

Understanding the Past, Predicting the Future: Causation, Not Intentional Action, Is the Root of Temporal Binding

Marc J. Buehner

Cardiff University

Psychological Science
 23(12) 1490–1497
 © The Author(s) 2012
 Reprints and permission:
sagepub.com/journalsPermissions.nav
 DOI: 10.1177/0956797612444612
<http://pss.sagepub.com>


Abstract

Temporal binding refers to a subjective shortening of elapsed time between actions and their resultant consequences. Originally, it was thought that temporal binding is specific to motor learning and arises as a consequence of either sensory adaptation or the associative principles of the forward model of motor command. Both of these interpretations assume that the binding effect is rooted in the motor system and, critically, that it is driven by intentional action planning. The research reported here demonstrates that both intentional actions and mechanical causes result in temporal binding, which suggests that intentional action is not necessary for temporal binding and that binding results from the causal relation linking actions with their consequences. Intentional binding is thus a special case of more general causal binding, which can be explained by a theory of Bayesian ambiguity reduction.

Keywords

causality, perceptual motor coordination, time perception

Received 6/20/11; Revision accepted 3/4/12

Correlation famously does not imply causation. Both allow people to predict the future, but only the latter affords explanation of the past (Cheng, 1997; Sloman, 2005). Understanding why things turned out the way they did allows humans to plan interventions in a flexible and goal-directed way (Buehner & Cheng, 2005; Pearl, 2000; Sloman, 2005). Previous research has shown that such goal-directed instrumental action results in *intentional* or *temporal binding*: Actions are tied to their consequences such that the perceived times of both are attracted together in conscious experience, with the outcomes of instrumental actions being perceived as occurring earlier than they actually did (Haggard, Clark, & Kalogeras, 2002).

This effect is robust and has been observed in various laboratories across a multitude of methods (Buehner & Humphreys, 2009; Engbert & Wohlschläger, 2007; Engbert, Wohlschläger, Thomas, & Haggard, 2007; Humphreys & Buehner, 2009, 2010; Moore & Haggard, 2007; Wohlschläger, Haggard, Gesierich, & Prinz, 2003). However, whether the effect reflects sensory adaptation (Kennedy, Buehner, & Rushton, 2009; Stetson, Cui, Montague, & Eagleman, 2006), an association-based advantage of a forward model of motor command (Haggard, Aschersleben, Gehrke, & Prinz, 2002; Haggard, Clark, & Kalogeras, 2002), a specific awareness of self- and other-generated actions via a mirror-neuron system (Wohlschläger, Haggard, et al., 2003), or is grounded in the Bayesian principles underlying causal

inference (Buehner, 2010; Eagleman & Holcombe, 2002) is still debated. The results reported here clearly support the latter interpretation: Event anticipation rooted in higher-level causal understanding selectively afforded early response to a target stimulus relative to expectation based on mere correlation or association—and this early response arose irrespective of whether the cause was an intentional action or a mechanical event.

The causal interpretation of temporal binding (Buehner, 2010; Buehner & Humphreys, 2009, 2010; Eagleman & Holcombe, 2002) is rooted in Hume's (1777/1888) treatise of causation: Causal relations cannot be perceived directly but are instead inferred from observable qualities, such as contingency and temporal and spatial contiguity—the closer two events follow one another in time and space, the more likely they are to be causally related, especially if the pairing occurs frequently. According to Bayes's theorem, this means that two causally related events are therefore also more likely to follow each other closely in space and time than two unrelated events are. Given

Corresponding Author:

Marc J. Buehner, Cardiff University, School of Psychology, Park Place,
 Cardiff CF10 3AT, Wales, United Kingdom
 E-mail: buehnerm@cardiff.ac.uk

that human time perception is inherently uncertain and noisy, it makes sense for the perceptual system to attempt to resolve ambiguities by drawing on prior experience of temporally contiguous cause-effect pairings and shifting estimates of cause and effect time toward one another. Critically, the nature of the cause is irrelevant, as long as the organism has prior knowledge or experience of the causal relation and its time course.

In contrast, competing approaches to temporal binding (Engbert & Wohlschläger, 2007; Engbert, Wohlschläger, et al., 2007; Haggard, Aschersleben, et al., 2002; Haggard, Clark, & Kalogeras, 2002; Moore & Haggard, 2007; Wohlschläger, Haggard, et al., 2003) are specific to the human motor system and require intentional action as the trigger of the subsequent event. According to the intentional account, the perceptual system binds together actions and effects, so that the association between them can be strengthened in order to facilitate reliable action-outcome learning (Haggard, Aschersleben, et al., 2002). A related variant of the intentional-binding account is the motor-adaptation perspective (Stetson et al., 2006). According to this perspective, binding is the result of cross-modal adaptation to temporal delays (see also Kennedy et al., 2009). Because of differences in how rapidly the neural system processes sensory information in different modalities, there is a binding problem, such that a single event might result in multiple, temporally discrete percepts (e.g., the sound of my hand slapping the table is processed faster than the visual and tactile feedback from that same event). It is argued that to enable the organism to function efficiently, the nervous system adapts to and resolves such discrepancies by shifting the perceptual streams accordingly. Thus, short delays between motor actions and outcomes could be adapted to and create a uniform near-simultaneous percept of an action and its immediate outcome.

Note that both variants of the intentional account are a special subset of the more general causal account: In all situations during which intentional binding occurs, there is a causal relation, such that the cause is an intentional action and the effect some sensory consequence. In a typical intentional-binding experiment, the outcome follows the action after a very short interval. Thus, the learner acquires a specific, short time frame, which then can be used to disambiguate or bias subsequent perception of when the respective events occurred and the temporal interval between them (see Humphreys & Buehner, 2009, 2010, for demonstrations of changes in time perception). Most reports of temporal binding (Engbert & Wohlschläger, 2007; Engbert, Wohlschläger, & Haggard, 2008; Engbert, Wohlschläger, et al., 2007; Haggard, Aschersleben, et al., 2002; Haggard & Clark, 2003; Haggard, Clark, & Kalogeras, 2002; Humphreys & Buehner, 2009, 2010; Moore & Haggard, 2007) have confounded intentionality and causality so that it is impossible to disentangle the competing approaches.

The experiments reported here afforded a genuine test of the intentional versus the causal approach to temporal binding by contrasting two types of causal actions (self, machine) against a noncausal observational condition. If intentionality is the critical component, only the self-causal condition should

afford temporal binding. If causality is at the root of temporal binding, then intentional action would simply be one particular type of cause among others (say, mechanical causation), which all result in mutual attraction of cause and effect. Thus both the machine- and self-causal conditions should produce more temporal binding than the observational baseline condition would.

Experiments 1 and 2

Experiments 1 and 2 used an event-anticipation paradigm: Participants had to predict when an LED on a response box would flash. The basis on which predictions were made was varied across three conditions (see Fig. 1). In the baseline condition, the target flash was preceded by a signal LED. In the self-causal condition, participants pressed a button to generate the target flash; in the machine-causal condition, a machine (separated from the rest of the experimental apparatus) pressed the button to generate the flash. The target LED always flashed after a fixed interval of 500, 900, or 1,300 ms following the predictor. According to the causal-binding hypothesis, the two causal conditions should elicit earlier anticipations of the target event compared with the baseline condition; according to all other accounts, only the self-causal condition should be privileged in this way.

It was essential that the machine be perceived as an autonomous mechanical causal agent. Connection to the main apparatus would have negated the machine's role as an independent causal agent and would have made the computer the common cause (Pearl, 2000) of both the machine's action and the target flash. This was a key weakness of earlier designs (Wohlschläger, Engbert, & Haggard, 2003; Wohlschläger, Haggard, et al., 2003), which contrasted self- and other-causal conditions against machine-causal conditions involving a rubber hand (see Fig. 1). In those studies, the rubber hand was attached to the same lever that subjects (or the experimenter) had to press in the causal conditions, but the lever was now pulled down by a solenoid (i.e., in the absence of intentional action). Critically, the scheduling of the lever press in these machine conditions was controlled by the same computer that generated the target tone. Thus, the computer was the common cause of both the lever depression and the tone. There was no causal link between the lever press and the subsequent tone in these conditions, and hence no binding would have been predicted from the causal perspective.

Experiment 2 was a replication of Experiment 1 with a small procedural change (see Method) to allow for an even stricter test of the causal hypothesis.

Method

Participants. Forty-six and 40 volunteers participated in Experiments 1 and 2, respectively. All were recruited from the Cardiff University subject pool and received £5 for participating.

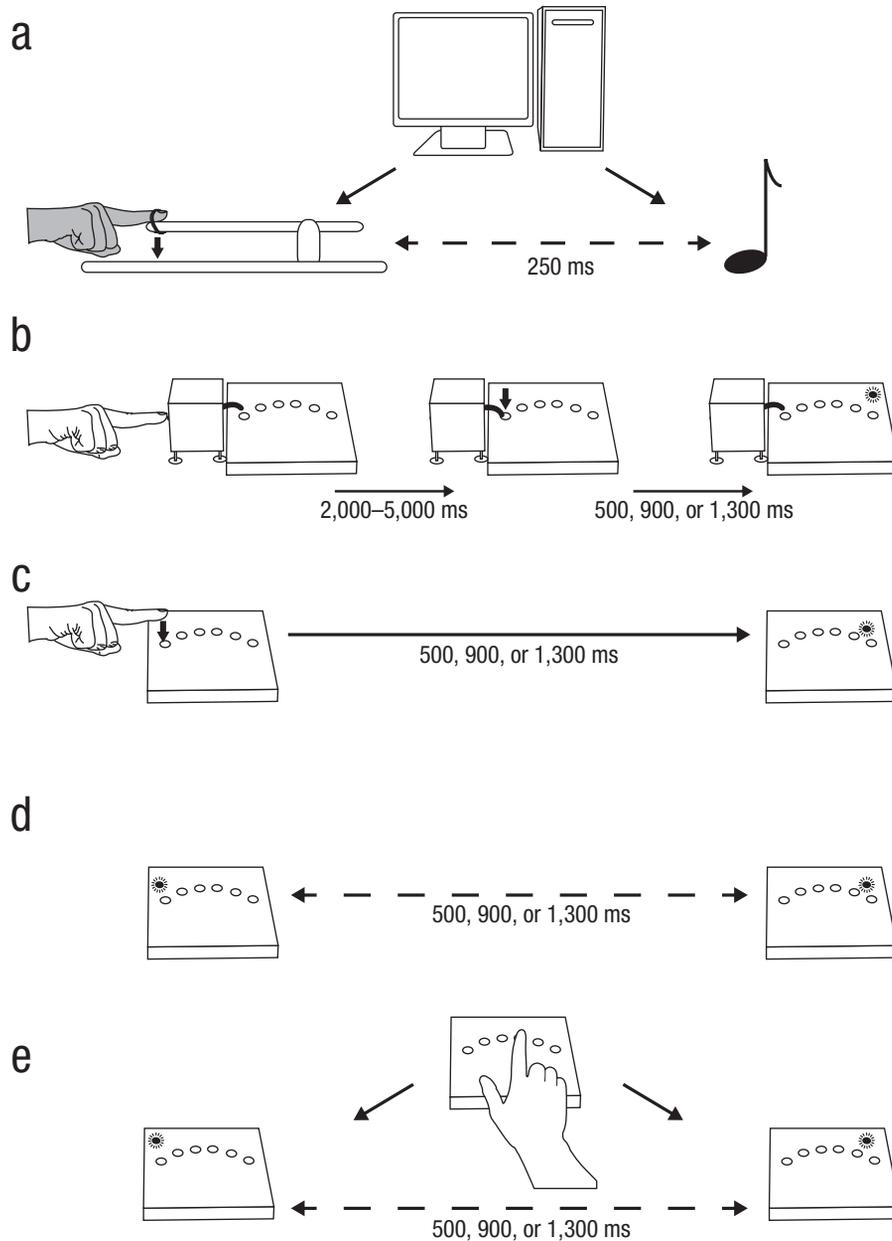


Fig. 1. Illustrations of (a) the experimental setup employed by Wohlschläger, Engbert, and Haggard (2003) and Wohlschläger, Haggard, Gesierich, and Prinz (2003) and (b–e) the setups used in the present experiments. In the machine-action conditions used in Wohlschläger et al.'s experiments (a), a lever was pulled down 250 ms before a tone was sounded. Both the movement of the lever and the sounding of the tone were controlled by a computer. In some of these experiments, a rubber hand was attached to the lever to maximize visual similarity to the self- or other-action conditions. In the machine-causal condition of the present experiments (b), participants supposedly initialized a machine at the start of each trial. After a random delay, the machine pressed a button on the left-hand side of a response box. After 500, 900, or 1,300 ms, a target LED flashed on the response box. The procedure in the self-causal condition of Experiment 1 (c) was similar, except that the participant pressed the button on the response box directly. In the baseline condition of Experiment 1 (d), a signal LED on the left-hand side of the response box flashed 500, 900, or 1,300 ms before the target LED flashed. In the baseline and self-causal conditions of Experiment 2 (e), participants were instructed to press a button in the center of the response box to begin each trial, after which the trials proceeded as in the corresponding conditions in Experiment 1 (only the baseline condition is illustrated here). In all conditions, participants had to press the right-most button on the response box at the time they expected the target LED to flash. Solid arrows indicate temporal intervals determined by genuine causal links. Dashed, double-headed arrows indicate merely correlational temporal intervals.

Apparatus. The experiment was conducted on an iMac computer running PsyScope X (Cohen, MacWhinney, Flatt, & Provost, 1993) with an ioLab Systems (www.iolab.co.uk) universal serial bus (USB) response box. A machine (Fig. 2) comprising a plastic casing, stopwatch display, switches, and solenoid-controlled lever was specifically built for the machine-causal condition.

Design and procedure. Participants were told they would partake in a stimulus-anticipation experiment, in which their task would always be to anticipate the flash of a target LED on the response box. Predictions were based either on a preceding flash of another LED (baseline condition) or on a preceding button press on the box by either the participant (self-causal condition) or a machine (machine-causal condition). Three prediction intervals were used (500, 900, 1,300 ms). Each interval was presented in a 20-trial block in each condition of Experiment 1 and in a 30-trial block in each condition of Experiment 2 to allow optimal prediction via experience-based learning. The order of the trial blocks was randomized within each condition. Participants performed all three conditions in random order.

In the baseline condition, a signal LED on the left side of the box flashed for 100 ms. Then, after the prediction interval, the target LED flashed (100 ms) on the right side of box. Participants were asked to press a button immediately below the target LED at the exact moment that they expected the target to flash. A trial ended as soon as the target appeared and the participant had made a response. The next trial followed after a randomly determined interval of 2,300 to 2,800 ms.

The procedure in the two causal conditions was identical to the procedure in the baseline condition, except that the target flash was triggered by pressing the left-most button on the box. In the self-causal condition, a message appeared on the screen asking participants to press the button to generate the target flash. Once they did so, the message disappeared and the target flashed after the relevant delay. The only difference

in the machine-causal conditions was that the button was pressed by a machine, which was completely separated from the rest of the experimental apparatus, and that the message on the screen asked participants to “initialize the machine.” Participants were told that the machine contained a random number generator, which would trigger a lever press after a random interval between 2 and 5 s following “initialization” via a switch on the front of the machine. This was demonstrated a few times to participants at the beginning of the machine-causal block: Participants pressed the switch on the front of the machine, which started a stopwatch display above it; after an interval of 2 to 5 s, the stopwatch halted and the lever went down and up again. Participants could then reinitialize the machine by pressing the switch on its front again, which started the stopwatch where it had stopped, and triggered a new cycle. In reality, the lever action was controlled by an infrared remote control hidden in the experimenter’s pocket.

The initiation and stopwatch procedure were used to make the setup more credible. It would have been impossible for a machine completely separated from the computer and other apparatus to press the button only during the relevant trials (and to wait for the next action until after a participant made his or her prediction response and was thus prepared for the next trial). In postexperimental debriefing, not a single participant reported having anthropomorphized the machine as having intentions; likewise, no participant had been aware of the deception that the machine’s action was controlled by a hidden remote.

Experiment 2 was identical to Experiment 1, except that each trial in the self-causal and baseline condition began with a prompt on the screen asking participants to press the green button (on the center of the response box) to begin the trial. Once participants pressed this button, the trials proceeded as in Experiment 1: Self-causal trials displayed the prompt asking participants to press the left-most button on the box to generate the target flash; baseline trials began after a randomly determined interval of 2,300 to 2,800 ms, after which the



Fig. 2. Two views of the machine and the response box used in the machine-causal conditions of Experiments 1 and 2. The signal and target LEDs are not visible here because they were situated underneath the opaque translucent plastic cover of the box. Each LED was placed centrally above the relevant button.

signal LED flashed, followed by the target LED after the relevant interval.

The rationale for these changes was to address concerns that the machine versus self distinction in Experiment 1 was not as clear as it could be: Participants were required to initiate the machine by pressing a switch on it. Although this action was not the immediate cause of the outcome that followed (which was, of course, the machine pressing the button on the response box), it nonetheless constituted an enabling condition, setting the causal mechanism in motion. Thus, there arguably could be a trace of intentional binding also affecting participants in the machine-causal condition. Amending both the baseline and self-causal conditions with similar enabling actions ruled out this concern. More specifically, adding an enabling action to the baseline condition altered its causal structure such that the enabling action served as a common cause for both the first and the second (target) flash (see Fig. 1). Consequently, there was a direct causal relation between the initializing action and the first LED flash at baseline, but only an indirect causal relation in the machine-causal condition.

Thus, if there is any trace of intentional binding between an earlier initializing action and a consequent event, more of it would be expected in this revised baseline condition than in the machine-causal condition. Note, however, that reliable prediction of the target event was only afforded relative to the

first flash, and not relative to the initializing action. The addition of the initializing action thus preserved the central contrast whereby the occurrence of a target event can be reliably predicted by a preceding event that either caused the target (machine and self-causal conditions) or merely correlated with it (baseline condition).

Results

Analyses of variance were conducted on log-transformed prediction errors to ensure normality. A significance level of .05 (two-tailed) was used for all tests. Event anticipation was operationalized by comparing each participant's median prediction error over the last 10 trials in each condition with the veridical target time in that condition. Figures 3 and 4 show a clear anticipation advantage for causality-based predictions. Statistical analysis of logged prediction errors confirmed a main effect of condition in both Experiment 1, $F(2, 90) = 6.98$, $\eta_p^2 = .13$, and Experiment 2, $F(2, 78) = 3.26$, $\eta_p^2 = .08$. Orthogonal Helmert contrasts further showed that correlation-based baseline predictions (mean raw error = 6.17 ms in Experiment 1 and -10.69 ms in Experiment 2) were significantly delayed relative to the two kinds of causality-based predictions, Experiment 1: $F(1, 45) = 11.06$, $\eta_p^2 = .20$, and Experiment 2: $F(1, 39) = 5.55$, $\eta_p^2 = .13$, which in turn did not differ from each other (self-causal condition: mean raw error = -26.68 ms in

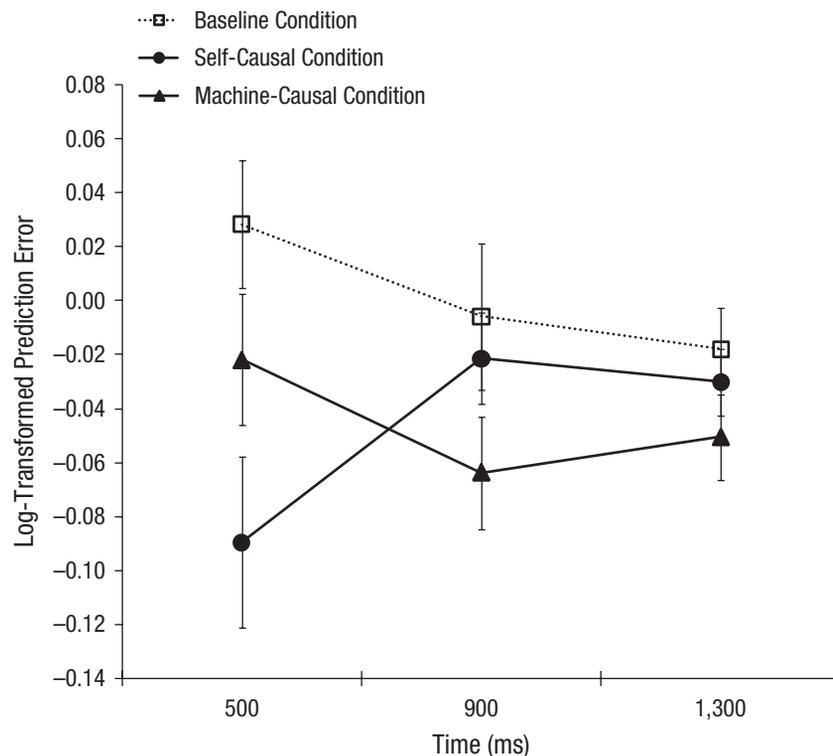


Fig. 3. Results from Experiment 1: mean log-transformed prediction error as a function of the interval between the signal and the target flash (500, 900, or 1,300 ms) and condition. Negative numbers indicate predictions that preceded the target flash. Error bars denote standard errors.

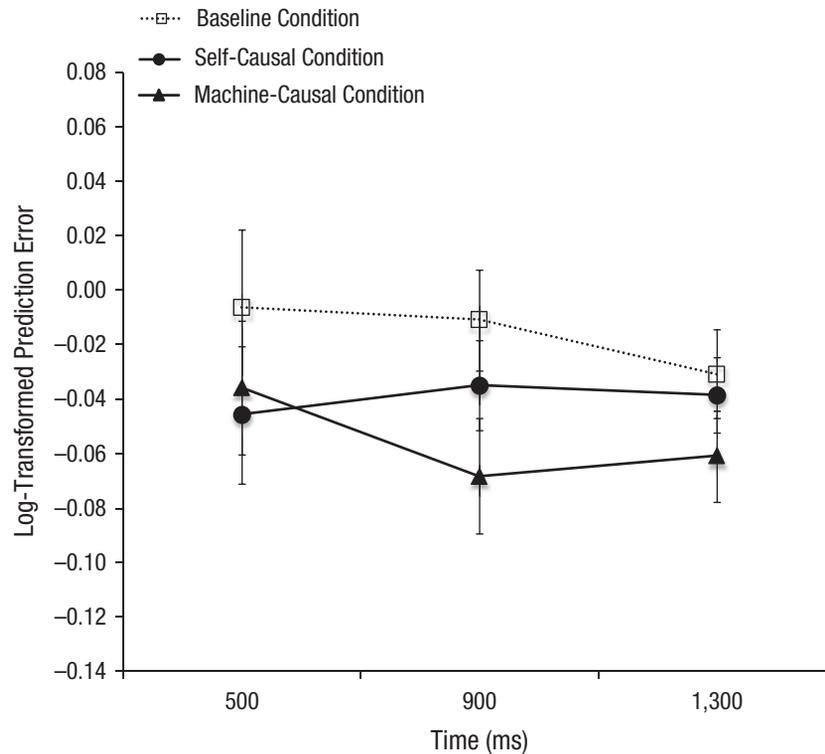


Fig. 4. Results from Experiment 2: mean log-transformed prediction error as a function of the interval between the signal and the target flash (500, 900, or 1,300 ms) and condition. Negative numbers indicate predictions that preceded the target flash. Error bars denote standard errors.

Experiment 1 and -28.98 ms in Experiment 2; machine-causal condition: mean raw error = -35.96 ms in Experiment 1 and -44.68 ms in Experiment 2), both F s < 1 . In Experiment 1, a Condition \times Time interaction, $F(4, 180) = 4.60$, $\eta_p^2 = .09$, suggested that the causality advantage at very short intervals (500 ms) was greater for self- than for machine causation, with a reversal of this pattern at slightly longer intervals (900 and 1,300 ms). However, this pattern did not hold up in Experiment 2, $F(4, 156) = 0.92$.

Discussion

The two experiments reported here showed clear evidence of temporal binding in the self-causal condition as well as in the machine-causal condition. More specifically, short-term event anticipation was boosted by causal understanding of the underlying predictive relation, and this finding was not limited to situations involving forward planning of motor control (Haggard, Clark, & Kalogeras, 2002) or intentional (self/other) agents (Wohlschläger, Haggard, et al., 2003).

These two experiments showed that people will predict that target events will occur earlier if they have a causal understanding of what brings the events about; mere correlation or association with a predictive signal did not yield such privileged processing, even though on a perceptual level, all three

conditions afforded the same regularities between cue and outcome. It is important to note that Experiment 2 confirmed that temporal binding in the machine-causal conditions could not have occurred as the result of any agency component or intentional attribution afforded by participants initializing the machine at the start of each trial.

These results thus provide novel but unequivocal evidence of temporal binding in the absence of intentional action. Earlier work from this lab has already demonstrated that intentional action on its own (in the absence of a causal relation) is not sufficient to produce temporal binding between actions and subsequent events (Buehner & Humphreys, 2009). The results reported here complement these findings by showing that intentional action is also not necessary, but can instead be replaced with nonintentional mechanical causation. Consequently, intentional binding is a special subset of causal binding, in which the cause in question constitutes an intentional action. Because the explanatory power lies with the causal relation rather than with the intentional action, the term *intentional binding* is a misnomer and should be replaced with *causal binding*. This interpretation is corroborated by the observation that reducing the causal efficacy of an intentional action to produce an outcome leads to a concomitant decrement in temporal binding (Moore, Lagnado, Dear, & Haggard, 2009).

The results of Experiment 1 suggested that at a very short interval (500 ms), intentional action might provide an additional boost over and above the causal binding afforded by causation in the absence of intention. Original reports of perceptual shifts following voluntary actions (Engbert & Wohlschläger, 2007; Engbert et al., 2008; Haggard, Aschersleben, et al., 2002; Haggard, Clark, & Kalogeras, 2002; Moore & Haggard, 2007) all involved very short intervals. In fact, Haggard, Clark, and Kalogeras (2002) found that the binding effect disappeared when action and outcome were separated by 650 ms (but see Humphreys & Buehner, 2009, 2010, for demonstrations of binding across much longer intervals). Thus, there is a possibility that causal binding receives an additional boost from intentional action if the consequence follows almost instantaneously. However, the failure to replicate this effect here in Experiment 2 casts doubt on this possibility. Furthermore, our earlier finding (Buehner & Humphreys, 2009) that intentional action in the absence of causality fails to elicit binding clearly shows that the explanation for temporal binding is causality and not intentionality (see also Moore et al., 2009, as discussed above).

It is important to note that the reported results cannot be explained by attentional modulation of time perception. If attentional load associated with planning and executing the self-causal trigger action drew resources away from the timing and anticipation task, then target anticipation in the self-causal condition should stand out compared with target anticipation in the baseline and machine-causal conditions, which did not require initial action. This was not the case. Alternatively, had the element of surprise associated with the uncontrollable signal LED (baseline condition) or machine action detracted from the main task of target anticipation, then these two conditions would have produced similar results, which was also not the case. Instead, predictions of self- and machine-caused targets gave rise to earlier responses to the anticipated event than correlationally based predictions of baseline targets did.

In conclusion, causal knowledge thus prepares people for the future not only at a high cognitive level (by informing their choices) but also on a low perceptual level (by increasing their awareness of what is to come). It is well established that many behaviors are based on habit and appear not to recruit deeper reflection on causal principles. Ironically, a strength of such association-based action is that it is thought to draw on fast, automatic, and nondeliberative processes (Sloman, 1996). The results presented here cast doubt on the adaptive value of action grounded in mere association and instead highlight once again that many findings previously thought to be supporting associative theories are better explained by causal models (for an example from animal learning, see Blaisdell, Sawa, Leising, & Waldmann, 2006). Specifically, causal binding arises despite causal and noncausal event sequences both having identical low-level predictability, and it cannot be explained by attentional differences. Furthermore, causal binding is insensitive to the intentional versus mechanical nature of the causal source, which rules out accounts of

intentional binding and forward planning of motor control (Haggard, Clark, & Kalogeras, 2002; Wohlschläger, Haggard, et al., 2003) in favor of a more general theory of causal binding (Buehner, 2010; Buehner & Humphreys, 2009, 2010; Eagleman & Holcombe, 2002).

Acknowledgments

I thank Dennis Price and Dennis Simmonds for constructing parts of the experimental apparatus, Lorraine Woods for creating Figure 1, Sophie Patterson and Hoda El-Gaby for acting as experimenters, and Ulrike Hahn for valuable comments on an earlier draft of the manuscript.

Declaration of Conflicting Interests

The author declared that he had no conflicts of interest with respect to his authorship or the publication of this article.

Funding

This research was funded by the Engineering and Physical Sciences Research Council and the Experimental Psychology Society.

References

- Blaisdell, A. P., Sawa, K., Leising, K. J., & Waldmann, M. R. (2006). Causal reasoning in rats. *Science*, *311*, 1020–1022.
- Buehner, M. J. (2010). Temporal binding. In C. Nobre & J. Coull (Eds.), *Attention and time* (pp. 201–211). Oxford, England: Oxford University Press.
- Buehner, M. J., & Cheng, P. W. (2005). Causal learning. In K. J. Holyoak & R. Morrison (Eds.), *The Cambridge handbook of thinking and reasoning* (pp. 143–168). Cambridge, England: Cambridge University Press.
- Buehner, M. J., & Humphreys, G. R. (2009). Causal binding of actions to their effects. *Psychological Science*, *20*, 1221–1228.
- Buehner, M. J., & Humphreys, G. R. (2010). Causal contraction: Spatial binding in the perception of collision events. *Psychological Science*, *21*, 44–48.
- Cheng, P. (1997). From covariation to causation: A causal power theory. *Psychological Review*, *104*, 367–405.
- Cohen, J. D., MacWhinney, B., Flatt, M., & Provost, J. (1993). PsyScope: An interactive graphic system for designing and controlling experiments in the psychology laboratory using Macintosh computers. *Behavior Research Methods, Instruments, & Computers*, *25*, 257–271.
- Eagleman, D. M., & Holcombe, A. O. (2002). Causality and the perception of time. *Trends in Cognitive Sciences*, *6*, 323–325.
- 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.

- Haggard, P., Aschersleben, G., Gehrke, J., & Prinz, W. (2002). Action, binding, and awareness. In W. Prinz & B. Hommel (Eds.), *Common mechanisms in perception and action: Attention and performance* (Vol. XIX, pp. 226–285). Oxford, England: Oxford University Press.
- 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.
- Hume, D. (1888). *An enquiry concerning human understanding* (L. A. Selby-Bigge, Ed.). Oxford, England: Clarendon Press. (Original work published 1777)
- Humphreys, G. R., & Buehner, M. J. (2009). Magnitude estimation reveals temporal binding at super-second intervals. *Journal of Experimental Psychology: Human Perception and Performance, 35*, 1542–1549.
- Humphreys, G. R., & Buehner, M. J. (2010). Temporal binding of action and effect in interval reproduction. *Experimental Brain Research, 203*, 465–470.
- Kennedy, J. S., Buehner, M. J., & Rushton, S. K. (2009). Adaptation to sensory-motor temporal misalignment: Instrumental or perceptual learning? *Quarterly Journal of Experimental Psychology, 62*, 453–469.
- Moore, J., & Haggard, P. (2007). Awareness of action: Inference and prediction. *Consciousness and Cognition, 17*, 136–144.
- Moore, J., Lagnado, D. A., Dear, D. C., & Haggard, P. (2009). Feelings of control: Contingency determines experience of action. *Cognition, 110*, 279–283.
- Pearl, J. (2000). *Causality: Models, reasoning, and inference*. Cambridge, England: Cambridge University Press.
- Sloman, S. A. (1996). The empirical case for two systems of reasoning. *Psychological Bulletin, 119*, 3–22.
- Sloman, S. A. (2005). *Causal models: How we think about the world and its alternatives*. New York, NY: Oxford University 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.
- Wohlschläger, A., Engbert, K., & Haggard, P. (2003). Intentionality as a constituting condition for the own self—and other selves. *Consciousness and Cognition, 12*, 708–716.
- Wohlschläger, A., Haggard, P., Gesierich, B., & Prinz, W. (2003). The perceived onset time of self- and other-generated actions. *Psychological Science, 14*, 586–591.