Temporal predictability enhances judgements of causality in elemental causal induction from both observation and intervention

Greville, William; Buehner, Marc

Published in:
The Quarterly Journal of Experimental Psychology

DOI:
10.1080/17470218.2015.1041535

Publication date:
2015

Citation for published version (APA):
Greville, W., & Buehner, M. (2015). Temporal predictability enhances judgements of causality in elemental causal induction from both observation and intervention. *The Quarterly Journal of Experimental Psychology*, *69*(4), 678-97. https://doi.org/10.1080/17470218.2015.1041535

Running head: Temporal Predictability

Temporal Predictability Enhances Judgments of Causality in Elemental Causal Induction from both Intervention and Observation

W. James Greville[1*], Marc J. Buehner[2]

[1]College of Medicine, Swansea University, SA2 8PP. E-mail: w.j.s.greville@swansea.ac.uk

[2]School of Psychology, Cardiff University, CF10 3AT. E-mail: buehnerm@cf.ac.uk

*Author to whom correspondence should be addressed.

Acknowledgements: This work was supported by a doctoral training grant awarded by the EPSRC.

Abstract

When the temporal interval or delay separating cause and effect is consistent over repeated instances, it becomes possible to predict when the effect will follow from the cause, hence *temporal predictability* serves as an appropriate term to describe consistent cause-effect delays. Greville & Buehner (2010) demonstrated that in instrumental action-outcome learning tasks, enhancing temporal predictability by holding the cause-effect interval constant elicited higher judgments of causality compared to conditions involving variable temporal intervals. Here, we examine whether temporal predictability exerts a similar influence when causal learning takes place through observation as well as through intervention by instrumental action. Four experiments consistently demonstrated that judgments of causality were higher when the temporal interval was constant than when it was variable, and that judgments declined with increasing variability. We further found that this beneficial effect of predictability was stronger in situations where the effect base-rate was zero (Experiments 1 and 3). The results therefore clearly indicate that temporal predictability enhances impressions of causality, and that this effect is robust and general. Factors that could mediate this effect are discussed.

Keywords: Causality, Predictability, Contiguity, Delay, Causal Learning

Temporal Predictability Enhances Judgments of Causality in Elemental Causal Induction from both Intervention and Observation

The capacity to infer causality allows humans and other intelligent agents to understand and interact with their environment. In order to learn the relationship between causes and effects, we rely upon a number of cues to causality (Buehner & Cheng, 2005; Einhorn & Hogarth, 1986), which are temporal order (effects must follow their causes and not vice versa), contingency (effects must follow their causes with a degree of consistency), and temporal contiguity (effects must follow their causes rapidly). With temporal order typically assumed as a necessity for causal learning, and the vast majority of research focusing on the role of contingency or statistical information, the role of temporal contiguity in causal learning is the least well characterised.

The majority of theories of causal learning, however, agree that the stronger the degree of temporal contiguity between two events, the stronger the perceived causal relationship between them will be, while conversely delays – a lack of contiguity – are detrimental to causal learning. There are exceptions, such as when contiguity is implausible given the situation (e.g. Buehner & McGregor, 2006), and there are cases where a lack of contiguity is no barrier to learning the causal relationship (e.g. Buehner & May, 2004). As a general rule though, contiguous causal relations are easier to detect and are judged as stronger than delayed causal relations, and the longer the delay or interval between cause and effect, the weaker this relationship will be perceived to be (Shanks, Pearson, & Dickinson, 1989; Wasserman, Chatlosh, & Neunaber, 1983).

Most research in contiguity focuses only on the extent of delay – whether intervals are short or long. However, causal relations are rarely experienced as one-off pairings; more usually we will experience multiple instances of cause and effect. Consequently, we also experience multiple cause-effect intervals. Focusing purely on the extent of delay then either assumes that delays are consistent over time, or simply takes the average delay across cause-effect pairings. In doing so, any putative influence of the variability of delays across multiple cause-effect pairings is ignored. Recent studies (Greville & Buehner, 2010; Lagnado & Speekenbrink, 2010) have attempted to address this oversight, examining the effect of variable delays in causal learning. The purpose of the current paper is to extend this line of research, and attempt to resolve some discrepancies between these studies.

Temporal Predictability

Griffiths and Tenenbaum (2009) point to the discovery of Halley's comet as a striking example of causal induction through the use of prior knowledge, theories, and in particular temporal predictability. Halley noted that comets observed in 1531, 1607, and 1682 had all taken remarkably similar paths across the sky, and using the principles of Newtonian physics, Halley inferred that the three comets previously observed were in fact one and the same following a regular solar orbit. Doubtless, Halley's theoretical understanding was crucial to this successful calculation. However, perhaps the most potent clue to this discovery was that the three comets had been observed approximately 76 years apart from one another in each case. In other words, there was a consistent temporal interval between the three appearances of the comet, that varied (in relative terms) minimally. It was this periodicity that allowed Halley to predict that the comet would return again in 1758 and indeed this prediction proved to be accurate, with Halley's comet visiting Earth every 76 years since. This facility of consistent timing, to enable predictions regarding the occurrence of future events and specifically when those events will occur, makes "temporal predictability" an apt term to describe such a feature.

Greville and Buehner (2010) recently demonstrated that temporal predictability enhanced judgments of causality in action-outcome learning. Participants investigated the extent to which they could make a triangle on a computer screen light up (the effect) by clicking on a button (the cause), in experiments based on earlier work by Shanks, Pearson and Dickinson (1989). The key manipulation in Greville and Buehner's studies was to control the variation of the temporal interval between the cause and effect. In certain conditions, the temporal interval was fixed, thus conferring maximal temporal predictability, and these conditions constantly received higher causal ratings than corresponding conditions with variable temporal intervals (but with equivalent mean delays overall). Furthermore, the greater the variability of the delay from one instance to the next, the lower the causal ratings provided by participants. Other explanations for this effect such as variations in action rate or the amount of time participants spent investigating the relationship were ruled out. Greville and Buehner concluded that temporal predictability enhances impressions of causality, whilst temporal variability impedes causal learning. Here, we extend this work by investigating whether the same holds true when learning from observing sequences of putative causes and effects rather than directly intervening to generate causal events.

Observation vs. Intervention

Instrumental learning tasks such as those used by Greville and Buehner (2010), trace their heritage to operant conditioning studies with animals (e.g. Ferster & Skinner, 1957). In such tasks, a putative causal link in the environment is actively investigated through the performance of an operant response (i.e. instrumental action) and monitoring the apparent consequences (action-outcome). Causal relations may, of course, also be uncovered by passively observing the occurrence of different stimuli (cue-outcome), analogously to Pavlovian or classical conditioning. Given the inherent similarities between instrumental and classical conditioning, it may be tempting to assume that Greville and Buehner’s findings in instrumental learning will generalize to observational situations. As a case in point, the introductory example of the discovery of Halley’s comet could well be regarded as a demonstration of temporal predictability facilitating causal learning from observation.

However, a great deal of recent research in causal learning has emphasized the special status of actions in intervention and the distinction between learning either from intervention or from observation (Lagnado & Sloman, 2004, 2006; Sloman & Lagnado, 2005; Waldmann, 1996, 2000; Waldmann & Hagmayer, 2005; Waldmann & Holyoak, 1992, 1997). Intervention – performing an instrumental action on a system to modify the value of a variable – affords different predictions compared to when the value of a variable is merely observed. By deliberately intervening on the environment, an organism can control the frequency or rate of their interventions, as well as their pattern or temporal distribution, intensity, strength and so forth. Simply put, patterns of intervention are self-governed, and choices can modulate the data that are received (Lagnado & Sloman, 2006). Learning from observation meanwhile may intuitively seem more difficult, since the occurrence of stimuli is beyond the control of the observer. Temporal predictability in particular might be easier to detect under instrumental rather than observational conditions. Under the former, one can produce meaningful or memorable action patterns or rhythms and then monitor the stream of outcomes to see if a similarly matching pattern occurs. This could be on as simple a level as comparing ratios of rates or frequencies (that is, comparing number of outcomes to number of actions) but could also involve more complex comparisons such as whether the specific timing of outcomes mirrors the pattern of actions (or to what degree the patterns have similar temporal distributions). Meanwhile, when learning through observation alone, one would have to wait for such meaningful (or at least recognizable) patterns to be generated by the environment or an alternative agent. Interventional learning thus may promote more directed hypothesis testing, as someone who repeatedly intervenes on a system is in a better position to test their own hypotheses than someone who merely observes the system. Sobel and Kushnir (2006) for instance demonstrated that people were better at learning causal models when they observed intervention data that they had generated themselves, compared to observing intervention data generated by another person. It is, therefore, far from a foregone conclusion that the facilitatory effect of temporal predictability observed in Greville and Buehner's (2010) experiments will manifest in observational learning.

Indeed, Young and Nguyen (2009) reported findings suggesting that temporal predictability does not facilitate causal learning from observation and may in fact exert the opposite influence. Young and Nguyen devised a task using a "first person shooter" game where participants were required to observe the occurrence of three candidate causes (characters firing weapons) and a single effect (an explosion), and then identify one of the candidates as the true cause of this effect. They found that when the temporal interval between the true cause and its effect was variable, participants correctly selected the true cause more often than when the interval was fixed. In other words, temporal *variability* rather than *predictability* promoted more accurate causal judgments. This undermined an earlier "predictability hypothesis" proposed by Young, Rogers and Beckman (2005) that temporal predictability should facilitate causal learning.

Conversely, Lagnado and Speekenbrink (2010; Experiment 1) found a beneficial effect of temporal predictability in observational causal learning. Their task required participants to judge whether the presence of certain features in bacteria (feelers, spots, or tail, indicated by pictorial stimuli) caused the outcome of stomach cramps (indicated by a flash of the screen). On a given trial, a bacterium was presented on the screen for 2s before disappearing. Each of the three features could be either present or absent on a given trial, with one (target) feature increasing the probability of the outcome. The variability of the delay was manipulated within-subjects, with delays on a given trial randomly sampled from a lognormal distribution either with a standard deviation of 0.1s (low delay) or 1 s (high delay). Lagnado and Speekenbrink found that causal ratings for the target feature were higher when the cause-effect delay was less variable. These results were consistent with Greville and Buehner's (2010) findings and suggest that temporal predictability is important in causal learning from observation as well as intervention.

The question then is why this discrepancy should exist between Young and Nguyen's (2009) results and those of Lagnado and Speekenbrink (2010), since both involve causal learning from observation. It is worth pointing out that both of these tasks differed somewhat from the typical causal learning task (Greville & Buehner, 2010; Shanks et al., 1989; Wasserman et al., 1983) in which participants investigate the putative relationship between a single candidate cause and a single candidate effect, and provide a numerical rating reflecting their assessment of the reliability or strength of the cause-effect relationship. Griffiths and Tenenbaum (2005) termed this decision "elemental causal induction" which is effectively a process of assessing the evidence for each of two hypotheses; one where there is a genuine causal link between putative cause and effect, and one where there is not. Lagnado and Speekenbrink's task instead involved assessing the relative contribution of three candidate causes in terms of increasing the likelihood of an outcome, while Young and Nguyen's tasks required participants to make a forced choice as to which of three candidates was the most likely cause. The hypothesis space in these tasks thus differed from that in the standard causal learning paradigm. In other words, although these tasks clearly involved causal decision making, they did not entail elemental causal induction.

The status of temporal predictability in causal learning from observation is therefore unclear. It may be the case that temporal predictability assists causal learning only in certain forms of causal decision making, but that it has no effect or even the opposite effect in other contexts. What is clear though is that the issue warrants further investigation, and the work presented in this article is a first step towards this. The following experiments attempted to definitively determine whether temporal predictability facilitates elemental causal induction in learning from observation.

Experiment 1

To clarify the role of temporal predictability in elemental causal induction, we implemented an observational analogue of the causal learning task used in Greville and Buehner (2010). Our approach was to effectively replicate Greville and Buehner's experiments, with the crucial distinction of requiring participants to simply observe a sequence of candidate causes and effects rather than generate them themselves by intervention. Greville and Buehner recorded the timing of all events during their experiments, and using this information it is therefore possible to simply "play back" the exact same sequence of actions and outcomes experienced by a previous participant, to be observed by a new participant in another experiment. We elected to use Experiment 2 from Greville and Buehner (2010) as the basis for our first experiment, as this experiment had the strongest effect of temporal predictability and largest sample size of the experiments in the series. This experiment included six non-contingent control conditions, in which there was no relation between action and outcome, and participants were easily able to distinguish between these and the "master" conditions in which the outcome was contingent on the action. For simplicity and economy, we therefore excluded the control conditions and presented only the event sequences from the master conditions in this observational study.

Method

Participants:

Thirty-three undergraduate psychology students from Cardiff University (5 male) with a median age of 19 years participated in exchange for partial course credit.

Design:

We exposed participants to sequences of causes and effects from six different experimental conditions in Greville and Buehner's (2010) Experiment 2. Two independent variables were employed; extent of delay (short: 3s or long: 6s) and variability of delay (zero, intermediate, or high), combining in a 2×3 within-subjects design to produce these six conditions. The key feature of interest was thus the temporal interval between cause and effect. In Greville and Buehner's experiments, the effect followed the cause with a probability of 75% i.e. $P(e|c) = 0.75$ and the effect never occurred independently of the cause i.e. $\underline{P}(e|\neg c) = 0$. If an effect was scheduled, it occurred after a temporal interval that was determined by the combination of the factors extent and variability of delay. If, for instance, variability was zero and the temporal interval was thus constant throughout the condition (i.e. maximal predictability), then the effect always occurred after either 3s (short delay conditions) or 6s (long delay conditions). If variability was intermediate, then the temporal interval could take values within a range of 3s about the 3s or 6s midpoint (i.e. 1.5-4.5s or 4.5-7.5s), while if variability was high the temporal interval could take values within a range of 6s (i.e. 0-6s or 3-9s), with all values equiprobable within the defined range. For example then, in a short delay/high variability condition, the temporal interval for any given cause effect pairing could range from 0-6s, while for a long delay/zero variability condition the temporal interval was always 6s.

Apparatus, Materials & Procedure:

Participants were tested in groups in a small computer lab, seated at individual workstations, separated by partitions. Each participant used a PC with a 19” LCD widescreen display and a standard mouse and keyboard to engage with the experiment, which was programmed in Python 2.4.

After being welcomed by the experimenter and giving their consent to participate, participants read on-screen instructions outlining the nature of the task, and then began the experiment. In each condition, an outline of triangle was presented in the centre of the screen, with images of a push-button beneath the triangle, and a pointing finger alongside this button. Every so often, according to the replayed schedule, the finger moved and pressed the button (which then illuminated for 250ms), before both button and finger returned to their previous state. This sequence constituted an instance of the candidate cause. If an effect was scheduled, the triangle lit up, also for 250ms.

The occurrence of causes and effects was simply a carbon copy of a recorded action and outcome schedule that was previously generated and experienced by a selected participant from Greville and Buehner’s (2010) Experiment 2. Occurrence of effects was therefore not determined anew using a probability schedule but instead matched the pattern in the recorded data. However it was necessary to ensure that the schedule that was presented consisted of useful evidence: The schedule must comprise sufficient pairings of cause and effect so that the statistical and temporal relationship between them is tangible. At the same time, if cause density is too high then the true causal relationship may be obscured, as it is necessary for the encountered data stream to contain periods where no causes are administered so that the baseline occurrence of the effect can be determined. Consequently, we decided to restrict ourselves to playing back only a subset of Greville & Buehner's original recorded data streams. Specifically, we discarded sets where the total number of instrumental actions performed by participants was in either the top or bottom quartile, thus excluding event streams containing too few or too many causal actions to provide meaningful data. This still left a total of 30 different data sets from the middle two quartiles. For each participant in Experiment 1, the computer randomly chose one set with replacement at the beginning of the session.

Participants engaged in six conditions as described above, presented in the same order as experienced by the previous participant (which was randomly determined in Greville and Buehner's original study), with each condition lasting 120 seconds. At the end of each, the screen cleared and participants were asked: "On a scale of 0–100, how effective was pressing the button at causing the triangle to light up?" Participants then typed in their rating, and progressed to the next condition. In total the experiment lasted around 20 minutes.

Results & Discussion

Figure 1 shows the mean causal ratings provided by participants for the six different conditions in Experiment 1. The maximally predictable conditions, where the temporal interval was invariant, received the highest ratings. Furthermore, ratings tended to decline as variability increased and temporal predictability was lost. The effect of delay is less apparent; while ratings were noticeably higher for short delays than long delays where variability was high, conditions with short and long delays received close to identical mean ratings where there was medium or zero variability.

For these and all subsequent analyses, we adopted an α-value of 0.05 and applied Greenhouse-Geisser correction to the degrees of freedom where appropriate. A 3×2 repeated measures ANOVA confirmed a significant main effect of delay variability on causal ratings, $F(1.798, 57.52) = 7.410$, $MSE = 434.7$, $p = .002$, $\eta_p^2 = .188$, with only the linear component reaching significance, $F(1,32) = 11.11$, $MSE = 504.5$, $p = .002$, $\eta_p^2 = .258$. No significant effect of delay extent was obtained, $F(1,32) = 0.546$, $MSE = 715.2$, $p = .465$, $\eta_p^2 = .017$, nor was there a significant interaction between variability and extent of delay, $F(1.995, 63.83) = 0.656$, $MSE = 474.1$, $p = .522$, $\eta_p^2 = .020$.

As the occurrence and timing of effects in the experimental paradigm was determined probabilistically, some small deviation of experienced values from programmed values is expected. It was thus important to verify that our results were not confounded by any unplanned variations in experienced cause-effect contingency or delay. Table 1 reports the means and standard deviations of occurrence rates of causes and effects[1], overall experienced $P(e|c)$, and overall experienced cause-effect delay, across participants, for all experimental conditions in Experiment 1, together with the causal ratings provided. Here (and in Experiment 3), $P(e|c)$ was calculated as the total number of effects over the total number of causes. Experienced $P(e|c)$ did not vary with either delay variability, nor delay extent, nor was there a significant interaction (all

---

[1] In Experiments 1, 3 and 4, observed data were derived from a free-operant paradigm and as such some variation in the rates of occurrence of causes (and effects) across experimental conditions was expected. However, cue and outcome density effects (e.g. Allan & Jenkins, 1983) whereby higher rates of event occurrence tend to elicit higher causal ratings, are typically limited to discrete trials procedures and not found using the free-operant procedure (Msetfi, Murphy, Simpson, & Kornbrot, 2005; Wasserman et al., 1983; Wasserman, Elek, Chatlosh, & Baker, 1993) and hence any such variations are not examined further in this paper.

*F*s<1.7). In this and all subsequent experiments, experienced delays were always a strong function of programmed delays as expected, and so we have not examined this further. Experienced delays did not vary as a function of delay variability, whilst delay extent and variability did not interact (both *F*s<0.3). We can thus be confident that causal ratings in this experiment were not influenced by unplanned variations in contingency or contiguity.

In sum, Experiment 1 thus produced results that are indicative of a facilitatory effect of temporal predictability on judgments of causality when learning from observation, consistent with Greville and Buehner's (2010) earlier findings from instrumental learning. Specifically, conditions with fixed temporal intervals were evaluated more favourably than those with variable temporal intervals and ratings declined as variability (and thus temporal uncertainty) increased. The implication is that the facilitatory effects of predictability seen in instrumental learning can indeed generalize to observational learning, at least in the particular case of elemental causal induction in a simple learning environment.

However, the lack of an effect of delay extent is rather surprising. Certainly, there is a plethora of studies in the literature that have previously demonstrated detrimental effects of delays in learning, both in human judgments of causality (Shanks et al., 1989) and conditioning in animals (Grice, 1948; Williams, 1976), and this is now a well-established finding. Robust and consistent effects of delay were also found in all of Greville and Buehner's (2010) instrumental studies. The failure to find an effect of delay extent here is therefore a cause for some concern. Prior research has demonstrated that the effects of delays may be mitigated, by prior knowledge (Buehner & May, 2002, 2004), experience (Buehner & May, 2003), or additional cues bridging the temporal gap (Young et al., 2005) or revealing hidden trial structure (Greville, Cassar, Johansen, & Buehner, 2013). Yet no such additional information was provided in the current study. What then could have attenuated the impact of delays?

One possibility is that the task was too easy for participants and they did not need to rely on contiguity as a cue. Note that Experiment 2 of Greville and Buehner (2010) included control conditions in order to provide a challenge for participants in terms of working out whether the effect was truly contingent on the cause, while their other experiments presented non-contingent “background” effects (i.e. deployed a non-zero base-rate) throughout the experimental conditions to achieve the same. This was done with the intention of prompting participants to make full use of the temporal information provided in making their causal decision. Since Experiment 1 here did not place similar demands on participants, it is possible that decisions were merely determined by predictability rather than a combination of predictability and contiguity.

This may seem strange given that in Greville and Buehner’s (2010) experiments, the effect of predictability was always subordinate to that of contiguity. It is however possible that the motivational significance of a contiguous outcome may be reduced in an observational learning task. Many normative theories analyze decision-making in terms of utility (Manski, 2000; Mongin, 1997), which is often characterized by a cost-benefit relation. The cost of making a response or an intervention is typically considered in terms of the effort expended by the animal in comparison to the animal’s energy budget (Caraco & Lima, 1987). Meanwhile, the benefit or subjective value conferred by a reward is strongly influenced by the delay until the receipt of that reward, as a vast body of literature on temporal discounting has made clear (e.g. Green & Myerson, 2004). In instrumental performance, contiguity is thus central in determining the utility of a particular response-outcome relation. In contrast, merely observing a cue incurs a negligible energy cost in comparison to performing an instrumental response. As such, contiguity may well have a diminished role in learning from observation. For example, although Buehner and May (2004) showed that expectation of a delay could mitigate its detrimental impact, according to the strong version of their knowledge mediation hypothesis, an expectation of a delayed mechanism should also result in a weaker perception of causality when events are contiguous, since the data is then inconsistent with mechanism beliefs. However, this finding was not obtained: when action and outcome were maximally contiguous, ratings were high regardless of whether contiguity was made plausible or implausible by the cover story. Thus the incompatibility of the expected mechanism was insufficient to negate the facilitatory effect of contiguity. Yet, in a Pavlovian analogue of Buehner and May’s (2002) grenade-launching task, Allan, Tangen, Wood and Shah (2003) found that ratings were consistently higher when delay and prior knowledge were congruent with each other; specifically, they found that non-contiguous cause-effect pairings received higher causal ratings than contiguous conditions if the participant expected a delayed relation. Consequently, it is possible that the importance of temporal contiguity on causal learning is higher in instrumental than in observational learning.

These concerns over the lack of an effect of delay extent should not however detract from the principal finding from the current study, that elemental causal induction through observation is facilitated by temporal predictability. Participants observing sequences of cues and outcomes obtained from performance of previous participants gave the strongest endorsement of causal effectiveness to those conditions with constant temporal intervals, in the same fashion as those participants who originally generated the data through instrumental responding. Even so, caution should be exercised before drawing firm conclusions from the results of this single study, given the lack of an effect of delay extent; replication and extension of this study is desirable. Accordingly, Experiment 2 examined the effect of temporal predictability on observational causal learning from stochastically (as opposed to human) generated data streams, while Experiments 3 and 4 compared instrumental and observational learning directly within the same experimental setting.

Experiment 2

Humans (and other animals) may be seen as intentional agents who perform naïve experiments and engage in hypothesis testing in order to uncover causal mechanisms. As such, they can intervene on the world in a structured manner in an attempt to elucidate meaningful patterns of events. We can also learn vicariously; that is, by observing the behaviour of others. However, many causal mechanisms are inaccessible to or independent of the behaviour of agents. One of the key benefits afforded by observational learning is that it allows onlookers to learn about causal systems on which they cannot directly intervene. At the same time, an important challenge for observational learning is that lack of control over stimulus delivery means there is no guarantee that events will be segregated into meaningful patterns. Causal inference in naturalistic systems, such as learning that the presence of clouds may cause rain or that forest fires may arise from an extended period of hot and dry weather, tends to be made from more haphazard distributions of events quite unlike the structured patterns resulting from the behaviour of organisms. A distinction can thus be made between patterns of events that might be emblematic of learning from one's own behaviour, learning from the behaviour of another, or learning by simply observing a stochastic pattern of events unfold.

The instrumental experiments of Greville and Buehner (2010) constitute learning by "doing"; the previous study meanwhile falls into the category of "watching it done" (Sobel, 2003). Though the participant observing the events sequences did not directly observe the previous participant performing the action, the event sequences were obtained from human performance. Even though participants in Experiment 1 did not know that the information they were processing had been generated by a previous participant, the sequences they observed most likely would have contained richer information than stochastic patterns of cue occurrence, as they were produced by an intentional agent engaging in hypothesis testing. Such information may include, for example, rapid successive action bursts, rhythms, and abstinence from intervening for extended periods. If learning through observation can truly be facilitated by temporal predictability, it needs to be demonstrated that predictability can facilitate induction from event sequences that more closely resemble those in naturalistic settings, where such characteristic patterns that might serve as useful diagnostic tools are absent. The goal of Experiment 2 therefore was to reduce the incidence of these structured patterns of cue presentation and see if the facilitatory effect of predictability obtained in the current experiment can be replicated with a more challenging causal induction task.

Accordingly, we utilized a similar observational variant of the elemental causal induction task closely based on the previous paradigm. The essential modification was that this time the distributions of cues and outcomes were not extracted from performance of previous human participants. Instead, the causal candidate occurred according to a stochastic rate process. The likelihood of obtaining patterns of cues resembling exploratory behaviour, such as successive action bursts or a long period of abstinence from intervening, is therefore reduced, and should thus appear more "natural" (or unintentional) to observers. In addition to substituting stimulus patterns generated from previous participants with stochastically created patters, Experiment 2 also included non-contingent background effects. This manipulation makes the task more challenging and provides a more strenuous test of the reliability of the predictability effect, as objective perception of predictability may be impaired by a non-contingent effect occurring between the cue and its programmed outcome.

Method

Participants

Thirty-three students from Cardiff University (12 male) with a median age of 21.5 years completed the experiment either freely with no compensation or to receive partial course credit. One participant self-reported as completely failing to understand the task, hence their data was discarded.

Design

The same 2×3 within-subjects design as for the previous experiment was again applied here. The factors delay extent (3s/6s) combined with delay variability (0s/3s/6s) provided six conditions, each lasting for two minutes, with participants providing a causal rating from 0-100 as the dependent measure.

Apparatus, Materials & Procedure

The experiment was carried out in the same location using the same equipment as for Experiment 1. The outward appearance and requirements of the task was identical to Experiment 1.

The first modification from the previous experiment was that the occurrence of cues or candidate causes was no longer obtained from pre-recorded data. Instead, each 2-minute trial was divided into a series of small segments during which there was a fixed probability of a cue being presented. Specifically, after every 500ms, there was a 1/6 chance of cue presentation. This created, on average, a rate of one cue every three seconds, which is in line with the approximate

20 responses per minute observed in the Greville & Buehner's (2010) instrumental studies and Experiment 1 of this paper. Following cue presentation, the outcome was delivered according to the appropriate probability schedule (once again set at 0.75) after the relevant temporal interval. The temporal intervals were likewise determined by the nominal delay and range of variation about this central point for a given condition. The delays and ranges used were identical to the previous experiment.

The second modification was the application of a base rate of background effects at a pseudo-random rate of one every ten seconds on average. In other words, the first background effect occurred at a randomly determined point between 0-10s into the condition, the second between 10-20s, and so on. These background effects occurred independently of the causal cues and their outcomes.

Results and Discussion

Figure 2 shows the mean of the causal ratings provided by participants for the six different conditions. As with the previous experiment, the condition with fixed short delays attracted noticeably higher ratings than all other conditions. The familiar effect of delay extent also is reinstated, with short-delay conditions receiving uniformly higher ratings than long-delay conditions. Ratings also appear to generally decline with increasing temporal interval range, though this is more pronounced with short than long delays.

A 2×3 repeated measures ANOVA found a significant main effect of delay extent, $F(1,31) = 12.73$, $MSE = 406.0$, $p = .001$, $\eta_p^2 = .291$ and delay variability, $F(1.785, 55.32) = 5.352$, $MSE = 351.9$, $p = .010$, $\eta_p^2 = .147$. The delay extent × delay variability interaction was not significant, $F(1.949, 60.42) = 0.169$, $MSE = 380.2$, $p = .840$, $\eta_p^2 = .005$. As in Experiment 1, only the linear component of the main effect of delay variability was significant, $F(1,31) = 7.805$, $MSE = 422.9$, $p = .009$, $\eta_p^2 = .201$.

Table 2 reports cause and effect rates, experienced $P(e|c)$, experienced delays, and causal ratings provided by participants, for each condition. In this experiment (and Experiment 4), background effects were ignored in terms of $P(e|c)$, which was instead calculated as the proportion of causes producing an effect; in other words, the total number of effects that were generated by the candidate cause divided by the total number of causes ignoring background effects. Analyses confirmed that $P(e|c)$ did not vary with either delay extent, or delay variability, nor was there a significant delay extent × delay variability interaction (all Fs < 1.78). There was no significant variation of the mean cause-effect interval experienced within a given condition either as a function of delay variability and or a delay extent × delay variability interaction (both Fs <1.08).In other words, there were no systematic unplanned variations in contingency or contiguity.

The most apparent difference between these results and those of Experiment 1 is the return of the familiar detrimental effect of delays on causal ratings. Indeed the effect is strong and robust, with shorter delays preferred to longer delays at each level of predictability. Causal ratings overall were lower than in the previous study, which is to be expected since the additional background effects inflate the value of $P(e|\neg c)$ and thus lower objective contingency as measured by ΔP (Allan, 1980). The most notable result in the wider context however is that a significant effect of temporal predictability has once again been obtained. A comparison of effect sizes reveals that although the influence of predictability was weaker here than in the previous experiment, and was subordinate to the influence of delay, the anticipated facilitatory effect was still evident, replicating the findings of Experiment 1 and of Greville and Buehner (2010).

Experiment 2 thus provides additional confirmation that predictability can facilitate causal induction in observational as well as instrumental learning. Furthermore, the predictability effect is maintained when observing patterns of events whose occurrence is governed by a probabilistic rate schedule as well as when observing those derived from intentional, exploratory behaviour. This finding thus completes a triplet of obtaining facilitatory effects of predictability in elemental causal induction tasks, whether learning from one's own actions (Greville & Buehner, 2010), learning by observing the information generated by someone else's actions (Experiment 1), or learning from identifying patterns in a stochastic process (Experiment 2).

Experiment 3

Experiments 1 and 2 reliably showed that temporal predictability facilitates causal learning from observational data. However, Experiment 1 failed to show the typical finding that (overall) delay impairs causal learning, a result that resurfaced in Experiment 2. Even though demonstrating the overall effect of delay was not the purpose of our experiments, failure to find it in Experiment 1 may instil doubt in the generalisability and robustness of our main finding – the facilitatory effect of temporal predictability. Experiments 3 and 4 address these concerns.

In Experiment 3 we attempted to replicate the conditions of Experiment 1 to determine whether the absence of the delay effect was a consistent phenomenon or just an anomaly of Experiment 1. As in Experiment 1, the distribution of candidate causes mirrored the causal actions of participants from an earlier procedure. However, rather than yoking the pattern of causes and effects from a dataset in a previous experiment, we opted to create a mixture of both instrumental and observational conditions in a single experiment, so that the presented sequence of events in observational conditions was obtained from the participant's own actions in a preceding instrumental condition. In the first experimental block, participants engaged in instrumental causal learning tasks, as with Greville and Buehner (2010). In the subsequent observational block, the exact same sequence of causes and effects as generated during instrumental conditions were played back to participants in corresponding observational conditions. This afforded us the opportunity to directly contrast instrumental and observational conditions within the same experiment, and, as each individual participant experienced the exact same sequences for matched instrumental and observational conditions, allowed us to see if there were any fundamental differences between learning from instrumental action or observation in elemental causal induction tasks.

Method

Participants

Thirty students from Cardiff University (7 male) with a median age of 19 years completed the experiment to receive course credit or £5 payment. Due to a computer error, three participants were unable to complete the experiment, and a further two participants failed to follow instructions and did not make any responses during some of the instrumental conditions, and so were excluded. Twenty-five participants in total thus contributed data to the analysis.

Design

Three independent variables were manipulated within subjects: mode of learning (instrumental vs observation), delay extent (3s vs 6s on average), and delay variability (0s, 3s and 6s maximal variation of a given interval), resulting in 12 different conditions, split into 2 blocks of 6, an instrumental block and an observational block. The instrumental conditions were presented first, in random order, and the observational conditions were presented subsequently, in a new random order. No indication was given during the observational conditions that the pattern of causes and effect was in any way linked to the previous instrumental conditions.

Apparatus, Materials & Procedure

The experiment was conducted in the same location using the same equipment as for the previous experiments, and was programmed using Python. The basic contingency judgement task in instrumental conditions was near-identical to the previous experiments, with the same layout and appearance, and participants were required to provide a causal rating from 0-100 at the end of each condition. The key differences were as follows: Firstly, the instructions were suitably modified to inform participants that the experiment would be divided into two blocks. For the instrumental conditions, participants were instructed to actively investigate the relationship between pressing a button and a shape lighting up, by pressing the button at certain periods, and then refraining from pressing it for certain periods, and monitoring the lighting up of the shape. In the subsequent observational block, participants were instructed to simply observe the sequence of candidate causes and effects. Secondly, in order to make the instrumental and observational blocks distinct, a square was used in place of a triangle during one block, and a different button style was used in each block. The shape and button style used during each block was counterbalanced across participants.

Results & Discussion

Figure 3 shows participants' mean causal ratings for all conditions in Experiment 3 and shows that, as in previous experiments, conditions where the cause-effect delay was predictable received the highest ratings. A 2×2×3 within subjects ANOVA confirmed a significant main effect of delay variability, $F(1.986, 47.66) = 13.45$, $MSE = 459.6$, $p < .001$, $\eta_p^2 = .359$. Neither the effect of learning mode (instrumental vs observation), $F(1,24) = 1.503$, $MSE = 441.3$, $p = .232$, $\eta_p^2 = .059$, nor the effect of delay extent, $F(1,24) = 1.199$, $MSE = 458.1$, $p = .284$, $\eta_p^2 = .048$, and none of the possible interactions (all Fs<2.3) were significant. For delay variability, only the linear component of the main effect was significant, $F(1,24) = 27.75$, $MSE = 440.7$, $p < .001$, $\eta_p^2 = .536$. The possibility of order effects (as conditions were blocked such that the instrumental conditions always preceded the observational conditions) was not examined as there was no effect of learning mode.

Table 3 details rates of causes and effect, actual $P(e|c)$, experienced delays, and causal ratings as for the previous experiments. Note that, with the exception of causal ratings, separate values are not detailed for instrumental and observational conditions as the exact same sequences were replayed during observational sequences as were generated in instrumental conditions. $P(e|c)$ did not vary across conditions (all Fs<1.3). There was no significant effect of delay variability on experienced delays and no significant delay extent × variability interaction (both Fs<2.9). The main effect of delay variability on ratings was therefore not confounded with unplanned differences in contiguity or contingency between conditions.

In sum, Experiment 3 replicated the findings of Experiment 1, both in terms of finding a significant facilitatory influence of temporal predictability, and finding no significant influence of delay. Since the detrimental effect of delays on causal judgment is so well established, we are thus left with the task of explaining the lack of delay effect in these particular experimental preparations. Earlier, we postulated that a lack of background effects might be responsible, and indeed when background effects were introduced in Experiment 2, the familiar main effect of delay resurfaced. Experiment 4 therefore attempted to verify this supposition.

Experiment 4

Experiment 4 used a similar procedure to Experiment 3, employing two blocks of conditions, instrumental and observational. The key differences between Experiments 3 and 4 were as follows: Firstly, Experiment 4 introduced non-contingent background effects, occurring at a rate of 1 every 10s on average in the same manner as for Experiment 2. Secondly, the intermediate (3s) level of predictability was removed to streamline the experiment and, given that the effect of predictability has been established as linear, to provide a straightforward and direct comparison between predictable and unpredictable conditions.

Method

Participants

Twenty-nine undergraduate students from Cardiff University (10 male) with a median age of 19 years completed the experiment to receive course credit.

Design

Three independent variables were manipulated, mode of learning (instrumental vs observation), delay extent (3s vs 6s on average), and delay variability (0s vs 6s maximal variation of a given interval), resulting in eight different conditions, split into two blocks of four, an instrumental block presented first and a subsequent observational block. Conditions within a given block were presented in random order.

Apparatus, Materials & Procedure

The experiment was conducted in the same location using the same equipment for all previous experiments. The basic experimental structure involving two separate blocks was identical to Experiment 3. Background effects occurred according to a preprogrammed schedule of 1 every 10s on average as for Experiment 2.

Results & Discussion

Figure 4 shows mean causal ratings in Experiment 4. It is apparent that conditions involving short delays (solid lines) were judged as more causally effective than conditions involving long delays (dashed lines). Meanwhile there is a general trend that ratings declined as predictability decreased, but this decline was much more marked in some circumstances compared to others.

A 2×2×2 within subjects ANOVA confirmed a significant effect of delay extent, $F(1,28) = 5.558$, $MSE = 450.3$, $p = .026$, $\eta_p^2 = .166$, and a marginally significant effect of delay variability, $F(1,28) = 4.025$, $MSE = 706.1$, $p = .055$, $\eta_p^2 = .126$. There was no significant effect of presentation mode ($p > 0.1$). There was, however, a significant three-way interaction, $F(1,28) = 4.705$, $MSE = 329.8$, $p = .039$, $\eta_p^2 = .144$. As shown in Figure 4, the interaction between extent and variability followed different patterns in instrumental and observational conditions. Specifically, in instrumental conditions, the effect of predictability was most marked for long delay conditions, whereas in observational conditions the effect of predictability was most marked in short delay conditions.

Cause and effect rates, mean experienced $P(e|c)$ and delays, and mean causal ratings are detailed in Table 4. Experienced values were once again consistent with programmed values across conditions, with no significant effects of either delay extent, delay variability or the extent by delay interaction on experienced $P(e|c)$, and no effect of delay variability or the extent by delay interaction on experienced delay (all ps > 0.1).

Experiment 4 has thus demonstrated the expected familiar effect of delay extent, as in Experiment 2. This lends weight to the idea that it is the absence of background effects in Experiments 1 and 3 that is responsible for the absence of the effect of delay. It may be the case that when temporal regularity of cause-effect intervals is readily identifiable, as would be the case with a base outcome rate of zero, then contiguity is not necessarily a strong cue to causality and may be secondary to predictability.

The effect of predictability in Experiment 4 was however only marginal. This, coupled with a smaller effect size for predictability in Experiment 2 suggests that, although predictability does exert a facilitatory influence, it may be a less powerful cue to causality when there is a background rate of non-contingent outcomes. Presumably, it becomes more difficult to identify regularity in the timing of causes and effects against a backdrop of noise (i.e. effects that are produced by other causes) as a non-contingent outcome may be misattributed to the candidate cause, and thus introduce variability to the perceived cause-effect interval. However, further between-experiment analyses (the results of which we have not reported for the sake of brevity) found no interaction between the variables experiment and predictability either when comparing Experiments 1 and 2, or Experiments 3 and 4, indicating that the influence of predictability is relatively stable.

The three-way interaction was a rather more unexpected finding. It appears that, in instrumental conditions, predictability only mattered if delays were long; that is, if delays were long and unpredictable, then ratings were low, otherwise ratings were high. Conversely for observational conditions, predictability appeared to matter only when delays were short; that is, if delays were short and predictable then ratings were high, otherwise ratings were low. Hence, although there was no significant main effect of presentation mode alone, the interaction suggests that learning from instrumental activity generally produced stronger impressions of causality. These results are somewhat different from the most comparable experiment of Greville and Buehner (2010), concerning conditions of 2-minutes duration in Experiment 3. In those conditions, which were instrumental, the general pattern was in fact rather more similar to the observational conditions in the current experiment; ratings tended to be low, unless both contiguity and predictability were high. In other words, at longer delays, predictability did not help all that much, but at shorter delays, predictability enhanced impressions of causality. The precise underlying reasons for this difference are not immediately obvious from either a theoretical or methodological standpoint; it is possible for instance that practice effects may have played a role. Future research may seek to further explore the precise nature of the interaction between these factors. In summary though, it appears that, depending on the way in which the learner engages with the causal relation in question, a lack of either contiguity or predictability may impair causal impressions, and a lack of both will impair causal impression regardless.

For present purposes though, the central issue was to determine whether temporal predictability facilitates elemental causal induction when learning from observation. The results of the four experiments contained herein strongly indicate that this is indeed the case (although there may be other factors that mediate this facilitatory influence). We have also addressed a secondary finding pertaining to the absence of an influence of contiguity, and our results suggest that the influence of contiguity may be secondary to that of predictability when $P(e|\neg c)$ is zero.

General Discussion

We set out to resolve the question of whether temporal predictability is as effective in facilitating elemental causal induction from observation as from intervention. Our focus has been on a particular type of causal learning problem, that of assessing the relationship between a single binary causal candidate and a single binary effect, which Griffiths and Tenenbaum (2005) termed elemental causal induction. Creating an observational analogue of Greville and Buehner's (2010) instrumental studies, we obtained facilitatory effects of temporal predictability consistent with those found with instrumental learning, both when the patterns of cue occurrence were based on action sequences of previous participants (Experiment 1), on participants' own previous actions (Experiments 3 and 4) and also when based on a random rate-based process (Experiment 2). The results confirm that predictability facilitates causal learning, at least with respect to the special case of elemental causal induction.

This work has also yielded an unexpected finding; in the absence of non-contingent background effects, longer cause-effect delays did not attenuate causal judgments, provided that these delays were consistent. The detrimental effect of delays on causal learning is of course strongly established (e.g. Shanks et al., 1989; Wasserman et al., 1983). A potential explanation for the present finding then is that it is not delay per se that is harmful for causal learning but rather the uncertainty created by delay, a thesis that has recently been garnering support. For instance, Lagnado and Speekenbrink (2010; Experiment 2) have demonstrated that it was specifically the probability of the occurrence of an intervening event rather than the extent of delay per se that was most detrimental to causal ratings. Conversely though, Greville and Buehner (2010; Experiment 1) investigated the impact of varying the level of background effects in an instrumental learning task and found that the extent of background effects did not mediate the effect of delay. However, there were important differences in the procedures used by Lagnado and Speekenbrink, and Greville and Buehner, which might also have led to differential results concerning the impact of intervening events. Future research could systematically investigate the role of intervening events in determining the effect of delay, and whether this role changes depending on the nature of the causal learning task or whether learning takes place through observation or intervention.

The overall implication from all four experiments – that predictability fosters causal impressions – notwithstanding, the effect size of predictability was weaker in Experiments 2 and 4. Why might this be the case? As described earlier, action patterns produced by an intentional agent tend to contain richer information than stochastic patterns as the agent produces conspicuous action patterns to enable hypothesis testing, whereas such patterns are unlikely to arise from purely stochastic processes. This may mean that statistical and temporal information is easier to detect in the former rather than the latter case. For example, if a participant performs a rapid series of successive causal actions, and both contingency and temporal predictability are strong, then this will be followed by a corresponding series of successive outcomes. Thus in situations where the learner has the ability to produce (or observe) conspicuous action patterns, temporal predictability may be a more effective cue to causality than where such patterns cannot be produced. Future research could explore this issue further by directly contrasting stochastic with agent-generated patterns within the same experiment.

While this would account for a weaker effect of predictability in Experiment 2, it does not account for the same in Experiment 4, where cue and outcome patterns were self-generated. However, if we are willing to endorse the idea that recognizing complementary patterns in the temporal distribution of causes and effects may serve as a strong cue to causality, it is easy to see how this might be obscured by the presence of non-contingent background effects; if a background effect occurs during a sequence of contingent effects, it breaks up the sequence. The diagnostic value of predictability is thus reduced as even if predictability is strong, a conspicuous series of causal actions may not necessarily be followed by a similarly conspicuous pattern of outcomes due to the occurrence of background effects. Thus, since background effects were present in both Experiments 2 and 4, it follows that the effect of temporal predictability should be weaker in these experiments.

The utility of predictability as a cue to causality may similarly be obscured or made redundant in other forms of causal decision making, for instance in tasks such as those of Young and Nguyen (2009). One of the difficulties involving causal learning with delays is that competing agents can come between the cause and the outcome. This is particularly true in a task such as Young and Nguyen's, involving choice between multiple candidates, since the non-causal competing candidates can, coincidentally, be more contiguous with a particular effect than the true cause, and thus precipitate incorrect selection of a non-causal candidate as the target. The longer the delay, the more likely this is to occur, and this is particularly true for a constant, long-delay causal candidate: Whilst for a variable-long-delay, there is the possibility on any given trial that there may be a contiguous pairing of the cause and effect, this cannot occur with fixed-long-delays. Consequently, the utility of temporal predictability as a cue to causality may be diminished when the task requires identifying a cause amongst a set of candidates, particularly when reasoners can control how much information they sample: An early observation of a contiguous pairing may lead to a termination of the search.

Theoretical Implications

The results of this paper are consistent with a "temporal predictability hypothesis" that was first outlined by Young, Rogers and Beckmann (2005), and which provided the impetus behind Greville and Buehner's (2010) studies. According to this framework, temporal predictability enhances impressions of causality by allowing learners to predict when, not just whether, an effect will follow from its putative cause. Young et al. argued that the key difference between the predictability hypothesis and earlier knowledge-based accounts (Tenenbaum & Griffiths, 2003) is that the latter rely on top-down models to explain the impact of temporal information, but are silent as to how this top-down knowledge is acquired in the first place. A more recent account that may however subsume the predictability hypothesis is that of theory-based causal induction (Griffiths & Tenenbaum, 2009), which proposes that people form intuitive top-down theories about causal mechanisms which subsequently affect the processing and interpretation of bottom-up causal data. This account would anticipate temporal predictability to facilitate causal induction if it is assumed that people have the a priori notion that causal mechanisms should be temporally predictable and unfold in a consistent manner over time. Whether of course people do have such a priori notions with regard to temporal predictability, and where such notions come from in the first place, remains an open question. What may be more likely is the converse; that people have the a priori notion that temporal regularity is unlikely to happen due to chance. It may then be that the experience of predictability facilitates causal learning by means of "coincidence detection" (Griffiths & Tenenbaum, 2007). In an exploratory task such as elemental causal induction, where a participant is actively investigating a putative causal relation, temporal predictability might simply serve to make a causal relation easier to detect through observing conspicuous patterns of events. If the effect keeps happening at the same point in time after the occurrence of the candidate cause, this reflects a situation that is statistically unlikely to happen if the effect is occurring due to stochastic background processes, and thus constitutes evidence for the existence of a causal relation between the candidate cause and effect.

Associative accounts of causality judgment are also able to accommodate predictability effects. The role of predictability from this perspective depends on whether experiencing a set of consistently delayed cause-effect pairings accrues more or less associative strength between cause and effect than an inconsistent set of short and long delays centred about the same mean. The degree of associative strength accrued would be determined by the function with which acquired associative strength declines as delay between cause and effect increases. A negatively accelerated function has been identified as the most likely function relating delay and associative strength (see, e.g., Chung, 1965). However such a function predicts that temporally variable causal relations result in stronger associations than temporally predictable causal relations (see Greville and Buehner, 2010). An associative perspective therefore struggles to explain a facilitatory effect of temporal predictability in causal learning unless there is an a priori reason to suspect a different function governs the relationship between associative strength and delay..

Research on the role of temporal predictability in causal learning is still in its infancy. Prior studies (Greville & Buehner, 2010; Lagnado & Speekenbrink, 2010) have indicated that temporal predictability tends to facilitate causal learning in elemental induction tasks. However other research (Young & Nguyen, 2009) demonstrate that this predictability advantage does not necessarily extend to all causal learning scenarios. The work in this paper has concluded that temporal predictability enhances judgments of causality in elemental causal induction. Further research should attempt more precisely identify situations in which the predictability hypothesis holds and in which it does not. Furthermore, theoretical accounts for the predictability effect are currently still at the abstract level; the development of computational explanations for this effect would be desirable, and augmentation of existing learning algorithms in order to encapsulate temporal predictability may enhance the ability of such algorithms to model causal learning in real time.

References

Allan, L. G. (1980). A note on measurement of contingency between two binary variables in judgment tasks. *Bulletin of the Psychonomic Society, 15*(3), 147-149.

Allan, L. G., Tangen, J. M., Wood, R., & Shah, T. (2003). Temporal contiguity and contingency judgments: A Pavlovian Analogue. *Integrative Physiological and Behavior Science, 31*(2), 205-211.

Buehner, M. J., & Cheng, P. W. (2005). Causal learning. In K. J. Holyoak & R. Morrison (Eds.), *Handbook of Thinking & Reasoning* (pp. 143-168). Cambridge, England: Cambridge University Press.

Buehner, M. J., & May, J. (2002). Knowledge mediates the timeframe of covariation assessment in human causal induction. *Thinking and Reasoning, 8*(4), 269-295.

Buehner, M. J., & May, J. (2003). Rethinking temporal contiguity and the judgement of causality: effects of prior knowledge, experience, and reinforcement procedure. *Quarterly Journal of Experimental Psychology, 56A*(5), 865-890.

Buehner, M. J., & May, J. (2004). Abolishing the effect of reinforcement delay on human causal learning. *Quarterly Journal of Experimental Psychology, 57B*(2), 179-191.

Buehner, M. J., & McGregor, S. (2006). Temporal delays can facilitate causal attribution: Towards a general timeframe bias in causal induction. *Thinking and Reasoning, 12*(4), 353-378.

Caraco, T., & Lima, S., L. . (1987). Survival, energy budgets, and foraging risk. In M. L. Commons, A. Kacelnik & S. J. Shettleworth (Eds.), *Quantitative Analyses of Behavior* (Vol. IV: Foraging, pp. 1-21). Hillsdale, NJ: Erlbaum.

Einhorn, H. J., & Hogarth, R. M. (1986). Judging probable cause. *Psychological Bulletin, 99*(1), 3-19.

Ferster, C. B., & Skinner, B. F. (1957). *Schedules of reinforcement*. New York: Appleton-Century-Crofts.

Green, L., & Myerson, J. (2004). A discounting framework for choice with delayed and probabilistic rewards. *Psychological Bulletin, 130*(5), 769-792.

Greville, W. J., & Buehner, M. J. (2010). Temporal predictability facilitates causal learning. *Journal of Experimental Psychology: General, 139*(4), 756–771.

Greville, W. J., Cassar, A. A., Johansen, M. K., & Buehner, M. J. (2013). Structural awareness mitigates the effect of delay in human causal learning. *Memory & Cognition, 41*(6), 904-916.

Grice, G. R. (1948). The relation of secondary reinforcement to delayed reward in visual discrimination learning. *Journal of Experimental Psychology, 38*(1), 1-16.

Griffiths, T. L., & Tenenbaum, J. B. (2005). Structure and strength in causal induction. *Cognitive Psychology, 51*(4), 334-384.

Griffiths, T. L., & Tenenbaum, J. B. (2007). From mere coincidences to meaningful discoveries. *Cognition, 103*(2), 180-226.

Griffiths, T. L., & Tenenbaum, J. B. (2009). Theory-based causal induction. *Psychological Review, 116*(4), 661–716.

Lagnado, D. A., & Sloman, S. A. (2004). The advantage of timely intervention. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 30*(4), 856-876.

Lagnado, D. A., & Sloman, S. A. (2006). Time as a guide to cause. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 32*(3), 451-460.

Lagnado, D. A., & Speekenbrink, M. (2010). The influence of delays in real-time causal learning. *The Open Psychology Journal, 3*(2), 184-195.

Manski, C. F. (2000). Economic Analysis of Social Interactions. *The Journal of Economic Perspectives, 14*(3), 115-136.

Mongin, P. (1997). Expected utility theory. In J. Davis, W. Hands & U. Maki (Eds.), *Handbook of Economic Methodology* (pp. 342-350). London: Edward Elgar.

Msetfi, R. M., Murphy, R. A., Simpson, J., & Kornbrot, D. E. (2005). Depressive realism and outcome density bias in contingency judgments: The effect of the context and intertrial interval. *Journal of Experimental Psychology: General, 134*(1), 10-22.

Shanks, D. R., Pearson, S. M., & Dickinson, A. (1989). Temporal contiguity and the judgment of causality by human subjects. *Quarterly Journal of Experimental Psychology, 41B*(2), 139-159.

Sloman, S. A., & Lagnado, D. A. (2005). Do we "do"? *Cognitive Science, 29*, 5-39.

Sobel, D. M., & Kushnir, T. (2006). The importance of decision making in causal learning from interventions. *Memory & Cognition, 34*(2), 411-419.

Tenenbaum, J. B., & Griffiths, T. L. (2003). Theory-based causal inference. In S. Becker, S. Thrun & K. Obermayer (Eds.), *Advances in neural information processing systems* (Vol. 15, pp. 35-42). Cambridge, MA: MIT Press.

Waldmann, M. R. (1996). Knowledge-based causal induction. In D. R. Shanks, K. J. Holyoak & D. L. Medin (Eds.), *Causal Learning* (Vol. 34, pp. 47-88). San Diego, CA: Academic Press.

Waldmann, M. R. (2000). Competition among causes but not effects in predictive and diagnostic learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 26*(1), 53-76.

Waldmann, M. R., & Hagmayer, Y. (2005). Seeing versus doing: Two modes of accessing causal knowledge. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 31*(2), 216-227.

Waldmann, M. R., & Holyoak, K. J. (1992). Predictive and diagnostic learning within causal models: Asymmetries in cue competition. *Journal of Experimental Psychology: General, 121*(2), 222-236.

Waldmann, M. R., & Holyoak, K. J. (1997). Determining whether causal order affects cue selection in human contingency learning: Comments on Shanks and Lopez (1996). *Memory & Cognition, 25*(1), 125-134.

Wasserman, E. A., Chatlosh, D. L., & Neunaber, D. J. (1983). Perception of causal relations in humans: Factors affecting judgments of response-outcome contingencies under free-operant procedures. *Learning and Motivation, 14*(4), 406-432.

Wasserman, E. A., Elek, S. M., Chatlosh, D. L., & Baker, A. G. (1993). Rating causal relations: Role of probability in judgments of response-outcome contingency. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 19*(1), 174-188.

Williams, B. A. (1976). The effects of unsignalled delayed reinforcement. *Journal of the Experimental Analysis of Behavior, 26*(3), 441-449.

Young, M. E., & Nguyen, N. (2009). The problem of delayed causation in a video game: Constant, varied, and filled delays. *Learning and Motivation, 40*, 298–312.

Young, M. E., Rogers, E. T., & Beckmann, J. S. (2005). Causal impressions: Predicting when, not just whether. *Memory & Cognition, 33*(2), 320-331.

Table 1

Rates and timings of causes and effects in Experiment 1. Mean values are shown with standard deviations in parentheses.

	delay					
	3s			*6s*		
	range of temporal intervals					
	0s	*3s*	*6s*	*0s*	*3s*	*6s*
total causes	34.70 (18.62)	31.18 (14.24)	36.42 (13.03)	41.88 (15.89)	30.64 (16.08)	31.33 (19.09)
total effects	25.58 (13.16)	22.27 (10.56)	26.94 (10.52)	31.94 (12.46)	22.58 (12.23)	23.33 (15.99)
actual P(e|c)	0.745 (0.0776)	0.709 (0.0622)	0.732 (0.111)	0.769 (0.0727)	0.728 (0.0871)	0.701 (0.110)
experienced delay in ms	3000 (0)	2977.61 (261.84)	2997.39 (303.98)	6000 (0)	5996.42 (153.94)	6043.55 (629.17)
causal rating	54.61 (32.38)	46.21 (26.87)	45.42 (24.01)	54.48 (28.98)	45.73 (33.22)	37.61 (25.63)

Table 2

Rates and timings of causes and effects in Experiment 2. Mean values are shown with standard deviations in parentheses.

	delay					
	3s			*6s*		
	range of temporal intervals					
	0s	*3s*	*6s*	*0s*	*3s*	*6s*
total causes	39.06 (5.52)	38.06 (5.72)	39.38 (6.31)	39.00 (6.55)	39.72 (4.94)	40.17 (5.32)
total effects	29.81 (4.82)	28.69 (5.42)	29.25 (5.38)	29.81 (5.61)	30.67 (4.62)	29.88 (4.19)
actual P(e|c)	0.763 (0.0617)	0.751 (0.0639)	0.743 (0.0663)	0.764 (0.0647)	0.772 (0.0602)	0.746 (0.0653)
experienced delay in ms	3000 (0)	2913.38 (193.46)	2957.38 (375.33)	6000 (0)	5966.03 (122.15)	5993.94 (305.51)
causal rating	52.81 (24.39)	44.78 (24.88)	41.00 (24.53)	40.16 (23.76)	35.66 (24.41)	31.66 (21.06)

Table 3

Rates and timings of causes and effects in Experiment 3. Mean values are shown with standard deviations in parentheses.

	delay					
	3s			*6s*		
	range of temporal intervals					
	0s	*3s*	*6s*	*0s*	*3s*	*6s*
total causes	52.68 (28.56)	51.04 (27.94)	43.56 (26.09)	52.16 (44.48)	49.12 (35.11)	42.80 (25.69)
total effects	38.48 (20.72)	38.52 (22.29)	32.24 (19.95)	41.12 (36.07)	36.52 (26.38)	32.04 (19.35)
actual P(e|c)	0.736 (0.0744)	0.751 (0.0871)	0.737 (0.0804)	0.776 (0.0657)	0.748 (0.0809)	0.740 (0.0781)
experienced delay in ms	3000 (0)	3020.99 (164.58)	2911.37 (418.54)	6000 (0)	5946.63 (169.11)	6058.96 (413.92)
causal rating (instrumental)	65.24 (20.89)	60.36 (26.20)	49.16 (26.93)	66.00 (24.39)	51.04 (25.24)	43.08 (25.23)
causal rating (observational)	60.24 (21.46)	53.00 (26.45)	46.08 (18.23)	58.16 (18.60)	50.80 (27.50)	48.76 (22.31)

Rates and timings of causes and effects in Experiment 4. Mean values are shown with standard deviations in parentheses.

	delay			
	3s		*6s*	
	range of temporal intervals			
	0s	*6s*	*0s*	*6s*
total causes	68.00 (62.06)	57.07 (45.73)	58.45 (62.24)	57.66 (63.80)
total effects	49.48 (46.36)	42.00 (33.85)	44.24 (46.91)	42.51 (45.82)
actual P(e|c)	0.732 (0.0946)	0.741 (0.0700)	0.753 (0.0750)	0.742 (0.114)
experienced delay in ms	3000 (0)	3101.10 (398.44)	6000 (0)	6013.28 (547.76)
mean causal rating (instrumental)	57.48 (29.96)	55.72 (22.52)	56.24 (27.63)	42.00 (27.20)
mean causal rating (observational)	58.14 (19.85)	48.03 (24.80)	48.38 (26.39)	46.48 (21.24)

Figure Captions

Figure 1: Mean Causal Ratings for all conditions in Experiment 1

Figure 2: Mean Causal Ratings for all conditions in Experiment 2

Figure 3: Mean Causal Ratings for all conditions in Experiment 3

Figure 4: Mean Causal Ratings for all conditions in Experiment 4

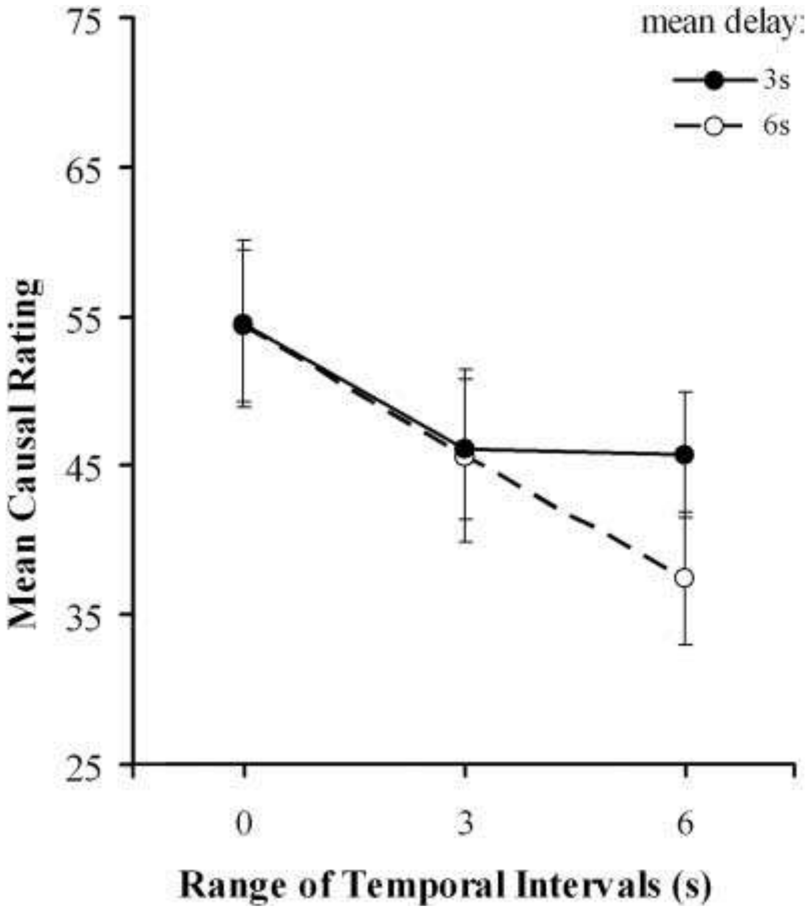
75
65
55
45
35
25
Mean Causal Rating
mean delay:
3s
6s
0
3
6
Range of Temporal Intervals (s)

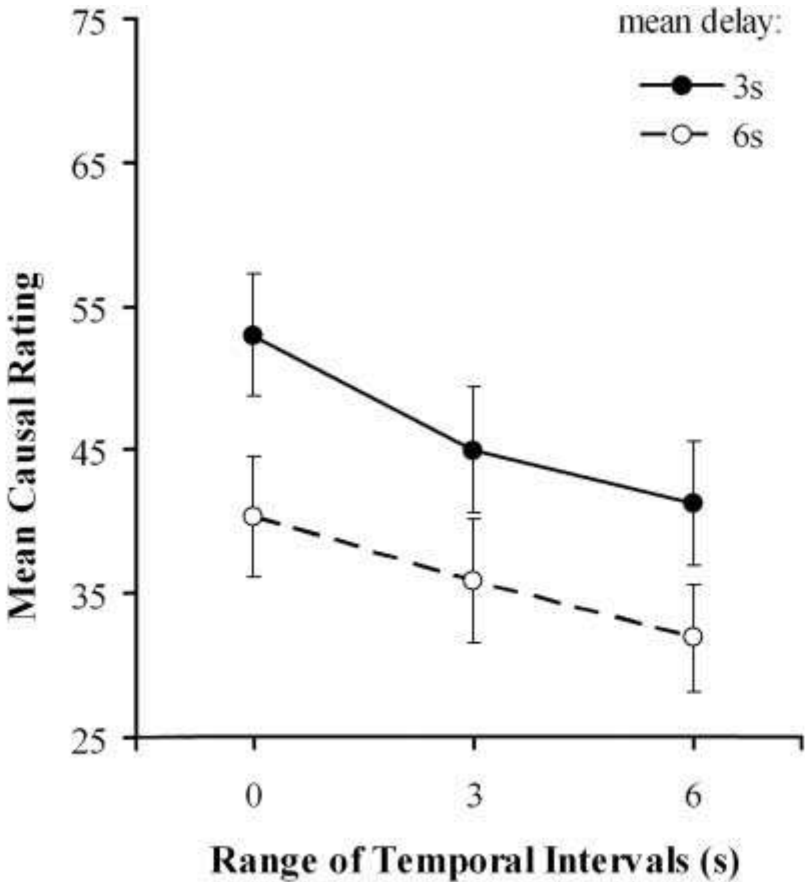

mean delay:
3s
6s
75
65
55
45
35
25
Mean Causal Rating
0
3
6
Range of Temporal Intervals (s)

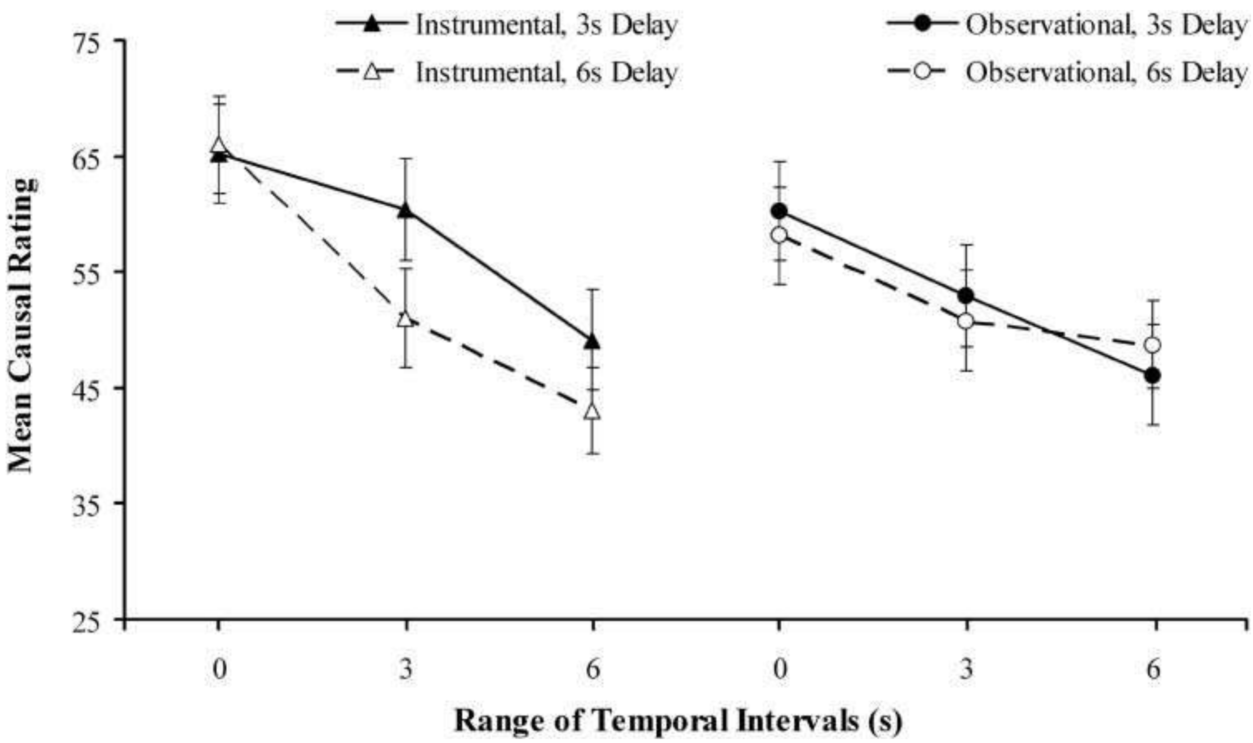
Instrumental, 3s Delay
Instrumental, 6s Delay
Observational, 3s Delay
Observational, 6s Delay
75
65
55
45
35
25
Mean Causal Rating
0
3
6
0
3
6
Range of Temporal Intervals (s)

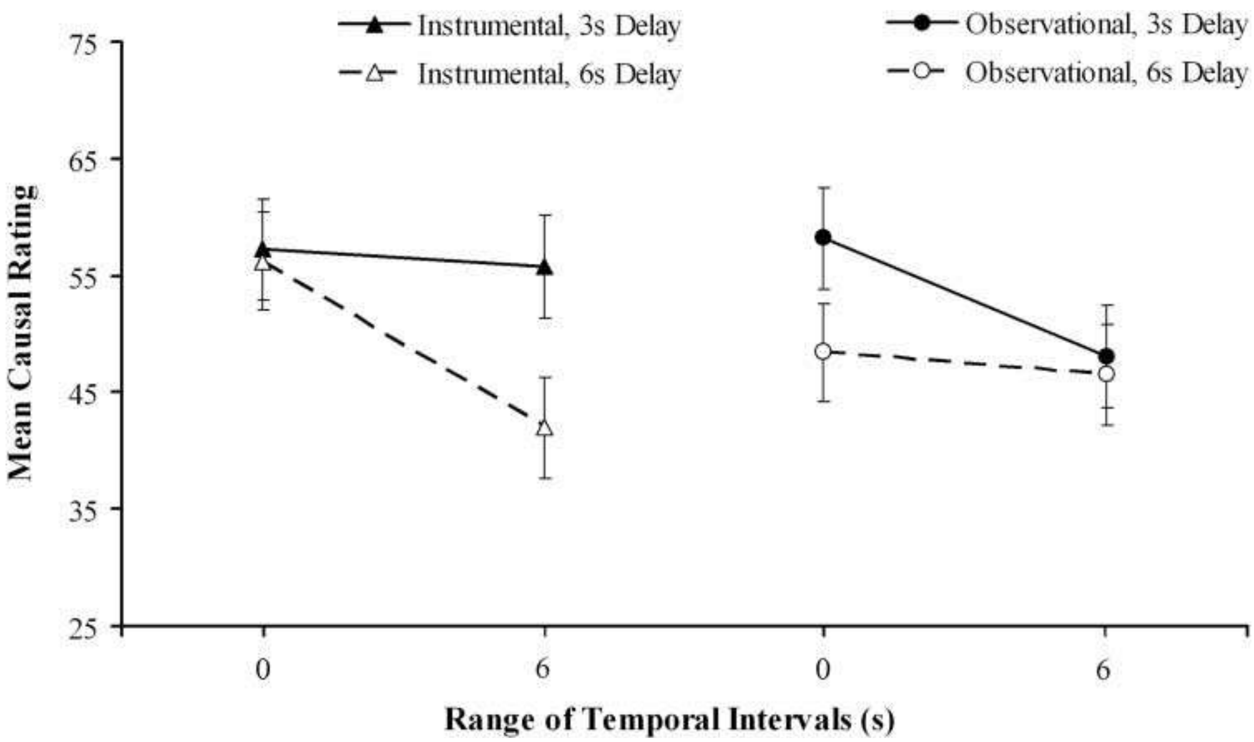
Instrumental, 3s Delay
Instrumental, 6s Delay
Observational, 3s Delay
Observational, 6s Delay
Mean Causal Rating
75
65
55
45
35
25
0
6
0
6
Range of Temporal Intervals (s)