Cognition 146 (2016) 58–66

Contents lists available at ScienceDirect

Cognition

journal homepage: www.elsevier.com/locate/COGNIT

Time reordered: Causal perception guides the interpretation of temporal order

Christos Bechlivanidis *, David A. Lagnado

Department of Experimental Psychology, UCL, 26 Bedford Way, London WC1H0DS, UK

A R T I C L E I N F O

Article history:
Received 12 December 2014
Revised 30 August 2015
Accepted 1 September 2015

Keywords:
Causality
Causal perception
Time perception
Temporal order
Perceptual illusion
Michottean launching

A B S T R A C T

We present a novel temporal illusion in which the perceived order of events is dictated by their perceived causal relationship. Participants view a simple Michotte-style launching sequence featuring 3 objects, in which one object starts moving before its presumed cause. Not only did participants re-order the events in a causally consistent way, thus violating the objective temporal order, but they also failed to recognise the clip they had seen, preferring a clip in which temporal and causal order matched. We show that the effect is not due to lack of attention to the presented events and we discuss the problem of determining whether causality affects temporal order at an early perceptual stage or whether it distorts an accurately perceived order during retrieval. Alternatively, we propose a mechanism by which temporal order is neither misperceived nor misremembered but inferred "on-demand" given phenomenal causality and the temporal priority principle, the assumption that causes precede their effects. Finally, we discuss how, contrary to theories of causal perception, impressions of causality can be generated from dynamic sequences with strong spatiotemporal deviations.

1. Introduction

Imagine watching a long line of dominos falling one after another. Apparently, each domino's fall causes the fall of the next one in the line. But suddenly one domino falls early, before being touched by the previous domino (Fig. 1). Would you notice that the domino had fallen prematurely?

Whether or not causal impressions can influence the experienced temporal order depends on two questions: (1) Can the perception of temporal order be influenced by information other than the order of the percepts themselves, and (2) do causal impressions possess those features necessary to influence presumably lower level percepts?

For some philosophers (Hoerl, 2013; Mellor, 1985; Phillips, 2014; Soteriou, 2010), the answer to the first question is negative: according to the mirroring theory, the order of our experiences mirrors or inherits the temporal structure of the environment. Thus, to experience event A happening before B, we must be exposed to that particular temporal order, even if it is illusory, such as when lightning is seen before the thunder is heard. Others, however, have argued for a more constructed view of temporal order (Dainton, 2010; Grush, 2007), i.e. for temporal order as a second-order judgement.

From a psychological perspective, the prior entry effect, the finding that attended stimuli are perceived earlier than unattended ones (Spence & Parise, 2010; Titchener, 1908) might seem at odds with the mirroring theory. This is because in prior entry the order of presentation does not match the perceived order. However, what matters to mirroring theory is the subjective order of presentation. What determines the experienced order is the time when successive stimuli (accelerated or delayed by attention) reach awareness, consistent with the mirroring theory of time perception (Vibell, Klinge, Zampini, Spence, & Nobre, 2007).

Multisensory integration, on the other hand, suggests that the experienced temporal order is in fact malleable: when two successive bimodal stimuli are assumed to originate from the same source, the perceived timing of each stimulus is shifted so that the two events are experienced as being simultaneous (King, 2005; Spence & Squire, 2003). Unlike the domino example, however, the order here is collapsed rather than reversed. Nevertheless, Stetson and colleagues (Stetson, Cui, Montague, & Eagleman, 2006) have combined multisensory integration with sensory adaptation to show that an initial adaptation to a short delay between an action and an outcome leads to an illusory experience of effects preceding their causes when that delay is subsequently reduced (see also Heron, Hanson, & Whitaker, 2009; Rohde, Scheller, &

* Corresponding author.
E-mail address: ucjtcbe@ucl.ac.uk (C. Bechlivanidis).

http://dx.doi.org/10.1016/j.cognition.2015.09.001
0010-0277/

Fig. 1. The temporal order (the 3/3 domino falls before the 2/2 touches it) does not match the assumed causal order (the 2/2 domino causes the fall of the 3/3). How do we resolve this incongruence?

Ernst, 2014). However, the bimodal nature of the stimuli and the requirement for sensory integration introduce additional complications and allow for multiple interpretations. It might, thus, be the case that the order distortions observed in multisensory integration happen exactly at the stage when signals from the various modalities are integrated to produce a coherent experience, whereas the order of unimodal stimuli is determined solely by the order of presentation, perhaps modulated by attention.[1]

Returning to our second question about the role of causality in perception, several recent studies show that judgments of spatial relations (Scholl & Nakayama, 2004), size (Buehner & Humphreys, 2010) trajectory (Kim, Feldman, & Singh, 2013) and, more relevantly, temporal duration (Buehner, 2012; Eagleman & Holcombe, 2002; Schutz & Kubovy, 2009) are sensitive to impressions of causality. For example, Buehner and colleagues (Buehner, 2012; Buehner & Humphreys, 2009) have shown that the temporal distance between the onset of successive events appears to shrink when the events are believed to be causally related. In this case, however, causal beliefs lead to quantitative shifts, whereas reversing the order of events requires a stronger qualitative change.

Nevertheless, Bechlivanidis and Lagnado (2013) show that recently acquired causal knowledge can switch temporal order judgements. In their study participants played a computer-based puzzle game which required learning a novel causal relation between two events. After training, participants observed events happening in an order that violated the causal order of the learned relationship. When asked to report the order they saw, the majority preferred the order that matched their acquired causal beliefs, thus distorting the objective order of presentation.

Does this reordering effect depend on recently learned causal rules or does it generalize to any case where strong causal beliefs are present? Similarly, does it depend on inference or can instances of causal perception (Michotte, 1963) also result in the reordering of events? Finally, did participants in our earlier studies simply report the most plausible order of events guided by learned causal rules or did they actually perceive a different order to that presented to them?

[1] Another set of findings that are commonly debated in the philosophy of time literature are collectively known as postdiction effects (see Shimojo, 2014 for a review). In postdiction a perception of some event appears to be influenced by events that temporally follow. However, given that these phenomena don't result in distortions of order but rather of the properties of events or objects, we will not be discussing them here.

2. Experiment 1

To evaluate the sensitivity of temporal order judgments to contradictory causal impressions, we modified the classic Michottean launching sequence (Michotte, 1963) by adding a third object. Participants observed a three-object pseudo-collision where, critically, the third object starts moving before the second object, i.e. the effect takes place before its presumed cause (Fig. 2a). Following presentation, participants were asked to report the temporal order of the events they had just witnessed.

Furthermore, to ensure that the order of events is perceptually distinguishable, as well as using relatively low speeds and long delays between critical events, we included a control condition in which object A is not present (Fig. 2b). In this case, we expected participants to report the veridical order of events since in the absence of A, there is no clear causal direction.

2.1. Participants and materials

The experiment was programmed in Adobe Flex 4.6 and conducted over the Internet using Amazon Mechanical Turk (both experiments presented here can be seen at http://goo.gl/4noAmR). While perceptual experiments are traditionally conducted in the controlled environment of the lab, we took a number of precautionary measures to ensure the consistent presentation of stimuli – see Appendix A for more details.

We recruited 60 participants in total but one participant was excluded from the analysis for providing a nonsensical answer to the order question (see Design and procedure), i.e. not identifying correctly the first object that started moving. Of the remaining 59 participants, 39 were male and 20 were female. The mean age was 32.39 years (SD = 9.96). Each participant was paid $0.50.

2.2. Design and procedure

The 59 participants were randomly assigned to one of two conditions that differed only in the displayed clip, resulting in 29 participants in condition 1 and 30 in condition 2.

After completing the calibration section (see Appendix A) participants were welcomed to the experiment and were asked for some simple demographic data. They were then informed that they would watch a short movie clip and answer some questions about

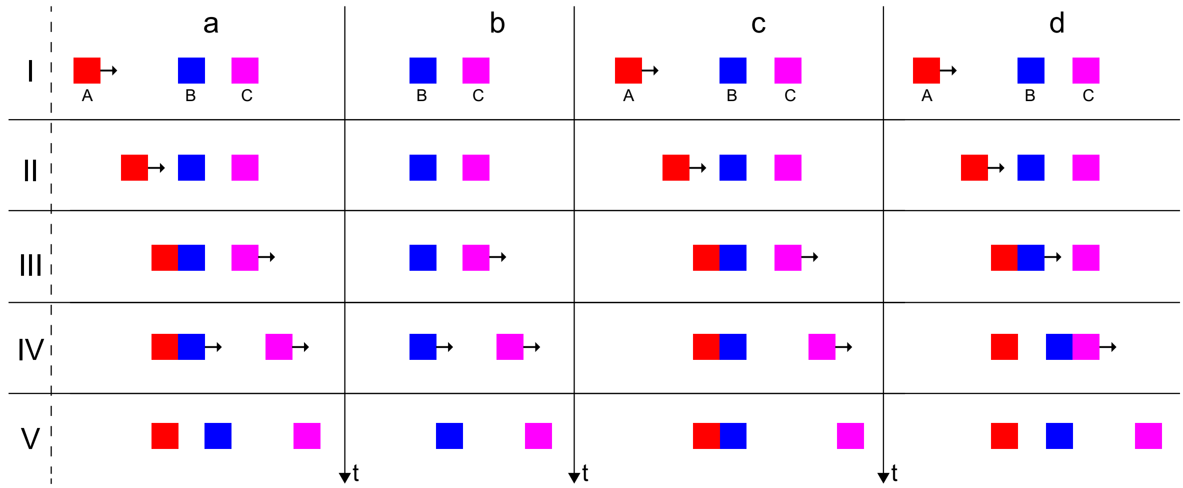

Fig. 2. The sequences used in the two experiments: (a) Object A approaches B (I–II) and stops next to it at which point object C starts moving (III). After 350 ms object B starts moving (IV) and stops to the left of object C's original position (V); (b) identical to sequence (a) without object A; (c) identical to (a) but object B remains stationary throughout; (d) realistic 3-object collision shown during the review question in Experiment 2 (the arrows show the direction of movement and were not visible in the experiments).

it. They were also asked to be as focused as possible since the clip would be displayed only once.

Participants saw the clip a single time. In condition 1 ("A present") that is shown in Fig. 2a, three 8×8 mm squares fade in slowly (2 s). Object A is located 35 mm to the left of object B and object C is located 16 mm to the right of object B. The squares remained static for another 2 s. Then object A starts moving to the right towards B at a speed of 30 mm/s, i.e. relatively slowly compared, for example, to the 300–400 mm/s in Michotte's original studies. Object A stops adjacent to B and, immediately after, object C starts moving also at 30 mm/s. After 350 ms object B starts moving to the right at the same speed and stops to the left of object C's original position. The clip ends immediately after C travels 35 mm. The clip was designed to be as similar as possible to a normal 3-object collision with the exception of the order of events between B and C.

In condition 2 ("A absent") the clip that was shown was exactly the same but object A was not present (Fig. 2b). Specifically, in this condition, object C starts moving to the right at 30 mm/s and 350 ms later B moves also to the right and stops next to C's original position. C travels for 35 mm and the clip ends.

For each condition there were two similar versions of the target clip that differed only in the objects' colours and the direction of movement. In the "normal" version, the colours were as shown in Fig. 2a, i.e. red: RGB(255,0,0), blue: RGB(0,0,255) and purple: RGB(212,7,171) and the direction of movement was left-to-right as described above. In the "mirrored" version the colours were A = purple, B = red, C = blue and the direction of movement was right-to-left, meaning that the initial position of objects was mirrored compared to the "normal", i.e. A starts to the right of B and C to the left of object B. Participants in each condition were randomly assigned to one of the two clip versions.

After watching the clip a single time, participants were shown the initial configuration of the objects (i.e. Fig. 2a-I or b-I) and were asked to place the events in the order that they saw them. To do this they had to drag-and-drop the event sentences "The red square started moving" (only in condition 1), "The blue square started moving" and "The purple square started moving" from their initial container to another box. The order of appearance of the sentences was randomised for each participant. Then participants were asked to indicate their confidence to the selected order by dragging a slider on a scale that was labelled "Not at all confident" to the left and "Very confident" to the rightmost position.

In the next screen, the initial object configuration was shown again and participants were asked for their causal impressions for all possible object pairs (six in condition 1 and two in condition 2). These were expressed by dragging on a slider labelled "Completely Disagree", "Neutral" and "Completely Agree" next to statements of the form "The X square made the Y square move", were X and Y were colour pairs (e.g. "The red square made the blue square move"). Finally, participants were asked for any comment they had regarding the experiment and they were thanked for participating.

2.3. Results

Fig. 3, shows the proportion of participants that reported the objective temporal (A–C–B) or the causal order (A–B–C) of events (we collapsed the normal and mirrored versions of the clips since no difference was observed). The overwhelming majority (82.76%) preferred the causal order when A was visible while a similar majority (83.33%) preferred the objective temporal order when A was absent, despite the fact that in both conditions the behaviour of objects B and C was identical ($\chi^2(1, N = 59) = 25.77$, $p < .001$). Furthermore, participants in both conditions were very confident in the order they reported, with mean confidence ratings 78.76/100 (SD = 24.62) for condition 1 and 73.63/100 (SD = 21.05) for condition 2.

The direct causal judgments (Fig. 4) show that when A was present participants thought that it caused B to move (88.45%) and also that B caused C to move (77.28%). Participants were relatively indecisive about the A–C relationship (51.76%) but given the strong endorsement of the A–B and B–C relationships, one can assume that those endorsing it probably referred to the indirect A–C relationship, through B. The judgments for the inverse relationships were, as expected, very low. The C–B relationship is significantly higher than the C–A relationship ($t(28) = 2.305$, $p < 0.05$) and approaches significance compared with B–A ($t(28) = 2.007$, $p = 0.054$) but this is driven by those few participants who reported the correct temporal order of the sequence and any significant difference goes away if these participants are excluded.

Similarly, when A was not present (condition 2) and, thus, when the majority of participants reported the objective temporal order of events, the causal judgments were far weaker. The strongest causal belief is in C making B move. In fact significantly more participants endorsed the C–B causal relationship in condition 2 compared to condition 1 ($t(57) = 4.837$, $p < 0.001$). Regarding the

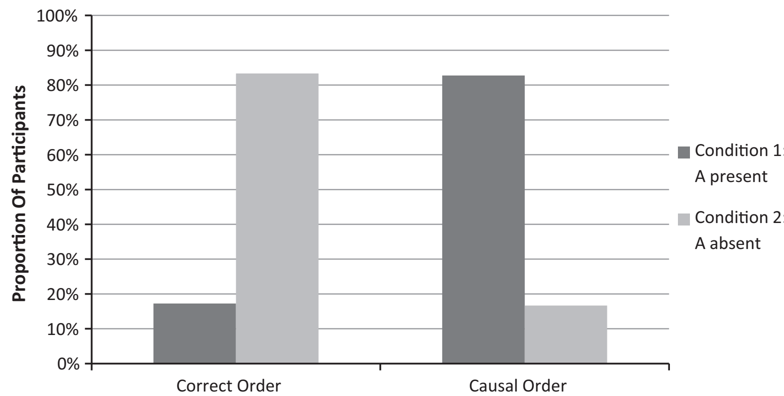

Fig. 3. Proportion of participants that reported the correct (temporal) or the causal order of events in each condition.

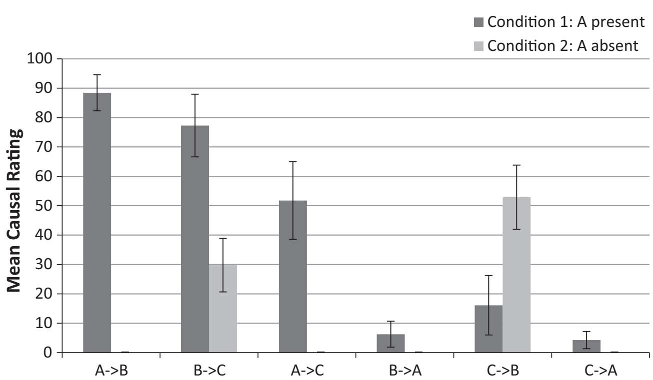

Fig. 4. Mean causal judgments for each object pair per condition (in condition 2, object A was not visible, so there are no ratings involving it). A value of 50 corresponds to neutrality – lower ratings indicate disagreement and higher rating indicate agreement with the causal statement (error bars represent 95% confidence intervals).

open-ended question at the end of the experiment, the majority of participants from both conditions left no comments or had nothing relevant to report.

2.4. Discussion

The results indicate that impressions of phenomenal causality can modify the perceived temporal order of events. The overwhelming majority of participants in the first condition reported the order that matched their causal impression, despite the fact that the actual temporal order was clearly perceivable, as both the long delay (350 ms) between the events and the veridical ordering in condition 2 indicate. The causal basis of the reordering effect is further demonstrated by the strong endorsement of the statement according to which B made C move.

Before attempting to explain and evaluate this finding, we need to address some potential confounds, especially related to attentional issues. First, the sequence becomes visually simpler in the absence of object A, therefore the erroneous temporal order reported in condition 1 could be attributed to the relative increase of perceptual load. Similarly, since motion and especially the onset of motion are known to attract attention (Abrams & Christ, 2003; Hillstrom & Yantis, 1994), perhaps participants' attention is drawn to object C when it starts moving, thus completely missing B's behaviour. Support for the latter comes from studies of inattentional blindness (Simons & Chabris, 1999) where participants fail to notice salient features of the environment that don't receive their focused attention. Thus, participants might not report the order that they actually see but rather the most plausible order given the lack of information due to perceptual overload or split attention. In other words, if participants miss part of the action, it makes sense to assume a causal relationship between events and thus report the order that matches this relationship.

A more trivial explanation is that participants perceive all the events in the order in which they were presented but, given the similarity to a normal 3-object collision, they assume that the presentation was distorted due to some computer error. In this case, the observed results can be explained by some form of response bias: participants do not report the order they experience, but the order which, in their view, the experimenter aimed to present. Despite the fact that no participant mentioned any abnormalities in the clip presentation, Experiment 2 will investigate this option, as well as the other potential confounds mentioned above by using a different control condition and an alternative, stricter dependent measure.

3. Experiment 2

We again presented the 3-object sequence of Experiment 1 (Fig. 2a) but instead of asking participants for an explicit ordering of the events, we presented the same sequence again side-by-side with a canonical collision sequence, i.e. a sequence in which the order of events is congruent with their causal relationships (Fig. 2d). After watching each of these sequences participants were asked to identify which of the two they saw earlier.

In the second condition of this between-group experiment, we presented participants with a very similar sequence that differed only in that object B remains stationary throughout (Fig. 2c). We hypothesised that the lack of motion would diminish the causal link between A and B as well as between B and C. In the absence of a causal interpretation, participants would be better at identifying the sequence they saw when asked to choose between that and a realistic collision. If this is the case, we will have evidence that the reordering effect observed in Experiment 1 and in the first condition of this experiment cannot be explained by lack of attention to B's behaviour.

3.1. Participants and materials

The experiment was programmed in Adobe Flex 4.6 and conducted over the Internet using Amazon Mechanical Turk. As in Experiment 1, we recruited 60 participants in total. Two were excluded from the analysis because in the critical question they did not watch one of the two sequences that they were asked to

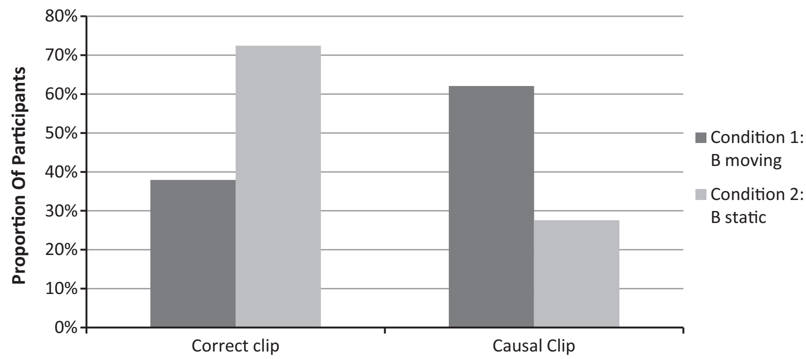

Fig. 5. Proportion of participants that selected the correct clip (the one they saw) or the canonical collision clip in each condition.

choose from so their answers were in fact random (see Design and procedure). Of the remaining 58 participants, 31 were male and 27 were female. The mean age was 34.57 years (SD = 12.12). Each participant was paid \$0.50. The same precautionary measures as in Experiment 1 were taken (see Appendix A).

3.2. Design and procedure

The 58 participants were randomly assigned to one the two conditions that differed only in the critical clip, resulting in 29 participants in each condition. The introductory screens were the same as in Experiment 1 and participants were asked to pay attention to the clip that would be shown a single time.

In condition 1 (“B moving”) the clip was identical to the clip shown in the first condition of Experiment 1 (Fig. 2a). In condition 2 (“B static”) the clip was similar with the exception that B remained static throughout the sequence (Fig. 2c). So, object A approaches from the left at 30 mm/s and stops next to B at which point C starts moving at the same speed and direction. The clip ends immediately when object C covers 35 mm. As in Experiment 1 there were two versions of each clip, one with the colours being red, blue and purple and direction left-to-right as in Fig. 2 and another version were the colours were shuffled (A = purple, B = red, C = blue) and the direction of movement was right-to-left.

After watching the clip, participants proceeded to the “review” screen in which two clips were displayed side-by-side. One of the clips was the critical clip that they had just seen and the other was a clip featuring a realistic three-object collision: Object A approaches from the left at 30 mm/s and stops next to B at which point B starts moving to the right and stops next to C, following which C starts moving to the right (Fig. 2d). So, in condition 1 the participants had to choose between two clips that differed only in the order in which B and C started moving (Fig. 2a vs d), while the difference in condition 2 was mainly whether object B moved or not (Fig. 2c vs d). Below each clip there was a “play” button and participants were allowed to watch each clip as many times as they wanted before reporting which of the two they had seen in the previous screen. The position of the clips (left–right) was counterbalanced between participants. Then participants were asked to indicate their confidence by dragging a slider on a scale that was labelled “Not at all confident” to the left and “Very confident” to the rightmost position. Participants were asked for direct causal judgments for each pair of objects in the clip, as in Experiment 1, and, finally, they were asked to report any general comments they had about the experiment.

3.3. Results

The proportion of participants that correctly identified the clip they saw was 37.93% for condition 1 and 72.41% for condition 2, as shown in Fig. 5.[2] These two conditions were significantly different: $\chi^2(1, N = 58) = 6.97$, $p = .008$. The order in which the clips were seen (i.e. whether the correct clip was seen first or second) made no difference to the selected clip: $\chi^2(1, N = 58) = 0.16$, $p = .684$. As in Experiment 1, participants were confident in their response: the mean confidence rating was 74.10/100 (SD = 23.15) for condition 1 and 81.21/100 (SD = 22.40) for condition 2.

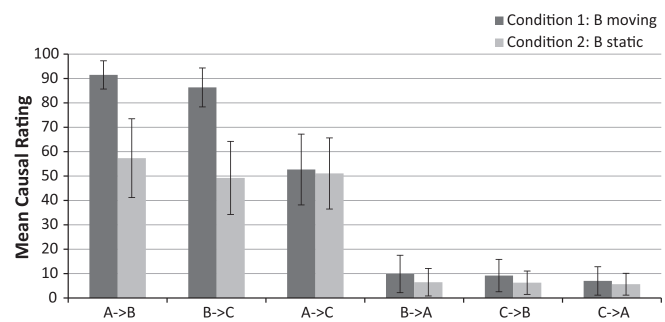

Fig. 6. Mean causal judgments for each object pair per condition. A value of 50 corresponds to neutrality – lower ratings indicate disagreement and higher rating indicate agreement with the causal statement (error bars represent 95% confidence intervals).

Regarding the causal ratings, for condition 1 they are almost identical to the respective ratings in Experiment 1: Participants agree strongly that A caused B to move and that B caused C to move, while being relatively neutral in the indirect A–C relationship and giving very low ratings to the inverse relationships. For condition 2, since B did not move at all, these causal questions are rather ambiguous and the answers participants gave reflect this ambiguity by being around the midpoint mark for all relationships with compatible temporal order (A–B, B–C and A–C). In any case there does not seem to be a prevalent causal perception in condition 2 (see Fig. 6). As in Experiment 1, the answers to the open-ended question at the end of the experiment did not reveal any systematic patterns.

3.4. Discussion

Compared to Experiment 1 the reordering effect was less pronounced, but perhaps more impressive given the different measure that we used. Participants in condition 1 saw a clip featuring relatively slow moving objects (30 mm/s compared to 300 mm/s in Michotte’s original experiments) in which object C moves 350 ms before object B, but failed to identify the clip they saw, choosing instead with high confidence a clip in which B moves earlier,

[2] Again, we collapsed the responses to the normal and the mirrored versions of the clips since no difference was observed.

and, most critically, appears to be launching C. The fact that when asked to report the order of events (Exp. 1) rather than identify the clip they saw (Exp. 2), participants show an even stronger preference for the causal order can be explained, in our view, by the nature of the measure. It is likely that some people do detect some irregularity in the clip when they first experience it, some deviation from an ideal collision clip, but they still don't identify the deviation to be the order of the events. Thus, when asked to choose between that clip and a realistic collision, these participants might prefer the deviant one, the one they actually saw, not because of the order, as the results of Experiment 1 show, but because it is the one that does not "look right". This hypothesis is supported by the direct causal judgments that remained roughly the same between the two experiments.

The most important finding from this experiment is that the reordering effect cannot be explained by split attention. When B remains stationary in condition 2 the majority of participants detect it and thus are able to correctly identify the clip they saw. This means that in condition 1, where B does move towards C, albeit late, its motion is in fact noticed and the subsequent reordering does depend on that detection. All the events that take place in the sequence are actually registered by the perceptual system and all are necessary for the causal impression to be formed; events need to be seen in order to be reordered.[3]

In addition, the current findings speak against the response bias interpretation that we put forward when discussing Experiment 1. As a reminder, we raised the possibility that participants hypothesise a computer-related error and thus provide the order that the experimenter, in their view, intended to present. Apart from the fact that no participant mentioned any such error in the open-ended question, the current experimental design provides two additional reasons to reject this hypothesis: first it is hard to imagine why participants would hypothesise a computer glitch when object B moves late but not when it remains stationary, such that the difference in the responses of the two groups can be accounted for. Second, given that the same clip is presented side-by-side with the canonical 3-object collision, even if there was the hypothesis of a glitch during the initial presentation it should be discarded at this stage, concluding that the presumed error was, after all, intentional. In other words, if participants experience a clip with a glitch and then have to select between the same clip and one without a glitch, it makes little sense to choose the latter. By the same token, we can rule out the possibility that participants in Experiment 1 perceive the objective order of events but distrust their perception and report the canonical order: the experienced version is available in Experiment 2 and subsequent viewings should strengthen the initially veridical impression resulting in the report of the objective order.[4]

4. General discussion

In two experiments we have provided evidence that phenomenal causation can influence the perception of temporal order. The causal reordering effect (Bechlivanidis & Lagnado, 2013) seems to generalize to situations where strong causal beliefs are present, irrespective of whether the causal links are recently learned or directly perceived. Moreover, the effect is not due to limitations of the perceptual system, since the relative onset of the events that are reordered is clearly perceivable in the absence of causal incongruences. Our results appear to discredit the mirroring view of order perception (Hoerl, 2013; Mellor, 1985; Phillips, 2014), since the order in which events are experienced does not match the order in which they occur. However, in order to better evaluate this conclusion, we need to determine the mechanism responsible for the erroneous temporal order judgement.

Table 1
Summary of the mechanisms through which a causal representation (schema) might influence the experienced temporal order resulting in the report of the causal rather than the objective temporal order of events.

Perceived B's behaviour?	Perceived B moving after C?	Remembered perceived order?	Role for causal schema
No	N/A	N/A	Fill in missing information
Yes	No	Yes	Influence perception
Yes	Yes	No	Distort working memory representation
Yes	No	No	Allow the inference of temporal order

4.1. Unattended, misperceived, misremembered or ignored?

Although we have produced strong evidence for the causal basis of the reordering effect, there are multiple hypotheses that can account for the effect, each postulating a different role for causal knowledge. Specifically, the collective evidence presented here and earlier (Bechlivanidis & Lagnado, 2013) suggests that participants' order judgements are influenced by stored schemata, learned representations of interactions (Zacks, Speer, Swallow, Braver, & Reynolds, 2007), causal interactions in particular (Sanborn, Mansinghka, & Griffiths, 2013; Weir, 1978; White, 2006, 2014). In the context of the current experiments, the causal schema would correspond to some abstract representation of successive collisions as experienced when balls collide or dominos fall. How does such a stored representation distort the objective order of events?

We have identified four potential roles for the causal schema[5]: as a source for filling in missing information, as a factor influencing perception, as a factor distorting working memory representations and as a premise in an inference of temporal order. As shown in Table 1, the role of the causal schema depends on whether participants notice all events that take place, whether they perceive the objective order of events at least at an early stage and whether the perceived order leaves an adequately strong trace in memory.

According to the first mechanism, participants do not attend to some of the key events that take place (in this case the behaviour of object B) because their attention is diverted to some other event (in this case the motion of object C). Subsequently the stored schema of colliding objects is used to fill in the incomplete percept, thus leading participants to report the causal order of events.[6]

As discussed, such mechanism might account for the results of Experiment 1 as well as those of earlier experiments (Bechlivanidis & Lagnado, 2013) but it is undermined by Experiment 2 which showed that participants do in fact perceive the behaviour of the middle object B: when we presented participants with a clip in which B remained stationary, the majority correctly identified it as the experienced clip against the canonical causal clip. One might insist, however, that there is a way for participants to select the causal clip when B is moving albeit late, and

[3] An alternative attention-based explanation will be discussed in Section 4.

[4] Note also, that, as described in the results section, even among the participants who reviewed the reordered clip first, the majority (64.7%) still chose the canonical clip with high confidence (86.09/100).

[5] We thank an anonymous reviewer for raising some of these issues, and helping us to clarify our position.

[6] We note, however, that if this was indeed the case one would still have to explain why participants ignore the fact that they see object C moving before B does.

select the presented clip when B remains stationary, without actually perceiving object B's behaviour. Even if the motion of object C draws attention away from B, participants might still observe B's final location and conclude that there was no collision if B remained in its original location and that there was a collision if B appears displaced. However, this strategy was not viable in the particular experimental design, since all presented clips end immediately after C completes its movement,[7] leaving no time to observe the final configuration of the objects. An even more complicated strategy, where one stops attending to C while it is still moving in order to observe B's final location, while possible, seems too far-fetched in our view, especially in a single-shot experiment. Future studies employing eye-tracking methods might provide more details regarding participants' attention, by revealing, for example, individual differences that potentially influence reported judgements. Nevertheless, the current experimental design suffices, in our view, to rule out lack of attention as the driving force for the observed reordering effect.

Assuming, therefore, that participants perceive all the events that take place, the question becomes whether the perceived order of those events is veridical or not, whether participants misperceive the order or, alternatively, whether they misremember an accurately perceived order. More specifically, given that the objective order of presentation was ACB and the order that the majority of participants reported was ABC, it is likely that a stored representation of a canonical collision intervened. Did this intervention take place online or offline: on the pathway from the retina to the working memory or on the way from working memory to the expressed behaviour? In other words is the representation stored in working memory already distorted (ABC) or does the causal schema operate on an initially veridical representation (ACB) stored in memory?

The presented experiments do not distinguish between these two mechanisms. We note, however, that it may generally be the case that such a distinction is empirically undecidable (Dennett, 1991; Dennett & Kinsbourne, 1992). Dennett rejects this distinction as "a difference that makes no difference" (Dennett, 1991, p. 132) and, moreover, argues that deciding between the two alternatives presupposes the endorsement of a Cartesian picture of the mind with a spatiotemporal locus of consciousness. Other authors maintain that even if such distinction is impossible with today's empirical methods, it is not theoretically unattainable (Block, 1992; Clark, 1992), if one, for example, was able to isolate and examine the contents of the working memory at various time points.

Both the misperception and the misremembrance interpretations assume that some spontaneous temporal order judgement is made and is influenced at different stages by top-down causal assumptions. A perhaps more radical suggestion is that no spontaneous temporal order judgement takes place, perceptually or inferentially.

There is strong evidence that humans from a very young age assume the temporal priority principle, the belief that causes precede their effects (Bullock & Gelman, 1979; Rankin & McCormack, 2013). Thus, since a causal representation has embedded temporal order information, it might be the case that causality drives directly the behaviour we have observed. In other words, either participants don't make a spontaneous temporal order judgment or if they do that judgement is inconsequential. The only efficacious representation that results from observing the target sequence is a causal representation. Subsequently, when participants are asked to provide the order of events, they infer that since B caused C and since causes precede their effects, then B occurred before C. It is interesting to note that representing the temporal order through a causal relationship contrasts with philosophical theories of time that argue for an isomorphic representation of temporal order, i.e. argue that the representation of order is itself an ordered set of representations (Mellor, 1985).

4.2. Generation of casual impressions from deviant stimuli

Up to this point we have discussed a number of alternative mechanisms explaining how a causal schema might influence temporal order judgements. How is the particular schema activated, though, and, more fundamentally, why is an impression of causality generated from the stimuli we have used?

The issue arises because our target sequence features extreme deviations from normal collisions, both temporal and spatial. Apart from the order of events, there is both a 350 ms temporal delay between object A stopping and object B starting to move and a 16 mm spatial gap between object B and object C (Fig. 2a). According to current models of causal perception (see Rips, 2011 for a review) such deviations should not lead to impressions of causality. For Michotte and the proponents of his theory (Michotte, 1963; Scholl & Tremoulet, 2000) a modular input analyser should not output a causal impression when, for example, delays exceed 150–200 ms (Fugelsang, Roser, Corballis, Gazzaniga, & Dunbar, 2005; Michotte, 1963; Straube & Chatterjee, 2010; Yela, 1952). Although schema-matching models (Sanborn, Griffiths, & Shiffrin, 2010; Weir, 1978; White, 2006, 2014), are more flexible by allowing any previously experienced sequence to drive causal perception, our target sequence is most likely an example of what White described as "stimuli that are unrepresentative of real interactions between objects in ways other than incompleteness" and thus should "not give rise to visual impressions of causality because they would not be matched against any schema" (White, 2006, p. 179).

White's (2014) more recent approach may provide a solution to this conundrum by specifying 14 low level cues to causality. Our target sequence contains 11 of those cues, though not all of them are instantiated by the appropriate objects. So, for example, there is "contact between actor and object" (cue 5) and "property transmission" (cue 8) despite the fact that upon contact between objects A and B, the momentum is phenomenally transmitted to a third object C.

Nevertheless, if the presence of those cues in the sequence as a whole triggers the search for a causal interpretation then ignoring spatiotemporal deviations and more interestingly reordering the events will result in the simplest representation of the observed sequence (Chater & Vitányi, 2003; Lombrozo, 2007). Rather than object A launching C from distance and object B moving spontaneously or being pulled by C, A will be represented as launching B followed by B launching C. The latter interpretation is clearly simpler by involving two instances of a single type of causal relationship and furthermore by matching a causal schema, similar to a queue of dominos falling.

In sum, we propose that the detection of abstract low level cues to causality (White, 2014) in a sequence as a whole triggers the search for familiar causal representations (schemata). If compatible schemata matching the sensory data are unavailable, the simplest among similar representations will be selected. Subsequently, the causal representation influences the experienced temporal order, at the time of perception or retrieval, or alternatively it completely overrides the need to spontaneously generate a temporal order judgement, since the order of events is implicitly represented in the causal relationship.

[7] In fact, one cannot tell whether C would stop when the clip ends or whether it would continue its motion.

Appendix A. Ensuring consistent presentation in a web-based experiment

Both experiments were developed with Adobe Flex 4.6 and conducted over the Internet. Subjects were recruited using Amazon Mechanical Turk. While the validity of using Mechanical Turk as a subject pool and its comparability to traditional lab based methods has been demonstrated in multiple occasions (Buhrmester, Kwang, & Gosling, 2011; Crump, McDonnell, & Gureckis, 2013; Gosling, Vazire, Srivastava, & John, 2000; Mason & Suri, 2012), here we went a step further by running a real-time perceptual experiment featuring dynamic sequences over the web. Such experiments are usually conducted in controlled conditions in order to ensure the uniformity of the presented stimuli (although see Germine et al., 2012; Hecht, Oesker, Kaiser, Civelek, & Stecker, 1999). So, how can we achieve a similar level of uniformity given that we have no control over the hardware or software (i.e. browser) used to display our stimuli?

Our approach was twofold: first, we took special measures to limit the potential stimulus variability prospectively and then, we recorded a number of variables while the experiment was running. More specifically, the main source of variability for online applications is both hardware and software related and may result in deviations both in the size of the presented objects and in the temporal duration of events. In terms of software, we chose the Flex SDK that allows for a minimum level of control over the timing of the stimuli while it compiles into the SWF file format and targets Flash Player making it browser-independent.

A.1. Temporal variability

Regarding the timing of the stimuli, although for the particular experiments small deviations would be acceptable, we recorded the minimum effective frame rate (fps)[8] while the critical sequence was presented as well as the effective delay between the onset of motion of objects C and B.

There were minimal deviations: While the target and the maximum framerate was 30 fps, the average minimum framerate was 29.70 fps (SD = 0.36) in Experiment 1 and 29.63 fps (SD = 0.39) in Experiment 2. Similarly, with the intended delay between the movements of objects C and B being 350 ms, the average recorded delay was 351.03 ms (SD = 5.22) in Experiment 1 and 351.00 (SD = 5.13) in Experiment 2.

A.2. Size variability: calibration section

Turning now to the issue of enforcing a consistent size of objects, the main problem is the variable monitor sizes and especially the variable ppi (pixels per inch). Although the area where the clip sequences were displayed measured 980×400 pixels, a pixel has varying dimensions depending on the exact ppi and this would result in variable object sizes. Without any way to access the actual ppi value, especially from within a web browser, we had to resort to more practical solutions.

Before proceeding to the actual experiment every participant had to go through a calibration session (programmed using DHTML+Ajax). This involved using an optical disc (CD, DVD, etc), a credit card or a dollar note (all participants were from the US) in order to match the size of the respective virtual object that appeared on screen. The participant would place, for example, her credit card on her screen where a virtual credit card was displayed and would use the provided controls to increase or decrease the size of the virtual card so that the virtual and the actual cards matched in size (the same would apply if using a dollar note or an optical disc).

Given that the size of those particular objects is standard, we then compared the size of the actual object against that of the virtual object as set by the participant, in order to derive the effective ppi. This value was then used to define the size of the objects and the distances in our experiments, thus ensuring consistency of stimuli among participants.

After the calibration section and before proceeding to the actual experiment, participants had to answer two further questions. The first displayed a horizontal line on the screen and required from the participants to use a physical ruler in order to measure it and then input their measurement. The size of the line was dependent on the ppi value derived during the calibration section. Participants had to input the correct width in order to proceed to the experiment, otherwise they were directed back to the calibration screen.

The second question asked participants to report their approximate physical distance from their monitor. The two options were: "I can more or less touch the screen if I extend my arm" and "I am further away from the screen". A single participant in Experiment 1 and no participants in Experiment 2 picked the second option.

A.3. Restricting participation

With online studies there is always the risk that some participants will attempt to do the experiment more than once. To prevent that from happening we advised participants against attempting to redo the experiment and we used a combination of their IP address and their Mechanical Turk Worker ID. The IP was recorded on our server after the initial instructions and before presenting any stimuli. Participants were warned against using the back button on their browser or refreshing the page as this would result in the termination of the experiment. Participants with an IP address that was already recorded on our server were not allowed to participate.

Of course, most Internet Service Providers assign dynamic IP addresses therefore it is possible that the IP of a participant would change (possible but improbable especially within a short period of time). However, all the experiments reported here constituted a single "Batch" in Mechanical Turk, meaning that no subject with the same Worker ID could participate twice.

Thus in order for someone to participate twice in our experiments he or she would have to ignore our request, maintain two or more Mechanical Turk accounts and use computers in different networks or find a way to renew their external IP address.

A.4. Repeating the analysis

Despite the measures that we took to limit the variability in the size of objects and the timing of events, we repeated the reported analyses after removing participants who deviated from the intended values. Specifically, in a second analysis, we excluded participants who stated that they were seated far from the monitor (1 in Exp. 1, 0 in Exp. 2) or for whom the minimum frame rate dropped below 29 fps at any point during the critical sequence or for whom the recorded delay deviated by more than 10% (i.e. 35 ms) from the 350 ms target (10 in Exp. 1 and 6 in Exp. 2). The results of this stricter analysis remained the same for both experiments.

[8] By framerate, here, we refer to the number of frames per second that our software attempted to draw on the screen. Given that the maximum framerate that we used was 30 fps and the refresh rate of modern displays is at least 50 Hz, we were confident that any variability in the refresh rate would not affect our stimuli.

References

Abrams, R. A., & Christ, S. E. (2003). Motion onset captures attention. *Psychological Science, 14*(5), 427–432. http://dx.doi.org/10.1111/1467-9280.01458.

Bechlivanidis, C., & Lagnado, D. A. (2013). Does the "why" tell us the " when"? *Psychological Science, 24*(8), 1563–1572. http://dx.doi.org/10.1177/0956797613476046.

Block, N. (1992). Begging the question against phenomenal consciousness. *Behavioral and Brain Sciences, 15*, 205–206. http://dx.doi.org/10.1017/S0140525X00068266.

Buehner, M. J. (2012). Understanding the past, predicting the future: Causation, not intentional action, is the root of temporal binding. *Psychological Science, 23*(12), 1490–1497. http://dx.doi.org/10.1177/0956797612444612.

Buehner, M. J., & Humphreys, G. R. (2009). Causal binding of actions to their effects. *Psychological Science, 20*(10), 1221–1228. http://dx.doi.org/10.1111/j.1467-9280.2009.02435.x.

Buehner, M. J., & Humphreys, G. R. (2010). Causal contraction: Spatial binding in the perception of collision events. *Psychological Science, 21*(1), 44–48. http://dx.doi.org/10.1177/0956797609354735.

Buhrmester, M., Kwang, T., & Gosling, S. D. (2011). Amazon's Mechanical Turk: A new source of inexpensive, yet high-quality, data? *Perspectives on Psychological Science, 6*(1), 3–5. http://dx.doi.org/10.1177/1745691610393980.

Bullock, M., & Gelman, R. (1979). Preschool children's assumptions about cause and effect: Temporal ordering. *Child Development, 50*(1), 89–96. http://dx.doi.org/10.2307/1129045.

Chater, N., & Vitányi, P. (2003). Simplicity: A unifying principle in cognitive science? *Trends in Cognitive Sciences, 7*(1), 19–22. http://dx.doi.org/10.1016/S1364-6613(02)00005-0.

Clark, A. (1992). Experiential facts? *Behavioral and Brain Sciences, 15*, 207–208. http://dx.doi.org/10.1017/S0140525X00068266.

Crump, M. J. C., McDonnell, J. V., & Gureckis, T. M. (2013). Evaluating Amazon's Mechanical Turk as a tool for experimental behavioral research. *PLoS ONE, 8*(3), e57410. http://dx.doi.org/10.1371/journal.pone.0057410.

Dainton, B. (2010). Temporal consciousness. In E. N. Zalta (Ed.), *The Stanford encyclopedia of philosophy (Spring 2014 edition)*. <http://plato.stanford.edu/archives/spr2014/entries/consciousness-temporal/>.

Dennett, D. C. (1991). *Consciousness explained*. UK: Penguin.

Dennett, D. C., & Kinsbourne, M. (1992). Time and the observer: The where and when of consciousness in the brain. *Behavioral and Brain Sciences, 15*, 183–247. http://dx.doi.org/10.1017/S0140525X12000477.

Eagleman, D. M., & Holcombe, A. O. (2002). Causality and the perception of time. *Trends in Cognitive Sciences, 6*(8), 323–325. http://dx.doi.org/10.1016/S1364-6613(02)01945-9.

Fugelsang, J. A., Roser, M. E., Corballis, P. M., Gazzaniga, M. S., & Dunbar, K. N. (2005). Brain mechanisms underlying perceptual causality. *Cognitive Brain Research, 24*(1), 41–47. http://dx.doi.org/10.1016/j.cogbrainres.2004.12.001.

Germine, L., Nakayama, K., Duchaine, B. C., Chabris, C. F., Chatterjee, G., & Wilmer, J. B. (2012). Is the Web as good as the lab? Comparable performance from Web and lab in cognitive/perceptual experiments. *Psychonomic Bulletin and Review, 19*(5), 847–857. http://dx.doi.org/10.3758/s13423-012-0296-9.

Gosling, S. D., Vazire, S., Srivastava, S., & John, O. P. (2000). Should we trust web-based studies? A comparative analysis of six preconceptions about internet questionnaires. *The American Psychologist, 59*(2), 93–104. http://dx.doi.org/10.1037/0003-066X.59.2.93.

Grush, R. (2007). Time and experience. In T. Müller (Ed.), *The philosophy of time*. Frankfurt: Klosterman. <http://mind.ucsd.edu/papers/time&exp/time&exp.pdf>.

Hecht, H., Oesker, M., Kaiser, A., Civelek, H., & Stecker, T. (1999). A perception experiment with time-critical graphics animation on the World-Wide Web. *Behavior Research Methods, Instruments, & Computers, 31*(3), 439–445. http://dx.doi.org/10.3758/BF03200724.

Heron, J., Hanson, J. V. M., & Whitaker, D. (2009). Effect before cause: Supramodal recalibration of sensorimotor timing. *PLoS ONE, 4*(11), e7681. http://dx.doi.org/10.1371/journal.pone.0007681.

Hillstrom, A. P., & Yantis, S. (1994). Visual motion and attentional capture. *Perception and Psychophysics, 55*(4), 399–411. http://dx.doi.org/10.3758/BF03205298.

Hoerl, C. (2013). A succession of feelings, in and of itself, is not a feeling of succession. *Mind, 122*(486), 373–417. http://dx.doi.org/10.1093/mind/fzt070.

Kim, S.-H., Feldman, J., & Singh, M. (2013). Perceived causality can alter the perceived trajectory of apparent motion. *Psychological Science, 24*(4), 575–582. http://dx.doi.org/10.1177/0956797612458529.

King, A. J. (2005). Multisensory integration: Strategies for synchronization. *Current Biology, 15*(9), R339–R341. http://dx.doi.org/10.1016/j.cub.2005.04.022.

Lombrozo, T. (2007). Simplicity and probability in causal explanation. *Cognitive Psychology, 55*(3), 232–257. http://dx.doi.org/10.1016/j.cogpsych.2006.09.006.

Mason, W., & Suri, S. (2012). Conducting behavioral research on Amazon's Mechanical Turk. *Behavior Research Methods, 44*(1), 1–23. http://dx.doi.org/10.3758/s13428-011-0124-6.

Mellor, D. H. (1985). *Real time*. Cambridge: Cambridge University Press.

Michotte, A. (1963). *The perception of causality*. London: Methuen & Co Ltd..

Phillips, I. B. (2014). Experience of and in time. *Philosophy Compass, 9*(2), 131–144. http://dx.doi.org/10.1111/phc3.12107.

Rankin, M. L., & McCormack, T. (2013). The temporal priority principle: At what age does this develop? *Frontiers in Psychology, 4*. http://dx.doi.org/10.3389/fpsyg.2013.00178.

Rips, L. J. (2011). Causation from perception. *Perspectives on Psychological Science, 6*(1), 77–97.

Rohde, M., Scheller, M., & Ernst, M. O. (2014). Effects can precede their cause in the sense of agency. *Neuropsychologia, 65*, 191–196. http://dx.doi.org/10.1016/j.neuropsychologia.2014.10.011.

Sanborn, A. N., Griffiths, T. L., & Shiffrin, R. M. (2010). Uncovering mental representations with Markov chain Monte Carlo. *Cognitive Psychology, 60*(2), 63–106. http://dx.doi.org/10.1016/j.cogpsych.2009.07.001.

Sanborn, A. N., Mansinghka, V. K., & Griffiths, T. L. (2013). Reconciling intuitive physics and Newtonian mechanics for colliding objects. *Psychological Review, 120*(2), 1–77. http://dx.doi.org/10.1037/a0031912.

Scholl, B. J., & Nakayama, K. (2004). Illusory causal crescents: Misperceived spatial relations due to perceived causality. *Perception, 33*(4), 455–469. http://dx.doi.org/10.1068/p5172.

Scholl, B. J., & Tremoulet, P. D. (2000). Perceptual causality and animacy. *Trends in Cognitive Sciences, 4*(8), 299–309. http://dx.doi.org/10.1016/S1364-6613(00)01506-0.

Schutz, M., & Kubovy, M. (2009). Causality and cross-modal integration. *Journal of Experimental Psychology: Human Perception and Performance, 35*(6), 1791–1810. http://dx.doi.org/10.1037/a0016455.

Shimojo, S. (2014). Postdiction: Its implications on visual awareness, hindsight, and sense of agency. *Frontiers in Psychology, 5*, 1–19. http://dx.doi.org/10.3389/fpsyg.2014.00196.

Simons, D. J., & Chabris, C. F. (1999). Gorillas in our midst: Sustained inattentional blindness for dynamic events. *Perception, 28*, 1059–1074. <https://www.cnbc.cmu.edu/~behrmann/dlpapers/Simons_Chabris.pdf>.

Soteriou, M. (2010). Perceiving events. *Philosophical Explorations, 13*(3), 223–241. http://dx.doi.org/10.1080/13869795.2010.501904.

Spence, C., & Parise, C. (2010). Prior-entry: A review. *Consciousness and Cognition, 19*(1), 364–379. http://dx.doi.org/10.1016/j.concog.2009.12.001.

Spence, C., & Squire, S. (2003). Multisensory integration: Maintaining the perception of synchrony. *Current Biology, 13*(13), R519–R521. http://dx.doi.org/10.1016/S0960-9822(03)00445-7.

Stetson, C., Cui, X., Montague, R. P., & Eagleman, D. M. (2006). Motor-Sensory recalibration leads to an illusory reversal of action and sensation. *Neuron, 51*(5), 651–659. http://dx.doi.org/10.1016/j.neuron.2006.08.006.

Straube, B., & Chatterjee, A. (2010). Space and time in perceptual causality. *Frontiers in Human Neuroscience, 4*, 1–10. http://dx.doi.org/10.3389/fnhum.2010.00028.

Titchener, E. B. (1908). *Lectures on the elementary psychology of feeling and attention*. New York: Macmillan.

Vibell, J., Klinge, C., Zampini, M., Spence, C., & Nobre, A. C. (2007). Temporal order is coded temporally in the brain: Early event-related potential latency shifts underlying prior entry in a cross-modal temporal order judgment task. *Journal of Cognitive Neuroscience, 19*(1), 109–120. http://dx.doi.org/10.1162/jocn.2007.19.1.109.

Weir, S. (1978). The perception of motion: Michotte revisited. *Perception, 7*(3), 247–260. http://dx.doi.org/10.1068/p070247.

White, P. A. (2006). The role of activity in visual impressions of causality. *Acta Psychologica, 123*(1–2), 166–185. http://dx.doi.org/10.1016/j.actpsy.2006.05.002.

White, P. A. (2014). Singular clues to causality and their use in human causal judgment. *Cognitive Science, 38*(1), 38–75. http://dx.doi.org/10.1111/cogs.12075.

Yela, M. (1952). Phenomenal causation at a distance. *Quarterly Journal of Experimental Psychology, 4*(4), 139–154. http://dx.doi.org/10.1080/17470215208416612.

Zacks, J. M., Speer, N. K., Swallow, K. M., Braver, T. S., & Reynolds, J. R. (2007). Event perception: A mind/brain perspective. *Psychological Bulletin, 133*(2), 273–293. http://dx.doi.org/10.1037/0033-2909.133.2.273.