Title: Effects can precede their cause in the sense of agency

Authors: Marieke Rohde[1,2], Meike Scheller[1,] and Marc O. Ernst[1,2]

Affiliations: [1]Department of Cognitive Neuroscience

University of Bielefeld,

Universitätsstr. 25, W3-240,

33615 Bielefeld, Germany

[2]Cognitive Interaction Technology (CITEC) Center of Excellence,

University of Bielefeld,

Inspiration 1,

33619 Bielefeld, Germany

Email addresses: marieke.rohde@uni-bielefeld.de

meike.scheller@uni-bielefeld.de

marc.ernst@uni-bielefeld.de

Correspondence: Marieke Rohde

Department of Cognitive Neuroscience, University of Bielefeld

Universitätsstr. 25, W3-246,

33615 Bielefeld, Germany

Phone: +49-521-106 5703

Email: marieke.rohde@uni-bielefeld.de

**Abstract**

The sense of agency, i.e., the feeling that one's action is the cause of an external sensory event, involves causal inference based on the predicted sensory outcome of a motor act. Here, we investigated whether this inference process faithfully implements the physical principle that a cause (motor act) temporally precedes its effect (external sensory feedback). To this end, we presented participants with visual flashes that were temporally offset from voluntary button presses, including scenarios where the flash occurred shortly before the press. Participants then judged their experience of agency. As expected, cause-effect order is an important cue for this task: participants were far more likely to report agency for temporally lagging flashes than for leading flashes, even if very long sensory delays also disrupted the sense of agency (Experiment 1). This suggests that the temporal order between action and sensation is the dominant temporal cue for agency. However, when participants judged whether they had caused a first flash that occurred before the button press or a second flash that occurred afterwards, the temporal threshold for rejecting leading first flashes was relaxed proportionally to the delay of the second flash (Experiment 2). There was competition between different sensorimotor timing cues (temporal order favored the second flash and temporal proximity favored the first flash), and participants' tolerance for cause-effect inversions was modulated by the strength of the later, conflicting cue. We conclude that the perceived order of action and sensation is not used in a winner-take-all fashion in inference of agency. Instead, a probabilistic negotiation of the different timing cues favoring different flash events takes place postdictively, after presentation of the second flash.

**Keywords**

Sense of agency, motor prediction, temporal processing, perception and action.

## Introduction

Humans experience a sense of agency (SoA) when sensory events are in agreement with the expected outcome of an action (Haggard & Chambon, 2012). An important cue for this kind of perceptual causal inference is the relative timing of action and sensation (Blakemore, Wolpert & Frith, 1998). If an action is self-initiated, it has to happen before the sensory consequence (temporal *priority principle*; Wegner & Wheatley, 1999). Additionally, the SoA decreases with an increasing sensory delay between action and effect (Dewey & Carr, 2013; Ebert & Wegner, 2010; Farrer, Valentin & Hupé, 2013; Haggard, Clark & Kalogeras, 2002; Moore, Wegner & Haggard, 2009; Sato & Yasuda, 2005; Weiss, Tsakiris, Haggard & Schütz-Bosbach, 2014). Generally, to be contingent with a causal interpretation, a sensory event has to occur relatively shortly after the action. This temporal *proximity principle* can be modulated by context (e.g., top-down beliefs about delayed causation; Humphreys & Buehner, 2009) and by changes in the temporal action-sensation statistics (e.g., Ebert & Wegner, 2010; Moore et al., 2009). Yet, it temporally constrains the SoA in most situations.

Most current models of the mechanisms of the SoA assume two stages of processing (e.g., Balslev, Cole & Miall, 2007; Blakemore et al., 1998; Kawabe, Roseboom & Nishida, 2013). A first sensorimotor integration stage compares different incoming streams of sensory information (touch, proprioception, visual or auditory feedback) to a prediction of sensory feedback from the history of motor commands (efference copy) and prior knowledge about the task, in order to assess temporal and spatial coherence of multisensory event information. The fact that we cannot tickle ourselves has for instance been explained with such a model. A self-produced tickle sensation is fully expected from motor output and is thus causally assigned to oneself, which makes it non-ticklish (Blakemore, Frith & Wolpert, 1998). Such low-level comparison includes an assessment of temporal coherence: For instance, if a delay is inserted between a tickling action and the corresponding touch sensation (by robotic manipulation), participants perceive their own tickling as ticklish. The delay leads to a discrepancy between the actual and expected sensory inputs and the brain does not attribute the sensation to the corresponding own tickling action (Blakemore et al., 2000). In such models, the result of the comparison operation is assumed to be fed forward for cognitive and conceptual evaluation of the action in the second stage of processing where the agency is assigned (cf., Synofzik, Vosgerau & Newen, 2008).

At least in a naïve interpretation, these models would predict a bottom-up processing of sensorimotor coherence and thus a strict application of the priority principle in the SoA. This prediction is supported by both intuition, as the rule that a cause precedes its effect is one of the most fundamental laws of physics, as well as by existing empirical evidence: Weiss et al. (2014) have shown a strong correspondence between subjective agency judgments and corticospinal activity related to motor preparation that also depended on visuomotor delays. Also, in an audiomotor temporal recalibration study, the threshold for perceiving agency was shown to shift with the threshold for perceiving simultaneity (Timm, Schönwiesner, SanMiguel & Schröger, 2014).

In order to test whether the temporal priority principle is really implemented as a fixed, bottom-up constraint for the SoA, we here investigated how humans negotiate conflicting temporal priority and temporal proximity cues when explicitly judging agency. To determine the relative importance of perceived order and perceived proximity as temporal cues for agency in Experiment 1, we asked participants to perform either agency judgments (AJs) or relative timing judgments (temporal order judgments: TOJ, simultaneity judgments: SJ) after exposure to a visual flash that occurred around the time of a voluntary button press. Contrary to other studies, we measured these judgments for both leading and lagging visual stimuli. This allowed us to characterize the temporal window of perceived agency. We observed a strong asymmetry around the Point of Subjective Simultaneity (PSS) that indicates dominance of the priority principle over the proximity principle. When judging agency, perceptible visual lags are tolerated, but visual stimuli that are perceived to occur before the press are strictly rejected.

In Experiment 2, we then studied how priority and proximity cues are negotiated in the case of a conflict, and whether a later second flash can postdictively alter the temporal processing of a leading first flash. Participants had to discriminate, which of two flashes (one before and one after their button press) they had caused. By varying the delay between the two flashes and their timing with respect to the action, we modulated the conflict between the priority cue (Is the order of first flash and the press contingent with agency?) and the proximity cue (How close in time is the second flash to the press?). We observed that the size of the conflict influences how strictly the priority principle is applied when judging agency for the first flash. That is, if the second flash occurred late, participants sometimes reported agency for a flash that occurred noticeably before their press. This shows that the priority principle is negotiated probabilistically in the inference of agency and can be relaxed postdictively, so as to accommodate the later conflicting cue.

## Experiment 1: The temporal window of perceived agency

In the experimental group (agency group), participants were asked to discriminate whether a flashed disk stimulus that was timed randomly around a button press was feedback to their own button press or feedback to a recorded press of a previous participant. From the literature (e.g., Farrer et al., 2013; Timm et al., 2014) we expected to find an asymmetrical window of perceived agency that contains lagging, but not leading visual stimuli relative to the PSS. In the control group (simultaneity group), participants had to rate whether the flash occurred at the same time as the press or not. We expected this window of perceived simultaneity to be narrower on the side of lagging visual stimuli than the window of perceived agency, but the same size for leading visual stimuli.

### Method

**Participants.** 20 participants (age range 18-42, 15 female, all right-handed by self-report) took part and received a small monetary compensation (6 €/h). Another four measurements had to be discarded due to technical device failure or participants disobeying the experimenter's instructions (two in the agency group, two in the simultaneity group). Participants gave written consent and were naïve with respect to the experimental hypothesis. The experiment was conducted in agreement with the ethics standards laid out in the 1964 Declaration of Helsinki and was approved by the ethics committee of the Department of Medicine of the University of Tübingen (Germany).

**Materials and apparatus.** We used the same set-up as in previous research (Rohde & Ernst, 2013; Rohde, Greiner & Ernst, 2014). A PHANToM™ force-feedback device (SensAble Technologies Inc.) was used to render a virtual haptic button consisting of a simulated mass (m = 0.1 kg) on a 4 mm spring (spring constant k = 500 kg/s2). Below the spring, there was an additional dead-band of 4 mm. When entering the dead-band, a button press event was registered and there was a sudden noticeable decrease in force-feedback (haptic click; see Fig. 1A and Methods section in Rohde & Ernst, 2013). The participants reached out from a defined start position and pressed this button, at approximately 5 cm distance in darkness. The finger movement was tracked in real time in order to predict the timing of complete button compression from the early part of the finger trajectory. This prediction algorithm allowed us to present visual stimuli even before the complete button press occurred, i.e., visual stimuli were timed relative to the estimated time of compression. The visual stimuli were flashed disks of 1.5° visual angle that were projected into participants' field of view using a mirror and a CRT monitor mounted upside-down at the top of the set-up in the otherwise dark room. The disks were not spatially aligned but displayed in close proximity with the haptically rendered button (i.e., within a 10 cm radius). The inherent latency of this setup was 37±7 ms.

Perceptual judgments asked were: AJ ("did you or another participant cause the flash?"), SJ ("did the button press and the visual flash occur at the same time?"), and TOJ ("which occurred earlier: button press or visual flash?").

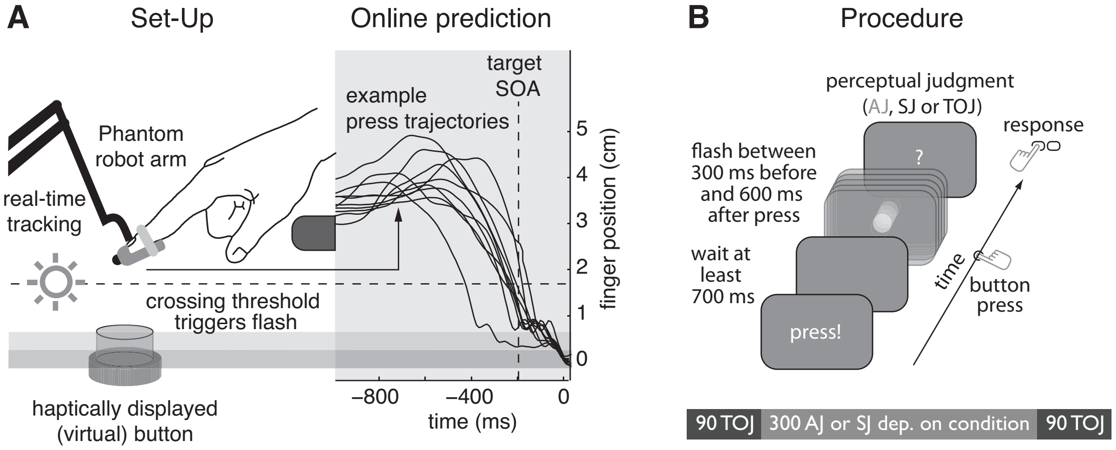


*Figure 1.* Setup and procedure. A: The set-up. Starting at the home position, participants reached out to press a button that was haptically rendered using a PHANToM™ force-feedback device. The height of the finger was analyzed online to estimate the timing of the upcoming button press, which allowed us to display visual stimuli also before the press. B: Procedure. Flashes were displayed around the time of the voluntary button press. The order of conditions is displayed as timeline at the bottom.

**Design and procedure.** Participants' right index finger was attached to the PHANToM™ force-feedback device. They practiced 20 button presses to get used to the task and to initiate the prediction algorithm. Each trial consisted of a go-signal, after which participants had to wait for at least 700 ms (in order to induce a voluntary movement) but not longer than 1500 ms before moving the finger from the home position and pressing the button (cf. Fig. 1A; timeline of a trial Fig. 1B top). Even if the flashed disks were sometimes displayed before the full compression of the button, participants had to complete the button press to move on to the perceptual judgment. If participants pressed the button too early or too late, trials were discarded (online) and repeated at the end of the block (21% of all trials).

Responses to the perceptual judgment tasks were given by pressing a key on a keyboard with the left hand.

At the beginning and the end of the experiment, participants performed 90 trials where they judged the temporal order between the button press and the visual flash (TOJ with answers ‘press first’ and ‘flash first’; timeline of the experimental conditions Fig. 1B, bottom). This was done to estimate a subjective baseline for the PSS, i.e., the asynchrony that corresponds to 50% “press first” responses. The timing of the visual flash was drawn from a uniform distribution between 300 ms before and 300 ms after the full compression of the button. From this, the PSS, was estimated (see Analysis).

There were ten participants in each of the experimental groups (agency group and simultaneity group). In the experimental blocks (recorded between the TOJ blocks) both groups were exposed to the same visual stimuli for 300 trials: Visual flashes were timed to occur randomly, drawn from a uniform distribution between 300 ms before and 600 ms after the estimated time of compression of the button press. Since there was some uncertainty in the temporal prediction procedure, flashes occasionally could occur more than 300 ms before the press.

Depending on the test group, participants had to perform different perceptual judgment tasks about these visual flashes. Participants in the agency group were instructed to report whether the presented flash was feedback to their own press or whether it was feedback to a recorded button press from a previous participant (yes/no AJ). Participants in the simultaneity group were instead instructed to rate the simultaneity of the press and the flash, which they were told would occur at a random time within the interval in which they had to press the button (yes/no SJ). The SJ condition was a temporal control task. As it is a yes/no task like the AJ task, it allows an analogous characterization of the limits of the window of perceived simultaneity and comparison of these to the limits of the window of perceived agency.

In the agency group, participants were instructed to move fast and not to modulate their press behavior (pressing late or slowly) to actively probe whether a flash occurred even without their press. We made sure that participants obeyed this instruction by testing for changes in the average time to respond to the go-signal.

The asymmetrical range of visual stimuli (300 ms before to 600 ms after the button press) was chosen because very early flashes might trigger a participant to press the button. However, it might also cause range biases (e.g., temporal recalibration; Rohde & Ernst, 2013; Stetson, Cui, Montague & Eagleman, 2006). Therefore, the correction for a subjective baseline (PSS) from the TOJs that were recorded at the beginning and the end of the experiment was important.

**Analysis.** The first 15 TOJ, SJ and AJ trials served as practice trials and were not part of the analysis. Additionally, trials were excluded if the flash was more than 425 ms earlier than the press, which indicates that participants might have missed the flash (<1% of TOJ trials and 3% of AJ/SJ trials).

The TOJ responses for each participant were analyzed by fitting a log-log psychometric function of the press-flash temporal discrepancies (Matlab Statistics Toolbox: glmfit.m with a log-log link function; cf. Fig. 2A, green line) to estimate the PSS:

$$\log(-\log(P(Y=1))) = \beta_1 + \beta_2 \cdot x \qquad \text{with} \qquad PSS = \frac{\log(-\log(0.5)) - \beta_1}{\beta_2} \tag{0}$$

Where *P(Y=1)* is the probability of answering ‘vision first’, and *x* is the temporal discrepancy between the full button compression and the flash. The fitted model parameters $\beta_1$ and $\beta_2$ describe the intercept and slope of the model. This asymmetrical psychometric function provided a better fit (Akaike Information Criterion) than the symmetrical cumulative Gaussian function often used in this kind of analysis. The PSS estimates were used as a subjective baseline for simultaneity for each participant (i.e., a temporal discrepancy of 0 corresponds to a participant’s PSS, a negative discrepancy represents a flash perceived before the press, a positive discrepancy a flash perceived after the press). The responses from the two TOJ blocks were pooled to estimate the PSS, as there were no significant differences between the two blocks when fitted individually (matched-sample t test pre vs. post: $t(19)=1.1$, $p= 0.274$). The pre-test TOJ responses of one participant had to be discarded, because he misunderstood the task in the first block (for this participant, the post-test PSS was used for baseline correction).

The SJ and AJ responses were analyzed by fitting a two-criterion window model to the responses (Cravo, Claessens & Baldo, 2011; Rohde et al., 2014; Ulrich, 1087; Yarrow, Jahn, Durant & Arnold, 2011):

$$P(Y=1) = g(x \mid \mu_V, \sigma_V) - g(x \mid \mu_M, \sigma_M) \tag{1}$$

where *P(Y=1)* is the probability of answering "yes", *x* is the temporal discrepancy between the full button compression and the flash, and $g(x|\eta,\sigma)$ denotes the cumulative probability function of a normal distribution with mean μ and standard deviation $\sigma$. This difference between two cumulative Gaussians describes a bell-shaped probability density function (example Fig. 2A, red and blue line) for parameter settings where $\mu_V < \mu_M$ and $\sigma_V$ and $\sigma_M$ are of a similar magnitude. $\mu_V$ delimits this temporal window on the vision-lead side of the range of stimuli; $\mu_M$ delimits it on the movement-lead side of the range of stimuli. $\sigma_V$ and $\sigma_M$ characterize the precision with which decisions are made on each side. This model can capture temporal asymmetries in the processing of temporal discrepancies (Rohde et al., 2014; Yarrow et al., 2011) and was fitted with a self-written maximum-likelihood-estimation algorithm implemented in Matlab 2010.

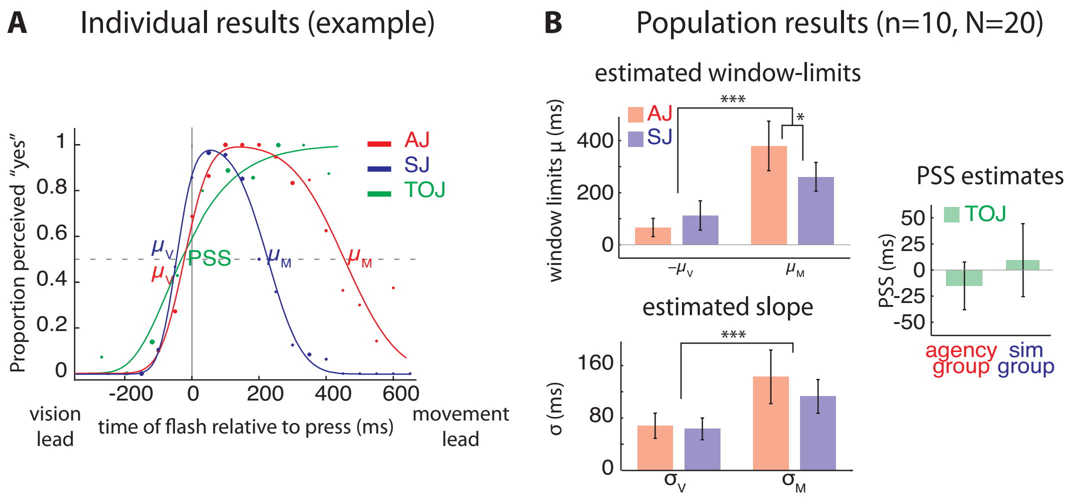


*Figure 2.* Results for Experiment 1. A: Results for two example participants, one from the agency group (AJ and TOJ) and one from the simultaneity group (SJ). B: Population results (parameter estimates; mean and confidence intervals). Window limits μ (top left) and slopes σ (bottom left) of the windows of perceived agency (red) and simultaneity (blue).

Note that $\mu_V$ are sign-inverted to indicate distance from PSS. Right: PSS estimates for both groups, pooled across pre-test and post-test.

Parameter estimates and differences between conditions on the population level are given as mean and 95% confidence intervals in the figures and the main text. Complementarily, null-hypothesis significance testing was conducted, using a mixed ANOVA (Matlab script mixed_between_within_anova, Johnson, 2010) with condition (agency vs. simultaneity) as between-participant factor and window side (vision-lead vs. movement-lead) as within-participant factor. As the emphasis was on an asymmetry around the PSS, μ parameters were sign-inverted on one side in the comparison (i.e., $-\mu_V$ was compared to $\mu_M$ to compare distance from the window mid-point; cf. Fig. 2B).

**Results**

The 2 x 2 comparison of the $\mu_V$ and $\mu_M$ limits of the windows of perceived agency (AJ task) and simultaneity (SJ task) revealed a strong asymmetry in time around the PSS for both conditions (Fig. 2A and B). There is a strong main effect of the factor window-side ($\mu_V$, $\mu_M$) in the mixed ANOVA comparison ($F(1,18)=38.3$, $p\ll0.001$). That such an asymmetry is already present in the window of perceived simultaneity (from the SJ task) can be explained by a previously reported compression of perceived duration under the assumption of agency ("intentional binding" Haggard et al., 2002), which distorts the perception of time asymmetrically (see also Rohde et al., 2014).

Importantly, there is a significant interaction of the factors window-side ($\mu_V$, $\mu_M$) and condition (agency vs. simultaneity; $F(1,18)=4.9$, $p=0.040$) but no main effect of the factor condition ($F(1,18)=1.7$, $p=0.208$). The temporal window of agency is 118±118 ms (mean and confidence interval) wider than the window of simultaneity on the movement-lead side of the range of discrepancies (two-sample t test comparing $\mu_{M,SJ}$ and $\mu_{M,AJ}$ $t(18)=2.1$, $p=0.049$), whereas the windows do not differ on the vision-lead side (two-sample t test comparing $\mu_{V,SJ}$ and $\mu_{V,AJ}$ $t(18)=1.4$, $p=0.187$). This shows that, when judging agency, participants tolerate small perceptible visual delays, which is in line with the literature (e.g., Farrer et al., 2013). Temporally leading flashes by contrast, which violate the possible cause-effect order, are judged much more strictly: The same criteria are used to assess simultaneity and agency.

The baseline PSS estimates from the TOJ task (Fig 2B right) do not differ between conditions (difference: 24±41 ms, mean and confidence interval).

Experiment 1 thus supports the hypothesis that the priority principle is strictly applied when judging agency, whereas small perceptible delays are tolerated. The 2 x 2 comparison of the $\mu_V$ and $\mu_M$ limits of the windows of both perceived agency and of perceived simultaneity shows that this is partially caused by a processing asymmetry relative to the PSS: The temporal asymmetry of agency perception is in part due to a compression of perceived duration for lagging visual stimuli. However, additionally, there is also tolerance of perceptible visual lags (Fig. 2B top left). Experiment 1 however does not yet address the question of the negotiation of conflicts between proximity and priority cues for the SoA. In Experiment 2, we put these temporal cues into conflict to address this question.

## Experiment 2: Agency Discrimination

In order to test for a possible negotiation of conflicting temporal priority and proximity cues we used an agency discrimination task (AD) in Experiment 2. Participants had to judge which one of two visual flashes that were presented sequentially around the time of the button press was a causal consequence of their action: the one presented before or the one after the press. By varying the interval between the flashes and their timing around the button press, we varied the conflict between priority cues and proximity cues: If the second flash occurs relatively late after the button press and the first one relatively shortly before the button press, the second flash is only weakly coupled to the button press in terms of temporal proximity, but is consistent in terms of temporal priority. Conversely, the first flash violates the priority principle, but is more strongly coupled to the button press in terms of temporal proximity. This conflict grows with the interval between the flashes. If the second flash occurs very late, would participants judge an earlier flash as related to their action because of its closer proximity to the button press, even if this conflicts with the required cause-effect order? This would speak against a bottom-up processing of temporal order information expected from comparator models.

## Method

**Participants.** Ten participants (age range 20-32, 6 female, all right-handed by self-report) took part in the experiment for which they received a small monetary compensation (6 €/h). They gave written consent and were naïve with respect to the experimental hypothesis. Two measurements had to be discarded due to technical device failure or participants disobeying the experimenter’s instructions. The experiment was conducted in agreement with the ethics standards laid out in the 1964 Declaration of Helsinki and was approved by the ethics committee of the Department of Medicine of the University of Tübingen (Germany).

**Materials and apparatus.** The apparatus was the same as in Experiment 1.

**Design and procedure.** As in Experiment 1, participants performed 90 TOJ trials at the beginning and the end of the experiment for a baseline PSS measure. During the main experiment, participants performed 600 AD trials (Fig. 3A). For this, two identical flashes were presented sequentially around the time of the button press. Participants were told that their button press caused only one of the two flashes, whereas the other one occurred randomly and independent of their action. They had to indicate which of the two flashes they believed was caused by their action. The timing of the first flash was drawn from a uniform distribution between 300 ms before and 200 ms after the button press. The second flash occurred at one of three fixed inter stimulus intervals (ISIs) after the first flash: 300 ms, 450 ms or 600 ms. Two hundred trials of each ISI were presented in randomized order between the two TOJ blocks (Fig. 3A bottom). If participants pressed the button too late or too early with respect to the go-signal, trials were discarded (online) and repeated (7% of all trials).

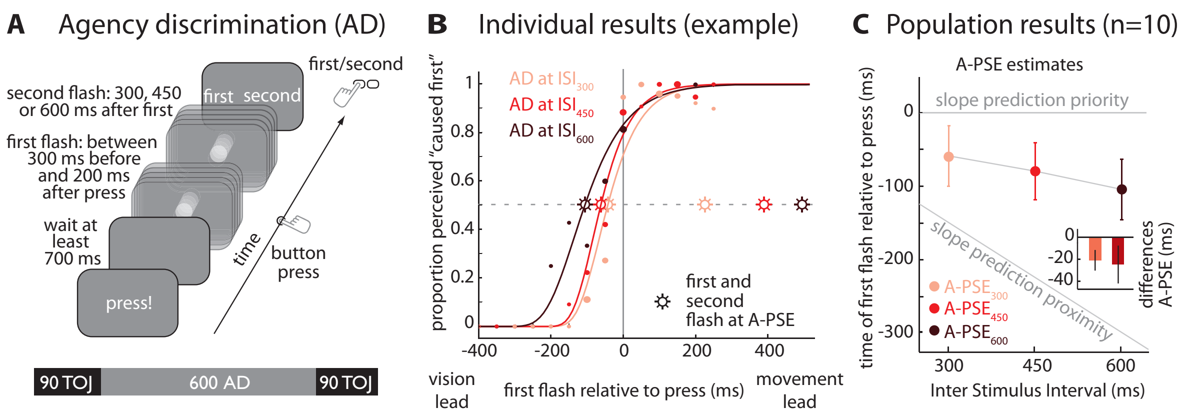


*Figure 3*. Procedure and Results Experiment 2. A: In AD trials, participants have to rate which of two stimuli was caused by their button press. B: Agency discrimination results for an example participant. C: Population results (A-PSE parameter estimates, mean and

confidence intervals). Horizontal grey line: prediction priority principle. Diagonal grey line: prediction proximity principle. The A-PSE for perceived agency for the first flash varies significantly and systematically with the ISI. Inlay: The differences in A-PSE between ISIs are small in magnitude but extremely reliable and present in all participants (bright red: difference $A\text{-}PSE_{450}$ - $A\text{-}PSE_{300}$; dark red: difference $A\text{-}PSE_{600}$ - $A\text{-}PSE_{450}$).

**Analysis.** The first 15 TOJ and AD trials served as practice trials and were not part of the analysis. Trials were excluded if the first flash was more than 425 ms earlier than the press (7% of trials). TOJ responses were analyzed as in Experiment 1.

AD responses were analyzed by modeling them as a log-log function (Eq. 1) of the temporal discrepancy between the first flash and the button press (Fig. 3B). Again, the asymmetrical log-log psychometric function provided a better fit than a symmetrical cumulative Gaussian function (Akaike Information Criterion). This analysis was performed for each participant and ISI to estimate the point of subjective equality for agency (A-PSE), i.e., the point at which participants respond with 50% probability that they had caused the first flash.

Differences in estimates of A-PSE between ISI conditions are given as mean and 95% confidence intervals in the figures and the main text. A weighting of the two competing cues (priority and proximity) was calculated as the proportional deviation from the two predictions. Complementarily, null-hypothesis significance testing was conducted, using a one-way repeated measure ANOVA with ISI (300, 450, 600 ms) as factor (Matlab script anova_rm, Salarian, 2008). This analysis was performed twice. Once, using the relative timing of the press and the first flash as a reference value (this tests for deviations from an exclusive use of the priority principle) and once using the relative timing of press and the midpoint of the ISI as reference (this tests for deviations from an exclusive use of the proximity principle). The resulting p-values were corrected for multiple measurements using the Bonferroni-Holmes method.

**Results**

By modulating the ISI across AD trials, we varied the size of the conflict between the competing proximity and priority cues. Based on the results from Experiment 1 (dominance of priority principle), we predicted that perceived agency would be maximal when the first flash and the button press are perceived at the same time, but would fall off sharply if it occurs earlier. Importantly, the timing of the later second flash should not influence the decision boundary (A-PSE) for having caused the first flash (prediction: horizontal grey line in Fig. 3C). On the other hand, if participants ignored the priority principle and relied on temporal proximity information only to assess agency, this would result in an A-PSE always in the center of the ISI. This prediction is indicated by the diagonal grey line in Fig. 3C. Finally, if A-PSEs fall of with an intermediate slope, this shows a negotiation between the two conflicting temporal cues.

The results (Figure 3B and C) demonstrate a small but highly significant influence of the timing of the second flash on participants' responses (rmANOVA testing for deviation from priority prediction: $F(2,9)=15.2$, $p \ll 0.001$; differences between A-$PSE_{300}$ and A-$PSE_{450}$: 21±9 ms; differences between A-$PSE_{450}$ and A-$PSE_{600}$: 25±17 ms; mean and confidence intervals, Fig. 3C inlay). The later the second flash occurs, the more likely it is that participants report having caused a first flash that occurred before their PSS, i.e., that is perceived to have happened before their button press. This contradicts our hypothesis that the priority principle strictly dominates the SoA. Calculating a weight between priority and proximity cues shows that the influence of the proximity cue on AD is on average 37%. With 63%, priority is the stronger temporal cue for judging agency (rmANOVA testing for deviation from proximity prediction: $F(2,9)=78.4$, $p \ll 0.001$) but it is relaxed on a trial-by-trial basis by 37% relative to the competing temporal information about the later flash. It should be noted that there is of course some inter-subjective variability in these weights. From ten participants the exact magnitude of the weights in the general population cannot be estimated with a high degree of certainty.

Other than in Experiment 1, there was a temporal drift of the PSS (Rohde & Ernst, 2013) over the course of the experiment as measured by the TOJ. That is, the PSS was higher before the experiment than in the end (pre-test PSS: 21±65 ms; post-test PSS: -56±51 ms; difference 77±27 ms; means and confidence intervals). This might have occurred because the distribution of first flashes was biased to the vision-lead side, and less widely distributed in time than the flashes in Experiment 1. To make sure that this drift does not compromise the result, the recalibrated PSS from the end block was used for baseline correction. Even with respect to the shifted post-test PSS, the A-PSEs for agency discrimination were still more negative, which indicates true tolerance for violations of the priority principle, i.e., agency over flashes that occurred perceptibly before the button press (Fig. 3C).

## General Discussion

When making perceptual decisions about agency, participants treat vision-lead temporal discrepancies (violations of cause-effect order) much more conservatively than movement-lead discrepancies (violations of temporal relatedness). This main result from Experiment 1 supports the hypothesis of a strict application of the priority principle in the sense of agency (SoA), if there are no other constraints. With competing cues, however, the priority principle can be relaxed to accommodate other causal cues, as demonstrated in Experiment 2. Participants' probability to report agency for a flash that they perceived to occur before the press depended on how likely they considered it, in terms of temporal proximity, that they caused the later second flash event. They thus tolerate violations of one of the most fundamental laws of physics (i.e., that effect comes before the cause) in the inference of agency to accommodate other, contextual factors associated with a later event. This probabilistic negotiation of temporal cues is performed postdictively and on a trial-by-trial basis. This seems, at least at first glance, not consistent with current two-stage models of agency perception that assume a low-level, bottom-up processing of the temporal coherence of motor predictions and sensory feedback in a first stage.

Bechlivanidis & Lagnado (2013) recently reported a similar result about a modulation of temporal order processing for the perception of external causality. In their study, contextual object information (object movement and collision) provided strong cues in favor of a causal link between events that violate the priority principle. Interestingly, this led participants to readjust their perceptual threshold for simultaneity (PSS) on a trial-by-trial basis: the conflict between the priority principle and the competing contextual cues was resolved by modulating the underlying perception of temporal order. That is, violations of the priority principle were perceptually accommodated so as to avoid experiencing a temporal inversion of cause and effect. The processing of temporal order information for time perception was influenced by causal inference. It is possible that the relaxation of the priority principle we observed here was accompanied by a similar trial-by-trial modulation also of perceived temporal order, as reported by Bechlivanidis & Lagnado (2013). In line with this result, our participants did not report the conscious experience of cause-effect inversions when debriefed, even though the shifts in their perceptual thresholds were well above their thresholds for detecting temporal discrepancies.

It is difficult to draw firm conclusions on the mechanism by which this postdictive reassessment of agency is performed. One possible explanation is that a later cognitive evaluation process overwrites the outcome of the low-level comparison module that is fed forward. Synofzik et al. (2008) have proposed such a two-stage model of agency to explain a partial dissociation between a non-conceptual (low-level) feeling of agency, which is associated with implicit markers such as 'intentional binding' (Haggard et al., 2002), and a conceptual sense of agency that drives explicit agency judgments. In their model, cognitive evaluation of different causal interpretations of a situation in the second stage can overwrite the outcome of the comparator module and cause dissociation between the feeling of agency and the conceptual sense of agency. It is possible that not just agency judgments, but also relative timing judgments can be overwritten in such a way. In this case, implicit measures of agency should not be affected by the postdictive modulation of explicit agency judgments we observed here.

A second possibility is that sensorimotor timing information is integrated over a longer period of time and is recurrently modulated by top-down processes. Neuroscientific evidence increasingly suggests that sensorimotor integration relies on much more complex processes than simple feed-forward processing of motor prediction and sensation. For instance, there appear to be distinct motor prediction processes for open-loop motor planning (forward model) and sensory feedback prediction (feedback model), as well as multiple and recurrent pathways for the processing and use of sensory feedback (for a review and a model see Perruchoud, Murray, Lefebre & Ionta, 2014). In such a framework, a postdictive modulation of perceived temporal order could also occur in a dedicated module for the processing of sensorimotor coherence that integrates information over a longer period of time and receives recurrent feedback from high-level nodes. In this case, implicit measures of agency should also be modulated by the latency of the second flash, just like explicit agency judgments. To settle this question, it would be necessary to test whether implicit measures of agency, such as the spatial analogue of temporal intentional binding (e.g., Beck, Wilke, Wirxel, Endres, Lindner & Giese, 2010), or neural measures, such as motor evoked potentials (e.g., Weiss et al., 2014) are modulated by the timing of later events as we observed it here for explicit agency judgments.

The results from this study show that, in contrast with most current models of agency perception, the mechanisms of the causal inference underlying the human SoA do not process temporal order information separately from and prior to other contextual causal cues. Judging agency involves a probabilistic negotiation of temporally separate visuomotor timing cues that support competing causal interpretations of the particular sensorimotor situation. This negotiation postdictively feeds back into the processing of temporal coherence and sometimes leads to the experience of agency in a situation where effect comes before the cause.

## Acknowledgements

The authors want to thank Alessandro Moscatelli for his help in developing the paradigm and Kielan Yarrow for providing the analysis script on which the MLE fitting routine is based. MR was funded by the DFG grant RO 4147/1-1 ELAPS: Embodied Latency Adaptation and the Perception of Simultaneity. MOE was funded by the FP7/2007-2013 project n8 601165 WEARHAP.

**References**


- Balslev, D., Cole, J., & Miall, R.C. (2007). Proprioception Contributes to the Sense of Agency during Visual Observation of Hand Movements: Evidence from Temporal Judgments of Action. *Journal of Cognitive Neuroscience, 19*(9), 1535–1541.
- Bechlivanidis, C., & Lagnado, D.A. (2013). Does the “Why” Tell Us the “When”? *Psychological Science, 24*(8), 1563-1572. doi: 10.1177/0956797613476046.
- Beck, T. F., Wilke, C., Wirxel, B., Endres, D., Lindner, A., & Giese, M. A. (2011). Me or Not Me: Causal Inference of Agency in goal-directed actions. *Journal of Vision*, *11*(11), 955–955. doi:10.1167/11.11.955.
- Blakemore, S.-J., Wolpert, D.M., & Frith, C.D. (1998). Central cancellation of self-produced tickle sensation. *Nature Neuroscience, 1*(7), 635–640.
- Blakemore, S.-J., Wolpert, D., & Frith, C. (2000). Why can’t you tickle yourself? *Neuroreport, 11*(11), R11–R16.
- 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: 10.1016/j.actpsy.2010.11.005.
- Dewey, J.A., & Carr, T.H. (2013). When dyads act in parallel, a sense of agency for the auditory consequences depends on the order of the actions. *Consciousness & Cognition, 22*(1), 155–166. doi: 10.1016/j.concog.2012.12.004.
- Ebert, J.P., & Wegner, D.M. (2010). Time warp: Authorship shapes the perceived timing of actions and events. *Consciousness & Cognition, 19*(1), 481–489. DOI: 10.1016/j.concog.2009.10.002.

– Farrer, C., Valentin, G., & Hupé, J.M. (2013). The time windows of the sense of agency. *Consciousness & Cognition, 22*(4), 1431–1441. doi: 10.1016/j.concog.2013.09.010.

– Haggard, P., & Chambon, V. (2012). Sense of agency. *Current Biology, 22*(10), R390–R392. doi: 10.1016/j.cub.2012.02.040.

– Haggard, P., Clark, S., & Kalogeras, J. (2002). Voluntary action and conscious awareness. *Nature Neuroscience, 5*(4), 382–385.

– Humphreys, G.R., & Buehner, M.J. (2009). Magnitude estimation reveals temporal binding at super-second intervals. *Journal of Experimental Psychology: Human Perception & Performance, 35*(5), 1542–1549. doi: 10.1037/a0014492.

– Johnson, M. (2010). Mixed (Between/Within Subjects) ANOVA (http://www.mathworks.com/matlabcentral/fileexchange/27080-mixed--between-within-subjects--anova/content/mixed_between_within_anova.m). *Matlab Central File Exchange*. Retrieved August 20, 2013.

– Kawabe, T., Roseboom, W., & Nishida, S. (2013). The sense of agency is action–effect causality perception based on cross-modal grouping. *Proceedings of the Royal Society B: Biological Sciences, 280*(1763), 20130991. doi: 10.1098/rspb.2013.0991.

– Moore, J.W., Wegner, D.M., & Haggard, P. (2009). Modulating the sense of agency with external cues. *Consciousness & Cognition, 18*(4), 1056–1064. doi: 10.1016/j.concog.2009.05.004.

– Perruchoud, D., Murray, M. M., Lefebvre, J., & Ionta, S. (2014) Focal dystonia and the Sensory-Motor Integrative Loop for Enacting (SMILE). *Frontiers in Human Neuroscience, 8*, 458. doi:10.3389/fnhum.2014.00458

- Rohde, M., & Ernst, M.O. (2013). To lead and to lag - forward and backward recalibration of perceived visuomotor simultaneity. *Frontiers in Psychology, 3*, 599. doi: 10.3389/fpsyg.2012.00599.
- Rohde, M., Greiner, L., & Ernst, M.O. (2014). Asymmetries in visuomotor recalibration of time perception: Does causal binding distort the window of integration? *Acta Psychologica, 147*, 127–135. doi: 10.1016/j.actpsy.2013.07.011.
- Salarian, A. (2008). Repeated Measures ANOVA (http://www.mathworks.com/matlabcentral/fileexchange/22088-repeated-measures-anova). *Matlab Central File Exchange*. Retrieved October 12, 2012.
- Sato, A., & Yasuda, A. (2005). Illusion of sense of self-agency: discrepancy between the predicted and actual sensory consequences of actions modulates the sense of self-agency, but not the sense of self-ownership. *Cognition, 94*(3), 241–255.
- 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* (1), 219–239.
- Timm, J., Schönwiesner, M., SanMiguel, I., & Schröger, E. (2014). Sensation of agency and perception of temporal order. *Consciousness & Cognition, 23*, 42–52. doi: 10.1016/j.concog.2013.11.002.
- Ulrich, R. (1987). Threshold models of temporal-order judgments evaluated by a ternary response task. *Perception & Psychophysics, 42*(3), 224–239.
- Wegner, D.M., & Wheatley, T. (1999). Apparent mental causation: Sources of the experience of will. *American Psychologist, 54*(7), 480–492.

- Weiss, C., Tsakiris, M., Haggard, P., & Schütz-Bosbach, S. (2014) Agency in the sensorimotor system and its relation to explicit action awareness. *Neuropsychologia*, *52*, 82-92.
- Yarrow, K., Jahn, N., Durant, S., & Arnold, D.H. (2011). Shifts of criteria or neural timing? The assumptions underlying timing perception studies. *Consciousness & Cognition, 20*, 1518–1531. doi: 10.1016/j.concog.2011.07.003.