless of whether their action was planned), and they gradually learn to

give more weight to cues that are better signals of the action (such as, for instance, haptic cues) while they progressively give less weight to cues such as auditory feedback, which in this case is a totally unreliable cue.

Also of interest is the effect observed in Experiment 2 with the estimated time of decision as the dependent variable (Experiment 2). In this case, participants shifted their estimated time of decision toward the outcome, as we expected, but quite surprisingly, they even shifted it beyond the action. That is, they reported that their decision had occurred after they had already acted. Given that the subjective time of decision should reside, at least in principle, before the action, we are using the term *superbinding* to refer to this effect. Importantly, just as participants learned to reduce their binding effect for the time of actions in Experiment 1, they also learned to reduce this superbinding effect for the time of decisions in Experiment 2. Thus, superbinding is also subject to learning. In the immediate condition participants learned in just a few trials to locate their decisions before their actions. In the delayed condition, however, they needed more time to progressively locate their decisions away from the outcomes and closer to their actions. Even so, by the end of the session they were still estimating that they had decided after they had already acted. This finding adds support to the idea that the time of decision is inferred from the apparent time of action, which in turn is strongly affected by the time of its consequences. Nevertheless, the shape of the learning curve suggests that if training were extended, then participants would probably learn to estimate their decisions as occurring before their actions, just as they do in the immediate and in the baseline conditions. In other words, this superbinding effect should probably vanish with a more prolonged learning experience. What this experiment clearly shows is that decisions are not directly perceived, nor are they inferred to occur at a fixed interval prior to actions. Instead, participants progressively learn to locate their decisions before their actions.

5.2. How do people learn to adjust their inferences for actions and decisions?

It is worth noting that with respect to decisions (i.e. Experiment 2) there are no objective parameters to learn from, that is, no parameters that require adjustment. In other words, there is nothing that can be taken as an objective cue for learning about the actual time of decisions. In principle, if no cues exist that can facilitate learning, then learning should not take place, and adjustments would not make sense. Even so, participants in this experiment are learning to locate their decisions before their actions, and they are modifying their estimates with experience. We believe that they are quite possibly learning to improve their estimations of the time of actions, as shown in Experiment 1, and then they progressively learn to also subtract some time from the estimated time of action, probably assuming that decisions must precede actions. As already noted, the action is possibly the most reliable cue that can be used to infer that a decision has been made and therefore the question of how people learn to infer their decisions can be substituted, at least in part, by the question of how people learn to infer their actions.

James (1890) Ideomotor Theory posits that associations are formed between actions and effects, and once they are formed the action will prime the mental representation of the effect, just as the observation of the effect will also activate the mental representation of the action. Consistent with this view, Ebbinghaus (1885) had also suggested that associations between mental representations of events were bidirectional, so that each event could activate the representation of the other one in either direction. Although early 20th century research was not conclusive about the bi-directionality of associations and the idea was generally abandoned (see Asch & Ebenholtz, 1962; Ekstrand, 1966, for review), more recent research has shown that associations can indeed activate the mental representation of either event upon observation of the other one (Arcediano et al., 2003; Arcediano, Escobar, & Miller, 2005; Dignath, Pfister, Eder, Kiesel, & Kunde, 2014; Gerolin & Matute, 1999).

Recent extensions of James' Ideomotor Theory (Dignath et al., 2014; Elsner & Hommel, 2001), as well as the Temporal Coding Hypothesis developed by Miller and his colleagues in the domain of associative learning with animals (Matzel et al., 1988; Savastano & Miller, 1998) and humans (Arcediano et al., 2003; Arcediano et al., 2005), suggest that associations encode not only the representation of the associated events but also their timing relationship. That is, if a given temporal relationship exists between the two associated events, each of these events will prompt the mental representation of the other at the time when it should occur or should have occurred (Dignath et al., 2014; Matzel et al., 1988; Savastano & Miller, 1998). This view predicts that whenever an action is performed its effect will be predicted based on previous learning; and that whenever the effect is observed, the occurrence of the action will be inferred at its usual time. This proposal is also consistent with the idea that action-effect binding might be the result of people using the consequences of their actions as temporal markers in order to reconstruct the time of their actions (and thus their decisions as well) (Banks & Isham, 2011; Friedman, 1990).

These two processes (predicting the second event from the first and inferring the first event from the second) have often been regarded in the literature as two alternative hypotheses, but they are perfectly compatible. Indeed, the two of them have been shown to play a critical role, and it has been suggested that there is both a predictive and a postdictive (inferential) component in the subjective estimation of actions and in the action-outcome binding effect (Chambon, Sidarus, & Haggard, 2014; Moore & Haggard, 2008; Moore et al., 2009; Wenke, Fleming, & Haggard, 2010). We do not want to mean that these predictions and inferences need be perfectly explicit, conscious, and verbalized. However, because the action is associated with its consequence (which includes timing), performing the action involves the prediction (whether explicit or incidental) that its consequence will occur at its usual time. Additionally, because of the bidirectionality of the association, the representation of the action (including its time of occurrence) can also be backwardly activated upon observation of its consequence.

All of the above suggests that continuous adaptation (i.e. learning) takes place when people infer the occurrence of their actions and decisions from consequences, a view that is strongly supported by the present findings. According to the most widely accepted theories of associative learning, learning consists of reducing prediction errors (Rescorla & Wagner, 1972). Combining traditional learning theories with the proposals described above concerning temporal coding and bi-directionality of associations (Dignath et al., 2014; Matzel et al., 1988; Savastano & Miller, 1998), it seems reasonable to expect that when participants perform an action they will

automatically predict (whether explicitly or not) that a consequence will occur at a given time (as a function of their previous experience). We should also expect that if the consequence does not occur at the expected time, then learning will consist of adjusting those prediction errors until they are minimized. Moreover, from the occurrence of the consequence, people should also adjust the time at which they believe the action should have occurred, so that, if the outcome is moved forward, then the action should also have been moved forward.

It is currently a matter of debate as to whether the subjective action-outcome interval becomes shortened in the action-outcome binding paradigm or whether the whole subjective interval is not shortened but instead displaced. In other words, it might be that action-outcome binding effects are due to a reduction in the subjective interval between actions and outcomes, but it could also be that when the outcome is delayed on a given trial, then the perception of the entire interval is shifted in that direction so that the estimated time of action will preserve the length and synchrony of the original interval (Heron et al., 2009). We believe our results are more consistent with this second view. The estimation of the action (and decision) tended to be shifted in the direction of the outcome or in the opposite direction as a function of whether the outcome occurred later or sooner than the time it was expected to occur. The estimated time of action and decision were also moved in the forward direction or in the backward direction as a function of whether learning was in its earlier or more advanced stages. Initially, the estimations tended to be displaced toward the outcome (which occurred surprisingly late), but then they moved away from the outcome as participants learned that the tone was not a reliable cue for the occurrence of action. Thus, it was not the case that the interval became shorter, it just moved in one or the other direction.

A similar proposal has been advanced by different versions of the Comparator Theory in the area of sensorimotor research (Frith, 2005; Frith, Blakemore, & Wolpert, 2000; Miall & Wolpert, 1996; Synofzik, Vosgerau, & Newen, 2008; Wolpert, Ghahramani, & Jordan, 1995), according to which, people compute the difference between the sensory input that they expect when performing the action and the one they receive, and they use the output of this sensorimotor comparator process to realign the temporal interval (see also Heron et al., 2009). Because in the present experiments we presented the consequences of the action at two different intervals rather than just one, it was impossible for participants to accurately predict whether a tone would follow the action (and decision) immediately or whether it would be delayed. Under those uncertain circumstances, we suggest that participants are likely to have minimized those errors by encoding some averaged prototype of the intervals rather than (or in addition to) the correct ones. Participants could then use this averaged interval, both when they predict the outcome of individual trials and when they infer the time at which their action (or decision) occurs. In other words, they probably predicted the consequences of their actions at some intermediate point (e.g., 200 ms) and then, when they observed the consequences of their actions, they possibly inferred that the actions had occurred at some intermediate time before the consequences (e.g., 200 ms). We believe this might be important in the development of the antibinding effect; on trials in which feedback occurred immediately after the action, participants would still infer that their action should have occurred earlier, thereby producing the antibinding effect. This view assumes that the whole interval moves forward (or backward) rather than becoming shortened in binding effects (see Heron et al., 2009).

Learning is, by definition, continuous, and should always correct errors. This means that one should expect that if changes are introduced in an already learned situation, then new errors could occur, and then relearning will need to take place again. Our results suggest that before starting the experiment, participants already have some expectations about the standard interval between their action and their consequences, due to their previous experience with the experimental context (typically a computer). This is consistent with Cravo et al.'s (2013) proposal that pre-experimental biases play a crucial role in the observation of binding effects with delayed outcomes, an effect that we replicated here during early trials. If those outcomes are presented later (or sooner) than expected, then participants assume that the action must have also occurred later (or sooner). Our results show that those expectancies, or pre-experimental biases, need to be learned and are continuously adjusted for each action-outcome pair. By the same reasoning, this adjustment does not always consist of shifting the estimation of the action toward the outcome (binding); sometimes it consists of shifting the estimation of the action away from the outcome (antibinding), and sometimes decisions can even be reported to occur after their actions (superbinding). In any case, and as learning proceeds, those errors tend to be corrected and people gradually learn to adjust the estimation of their actions to the occurrence of the actual actions. As they learn to adjust the estimation of their actions, they also learn to adjust the estimation of their decisions so that they appear to occur before their actions.

Acknowledgements

Support for this research was provided by Grant PSI2016-78818-R AEI/FEDER from Agencia Estatal de Investigación of the Spanish Government (AEI) and Fondo Europeo de Desarrollo Regional (FEDER), as well as Grant IT955-16 from the Basque Government, both awarded to HM. The funders had no role in study design, data collection and analysis, decision to publish, or preparation of the manuscript.

Author contributions

Conceived and designed the experiments: HM, CPC, PG. Developed the software: PG. Conducted the experiments: CPC. Analyzed the data: CPC, PG, HM. Contributed reagents/materials/analysis tools: HM. Wrote the paper: HM, CPC, PG.

Competing financial interests

The authors declare no conflict of interest in the publication of these experiments.

Data availability

The data file for these experiments is available at the Open Science Framework, <https://osf.io/8zxbk/>.

Appendix A. Supplementary material

Supplementary data associated with this article can be found, in the online version, at <http://dx.doi.org/10.1016/j.concog.2017.09.009>.

References

- Abdoud, H., Schultz, W. H., & Zeitlin, V. (2006). *Superlab, stimulus presentation software (version 4.0)*. San Pedro, CA: Cedrus Corp.
- Alquist, J. L., Ainsworth, S. E., & Baumeister, R. F. (2013). Determined to conform: Disbelief in free will increases conformity. *Journal of Experimental Social Psychology*, *49*(1), 80–86. <http://dx.doi.org/10.1016/j.jesp.2012.08.015>.
- Arcediano, F., Escobar, M., & Miller, R. R. (2003). Temporal integration and temporal backward associations in human and nonhuman subjects. *Learning & Behavior*, *31*, 242–256. <http://dx.doi.org/10.3758/BF03195986>.
- Arcediano, F., Escobar, M., & Miller, R. R. (2005). Bidirectional associations in humans and rats. *Journal of Experimental Psychology: Animal Behavior Processes*, *31*, 301–318. <http://dx.doi.org/10.1037/0097-7403.31.3.301>.
- Asch, S. E., & Ebenholtz, S. M. (1962). The principle of associative symmetry. *Proceedings of the American Philosophical Society*, *106*, 135–163.
- Banks, W. P., & Isham, E. A. (2009). We infer rather than perceive the moment we decided to act. *Psychological Science*, *20*(1), 17–21. <http://dx.doi.org/10.1111/j.1467-9280.2008.02254.x>.
- Banks, W. P., & Isham, E. A. (2011). Do we really know what we are doing? Implications of reported time of decision for theories of volition. In W. Sinnott-Armstrong, & L. Nadel (Eds.), *Conscious will and responsibility* (pp. 47–60). New York: Oxford University Press.
- Baumeister, R. F., Masicampo, E. J., & DeWall, C. N. (2009). Prosocial Benefits of Feeling Free: Disbelief in Free Will Increases Aggression and Reduces Helpfulness. *Personality and Social Psychology Bulletin*, *35*(2), 260–268. <http://dx.doi.org/10.1177/0146167208327217>.
- Bechlivanidis, C., & Lagnado, D. A. (2013). Does the “why” tell us the “when”? *Psychological Science*, *24*(8), 1563–1572. <http://dx.doi.org/10.1177/0956797613476046>.
- Bechlivanidis, C., & Lagnado, D. A. (2016). Time reordered: Causal perception guides the interpretation of temporal order. *Cognition*, *146*, 58–66. <http://dx.doi.org/10.1016/j.cognition.2015.09.001>.
- Buehner, M. J. (2012). Understanding the past, predicting the future: Causation, not intentional action, is the root of temporal binding. *Psychological Science*, *23*(12), 1490–1497. <http://dx.doi.org/10.1177/0956797612444612>.
- Buehner, M. J., & Humphreys, G. R. (2009). Causal binding of actions to their effects: Research article. *Psychological Science*, *20*(10), 1221–1228. <http://dx.doi.org/10.1111/j.1467-9280.2009.02435.x>.
- Buehner, M. J., & May, J. (2002). Knowledge mediates the timeframe of covariation assessment in human causal induction. *Thinking & Reasoning*, *8*(4), 269–295. <http://dx.doi.org/10.1080/13546780244000060>.
- Buehner, M. J., & May, J. (2003). Rethinking temporal contiguity and the judgement of causality: Effects of prior knowledge, experience, and reinforcement procedure. *The Quarterly Journal of Experimental Psychology. A, Human Experimental Psychology*, *56*(5), 865–890. <http://dx.doi.org/10.1080/02724980244000675>.
- Buehner, M. J., & May, J. (2004). Abolishing the effect of reinforcement delay on human causal learning. *Quarterly Journal of Experimental Psychology Section B: Comparative and Physiological Psychology*, *57*(2), 179–191. <http://dx.doi.org/10.1080/02724990344000123>.
- Caspar, E. A., & Cleeremans, A. (2015). “Free will”: are we all equal? A dynamical perspective of the conscious intention to move. *Neuroscience of Consciousness*, *20*(1), niv009. <http://dx.doi.org/10.1093/nc/niv009>.
- Chambon, V., Sidarus, N., & Haggard, P. (2014). From action intentions to action effects: How does the sense of agency come about? *Frontiers in Human Neuroscience*, *8*, 320. <http://dx.doi.org/10.3389/fnhum.2014.00320>.
- Cravo, A. M., Haddad, H., Claessens, P. M. E., & Baldo, M. V. C. (2013). Bias and learning in temporal binding: Intervals between actions and outcomes are compressed by prior bias. *Consciousness and Cognition*, *22*(4), 1174–1180. <http://dx.doi.org/10.1016/j.concog.2013.08.001>.
- Desantis, A., Rousset, C., & Waszak, F. (2011). On the influence of causal beliefs on the feeling of agency. *Consciousness and Cognition*, *20*(4), 1211–1220. <http://dx.doi.org/10.1016/j.concog.2011.02.012>.
- Dignath, D., Pfister, R., Eder, A. B., Kiesel, A., & Kunde, W. (2014). Representing the Hyphen in Action-Effect Associations: Automatic Acquisition and Bi-Directional Retrieval of Action-Effect Intervals. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *40*(6), 1701–1712.
- Ebbinghaus, H. (1885). *Memory: A contribution to experimental psychology*. New York: Teachers College, Columbia University.
- Ekstrand, B. R. (1966). Backward associations. *Psychological Bulletin*, *65*, 50–64.
- Elsner, B., & Hommel, B. (2001). Effect Anticipation and Action Control. *Journal of Experimental Psychology*, *27*(1), 229–240. <http://dx.doi.org/10.1037/AM96-1523.27.1.229>.
- Fahle, M. W., Stemmler, T., & Spang, K. M. (2011). How Much of the “Unconscious” is Just Pre-Threshold? *Frontiers in Human Neuroscience*, *5*(120), 1–6. <http://dx.doi.org/10.3389/fnhum.2011.00120>.
- Friedman, W. J. (1990). *About time: Inventing the fourth dimension*. Cambridge, MA: MIT Press.
- Frith, C. D. (2005). The self in action: Lessons from delusions of control. *Consciousness and Cognition*, *14*, 752–770. <http://dx.doi.org/10.1016/j.concog.2005.04.002>.
- Frith, C. D., Blakemore, S. J., & Wolpert, D. M. (2000). Abnormalities in the awareness and control of action. *Philosophical Transactions of the Royal Society B*, *355*, 1771–1788. <http://dx.doi.org/10.1098/rstb.2000.0734>.
- Garaizar, P., Cubillas, C. P., & Matute, H. (2016). A HTML5 open source tool to conduct studies based on Libet’s clock paradigm. *Scientific Reports*, *6*(March), 32689. <http://dx.doi.org/10.1038/srep32689>.
- Gerolin, M., & Matute, H. (1999). Bidirectional associations. *Animal Learning & Behavior*, *27*, 42–49. <http://dx.doi.org/10.3758/BF03199430>.
- Guggisberg, A. G., & Mottaz, A. (2013). Timing and awareness of movement decisions: Does consciousness really come too late? *Frontiers in Human Neuroscience*, *7*, 1–11. <http://dx.doi.org/10.3389/fnhum.2013.00385>.
- Haering, C., & Kiesel, A. (2012). Mine is earlier than yours: Causal beliefs influence the perceived time of action effects. *Frontiers in Psychology*, *3*, 1–8. <http://dx.doi.org/10.3389/fpsyg.2012.00393>.
- Haering, C., & Kiesel, A. (2014). Intentional binding is independent of the validity of the action effect’s identity. *Acta Psychologica*, *152*, 109–119. <http://dx.doi.org/10.1016/j.actpsy.2014.07.015>.
- Haering, C., & Kiesel, A. (2015). Was it me when it happened too early? Experience of delayed effects shapes sense of agency. *Cognition*, *136*, 38–42. <http://dx.doi.org/10.1016/j.cognition.2014.11.012>.
- Haggard, P., Clark, S., & Kalogeras, J. (2002). Voluntary action and conscious awareness. *Nature Neuroscience*, *5*(4), 382–385. <http://dx.doi.org/10.1038/nn827>.
- Haggard, P., & Eimer, M. (1999). On the relation between brain potentials and the awareness of voluntary movements. *Experimental Brain Research*, *126*(1), 128–133. <http://dx.doi.org/10.1007/s002210050722>.
- Heron, J., Hanson, J. V. M., & Whitaker, D. (2009). Effect before cause: Supramodal recalibration of sensorimotor timing. *PLoS ONE*, *4*(11), e7681. <http://dx.doi.org/10.1371/journal.pone.0007681>.
- Isham, E. A., Banks, W. P., Ekstrom, A. D., & Stern, J. A. (2011). Deceived and distorted: Game outcome retrospectively determines the reported time of action. *Journal*

- of *Experimental Psychology: Human Perception and Performance*, 37(5), 1458–1469. <http://dx.doi.org/10.1037/a0023111>.
- James, W. (1890). *The principles of psychology*. New York: Holt.
- Lau, H. C., Rogers, R. D., & Passingham, R. E. (2007). Manipulating the experienced onset of intention after action execution. *Journal of Cognitive Neuroscience*, 19, 81–90.
- 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*, 106, 623–642. <http://dx.doi.org/10.1093/brain/106.3.623>.
- Matzel, L. D., Held, F. P., & Miller, R. R. (1988). Information and expression of simultaneous and backward associations: Implications for contiguity theory. *Learning & Motivation*, 19, 317–344.
- Miall, R. C., & Wolpert, D. M. (1996). Forward models for physiological motor control. *Neural Networks*, 9, 1265–1279. [http://dx.doi.org/10.1016/s0893-6080\(96\)00035-4](http://dx.doi.org/10.1016/s0893-6080(96)00035-4).
- Miller, R. R., & Barnet, R. C. (1993). The role of time in elementary associations. *Current Directions in Psychological Science*, 2, 106–111.
- Moore, J. W. (2016). What is the sense of agency and why does it matter? *Frontiers in Psychology*. <http://dx.doi.org/10.3389/fpsyg.2016.01272>.
- Moore, J. W., & Haggard, P. (2008). Awareness of action: Inference and prediction. *Consciousness and Cognition*, 17(1), 136–144.
- Moore, J. W., Lagnado, D., Deal, D. C., & Haggard, P. (2009). Feelings of control: Contingency determines experience of action. *Cognition*, 110(2), 279–283. <http://dx.doi.org/10.1016/j.cognition.2008.11.006>.
- Nolden, S., Haering, C., & Kiesel, A. (2012). Assessing intentional binding with the method of constant stimuli. *Consciousness and Cognition*, 21(3), 1176–1185. <http://dx.doi.org/10.1016/j.concog.2012.05.003>.
- Rescorla, R. A., & Wagner, A. R. (1972). A theory of Pavlovian conditioning: Variations in the effectiveness of reinforcement and nonreinforcement. In A. H. Black, & W. F. Prokasy (Eds.), *Classical conditioning II: Current research and theory* (pp. 64–99). New York: Appleton-Century-Crofts.
- Rigoni, D., Brass, M., Roger, C., Vidal, F., & Sartori, G. (2013). Top-down modulation of brain activity underlying intentional action and its relationship with awareness of intention: An ERP/Laplacian analysis. *Experimental Brain Research*, 229(3), 347–357. <http://dx.doi.org/10.1007/s00221-013-3400-0>.
- Rigoni, D., Brass, M., & Sartori, G. (2010). Post-action determinants of the reported time of conscious intentions. *Frontiers in Human Neuroscience*, 4. <http://dx.doi.org/10.3389/fnhum.2010.00038>.
- Rigoni, D., Kühn, S., Gaudino, G., Sartori, G., & Brass, M. (2012). Reducing self-control by weakening belief in free will. *Consciousness and Cognition*, 21(3), 1482–1490. <http://dx.doi.org/10.1016/j.concog.2012.04.004>.
- Savastano, H. I., & Miller, R. R. (1998). Time as content in Pavlovian conditioning. *Behavioural Processes*, 44, 147–162.
- Schneider, W. E., Eschman, A. A., & Zuccolotto, A. (2002). *E-prime user's guide*. Pittsburgh: Psychology Software Tools.
- Soon, C. S., Brass, M., Heinze, H.-J., & Haynes, J.-D. (2008). Unconscious determinants of free decisions in the human brain. *Nature Neuroscience*, 11(5), 543–545. <http://dx.doi.org/10.1038/nn.2112>.
- 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), 651–659.
- Synofzik, M., Vosgerau, G., & Newen, A. (2008). Beyond the comparator model: A multifactorial two-step account of agency. *Consciousness and Cognition*, 17, 219–239.
- Tobias-Webb, J., Limbrick-Oldfield, E. H., Gillan, C. M., Moore, J. W., Aitken, M. R. F., & Clark, L. (2017). Let me take the wheel: Illusory control and sense of agency. *The Quarterly Journal of Experimental Psychology*, 70(8), 1732–1746. <http://dx.doi.org/10.1080/17470218.2016.1206128>.
- Vohs, K. D., & Schooler, J. W. (2008). The value of believing in free will: Encouraging a belief in determinism increases cheating. *Psychological Science*, 19(1), 49–54. <http://dx.doi.org/10.1111/j.1467-9280.2008.02045.x>.
- Wegner, D. M. (2002). *The illusion of conscious will*. Cambridge, MA: MIT Press.
- Wenke, D., Fleming, S. M., & Haggard, P. (2010). Subliminal priming of actions influences sense of control over effects of action. *Cognition*, 115(1), 26–38. <http://dx.doi.org/10.1016/j.cognition.2009.10.016>.
- Wolpert, D. M., Ghahramani, Z., & Jordan, M. I. (1995). An internal model for sensorimotor integration. *Science*, 29, 1880–1882. <http://dx.doi.org/10.1126/science.7569931>.