Time Warp: Authorship Shapes the Perceived Timing of Actions and Events

Citation	Ebert, Jeffrey P., and Daniel M. Wegner. 2010. Time warp: Authorship shapes the perceived timing of actions and events. Consciousness and Cognition 19(1): 481-489.
Published Version	doi:10.1016/j.concog.2009.10.002
Accessed	February 19, 2015 10:53:07 AM EST
Citable Link	http://nrs.harvard.edu/urn-3:HUL.InstRepos:9967607
Terms of Use	This article was downloaded from Harvard University's DASH repository, and is made available under the terms and conditions applicable to Open Access Policy Articles, as set forth at http://nrs.harvard.edu/urn-3:HUL.InstRepos:dash.current.terms-of-use#OAP

(Article begins on next page)

RUNNING HEAD: Authorship and Time

Time Warp: Authorship Shapes the Perceived Timing of Actions and Events

Jeffrey P. Ebert* and Daniel M. Wegner

Harvard University

*Corresponding author. Address: Department of Psychology, Harvard University, 33 Kirkland St, Cambridge, MA 02143, USA. Email address: jebert@wjh.harvard.edu.

Abstract

It has been proposed that inferring personal authorship for an event gives rise to *intentional binding*, a perceptual illusion in which one's action and inferred effect seem closer in time than they otherwise would (Haggard, Clark, & Kalogeras, 2002). Using a novel, naturalistic paradigm, we conducted two experiments to test this hypothesis and examine the relationship between binding and self-reported authorship. In both experiments, an important authorship indicator – consistency between one's action and a subsequent event – was manipulated, and its effects on binding and self-reported authorship were measured. Results showed that action-event consistency enhanced both binding and self-reported authorship, supporting the hypothesis that binding arises from an inference of authorship. At the same time, evidence for a dissociation emerged, with consistency having a more robust effect on self-reports than on binding. Taken together, these results suggest that binding and self-reports reveal different aspects of the sense of authorship.

Time Warp: Authorship Shapes the Perceived Timing of Actions and Events

A central problem agents face is determining which events they have caused, and which they have not. When a room goes dark, it could be because one flipped the light switch, but it could also be because someone else in the room clapped, or because somewhere else a butterfly flapped its wings (straight into a power line, perhaps). It has been proposed that to solve this problem of authorship, the mind must weigh a variety of authorship indicators and arrive at an appropriate causal inference (Wegner & Sparrow, 2004; Wegner, 2002). An important authorship indicator is the temporal contiguity between one's action and a subsequent event (Shanks, Pearson, & Dickinson, 1989), which is a special case of the more general principle that causality is perceived when one event is followed closely by another (Hume, 1739/1888; Michotte, 1963). When one's flip of the switch is followed immediately by the room's going dark, one can be more certain of having authored the darkness.

It is hardly controversial to say that perceived contiguity enhances perceived authorship, but there is now reason to believe that its less familiar converse is also true: perceived authorship enhances perceived contiguity. Evidence for this hypothesis comes from research on *intentional binding*, a perceptual illusion in which actions and their effects seem to occur especially close in time (Haggard, Clark, & Kalogeras, 2002). In an experiment based on Libet's time judgment paradigm (Libet, Gleason, Wright, & Pearl, 1983), participants were asked to press a key when they felt the urge to do so, and an auditory tone was played shortly after they acted. Participants were supposed to judge, in separate blocks, the time of either their keypress or the subsequent tone, by referring to a hand that rotated around a clock face. The main results showed that participants' judgments of when their action occurred were shifted forward in time (toward the tone), whereas judgments of when the tone occurred were shifted backward (toward the action),

relative to judgments made in baseline blocks in which action and tone occurred alone. Critically, these perceptual shifts did not occur when transcranial magnetic stimulation (TMS) induced the participant to make an involuntary keypress, suggesting that voluntary action is needed. Haggard et al. (2002) concluded that the perceived shifts in time result from the mind's attempt to construct a coherent conscious experience of our own agency, by binding our intentional actions to their effects.

Several experiments have provided additional support for the hypothesis that binding arises when personal authorship for an event is inferred. Some have focused on the "personal" aspect of this hypothesis, showing that binding occurs for one's own, voluntary actions, but not for involuntary movements (Engbert, Wohlschläger, & Haggard, 2008; Haggard & Clark, 2003) or the actions of others (Engbert et al., 2008; Engbert, Wohlschläger, Thomas, & Haggard, 2007). Others have focused on the "authorship" aspect, showing that the presence of authorship indicators (Moore, Wegner, & Haggard, in press) or the provision of control over events (Cravo, Claessens, & Baldo, 2009; Moore, Lagnado, Deal, & Haggard, 2008) enhances binding (see also Stetson, Cui, Montague, & Eagleman, 2006).

Still, many questions remain about the link between authorship and binding. Previous investigations into binding have not included a corroborating self-report measure of authorship, and without such a measure it is not possible to know for sure whether binding occurs alongside the experience of authorship, increasing and decreasing in response to the same indicators. The lack of a self-report measure has also precluded research into dissociations between what might be considered relatively implicit (binding) and explicit (self-report) measures of authorship (see Synofzik, Vosgerau, & Newen, 2008). One possible dissociation is that people readily report authorship for events that occur several seconds after their actions (Shanks et al., 1989) – and

indeed can claim responsibility for an event taking place years after their action – yet binding has been found only for action-event delays lasting a fraction of a second (Haggard et al., 2002; Stetson et al., 2006). Another open question is whether binding occurs for the kinds of actions performed, and events encountered, in everyday life, such as kicking a ball and watching it fly away. Nearly all previous investigations have based their methods on the Libet paradigm, asking participants to press a key and judge either the time of their keypress or a subsequent tone, which raises concerns about the external validity of the phenomenon.

We conducted two experiments to address validity concerns and explore the relationship between binding and self-reported authorship. Both experiments employed a novel "push/pull" paradigm that was designed to be relatively naturalistic. Rather than press a key, participants pushed or pulled on a joystick in response to images of everyday objects; rather than hear a tone, participants saw these objects move away or come closer. Unlike a keypress, the acts of pushing and pulling are imbued with bodily significance, implying both an orientation toward whatever is acted upon and an expected outcome of that action. Specifically, pushing corresponds to an avoidance orientation and is undertaken with the expectation that an object will go away, whereas pulling corresponds to an approach orientation and is undertaken with the expectation that an object will come closer (Cacioppo, Priester, & Berntson, 1993; Chen & Bargh, 1999). In both experiments, we manipulated whether an event was consistent with the participant's prior action (i.e., whether the object moved in the same direction as the action), both because action-event consistency is an important authorship indicator (Wegner & Sparrow, 2004), and because previous research had not directly examined its effect on binding. In addition to measuring binding, we asked participants to report their feelings of authorship. Our main hypothesis was

that action-event consistency would enhance both self-reported authorship and intentional

binding – but we were also interested in any differences the two measures might reveal.

<div align="center">Experiment 1</div>

Participants completed a series of trials on which they saw a picture of an everyday

object (e.g., an apple) and chose whether to push or to pull on a joystick in response. This action

was followed by a brief delay, after which the object appeared to move in a direction either

consistent or inconsistent with the action (the event). Participants then estimated the length of the

delay between their action and the object's movement, and these interval estimates served as the

measure of binding. Interval estimation has been used successfully in other investigations of

binding (Cravo et al., 2009; Engbert et al., 2008; Engbert et al., 2007; Moore et al., in press), and

this method was preferred because estimating intervals does not require participants to pay

attention to a clock, thus freeing them to pay attention to their actions and subsequent events. For

the measure of authorship, participants reported the degree to which they felt that their action had

caused the object to move.

Method

Participants. Twenty-five participants[1] completed the experiment after signing up

through the Harvard Study Pool website, which requested that only college students participate.

Participants were compensated with either $5 or 0.5 hours of credit.

Procedure. Participants were seated at a desk with an LCD monitor (set at 1024×768

resolution and a 60-Hz refresh rate), which was connected to a Dell OptiPlex computer with a

3.0 GHz Intel Pentium processor. The experiment was run in DirectRT 2006. A SideWinder

Precision Pro Joystick was mounted in front of the screen. Participants were asked to use their

dominant hand for pushing and pulling the joystick. To facilitate response entry for the

---

[1] Participant sex was not recorded for this experiment.

dependent measures, a keyboard was positioned to the left of the joystick for right-handed participants and to the right for left-handed participants.

First, participants completed a series of 12 practice trials on which they were cued either to push or to pull the joystick and asked to estimate the amount of time that passed between when they began to perform this action and when a flash appeared. Prior to the practice trials, participants were familiarized with the meaning of the term "millisecond." For instance, they were told that there are 1,000 milliseconds in one second, and that an eyeblink takes around 200 milliseconds. Each practice trial began with text prompting the participant how to act (e.g., "PULL joystick when ready"). When the participant acted, the screen went completely black for a delay of between 0 ms and 900 ms (including all and only integer multiples of 100 ms: 0 ms, 100 ms, 200 ms,…, 900 ms), after which a white flash appeared in the middle of the screen for 34 ms. Participants were asked to estimate the duration of this delay on a 10-point scale, ranging from 0 ms to 900 ms, in 100-ms increments. After they entered each estimate, participants were told the actual delay for that trial. These practice trials served multiple purposes. First, it was hoped that they would improve participants' comfort with, and accuracy at, estimating very brief intervals. Second, by exposing participants to 10 different delays, it was hoped that they would later be more likely to believe that the critical trials of the experiment also involved those 10 delays – whereas, in keeping with previous binding research and to facilitate comparison between conditions, only three different delays were used. Third, it was anticipated that performance on the practice trials could be used to screen for individuals who had difficulty estimating brief intervals or who did not take the task seriously.

Participants then completed 60 critical trials. On each trial, they were supposed to choose between pushing and pulling the joystick in response to a picture of an everyday object on the

screen. The 60 objects, which included such items as an apple, baseball, and pencil, were chosen because, in real life, each would be graspable and movable with a push or pull. Participants were instructed to try to have a 50-50 split between pushing and pulling – but to avoid simply alternating between the two. Otherwise, no instruction was given about how they should decide which action to perform. When the participant began to move the joystick, a delay of 100 ms, 400 ms, or 700 ms was initiated (Moore et al., in press), and for the duration of this delay the object remained stationary on the screen. After the delay, the object appeared to move either away from or toward the participant – half the time in the direction that was consistent with the participant's action, half the time in the direction that was inconsistent.

This apparent motion was accomplished with animations created in Macromedia Flash MX 2004. For each object, two 0.5-s animations were created, both beginning with an object identical in size to the still image shown to participants. In one, the object grew 625% in area over the course of 12 frames (for a frame rate of 24 fps), making it appear to come toward the participant; in the other, the object shrank an equivalent amount, to 16% of its initial area, making it appear to go away. After seeming to move for 0.5 s, the object remained still at its terminal position for an additional 0.5 s before the screen was cleared.

Participants then completed the binding and authorship measures. First they were asked: "How much time passed between when you first moved the joystick and when the object first started to move?" They gave their interval estimates on the same 10-point scale used for practice trials, with possible responses ranging from 0 ms to 900 ms, in 100-ms increments. Then they were asked to "indicate how much it felt like [their] moving the joystick caused the object on the computer screen to move." They reported their feelings of authorship on a 7-point rating scale anchored at 1, 4, and 7 with "not at all," "somewhat," and "very much." Because binding was of

primary interest and binding effects are relatively small, the binding measure was administered first to reduce the risk of interference from the authorship measure.

At the end of the session, participants were asked a question gauging their awareness of the predicted binding effect: "Do you feel that the time between when you moved the joystick and when the object moved was briefer when the object moved in the SAME direction as the joystick, briefer when the object moved OPPOSITE the direction of the joystick – or do you feel it made NO DIFFERENCE?" Then participants were debriefed and compensated for their participation.

The experiment had a 2 (action-event consistency: consistent, inconsistent) × 3 (delay: 100 ms, 400 ms, 700 ms) within-participants design. Which of the 60 objects appeared on consistent trials was counterbalanced across participants.

Results

Preliminary analyses. Overall, participants exhibited a high degree of accuracy in estimating intervals on the practice trials, and the mean correlation between estimated and actual intervals (*M* Fisher z = .73) was significantly greater than zero, $t(24) = 6.92$, $p < .001$. For three participants this correlation was negative, suggesting unusual difficulty with the interval estimation task, and these individuals were excluded from subsequent analyses. The remaining participants pulled in response to 64.1% of objects, a rate significantly higher than 50%, $t(21) = 4.57$, $p < .001$; however, participants' rates of pulling did not correlate with any of the significant effects of consistency reported for either dependent variable at any of the three delays, all ps > .289.

Data from the primary dependent measures were prepared for each participant by computing the means of the 10 interval estimates and 10 authorship ratings made in each of the

six cells of the design. Fig. 1a shows the means and standard errors for each of these interval estimate means, and Fig. 1b shows the means and standard errors for each of these authorship rating means.

Binding. A Consistency × Delay repeated-measures ANOVA was conducted on participants' mean interval estimates. The main effect of consistency was significant and in the predicted direction, with briefer interval estimates for consistent ($M = 413$ ms) than inconsistent ($M = 452$ ms) trials, $F(1, 21) = 9.19$, $p = .006$. Confirming that participants' estimates of delays were generally accurate, the main effect of delay was highly significant, $F(2, 42) = 101.51$, $p < .001$, with briefer estimates ($M = 272$ ms) at a delay of 100 ms, intermediate estimates ($M = 438$ ms) at 400 ms, and longer estimates ($M = 588$ ms) at 700 ms. The Consistency × Delay interaction was also significant, $F(2, 42) = 3.25$, $p = .049$. To examine this interaction, paired *t*-tests comparing consistent and inconsistent trials were carried out at each of the three delays. The effect of consistency was significant at delays of 100 ms (*M* difference = 48 ms), $t(21) = 3.43$, $p = .003$, and 400 ms (*M* difference = 58 ms), $t(21) = 2.86$, $p = .009$, but not at 700 ms (*M* difference = 12 ms), $t(21) = 0.73$, $p = .472$.

Were participants aware of the main effect of consistency? In response to the awareness question, eight participants reported that the delay felt briefer for consistent trials than for inconsistent, 12 reported no difference, and only one reported that the delay felt briefer for inconsistent trials.[2] This finding – an 8-to-1 ratio of participants who thought that delays were briefer for consistent, rather than inconsistent, trials – is interesting in itself, as it suggests that the effect of consistency on binding might have been great enough to be noticed by participants during the course of the procedure. But one might also wonder whether awareness was necessary for obtaining the effect of consistency. To address this question, a binding effect was computed

---

[2] Data for one participant were missing because the first participant was not asked this awareness question.

for each participant by subtracting mean interval estimates on consistent trials from mean estimates on inconsistent trials, and these binding effects were regressed on participants' reported awareness.[3] Not only was awareness unrelated to binding effects, $t(20) = -0.97$, $p = .345$, but the intercept of the model was positive and significant, $t(20) = 2.84$, $p = .010$, suggesting that the effect of consistency on binding would obtain even in the absence of any awareness (cf. Greenwald, Klinger, & Schuh, 1995).

Authorship. A Consistency $\times$ Delay repeated-measures ANOVA was conducted on participants' mean authorship ratings. The main effect of consistency was significant and in the predicted direction, with greater authorship reported for consistent ($M = 4.10$) than inconsistent ($M = 3.35$) trials, $F(1, 21) = 21.47$, $p < .001$. Confirming that temporal contiguity is an important authorship indicator, the main effect of delay was highly significant, $F(2, 42) = 21.07$, $p < .001$, with higher authorship ratings ($M = 4.24$) at a delay of 100 ms, intermediate ratings ($M = 3.72$) at 400 ms, and lower ratings ($M = 3.22$) at 700 ms. The Consistency $\times$ Delay interaction was also significant, $F(2, 42) = 7.56$, $p = .002$. To examine this interaction, paired t-tests comparing consistent and inconsistent trials were carried out at each of the three delays. Despite the significant interaction, the effect of consistency was significant at all three delays: 100 ms (M difference = 1.08), $t(21) = 5.95$, $p < .001$, 400 ms (M difference = 0.55), $t(21) = 2.84$, $p = .010$, and 700 ms (M difference = 0.60), $t(21) = 3.53$, $p = .002$.

Relationship between authorship and binding. To examine the relationship between consistency's effect on self-reported authorship and its effect on binding, an authorship effect was computed for each participant by subtracting mean authorship ratings on inconsistent trials from mean ratings on consistent (the opposite of the computation for the binding effect described earlier). The correlation between authorship and binding effects was positive and significant,

---

[3] Briefer when inconsistent was coded as -1, no difference as 0, and briefer when consistent as 1.

$r(21) = .45$, $p = .037$, indicating that participants who showed the greatest effect of consistency on self-reported authorship also showed the greatest effect on binding.

Discussion

These results support the hypothesis that intentional binding arises when personal authorship for an event is inferred. An important indicator of authorship, action-event consistency, was found to increase both self-reported authorship and binding, and these two effects were positively correlated. Action-event consistency had these effects even though, by design, how an object moved was not contingent upon how the participant acted, and objects were no more likely to move in a way that was consistent with the participant's action than they were to move in a way that was inconsistent – suggesting that the effect of consistency was driven by participants' prior expectations about how objects should react to pushing and pulling (Cacioppo et al., 1993; Chen & Bargh, 1999). It should be noted that the present test for binding was inherently relative, comparing interval estimates for consistent trials to those for inconsistent trials. These results thus do not exclude the possibility that participants experienced some amount of binding on inconsistent trials; after all, even on inconsistent trials, participants chose when and how to perform a voluntary action that caused an event to occur soon afterward.[4] What these results do show is that action-event consistency can *enhance* binding, which is especially noteworthy because consistency cannot be determined until after an event occurs, and only then by comparing the event to what one expected when performing the action – confirming that both predictive and retrospective authorship indicators affect binding (Moore & Haggard, 2008).

The present results also establish that intentional binding can occur in a more naturalistic paradigm than has been used in previous research. Participants performed the everyday actions

---

[4] Alternatively, participants may have experienced temporal repulsion between their actions and inconsistent events, similar to that reported by Haggard et al. (2002) for involuntary actions.

of pushing and pulling and then observed events that typically result from such actions: objects going away or coming closer. These results not only indicate that binding generalizes beyond the Libet paradigm, but also lend credence to the external validity of the phenomenon. They also confirm the usefulness of the push/pull paradigm for investigations of binding and suggest that future research might profitably employ the paradigm as an implicit measure of the sense of authorship (Moore et al., in press).

At the same time, this experiment revealed important differences between binding and participants' self-reports. First, consistency had a greater effect on self-reported authorship than on binding. Perhaps the process of making an explicit judgment about authorship leads to more rule-based consideration of authorship indicators (see Smith & DeCoster, 1999), or the mere act of conscious reflection serves to amplify the importance of salient indicators. Second, whereas consistency was found to enhance authorship at all three delays that were tested, it was found to enhance binding only at the two briefest delays. The latter findings are in keeping with previous research that has found evidence for binding only at brief delays between action and effect (Haggard et al., 2002; Stetson et al, 2006). One possible, albeit speculative, explanation for this feature of binding is that there exists a narrow window of a few hundred milliseconds during which one event (in this case, the effect) may influence the conscious perception of a prior event (one's action) (Choi & Scholl, 2006; Libet, Wright, Feinstein, & Pearl, 1979). In sum, although Experiment 1 showed that both binding and self-reported authorship are sensitive to the same authorship indicator, it also provided evidence for a dissociation between the two.

Experiment 2

The present experiment was conducted to verify the results of the first experiment. To provide purer measures of binding and authorship, the experiment was divided into two parts: a

block in which only the binding measure was administered, followed by a block in which only the authorship measure was administered. Removing the authorship measure from the first block carried the risk that participants would not consider the degree to which their actions were causing the events to occur, or even that they would fail to notice whether events were consistent with their actions, thereby decreasing the likelihood of finding an effect of consistency on binding. Two precautions were taken to reduce this risk. First, participants were explicitly given the goal of trying to make objects move in the same direction as their actions, a goal that should promote binding for consistent, but not inconsistent, events (Engbert & Wohlschläger, 2007). Second, on a random subset of trials participants were asked to indicate whether the object had moved in the same direction as their action; success on these catch trials necessitated paying attention to action-event consistency. Though asking participants to have a goal while acting and administering catch trials may strike some as artificial, we believed that their inclusion would help the procedure more closely approximate what happens in the real world when people perform actions, because actions are typically undertaken with a goal in mind, and their effects are evaluated with respect to this goal (Miller, Galanter, & Pribram, 1960).

Another procedural change is that participant sex was recorded. Although previous research has not suggested that there exist sex differences in binding, it was of interest whether participant sex would moderate any effects of action-event consistency.

Method

Participants. Fifty-four participants (31 female, 23 male) completed the experiment. As with the first experiment, it was requested that only college students participate, and participants were compensated with either $5 or 0.5 hours of credit.

Procedure. Participants completed a series of practice trials like those in the first experiment. For the critical trials, participants were instructed that their goal was to push or pull on the joystick and try to make the objects move in the same direction as their action, and that, just as in real life, they should push on the joystick if they wanted an object to go away and pull if they wanted it to come closer. As in the first experiment, they were asked to maintain a 50-50 split between pushing and pulling.

Participants completed a block of trials in which only the binding measure was administered, followed by a block in which only the authorship measure was administered. Each block consisted of 66 trials, six of which were catch trials. Participants were informed about the presence of these catch trials in advance ("on a few trials, you will be asked a different question: whether or not the object moved in the same direction as your action"). Aside from the presence of catch trials and the administration of only one dependent measure per block, trials were structured as they were in the first experiment. After completing both blocks, participants were asked a question gauging their awareness of the predicted binding effect, and then debriefed and compensated.

The experiment had a 2 (action-event consistency: consistent, inconsistent) × 3 (delay: 100 ms, 400 ms, 700 ms) within-participants design, with the binding and authorship measures administered in separate blocks. Which of the 132 objects appeared on consistent trials was counterbalanced across participants.

Results

Preliminary analyses. Overall, participants exhibited a high degree of accuracy in estimating intervals on the practice trials, and the mean correlation between estimated and actual intervals (*M* Fisher z = .62) was significantly greater than zero, $t(53) = 8.89$, $p < .001$. For seven

participants this correlation was negative, and as in the first experiment these individuals were excluded from subsequent analyses. The remaining participants pulled in response to 52.8% of objects during the binding block, a rate that did not differ significantly from 50%, $t(46) = 1.71$, $p = .094$, and to 54.5% of objects during the authorship block, which did differ significantly from 50%, $t(46) = 2.80$, $p = .007$; however, participants' rates of pulling in the authorship block did not correlate with their effects of consistency on authorship ratings, $r(45) = .05$, $p = .746$. Participants performed significantly above chance on the catch trials, both in the binding block – 93.6% correct, $t(46) = 18.12$, $p < .001$ – and in the authorship block – 95.0% correct, $t(46) = 19.91$, $p < .001$. Even if only those participants who responded correctly to all 12 catch trials (n = 31) were included in subsequent analyses, key results would remain statistically significant, so in the reported analyses no participants were excluded on the basis of catch trial performance. Participant sex was not found to interact with the consistency manipulation for either dependent variable, both ps > .150, and was not included in the reported analyses.

Data from the primary dependent measures were prepared for each participant by computing the means of the 10 interval estimates and 10 authorship ratings made in each of the six cells of the design. Fig. 2a shows the means and standard errors for each of these interval estimate means, and Fig. 2b shows the means and standard errors for each of these authorship rating means.

Binding. A Consistency × Delay repeated-measures ANOVA was conducted on participants' mean interval estimates. Replicating a key result of the first experiment, the main effect of consistency was significant and in the predicted direction, with briefer interval estimates for consistent ($M = 422$ ms) than inconsistent ($M = 435$ ms) trials, $F(1, 46) = 5.69$, $p = .021$. Confirming that participants' estimates of delays were generally accurate, the main effect of

delay was highly significant, $F(2, 92) = 89.12$, $p < .001$, with briefer estimates ($M = 338$ ms) at a delay of 100 ms, intermediate estimates ($M = 428$ ms) at 400 ms, and longer estimates ($M = 521$ ms) at 700 ms. The Consistency × Delay interaction was not significant, $F(2, 92) = 0.32$, $p = .726$, but paired t-tests at each of the delays suggested that the effect of consistency was strongest (M difference = 21 ms) at a delay of 100 ms, $t(46) = 1.82$, $p = .075$, intermediate (M difference = 12 ms) at 400 ms, $t(46) = 1.13$, $p = .265$, and weakest (M difference = 7 ms) at 700 ms, $t(46) = 0.56$, $p = .582$.

In response to the awareness question, 14 participants reported that the delay felt briefer for consistent trials than for inconsistent, 28 reported no difference, and only five reported that the delay felt briefer for inconsistent trials. As was done in the first experiment, a binding effect was computed for each participant, and these effects were regressed on reported awareness. Not only was awareness unrelated to binding effects, $t(46) = 0.51$, $p = .612$, but the intercept of the model was positive and significant, $t(46) = 2.10$, $p = .041$, suggesting that the effect of consistency on binding would obtain even in the absence of any awareness.

Authorship. A Consistency × Delay repeated-measures ANOVA was conducted on participants' mean authorship ratings. Replicating a key result of the first experiment, the main effect of consistency was significant and in the predicted direction, with greater authorship reported for consistent ($M = 4.77$) than inconsistent ($M = 3.01$) trials, $F(1, 46) = 59.18$, $p < .001$. Confirming once again that temporal contiguity is an important authorship indicator, the main effect of delay was highly significant, $F(2, 92) = 28.78$, $p < .001$, with higher authorship ratings ($M = 4.34$) at a delay of 100 ms, intermediate ratings ($M = 3.87$) at 400 ms, and lower ratings ($M = 3.47$) at 700 ms. The Consistency × Delay interaction was not significant, $F(2, 92) = 2.13$, $p = .124$. Not surprisingly, the effect of consistency was significant at all three delays: 100 ms (M

difference = 1.87), $t(46) = 7.59$, $p < .001$, 400 ms (M difference = 1.75), $t(46) = 7.27$, $p < .001$, and 700 ms (M difference = 1.66), $t(21) = 7.53$, $p < .001$.

Relationship between authorship and binding. As in the first experiment, an authorship effect was computed for each participant, and the relationship between these effects and participants' binding effects was examined. Unlike in the first experiment, these effects were not correlated, $r(46) = -.05$, $p = .744$.

Discussion

These results replicate the main effects of action-event consistency on binding and authorship found in Experiment 1, further supporting the hypothesis that binding arises when personal authorship is inferred. These effects were found even though the binding and authorship measures were administered in separate blocks – demonstrating that consistency can affect binding even when individuals are not asked to judge whether they authored an event. Participant sex was not found to moderate the effect of consistency on either dependent variable, which suggests that males and females are equally sensitive to this indicator of authorship. As in the first experiment, compared to its effect on binding, consistency's effect on self-reported authorship was stronger and relatively unaffected by the delay between action and event, supporting the idea that binding and self-reported authorship assess two dissociable modes of authorship (Moore et al., in press).

The present results differed from what was found in the first experiment in two important ways: consistency had a weaker effect on binding in the present experiment, and, unlike before, this binding effect did not correlate with the effect of consistency on self-reported authorship. In both cases, the separate administration of the binding and authorship measures in the present experiment may explain the difference. During the binding block, participants were not required

to pay attention to their feelings of authorship, which may have reduced the amount of binding experienced. Conversely, during the authorship block, participants did not make interval estimates prior to reporting their authorship, which may have decreased the weight given to perceived delay when participants rated their authorship. Indeed, examination of the authorship results for both experiments suggests that, compared to participants in the first experiment (Fig. 1b), participants in the present experiment (Fig. 2b) based their authorship ratings less on perceived delay and more on action-event consistency. The less individuals base their explicit authorship judgments on perceived delay, the weaker the correlation between binding and these judgments should be.

General Discussion

The present research tested the hypothesis that intentional binding, in which one's action and a subsequent event seem to occur especially close in time (Haggard et al., 2002), arises when personal authorship for the event is inferred. Providing support for this hypothesis, two experiments found that an important indicator of authorship – consistency between one's action and a subsequent event – enhanced both binding and self-reported authorship. Both experiments employed a relatively naturalistic paradigm, in which action-event sequences resembled those that participants might encounter in their everyday lives – pushing or pulling for a graspable object, and watching this object move away or come closer – boosting the case that intentional binding is an externally valid phenomenon.

It is intriguing that such disparate measures of the experience of authorship – asking participants to judge the length of the delay between action and event, versus asking them to rate the degree to which it felt like they caused the event – could yield similar results. At the same time, important differences between these measures emerged. Action-event consistency affected

self-reported authorship more than it did binding, and this difference was particularly stark at longer delays between action and event. Furthermore, when the measures were administered in separate blocks, the degree to which consistency affected a participant's self-reported authorship was unrelated to the degree to which it affected his or her binding. Taken together, these results suggest that binding and self-report measures may be capturing different aspects of the sense of authorship, with the former measuring an implicit, lower-level, reflexive feeling, and the latter measuring an explicit, higher-level, reflective judgment (Synofzik et al., 2008).

The possibility that binding may be used as an implicit measure of authorship (Moore et al. in press) holds promise for future research. Researchers could use binding to assess feelings of authorship that participants may be unwilling or unable to report (cf. Greenwald & Banaji, 1995). With proper validation, binding could prove particularly useful for research with clinical populations known to have distorted feelings of authorship (Haggard, Martin, Taylor-Clarke, Jeannerod, & Franck, 2003). In particular, by manipulating various authorship indicators and measuring their effect on binding, researchers might be able to determine which indicators these populations overweight or overlook when making low-level inferences about authorship (Ebert & Wegner, in press).

To be sure, more research is needed to fully understand the relationship between authorship indicators, intentional binding, and self-reported authorship. Nevertheless, the present findings add to accumulating evidence that tells the same story: The mind relies on a variety of indicators to make inferences about authorship, and inferring personal authorship for an event warps one's perception of time, such that one's action and inferred effect seem temporally closer than they otherwise would.

References

Cacioppo, J. T., Priester, J. R., & Berntson, G. G. (1993). Rudimentary determinants of attitudes: II. Arm flexion and extension have differential effects on attitudes. *Journal of Personality and Social Psychology, 65*, 5-17.

Chen, M., Bargh, J. A. (1999). Consequences of automatic evaluation: Immediate behavioral predispositions to approach or avoid the stimulus. *Personality and Social Psychology Bulletin, 25*, 215-224.

Choi, H., & Scholl, B. J. (2006). Perceiving causality after the fact: Postdiction in the temporal dynamics of causal perception. *Perception, 35*, 385-399.

Cravo, A. M., Claessens, P. M. E., & Baldo, M. V. C. (2009). Voluntary action and causality in temporal binding. *Experimental Brain Research, 199*, 95-99.

Ebert, J. P., & Wegner, D. M. (in press). Bending time to one's will. In W. Sinnott-Armstrong (Ed.), *Conscious will and responsibility: a tribute to Benjamin Libet*.

Engbert, K., & Wohlschläger, A. (2007). Intentions and expectations in temporal binding. *Consciousness and Cognition, 16*, 255-264.

Engbert, K., Wohlschläger, A., & Haggard, P. (2008). Who is causing what? The sense of agency is relational and efferent-triggered. *Cognition, 107*, 693-704.

Engbert, K., Wohlschläger, A., Thomas, R., & Haggard, P. (2007). Agency, subjective time, and other minds. *Journal of Experimental Psychology: Human Perception and Performance, 33*, 1261-1268.

Greenwald, A. G., & Banaji, M. R. (1995). Implicit social cognition: Attitudes, self-esteem, and stereotypes. *Psychological Review, 102*, 4-27.

Greenwald, A. G., Klinger, M. R., & Schuh, E. S. (1995). Activation by marginally perceptible ("subliminal") stimuli: Dissociation of unconscious from conscious cognition. *Journal of Experimental Psychology: General, 124*, 22-42.

Haggard, P., & Clark, S. (2003). Intentional action: Conscious experience and neural prediction. *Consciousness and Cognition, 12*, 695-707.

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

Haggard, P., Martin, F., Taylor-Clarke, M., Jeannerod, M., & Franck, N. (2003). Awareness of action in schizophrenia. *Cognitive Neuroscience and Neuropsychology, 14*, 1081-1085.

Hume, D. (1888). *A treatise on human nature* (L. A. Selby-Bigge, Ed.). London: Oxford University Press, (original work published 1739)

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.

Libet, B., Wright, E. W., Feinstein, B., & Pearl, D. (1979). Subjective referral of the timing for a conscious sensory experience: A functional role for the somatosensory specific projection system in man. *Brain, 102*, 193-224.

Michotte, A. (1963). *The perception of causality* (T. R. Miles & Elaine Miles, Trans.). New York: Basic Books.

Miller, G. A., Galanter, E., & Pribram, K.A. (1960). *Plans and the structure of behavior.* New York: Holt, Rhinehart, & Winston.

Moore, J., & Haggard, P. (2008). Awareness of action: inference and prediction. *Consciousness and Cognition, 17*, 136-144.

Moore, J. W., Lagnado, D., Deal, D. C., & Haggard, P. (2008). Feelings of control: Contingency determines experience of action. *Cognition, 110*, 279-283.

Moore, J. W., Wegner, D. M., & Haggard, P. (in press). Modulating the sense of agency with external cues. *Consciousness and Cognition*.

Shanks, D. R., Pearson, S. M., & Dickinson, A. (1989). Temporal contiguity and the judgement of causality by human subjects. *The Quarterly Journal of Experimental Psychology Section B, 41*, 139-159.

Smith E. R., & DeCoster, J. (1999). Associative and rule-based processing: A connectionist interpretation of dual process models. In S. Chaiken & Y. Trope (Eds.), *Dual process theories in social psychology* (pp. 323-336). New York: Guilford Press.

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*, 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.

Wegner, D. M. (2002). *The illusion of conscious will*. Cambridge, MA: MIT Press.

Wegner, D. M., & Sparrow, B. (2004). Authorship processing. In M. Gazzaniga (Ed.), *The cognitive neurosciences III* (pp. 1201-1209). Cambridge, MA: MIT Press.
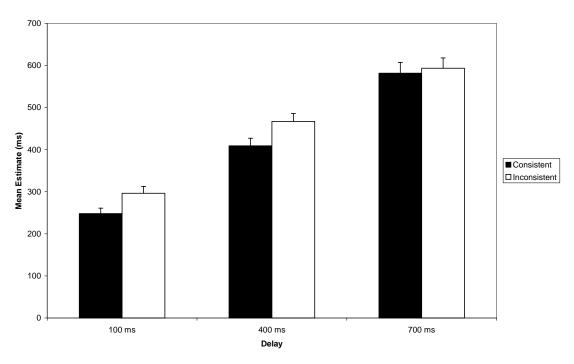
Author Note

Figure Captions

Figure 1. Displayed are a) mean interval estimates and b) mean authorship ratings by consistency and delay in Experiment 1. Error bars show SE across participants.
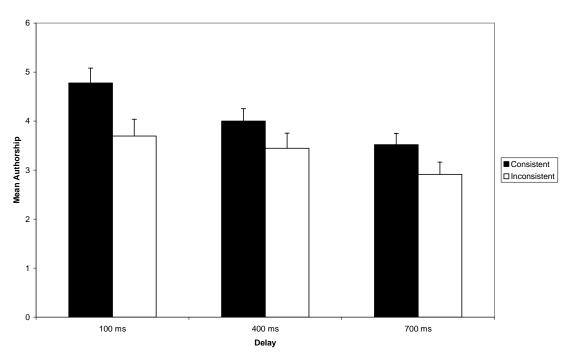
Figure 2. Displayed are a) mean interval estimates and b) mean authorship ratings by consistency and delay in Experiment 2. Error bars show SE across participants.
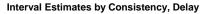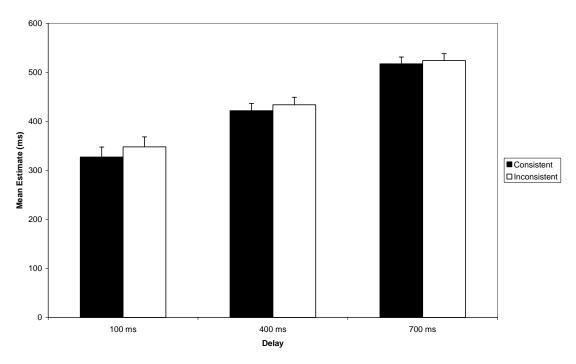
Figures 1a and 1b

Interval Estimates by Consistency, Delay

Self-Reported Authorship by Consistency, Delay

Figures 2a and 2b

Interval Estimates by Consistency, Delay

Self-Reported Authorship by Consistency, Delay