

Running head: TEMPORAL CONTIGUITY IN CAUSAL INDUCTION

Knowledge mediates the timeframe of covariation assessment in human causal induction

Marc J Buehner & Jon May

University of Sheffield, England, UK

Address for Proofs: Marc J Buehner
Department of Psychology
University of Sheffield
Western Bank
Sheffield S10 2TP
England, UK
Phone: +44-114-222-6520
Fax: +44-114-276-6515
email: marc.buehner@sheffield.ac.uk

Abstract

How do humans discover causal relations when the effect is not immediately observable? Previous experiments have uniformly demonstrated detrimental effects of outcome delays on causal induction. These findings seem to conflict with everyday causal cognition, where humans can apparently identify long-term causal relations with relative ease. Three experiments investigated whether the influence of delay on adult human causal judgments is mediated by experimentally induced assumptions about the timeframe of the causal relation in question, as suggested by Einhorn & Hogarth (1986). Causal judgments generally decreased when a delay separated cause and effect. This decrease was less pronounced when the thematic context of the causal relation induced participants to expect a delay. Experiment 3 ruled out an alternative explanation of the effect based on variations of cue and outcome saliencies, and showed that detrimental effects of delay are reduced even more when instructions explicitly mentioned the timeframe of the causal relation in question. Knowledge thus mediates the impact of delay on human causal judgment. Implications for contemporary theories of human causal induction are discussed.

“I find in the first place, that whatever objects are consider’d as causes or effects, are
contiguous;”

David Hume (1739/1978 p.75)

Understanding the causal structure of the world around us is one of the most fundamental cognitive capacities humans and other intelligent species have evolved. Although the study of causality has attracted scholars throughout the ages, David Hume’s “riddle of induction” (1739/1978) still puzzles cognitive psychologists and philosophers alike. Given that our sensory system is not susceptible to causal information, how can we ever infer that one thing causes another? Most of the current research adopts the Humean approach that causal strength must be computed from evidence available to our senses, such as the presence and absence of candidate causes and effects (see Shanks, Holyoak, & Medin, 1996 for a review). How this information is used to infer causality is still debated (for example see Cheng, 1997; Buehner & Cheng, 1997; Buehner, Cheng, & Clifford, in preparation; Lober & Shanks, 2000), but all computational approaches agree that the covariation between candidate cause and effect (or some transformation of it) is a crucial component in the inductive process. Because research has focused primarily on the question of what exactly is inferred, a very important aspect of causality has largely been overlooked: the temporal relation between causes and effects. In order to derive causal knowledge from covariation, an organism must first be able to successfully identify (in real-time) that two events have or have not co-occurred. At present, no computational approach addresses a solution to this problem. Experimental work most closely relevant to the problem has investigated the importance of temporal contiguity in causal induction (Michotte, 1963; and Shanks, Pearson, & Dickinson, 1989 are the seminal contributions) and has uniformly concluded that temporal contiguity is of utmost importance to causal induction. The reasons given as to why participants either failed to report a causal relation, or reported a substantially degraded relation when cause and effect were separated by only a few seconds, vary between accounts.

The Influence of Delay in Associative Learning Theory

According to an associationist interpretation of causal inference, causal learning is identical to associative learning. Judged causal strength reflects no more than the associative strength between candidate cause (equivalent to a conditioned stimulus or response) and effect (unconditioned stimulus) (Shanks & Dickinson, 1987). Every time a cause and an effect co-occur, the association between them is strengthened until it reaches the maximum strength the effect can support (i.e., asymptote), and every time the cause fails to produce the effect, the cause-effect association weakens.

Shanks and Dickinson (1987, p. 231), in a paper outlining the principles of an associationist account of human causal learning explained the importance of temporal contiguity as follows: “Contemporary accounts are usually silent about the actual interevent interval over which an association can be formed, but all argue that the size of the increment in associative strength accruing from a pairing decreases as the contiguity is degraded.” Shanks et al. (1989) confirmed these predictions: in their experiments, participants had to find out how strongly pressing the SPACE bar made a triangle light up on the computer screen. The apparatus was programmed with a .75 contingency; more specifically, 3 out of 4 responses produced the outcome. Shanks et al. varied the temporal delay between (reinforced) responses and the delivery of the outcome from 0s up to 16s. Causal judgments decreased substantially if action and outcome were separated by more than 2 seconds, so much that participants could no longer distinguish between contingent conditions (where their actions produced outcomes with a probability of .75) and non-contingent conditions (where their actions had no consequences

whatsoever). These experiments thus suggest that delays seriously impair causal learning.

This is not to say that associationism would claim that causal relations can only be learnt when cause and effect follow each other immediately. In a more recent theoretical paper Shanks wrote about his earlier experiments: “It should be emphasized (...) that much longer delays can certainly be tolerated in other situations. The slope of the contiguity function is likely to be highly task-specific.” (Shanks, 1993, p.323). This statement was clearly intended to account for findings such as the Garcia effect (Garcia, Ervin, & Koelling, 1966; Garcia & Koelling, 1966), which demonstrated that animals can bridge considerable time-spans in taste-aversion learning paradigms, and led theorists to propose that gustatory and olfactory cues are “biased” to be associated with internal malaise, “even when these stimuli are separated by long time periods” (Garcia et al., 1966, p.122).

Causal Power Theories

According to the causal power view (Ahn, Kalish, Medin, & Gelman, 1995; Bullock, Gelman, & Baillargeon, 1982) the crucial component of causal inference is knowledge about some causal power or mechanism linking cause and effect. The regular co-occurrence between a cause and an effect can only be rendered to reflect a causal relation, if the reasoner knows of a causal mechanism that explains how the cause could bring about the effect. In Michotte’s (1963) launching paradigm the stimuli created a visual illusion of two moving objects. Object A moves towards and collides with a stationary object B. B moves away either immediately after A’s impact (which results in a causal impression), or a few seconds after A’s impact (which fails to elicit causal

impressions). Because participants in Michotte's and similar studies presumably applied their naïve understanding of physics to the task, they evaluated contiguous launching events as reflecting a causal relation (A caused B to move away), but refrained from doing so if contiguity was disrupted by a delay (B moved on its own).

The causal power explanation for why previous experiments revealed contiguity as an important prerequisite in causal induction differs from the associationist explanation in one very important respect. Whereas associationism holds contiguity as an essential cue to causality, causal power theory does not make any such claims. Whether or not a particular action sequence gives rise to a causal impression depends on the observer's assumption about the potential mechanism linking cause and effect. If the assumed mechanism implies immediacy (as in Michotte's paradigm), only contiguous sequences should be interpreted as causal. If the reasoner assumes a delayed mechanism, delayed sequences should readily be judged as causal, whereas immediate sequences should fail to create a causal impression. According to causal power theory, there is nothing special about temporal contiguity. What matters are the assumptions about the causal mechanism that the reasoner brings to the task.

Einhorn and Hogarth's (1986) Knowledge Mediation Hypothesis

Einhorn and Hogarth (1986) argued that the main function of contiguity is to be an "important cue for directing attention to contingencies between variables, and such contingencies may then be considered as to their causal significance." (Einhorn & Hogarth, 1986 p.10). However, Einhorn and Hogarth also accounted for our intuitions about everyday causal inference, where people can recognize causal relations involving delays as follows:

When temporal and/or spatial contiguity is low (or temporal contiguity is erratic), inferring causality becomes more difficult. That is, in the absence of contiguity, relations are hard to develop, unless one uses intermediate causal models to link the events (...). For instance, the temporal gap between intercourse and birth requires some knowledge of human biology and chemistry to maintain links between those events (Einhorn & Hogarth, 1986, p.10).

This Knowledge Mediation Hypothesis of course borrows heavily from the causal power account described above. In fact, Einhorn and Hogarth (1986) go on to illustrate that in certain scenarios high temporal contiguity may conflict with other cues to causality, most notably covariation. This hypothesis fuses the explanations offered by associationism and causal power. Although it acknowledges temporal contiguity as a useful cue that helps to identify contingencies between causes and effects, it does not bestow an especially privileged role to it. People can overcome the need for temporal contiguity, if they know of some causal link, mechanism, or chain that takes time to unfold. In other words, the influence of time is mediated by prior knowledge. Whether or not a particular (contiguous or delayed) covariation will be judged causal depends on the assumptions about the causal mechanism (in particular the assumptions about the timeframe of this mechanism) that a reasoner brings to the task.

Distinguishing Knowledge Mediation from Associationism

Generating Predictions

For the purposes of generating predictions, the knowledge-mediation and causal power accounts do not differ. They are top-down approaches to causal reasoning: pre-existing mental concepts or structures determine how subsequent sensory experience

will be parsed and processed. It follows that identical sensory experiences could potentially be parsed differently, if the reasoner applies different assumptions to the task. Imagine that a person presses a button, and fifteen seconds later a green light illuminates. If the person assumes the connection between the button and the light to be an ordinary electric circuit, this delayed course of events would not qualify as a causal sequence. If, however, the person assumes the button triggers a timer-relay, as on a pedestrian crossing, a delay of fifteen seconds between pressing the button and the signal changing is within the expected range, and a causal connection is plausible. Because the assumptions about the causal mechanism were different in the two situations, the same sensory experience will give rise to a causal interpretation in the latter, but not the former, scenario. Logically, if a reasoner expects a delayed relation, immediate contingencies should not be attributed to the causal relation in question. If the person pressed the button and an instant later the traffic light turns green, one would most likely think that either the light would have changed anyway, or that someone else had pressed the button earlier.

The bottom-up nature of associationism stands in stark contrast to the top-down ideas of Knowledge-Mediation. In associationism, sensory experiences give rise to mental connections (associations) according to the laws of a specific learning mechanism, and every abstract idea can be reduced to be the result of associations (in fact, abstract ideas are thought to be nothing more than associations). Prior knowledge has no place in such an account, unless it is the result of previous associations itself (i.e. pre-existing associative strengths between a particular cue and outcome). What is not, possible, however, is that abstract ideas influence how sensory experience will be parsed. The

same covariational evidence should always give rise to the same causal impression.

While associationism acknowledges that different stimuli and effects may vary as to how far they can be separated by a delay for reasoners to still learn the connection successfully (Shanks, 1993), it cannot account for different interpretations of exactly the same evidence, as outlined in the previous paragraph.

Current Evidence does not Favour One Account over the Other

Knowledge-mediation in causal inference is well studied, but research has mainly focused on how prior knowledge determines whether a given covariation will give rise to causal impressions or not (e.g. White, 1995; Ahn et al., 1995; but see also Lien & Cheng, 2000), and how previously acquired category structures (Waldmann & Hagmayer, 1999) or assumptions about the direction of learning influence covariation assessment in causal induction (predictive or diagnostic learning, see Waldmann & Holyoak, 1992; Waldmann, 2000). However, whether and how beliefs about the timeframe of causal relations influence adults' assessment of delayed or contiguous contingencies is not yet clear.

A careful analysis of the predictions derived from the knowledge mediation and associationist accounts and the experimental materials employed in previous investigations reveals that current findings cannot distinguish between these two competing explanations. The reason for this lack of discrimination is that all these experiments employed scenarios where participants either expected contiguous cause-effect pairings, or used paradigms that lend themselves to such interpretations.

Michotte's (1963) and related perceptual causality experiments presented participants with images of colliding objects. Participants brought to the task their world

knowledge about the physical properties of impact and therefore expected immediate motion after impact. If this expectation was violated, the sequence was judged as non-causal. The same principle applies to Shanks et al.'s (1989) instrumental paradigm: "subjects in judgment studies such as ours assume that the word 'causes' in the experimental instructions means 'causes immediately'. After all, they presumably have considerable experience of the immediacy of cause-effect relations in such electrical devices as computers" (Shanks et al., 1989, p.155).

The finding that delays impair causal reasoning performance is thus compatible both with the Knowledge Mediation hypothesis and the associative learning approach, which denies knowledge mediation. What is needed to disentangle the predictions from both accounts and to allow a systematic investigation into the role of temporal contiguity in causal induction is a paradigm that explicitly addresses and manipulates participants' expectations about the timeframes of the causal relations in question. In addition, participants would have to be exposed to immediate as well as delayed cause-effect pairings. In other words, the design would comprise the factors of (experienced) Time and Knowledge/Assumptions (about the timeframe). Such studies would effectively constitute an experimental test of Einhorn & Hogarth's (1986) knowledge mediation hypothesis, which is, surprisingly, still lacking in the literature.

A New Paradigm

Schlottmann (1999) reported a developmental study that involved immediate and delayed causal relations, and also explicitly manipulated participants' knowledge about the timeframe of the causal mechanism in question. Her apparatus consisted of a "mystery box" which could contain one of two toys. The box had two holes on its top,

and the toys inside the box could be placed in such a way that balls dropped into the holes fell onto the toy. One of the toys made a bell ring immediately if a ball was dropped onto it (the weight of the ball made a see-saw swing and hit the bell), while the other made the bell ring a few seconds after the ball was dropped onto it (the ball rolled down a sloped runway to crash into the bell). There was always only one toy inside the box, and it was positioned in such a way that only balls dropped from one of the two holes fell onto it and subsequently made the bell ring. Participants (five to ten year old children and a control group of adults) were left to explore the box and the two mechanisms, and learned that one mechanism made the bell ring immediately, while the other one involved a delay. The children then had to predict how long each mechanism would take to make the bell ring. They were also asked to make diagnostic inferences (“Which mechanism is inside the box, the slow one or the fast one?”) after observing one ball being dropped in a hole, followed by the bell ringing either immediately or after a delay. Children of all age groups performed very accurately on these tasks, both when the fast and the slow mechanism were inside the box. This result indicates that they could use their knowledge of the different timeframes of the mechanisms to bridge temporal gaps.

The critical test in Schlottmann’s study (1999), however, pitted knowledge of mechanism and contiguity directly against each other and involved a forced choice between a delayed and a contiguous candidate cause. First, the experimenter put the slow mechanism inside the box and confirmed that the child understood that the slow mechanism was inside the box. The child did not know, however, under which hole the mechanism was placed. The experimenter then dropped a ball in one hole (unbeknownst to the child: the one with the slow mechanism under it), waited for a couple of seconds,

and then dropped a second ball in the other (ineffective) hole. The dropping of the balls was carefully timed so that the bell would ring immediately after the second ball was dropped. Children then had to indicate which ball they thought had made the bell ring. Results showed that contiguity was necessary at least for 5 and 7-year olds: they (erroneously) preferred to attribute causality to the second ball (contiguous cause), even though they explicitly knew that the operating causal mechanism involved a delay. Older children and the control group of adults correctly attributed causality to the first ball (noncontiguous cause). Schlottmann's results from the forced-choice task show that contiguity is a very important cue to causality at least for young children, but her results from the exploration phase also show that even young children could learn that the slow mechanism involved a delay, and could successfully use this knowledge to draw simple causal inferences of both predictive and diagnostic nature. In the complex forced-choice task they failed, however, presumably because they could not integrate knowledge of a delayed mechanism with the immediate perceptual feedback of contiguity they received through the second (noncausal) ball. Adults and older children learned that mechanism was superordinate to contiguity, and could thus successfully integrate the contiguous feedback with their knowledge of mechanism to come to the correct conclusion.

Schlottmann's (1999) study provides very specific support for the knowledge mediation hypothesis (Einhorn & Hogarth, 1986), but her paradigm has some limitations that make it hard to generalize from her results. In particular, her experiments fall outside the scope of associative learning theory. The exploration phase, during which participants learned about the two possible configurations of the apparatus was mostly driven by acquiring physical concepts and understanding how each of the two

mechanisms worked. Although the children experienced balls being dropped into the mystery box several times, one could hardly refer to these encounters as learning trials in an associative learning sense. Participants' feedback was mostly in the form of a Socratic dialogue with the experimenter (e.g. "The slow runway ball is dropped first, and when it is almost there, then the fast seesaw ball is dropped second", Schlottmann, 1999, p. 307). Associative-learning algorithms operate via error-correction, driven by whether the effect occurs or not, given the presence of a particular candidate cause, but are silent as to how feedback of such complex quality as in Schlottmann's study could be represented. Another limiting factor that is true for all causal learning paradigms in the mechanistic tradition is that they usually employ deterministic causal mechanisms. Both mechanisms in Schlottmann's study always made the bell ring, the only difference being how quickly they did so. Most causal relations, however, are probabilistic. One of the key advantages of covariation-based theories (including associative learning, but also the computational causal power approach, Cheng, 1997) is that they explicitly account for knowledge gained from probabilistic feedback. Finally, the dependent measures in Schlottmann's study were based on a forced-choice between two options. Systematic studies on the impact of delay (e.g. Shanks et al., 1989) used rating scales to probe causal judgments. Only the latter procedure, but not the former, is sensitive to gradual decreases in judged causal strength due to delays.

The motivation of the work presented here is to see if knowledge mediation effects as reported by Schlottmann (1999) can also be obtained in situations commonly used in causal and associative learning tasks. The procedure of the experiments reported in this paper is similar to the one developed by Shanks et al. (1989), in that it employs a

free-operant procedure. Participants were instructed that their goal was to find out how strongly a response (pressing a button) caused an event on the computer screen.

Reinforcement was probabilistic, and the dependent measures were answers on a rating scale. To see whether contiguity was necessary for adults, as suggested by the associative learning account, or whether it could be mediated by knowledge, as in Schlotzmann's study, we varied cause-effect delays and manipulated expectations about the length of delay.

Experiment 1

Based on a pre-experimental questionnaire (N=38) we identified two thematic scenarios that would elicit appropriate assumptions about delays. In the Light bulb scenario participants expected the cause c (flicking a switch) to produce the effect e (bulb lights up) immediately (Mdn=0s); in the Grenade scenario they expected c (pressing a FIRE! button on a grenade launcher) to produce e (observing an explosion several miles away) after a delay (Mdn=8s).

The experiments employed a free operant procedure (FOP) with two important modifications. First, we abolished the notion of pre-defined learning trials. In earlier studies (e.g. Shanks et al., 1989, Experiments 1 and 2) contingencies were defined relative to experimenter-defined learning trials. This usually means that only the first response within a specified time-bin (e.g. 1s) is recorded and subjected to the reinforcement schedule. Employing such arbitrary learning trials unnecessarily reduces the ecological validity of the procedure, as the participants are usually not informed about the restrictions of the underlying trial structure. It potentially also results in a discrepancy between the programmed and the objectively achieved cause-effect

contingency. If participants ever respond at a rate that is higher than the frequency of trial-spaces (e.g. more than once a second), a substantial proportion of their responses would not be subjected to the reinforcement schedule and thus be effectively unreinforced. Our experiments therefore did not involve any pre-defined learning trials, but employed a truly continuous paradigm instead.

Secondly, we programmed the apparatus to continue recording (and subjecting to the reinforcement schedule) responses made during cause-effect delays. The reason for these alterations is that a standard FOP may actually lower the experienced contingency by not reinforcing responses during the delay period. Lower causal ratings in delayed FOPs thus might simply reflect reduced contingency, rather than detrimental effects of delay (Buehner & May, under review)

Method

Participants

18 undergraduate students (6 male, 12 female, median age: 19) from the University of Sheffield participated either to fulfil a partial course requirement or to receive £ 2.

Design

We combined two thematic scenarios and five levels of causality to produce a 2 x 5 within subject design. Scenario had the levels Light bulb and Grenade. The five levels of causality involved three experimental and two control levels. The three experimental levels all shared the conditional probability $P(\text{etc})=.75$, but had different cause-effect delays (0, 2, and 5 s). In these three experimental levels the outcome never happened

unless the participant pressed the button ($P(e|c)=0$). There were no pre-defined learning trials, i.e. any button press triggered the effect with a probability of .75 after the relevant delay. Responses made during a delay period were also recorded and subjected to the reinforcement schedule. Each condition lasted 2 minutes. The two control levels involved 24 background outcomes, which were not influenced by the participant pressing or not pressing the key. These background outcomes were randomly scheduled within each 2-minute condition with the restrictions that (a) one background event occurred in each of 24 5-second intervals and (b) background events were separated by at least 500ms. One control level employed $P(e|c)=0$ (so that participants' actions were ineffective), the other $P(e|c)=.75$ (i.e. participants' actions had a 75% chance to produce the outcome, just as in the experimental condition) with an action-outcome delay of 0 seconds. In each scenario (Light bulb and Grenade) there were thus three experimental conditions with $P(e|c)=.75$ and 0, 2, or 5s delays, and two control conditions, one with $P(e|c)=.75$ and 0s delay, but an additional 24 background events, and one with $P(e|c)=0$ and 24 background events.

For the four levels with a programmed value of $P(e|c)=.75$, the mean value actually obtained was .76.

Materials and Procedure

We used a Macintosh 8500 computer, running Macromedia Director 7.0, to administer the experiment. Participants read instructions on the screen, telling them that in this experiment they had to learn to what extent their actions caused something to happen on the computer screen. They then proceeded to a specific description of the first scenario. The Light bulb scenario asked participants to imagine that a switch and light

bulb displayed on the screen were connected. It further instructed them about another switch in a different room invisible to them. Other persons might flick that switch sometimes, so that when the bulb lights up, it could be either due to the participant's or other people's action. The instructions for the Grenade scenario were structurally identical, but asked participants to imagine that a FIRE! Button operated a grenade launcher to fire shells into a training range, and that other people might also fire shells into the range (see Appendix).

In the Light bulb conditions, the computer displayed a drawing of a light bulb centrally against a grey background. 2cm below the light bulb was a white rectangular push-button labelled "Lightswitch" which inverted its colours when participants clicked on it with the mouse. An effect was represented as the bulb illuminating in yellow for 500ms accompanied by a whistling sound. In the Grenade scenario, the screen centrally displayed a rectangular viewing window (5 x 12 cm). The window displayed a view of the horizon with lines of schematic trees on the left and right side of the landscape. Approximately 2cm below the window was a blue rectangular push-button labelled "FIRE!". An effect was implemented as a 500ms display of a red and orange mushroom cloud in the centre of the landscape, accompanied by an explosion sound. To make the representations more realistic, a continually repeating 20-second sound loop of gunfire and battlefield noises provided background sound in the Grenade scenario. The Light bulb scenarios were accompanied by a continually looping 20-second piece of funky music. Each condition lasted 2 minutes. At the end of each condition, participants had to indicate whether or not there was a causal relation between clicking the button and the outcome by selecting one of two radio buttons. If a participant indicated and confirmed

that there was no causal relation, the answer was scored as 0. If participants indicated that there was a causal relation, they read the following prompt:

Nobody else is (turning on the light / causing explosions).

If you clicked the switch 100 times,

(how often would the bulb light up? / how many explosions would occur?)

Participants then had to enter a number between 1 and 100, confirm their entry and then proceeded to the next condition. After having worked on all five conditions in one scenario (Light bulb or Grenade), participants received instructions and worked on the five conditions for the other scenario. The order of conditions within a scenario was random, and the order of scenarios was counterbalanced between subjects. Participants worked individually and the experiment lasted about 30 minutes.

Results

Table 1 displays participants' mean ratings for all five levels of causality, broken down by scenario. Mean causal ratings in the experimental condition with $P(\text{elc})=.75$ and a 0s delay were close to the normative level for both the Light bulb(80.00) and the Grenade (81.56) conditions. As the temporal delay increased, ratings decreased in both scenarios. Results from the control conditions show that when $P(\text{elc})=.00$, causal ratings were substantially lower than in any of the three experimental conditions. However, introducing 24 “uncaused” events but maintaining $P(\text{elc})=.75$ with no delay only produced slightly lower ratings than the corresponding experimental conditions.

The main point of interest in this study was how the causal assessment of identical contingencies is affected by cause-effect-delays, and how knowledge mediates the

influence of delay, so further analyses exclude the control conditions. Inspection of Figure 1 suggests a main effect of delay, unmediated by knowledge about the timeframe of the causal relation. However, the heterogeneity of variance between the experimental conditions does not permit an ANOVA. It is worthwhile to look at the distribution patterns of causal ratings in the experimental conditions, to understand how these differences in variance came about. Figure 2 displays histograms of the causal ratings in the six experimental conditions. In both the Light bulb and Grenade scenarios ratings for the 0s delay conditions scattered closely around 80 and varied widely for the 2s delay conditions. For the 5s delay, ratings in the Light bulb scenario were also extensively distributed, but low ratings (0 to 10) are the largest category. However, in the grenade condition the 5s delay produced a bi-modal distribution of judgments: participants largely either rated the causal relation to be non-existent (0 -10) or to be perfectly deterministic (90-10), with five participants providing ratings in between those extremes.

To investigate whether the impact of delay is mediated by knowledge, we therefore computed the differences between participant's ratings in the 0s and 5s conditions of each scenario. The distribution of difference scores fulfils the key assumptions for parametric statistics and can thus be evaluated on a t test. The influence of delay was the same in the Light bulb ($M=29.83$, $SD=34.14$) and Grenade ($M=29.17$, $SD=41.27$) scenarios, $t(17)=.08$, n.s.

Discussion

In the Light bulb scenario participants expected an immediate causal link between their actions and the outcomes. Ratings in this scenario were extremely close to the implemented contingency when action and outcome followed each other immediately.

However, when an implausible delay separated cause and effect, participants gave lower causal ratings, even though the contingencies were equal across conditions.

In the Grenade scenario participants expected a delayed causal relation. According to Einhorn and Hogarth (1986) such expectations should bridge temporal gaps between causes and effects and render immediate cause-effect pairings noncausal. Contrary to this hypothesis, participants did not rate immediate contingencies lower under Grenade than Light bulb instructions; $Z=.35$, $p=.72$ on a Wilcoxon test. The immediate feedback of a high positive contingency in the no-delay conditions was sufficient to establish an impression of a causal link between participants' actions and the outcomes, despite prior expectations about delay. In the absence of this immediate feedback, causal ratings dropped just as much in the Grenade as in the Light bulb scenario. Experiment 1 thus demonstrated a main effect of delay on causal assessment of identical contingencies, but failed to demonstrate evidence for knowledge mediation.

Prior experience of contiguity and within-subjects design

The design employed in Experiment 1 may have been sub-optimal to study knowledge-mediation in assessment of delayed causal relations for a number of reasons. Because the experiment was conducted on a computer, the knowledge manipulation induced by the cover stories may not have been convincing enough. Participants were very aware that the computer controlled the presentation of outcomes. Every participant rated each combination of delay and contingency twice, once with Light bulb, once with Grenade instructions. It seems therefore especially plausible that participants noticed the structural identity between the problems. Another reason for the absence of knowledge mediation in Experiment 1 could be that (previously) experienced contiguity is a very

powerful cue to causality. Participants underestimate delayed contingencies more severely if they previously encountered contiguous as compared to delayed contingencies (Buehner & Hagmayer, in preparation, Buehner & May, under review). The logical way to circumvent both of these problems is to forfeit the economy of a within-subject design and collect data between-subjects instead.

Perceptual or causal judgments?

The dependent variable employed in Experiment 1 was based on a frequency estimate couched in a counterfactual question (e.g., “If you clicked the switch 100 times, how many explosions would occur?”), because earlier work (Buehner et al., in preparation) showed that judgments on a conventional rating scale (“How strongly do you think clicking the switch causes explosions?”) can reflect a conflation of causal strength and the reliability of the information provided. Such conflations are particularly problematic in designs where reliability and causal strength vary independently (e.g. Lober & Shanks, 2000, Buehner & Cheng, 1997).

In this study, however, reliability was not an issue, as there was only one contingency. The frequency estimate may not be the best way to assess causal judgements in the context of delays, since a participant who experiences a high contingency implemented with an implausible long action-outcome delay may think that the causal relation is weak, but still observe that effects are frequently preceded by candidate causes. Consequently, a high frequency estimate would be consistent with the perceptual quality of the learning experience. Because the question format used in Experiment 1 was ambiguous, some participants may have provided perceptual judgments while others gave actual causal ratings. This would explain the heterogeneity

of variance, including the bi-modal distribution in the experimental conditions of Experiment 1.

Another problem with the rating procedure in Experiment 1 has to do with the wording of the counterfactual question: “Nobody else is causing explosions”. While we intended it to be interpreted as a counterfactual, it may well be that some participants did not process the statement as a counterfactual, but as a true statement instead. Some participants may have taken the statement at face value, which of course would have violated what they learnt in the instructions, namely that another person is also trying to cause the effect. This in turn would have led them to attribute all occurrences of the effect to their button presses, as – according to their belief – no other plausible cause existed. The frequency estimate procedure is also potentially confusing with respect to the timeframe, both over which one has to imagine pressing the button 100 times and over which one has to imagine the occurrence of potential effects. Because this experiment dealt with causal relations elapsing in continuous time, frequency estimates are probably conceptually harder to represent than ratings on a scale. For all these reasons we decided to change the dependent variable used to probe causal estimates in the next two experiments to a standard rating scale.

Experiment 2

Experiment 2 was a between-subjects replication of Experiment 1, but used a rating scale to probe for causal ratings, and employed only the 0s and 5s experimental conditions only in order to obtain a larger sample in the maximally informative conditions.

Method

Participants

160 (116 female, 44 male, median age 18) visitors to the Department of Psychology, Sheffield participated as part of an Open Day Demonstration, run on 4 separate days. Participants were randomly assigned to the four conditions, but an effort was made to achieve approximately equal numbers in each cell. The Lightbulb/0s cell comprised 57, the Lightbulb/5s 36, the Grenade/0s 36, and the Grenade/5s 31 participants.

Design, Materials, and Procedure

The materials were identical to Experiment 1, except that only the 0s and 5s experimental conditions were included, and the question used to probe causal ratings was reworded. Because the experiment was run in large groups, all the sound effects (background music/gunfire and whistle/explosion) were removed in order to keep mutual disturbance at a minimum. The initial instructions now informed participants that they would solve one problem, lasting about two minutes. Subsequently, participants read the specific instructions relevant for their assigned scenario and then proceeded to the experiment. After having sampled evidence from the continuous paradigm for two minutes, participants had to rate causal strength on a scale from 0 to 100. In order to minimize any possibility that participants' ratings were influenced by the perceptual quality of the feedback, we labelled the extreme ends and midpoints of the scale in such a way that would encourage participants to provide ratings based on their causal beliefs, i.e. "0 means that clicking the switch has no influence on whether or not the effect occurs",

“50 means that clicking the switch moderately causes the effect”, and “100 means that clicking the switch strongly causes the effect”.

Results

Figure 3 displays mean ratings of causal strength in the four groups. As in Experiment 1, participants interpreted contiguous contingencies as highly causal in both the Light bulb ($M=74.1$, $SD=12.51$) and Grenade ($M=74.7$, $SD=12.56$) scenarios. The 5s delay generally resulted in lower ratings of causality; this finding was more pronounced in the Light bulb ($M=38.9$, $SD=27.44$) than the Grenade ($M=51.2$, $SD=20.82$) group. Visual inspection of Figure 4 reveals that the distributions of causal ratings were considerably more normal than in Experiment 1. Causal ratings in the 5s Grenade condition in particular did not follow a bi-modal distribution, but scattered mostly between 30 and 80, producing a flattened normal distribution.

An ANOVA with significance level .05 revealed main effects of Delay, $F(1,156)=95.272$, and Cover Story, $F(1,156)=4.556$, and a marginal Delay x Cover Story Interaction, $F(1,156)=3.811$, $p=.053$. The Delay x Cover Story interaction indicates that the impact of delay on causal ratings was mediated by assumptions about the timeframe of the causal mechanism. Planned t-tests with Welch-Satterthwaite adjusted degrees of freedom (Howell, 1997) confirmed the qualitative pattern described above: causal ratings were significantly lower in the Light bulb than in the Grenade group with 5s $t(64)=2.083$, $p<.05$, but not with 0s delay, $t(29)=.208$, n.s.

Discussion

Implicit assumptions about the timeframe of a causal relation determined how delayed contingencies were evaluated with respect to causal strength. Causal ratings no

longer followed a bi-modal distribution, a sign that the dependent variable used to probe participants' estimates of causal strength was much clearer and less ambiguous than in Experiment 1. The Delay x Cover Story interaction was straightforward to interpret: given a constant cause-effect contingency of .75, a 5s delay always impaired causal judgments relative to no delay; this detrimental effect of delay was less pronounced when participants thought a delay was plausible (Grenade instructions) than when they expected immediate cause-effect pairings (Light bulb instructions). This finding supports Einhorn and Hogarth's (1986) knowledge-mediation hypothesis: participants bridged temporal gaps between causes and effects better when they assumed a causal mechanism that took time to unfold, than when they contemplated an immediate mechanism. Such a top-down influence of abstract knowledge is incompatible with an associationist perspective on causal learning. Associationism is a strict bottom-up approach, and has no scope for instructional effects on learning; instead, what is learnt solely depends on the learning experience.

The learning experience in the Light bulb scenarios was different from the Grenade scenario, however, as the stimuli were dissimilar in these two conditions. One could therefore argue that this difference in the visual properties of the learning situation led to diverging evaluations of the delayed contingencies. Higher causal ratings in the Grenade than in the Light bulb scenario conditions involving a delay thus may be not the result of instructional effects, but merely the result of inconsistent stimulus properties between the scenarios. In associative learning theory, properties of the stimuli present in a learning situation are usually represented in learning parameters. In the Rescorla-Wagner learning rule (Rescorla & Wagner, 1972) $\Delta V_{CS} = \lambda_{CS} \cdot \lambda_{US} \cdot (\lambda_{US} - V)$, for

instance, the change of associative strength on a given learning trial ΔV is a function of the difference between the experienced outcome (Δ) and the expected outcome (ΔV), multiplied by the learning parameters α and β , which, respectively, represent the salience of the conditioned and unconditioned stimuli. From such a perspective, the stimuli in the Grenade scenario may have simply been more salient than those in the Light bulb scenario. Higher parameters in the Grenade than in the Light bulb scenario could thus account for higher ratings of causal strength in the former but not the latter scenario when the relation involved a 5s delay. A corollary from such an interpretation is that in the immediate conditions the Grenade scenario should likewise support higher estimates of causal strength than the Light bulb scenario. Our results do not confirm this, casting doubt on the associationist interpretation of the scenario effect. Nonetheless, differences between the visual properties of the learning experience could constitute a confounding factor that should be ruled out in order to strengthen the case for knowledge mediation.

Experiment 3

The goal of Experiment 3 was to test whether effects of knowledge mediation could also be obtained when stimulus properties are identical across conditions. To this end, we presented all participants with the same stimulus materials taken from the Grenade scenario of Experiments 1 and 2. To induce different assumptions about the timeframe of the causal relation in question, we employed two different instructions. One group of participants received instructions similar to those used for the Grenade conditions in Experiments 1 and 2: they were told their task was to test whether a Grenade launcher causes explosions in a training range several miles away from their post (Grenade-Delayed). Another group of participants was also told to check how strongly

pressing the FIRE! button causes explosions, but instructions stated that the FIRE! button was part of a remote control detonator; each time it was clicked, it emitted a radio signal directed to the site (Grenade-Immediate). The instructions stated explicitly that explosions should occur immediately in the latter, and after a few seconds in the former scenario (see Appendix). As in Experiment 2, we focused on the 0s and 5s conditions only, and collected data between subjects.

Method

Participants

Ninety-one volunteers participated in the study, 31 of which were enrolled in an undergraduate psychology course at the University of Sheffield and participated to fulfil a course requirement. The remaining 60 participants were visitors to the Department of Psychology, University of Sheffield, who participated during departmental Open Day demonstrations. Participants were randomly assigned to one of the four conditions, but an effort was made to achieve approximately equal numbers in all cells. The 0s-Grenade-Immediate cell comprised 22, the 0s-Grenade-Delay cell 23, the 5s-Grenade-Immediate 21, and the 5s-Grenade-Delay cell 25 participants.

Design, Materials, and Procedure

Design and Procedure were identical to Experiment 2. The only difference between Experiments 2 and 3 consisted in the stimulus materials employed. All participants saw FIRE! buttons and Explosions as stimuli, but instructions suggested an immediate causal link in one (Grenade-Immediate), and a delayed link in the other group (Grenade-Delayed). Full instructions for both groups are listed in the Appendix.

Results

Figure 5 displays mean ratings of causal strength in the four groups. Experiment 3 replicated the findings of Experiment 2. Participants evaluated a .75 contingency with immediate cause-effect pairings as strongly causal when the instructions suggested an immediate causal mechanism ($M=77.41$, $SD=14.26$); they interpreted the same immediate contingency as slightly less causal when instructions stated that the mechanism in question involves a delay ($M=69.83$, $SD=15.72$). Delayed contingencies were generally judged to be less causal, but the impact of the 5s delay was stronger in participants who expected immediate cause-effect pairings ($M=30.96$, $SD=21.13$), than in participants who assumed a delayed mechanism ($M=54.80$, $SD=15.72$).

An ANOVA with the factors Delay and Cover Story corroborated the qualitative pattern of results described above. It revealed a highly significant main effect of Delay, $F(1, 87)=47.9$, and a highly significant interaction between Delay and Cover Story $F(1, 87)=12.52$. Planned t-tests with Welch-Satterthwaite adjusted degrees of freedom (Howell, 1997) revealed that causal ratings for the 5s contingency were significantly higher in the Delay than the Immediate Instructions group, $t(43)=3.22$; the 0s contingency was evaluated as slightly less causal by participants receiving Delay Instructions compared to participants receiving Immediate Instructions, but the difference fell short of statistical significance, $t(43)=1.69$.

Discussion

The results of Experiment 3 showed that knowledge about the timeframe of causal relations as manipulated via instructions determined the extent to which a 5s delay impaired the assessment of a .75 contingency. When participants were told that the

causal mechanism between pressing a button and observing an explosion on the screen would take a few seconds to unfold, they bridged an experienced temporal gap between cause and effect much better than participants who were led to expect immediate pairings. This knowledge-mediation effect was obtained by merely changing the wording of a couple of sentences in the instructions. Because the visual properties were identical across conditions, it is not plausible to explain any differential impact of delay on causal ratings by differences in the learning experiences (i.e. by postulating different cue- and outcome-salencies for the two scenarios). In a related study (Buehner & May, under review), participants were presented with equally impoverished stimuli in all conditions (a triangle flashing on the computer screen, adapted from the original study by Shanks et al., 1989). Again, delayed contingencies were rated as significantly more causal when participants were alerted about the possibility of delays than when they were left ignorant of it. Such instructional effects fall outside the scope of associative learning theory and must be attributed to abstract knowledge (induced via instructions) driving the parsing of causal episodes.

In Experiment 3 the explicit reference to a specific cause-effect interval (“immediate” vs. “a few seconds”) afforded clearer and stronger results than either of the preceding experiments employing implicit manipulations. Overall, the pattern of results suggests that knowledge of both implicit (Experiment 2) and explicit nature (Experiment 3) influences the timeframe of covariation assessment in causal reasoning. The clearer and stronger pattern of results from Experiment 3 suggests that the extent of knowledge-mediation may be determined by how well established the knowledge-base is: explicit

instructions may be more reliable and thus exert stronger influence than implicit assumptions.

General Discussion

Our goal was to disentangle the predictions regarding cause-effect delays from causal power and associative learning theories. According to the former, the influence of delay crucially hinges on prior assumptions about the timeframe of the causal relation in question. According to the latter, instructions manipulating temporal assumptions should have no impact on the extent of this detrimental effect of delay. Overall, the pattern of results clearly contradicted predictions derived from a strict associative learning perspective. The extent to which action-outcome delays impaired causal ratings was significantly reduced when participants were led to believe that the causal relation in question involved a delay. Experiment 3 showed that this increased tolerance of delays cannot be attributed to variations in stimulus saliency. Instead, our experiments demonstrated that participants evaluate identical delayed covariations differently, depending on the temporal assumptions they bring to bear. Knowledge thus mediates covariation assessment in human causal induction.

Although the knowledge mediation effect we found for the assessment of delayed contingencies supports a causal power perspective, the absence of such an effect in the evaluation of immediate contingencies may pose a problem for this framework. With the exception of Experiment 3, immediate contingencies consistently received equally high causal ratings, regardless of whether participants expected immediate or delayed mechanisms. This finding is at variance with a strong interpretation of the knowledge-mediation hypothesis. Immediate cause-effect pairings should not give rise to causal

interpretations if the hypothesized mechanism calls for a delay; instead they should be judged as spurious. We thus have to be cautious in interpreting our results as uniformly supporting an unmodified causal power perspective. Be that as it may, associative learning theory as applied to causal judgment (see e.g. Dickinson, 2001; Shanks & Dickinson, 1987) seems to be incapable of accounting for the instructional effects we found on the assessment of delayed contingencies. It remains to be determined whether it could in principle be modified to account for knowledge mediation effects.

Is there a bottom-up contiguity bias?

The absence of a knowledge-mediation effect on the evaluations of immediate contingencies would imply that contiguity is more important than expectations about the timeframe of the causal relation, or that contiguity overrides temporal assumptions. All previous research on this topic (including Shanks et al., 1989; and Schlottmann, 1999) has shown that contiguity is a very powerful cue to causality. Schlottmann's results from the 5 and 7 year olds and our results from the contiguous conditions showed that experienced contiguity may even appear to override knowledge-based expectations about the timeframe of the causal relation in question, suggesting a bottom-up contiguity bias. It could be, for instance, that the influence of knowledge mediation is confined to evaluations of delayed contingencies, and that temporal assumptions do not affect causal judgment in case of immediate events. In other words, it is unclear as yet whether adults, when reasoning about complex probabilistic causal relations will always behave like 5-7 year olds, or whether they can in principle appreciate the necessity of delays. Naïve intuitions about everyday causal inference certainly suggest the latter.

One explanation for the strong effect of temporal contiguity in our experiments could be that the task was administered on a computer. As discussed earlier, people have very rich conceptions about (and, in most cases, a great deal of experience with) interactions with computers. Consequently, our participants must have been aware that the computer controlled all aspects of the stimulus display. We tried to minimize the impact of this problem by introducing a between-subjects design in Experiments 2 and 3, but apparently did not succeed completely. Even when solving only one 2-minute problem, participants in the Grenade scenario had no reason to believe that a delay between cause and effect was necessary. We managed to induce participants to think that a delay between their actions and the display of outcomes was plausible, but could not get them to assume it was essential.

Another possible reason for why our experiments failed to provide support for the strong claim of the Knowledge Mediation hypothesis could be a discrepancy in the amount of experience participants had with the (imagined) causal mechanisms. A typical undergraduate student probably is exposed to the Light bulb problem in the real world several times throughout the day. We are surrounded by electric light nearly everywhere we go, and we use it regularly. Consequently, we have a vast amount of experience with the simple causal connection between a light switch and the bulb lighting up. In contrast, it is unlikely that any of the participants ever had any direct experience of firing off grenades. As a consequence, participants' beliefs about the timeframe of the causal mechanism in question could have been more stable in the Light bulb than in the Grenade scenario. Interestingly, results from Experiment 3 did show a nonsignificant trend towards supporting the strong claim of the Knowledge Mediation hypothesis: causal

ratings for the 0s contingency were lower when participants were told the relation involved a delay than when they were told it was instantaneous. It may well be that despite the overall unfamiliarity with the Grenade scenario, an explicit mention of a specific time lag had a stronger and more stable impact on participants' performance than the implicit assumptions generated by the different scenarios in Experiments 1 and 2.

These explanations notwithstanding, it is important to stress that participants were "correct" (in a normative sense) in assigning high causal ratings to contiguous conditions in both the Light bulb and Grenade scenarios. After all, their actions really did cause the outcomes with a probability of .75, and the computer never produced the outcome unless they pressed the button. Because participants knew that the computer controlled stimulus display and feedback, it would have even been irrational for them to deny the existence of a causal link between their actions and the outcomes in the 0s Grenade scenario.

Participants would have had to explicitly disregard objective evidence that suggested a strong causal link in order to conclude that the 0s Grenade condition was non-causal. Be that as it may, in order to make an even more compelling case for knowledge mediation, future research will have to demonstrate its influence on contiguous contingencies as well. In order for such studies to succeed, it seems imperative to create a scenario where delays are indeed judged to be necessary, rather than merely plausible. Our results suggest that a Free-Operant Procedure on a computer may be ill-suited for these requirements. Real physical causal mechanisms, like the ones in Schlottmann's (1999) study are probably better suited for strong tests of Einhorn and Hogarth's (1986) hypothesis. Adult participants in her study easily learnt that a five second delay was necessary when the slow toy was inside the box.

Knowledge-based Causal Induction

We have demonstrated that different beliefs about the timeframes of causal relations result in different interpretations of identical covariations. Prior knowledge mediates causal inference in many other interesting ways. Michael Waldmann, for example, showed that different assumptions about causal models (predictive vs. diagnostic learning) likewise result in different interpretations of identical covariations (see Waldmann, 1996 for an overview). Both his work and the experiments presented here demonstrated an interaction between bottom-up (covariation assessment) and top-down (knowledge mediation) components in causal induction. Associationism is at a loss explaining such interactions, as it disallows any influence of knowledge (be it assumptions about structure or delay) beyond pre-existing associations. The causal power approach, although it accounts for knowledge mediation, suffers circularity: it cannot explain how knowledge is acquired in the first place. Cheng's (1997) power PC theory combines bottom-up and top-down components, and suggests that all aspects of causal knowledge can ultimately be derived from observation. Lien and Cheng (2000), for instance, demonstrated that humans are able to derive abstract categories of causal and non-causal entities from experienced covariations, and use this category knowledge when classifying novel objects as genuine or spurious causes. Power assumptions that distinguish causal from non-causal covariations thus are not shrouded in mystery or innateness but can themselves be inferred from covariation.

What lies ahead is to determine how people acquire assumptions about the timeframe of causal relations (other than extracting them from experimental instructions). Analogously to Lien and Cheng's findings, future research may show humans to be

capable of deriving surprisingly well-formulated temporal assumptions from statistical information. Once again David Hume can inspire our search: “In vain, therefore, should we pretend to determine any single event, or infer any cause and effect, without the assistance of observation and experience” (1777/1902 p.30).

References

- Ahn, W.-K., Kalish, C. W., Medin, D. L., & Gelman, S. A. (1995). The role of covariation vs. mechanism information in causal attribution. Cognition, *54*, 299-352.
- Buehner, M. J., & Cheng, P. W. (1997, August). Causal induction: The power PC theory versus the Rescorla-Wagner model. Paper presented at the Nineteenth Annual Conference of the Cognitive Science Society, Stanford, CA.
- Buehner, M. J., Cheng, P. W., & Clifford, D. (in preparation). From covariation to causation: A test of the assumption of causal power.
- Buehner, M. J., & Hagmayer, Y. (in preparation). Contingency, contiguity, and prior experience in causal learning.
- Buehner, M. J., & May, J. (under review). Temporal contiguity and the judgment of causality revisited: How Prior Knowledge, Experience, and Reinforcement Procedure contribute to the Assessment of Delayed Causal Relations. Manuscript under Review, Quarterly Journal of Experimental Psychology.
- Bullock, M., Gelman, R., & Baillargeon, R. (1982). The development of causal reasoning. In W. J. Friedman (Ed.), The developmental psychology of time (pp. 209-254). New York: Academic Press.
- Cheng, P. W. (1997). From covariation to causation: A causal power theory. Psychological Review, *104*(2), 367-405.
- Dickinson, A. (2001). Causal Learning: An associative analysis. Quarterly Journal of Experimental Psychology Section B- Comparative and Physiological Psychology, *54*(1), 3-25.

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

Garcia, J., Ervin, F. R., & Koelling, R. A. (1966). Learning with prolonged delay of reinforcement. Psychonomic Science, *5*(3), 121-122.

Garcia, J., & Koelling, R. A. (1966). Relation of cue to consequence in avoidance learning. Psychonomic Science, *4*(3), 123-124.

Howell, D. C. (1997). Statistical methods for psychology (4th ed.). Belmont, CA: Wadsworth.

Hume, D. (1739/1978). A treatise of human nature. In L. A. Selby-Bigge (Ed.), Hume's treatise. Oxford, England: Clarendon Press.

Hume, D. (1777/1902). An enquiry concerning human understanding. In L. A. Selby-Bigge (Ed.), Hume's Enquiries. Oxford, England: Clarendon Press.

Lien, Y. W., & Cheng, P. W. (2000). Distinguishing genuine from spurious causes: A coherence hypothesis. Cognitive Psychology, *40*(2), 87-137.

Lober, K., & Shanks, D. R. (2000). Is causal induction based on causal power? Critique of Cheng (1997). Psychological Review, *107*(1), 195-212.

Mendelson, R., & Shultz, T. R. (1976). Covariation and temporal contiguity as principles of causal inference in young children. Journal of Experimental Child Psychology, *22*(3), 408-412.

Michotte, A. E. (1963). The perception of causality (T. R. Miles, Trans.). London, England: Methuen & Co.

Rescorla, R. A., & Wagner, A. R. (1972). A theory of Pavlovian conditioning: Variations in the effectiveness of reinforcement and nonreinforcement. In A. H. Black &

W. F. Prokasy (Eds.), Classical Conditioning II: Current theory and research (pp. 64-99).
New York: Appleton-Century Crofts.

Schlottmann, A. (1999). Seeing in happen and knowing how it works: How children understand the relation between perceptual causality and underlying mechanism. Developmental Psychology, 35(5), 303-317.

Shanks, D. R. (1993). Human Instrumental Learning - a Critical-Review of Data and Theory. British Journal of Psychology, 84, 319-354.

Shanks, D. R., & Dickinson, A. (1987). Associative Accounts of Causality Judgment. In: G.H. Bower (Ed.) The Psychology of Learning and Motivation-Advances in Research and Theory (Vol. 21, pp. 229-261). San Diego, CA: Academic Press.

Shanks, D. R., Holyoak, K. J., & Medin, D. L. (Eds.). (1996). The Psychology of Learning and Motivation (Vol. 34): Causal Learning. San Diego, CA: Academic Press.

Shanks, D. R., Pearson, S. M., & Dickinson, A. (1989). Temporal Contiguity and the Judgment of Causality By Human Subjects. Quarterly Journal of Experimental Psychology Section B- Comparative and Physiological Psychology, 41(2), 139-159.

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

Waldmann, M. R., & Hagmayer, Y. (1999). How categories shape causality. In M. Hahn & S. C. Stoness (Eds.), Proceedings of the Twenty-first Annual Conference of the Cognitive Science Society (pp. 761-766). Mahwah, NJ: Erlbaum.

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.

White, P. A. (1995). Use of prior beliefs in the assignment of causal roles: Causal powers versus regularity-based accounts. Memory and Cognition, 23(2), 243-254.

Appendix

General Instructions for Experiment 1

In this experiment you have to evaluate the extent to which your actions can cause something to happen. There will be a button on the computer screen and your task is to observe whether clicking it causes something to happen on the screen.

You can choose at any time whether or not to click the button. You can click it as often or as little as you like. However, because of the nature of the task it is to your advantage to click it some of the time and not to click it some of the time.

The effectiveness of you clicking the button stays the same within a particular condition but may well vary between problems.

At the end of each problem you will be asked whether clicking the button causes the outcome and if so, how strongly it causes the outcome.

You will work on two different scenarios with five problems per scenario, each lasting about two minutes.

Specific Instructions for Light bulbscenario

In the upcoming five problems there will be a light bulb and a lightswitch on the screen. Your task is to judge the extent to which clicking the switch causes the bulb to light up.

Imagine that the light bulb is connected to the switch you can click on and to a switch in another room that other persons can flick without you being aware of it. Thus, if the bulb lights up, it may be because you clicked the switch or because a person in the other room flicked the second switch.

You can choose at any time whether or not to click the switch. You can click it as often or as little as you like. However, because of the nature of the task it is to your advantage to click it some of the time and not to click it some of the time.

You will work on five different problems with the light bulb, each lasting for 2 minutes. The relationship between your clicking the switch and the bulb lighting up will be constant within each problem but may well differ from one problem to the next.

At the end of each problem you will be asked whether and how strongly clicking the lightswitch makes the bulb light up.

Specific Instructions for Grenade scenario

In the upcoming five problems you will view a military training range from a command post several miles away. Your task is to find out whether clicking on a "FIRE!" button produces explosions in the range.

Imagine that the "FIRE!" button operates a grenade launcher which is situated in your post and fires grenades into the training range. When a grenade you've fired hits the training range, you will see an explosion. However, an officer in another post is also firing into the range. Thus if you see an explosion, it may be because you clicked the "Fire!" button or because the second officer in the other post launched a grenade.

You can choose at any time whether or not to click on "Fire!". You can click it as often or as little as you like. However, because of the nature of the task it is to your advantage to click it some of the time and not to click it some of the time.

You will work on five different problems with the grenade launcher, each lasting for 2 minutes. The relationship between your clicking "Fire!" and explosions in the range will be constant within each problem but may well differ from one problem to the next.

At the end of each problem you will be asked whether and how strongly clicking "FIRE!" causes explosions in the range.

Coverstory for Grenade-Immediate Scenario in Experiment 3

You are a military officer on an army training site. Today's training objective is to test some weapons used in your unit. In particular, your task is to find out whether triggering a "FIRE!" button produces explosions in the training range.

In order to find out whether the equipment works correctly, you and a colleague have positioned yourselves in two individual command posts, which allow a perfect view of the training range a few miles away. You each have a remote control detonator; when you click "FIRE!", it sets off a mine in the training range. The detonator emits a radio signal directed from your post to the site, so your clicking of the "FIRE!" button should produce an explosion right away.

Remember that there is also another officer, who is also testing his explosives at the same time. Thus, if you observe an explosion, it may be because you clicked the "FIRE!" button or because the other officer caused an explosion.

You can choose at any time whether or not to click on "FIRE!". You can click it as often or as little as you like. However, because of the nature of the task it is to your advantage to click it some of the time and not to click it some of the time.

You will work on the task for 2 minutes.

The relationship between your clicking "Fire!" and explosions in the range will be constant throughout the task.

At the end of the problem you will be asked how strongly clicking your "FIRE!" button causes explosions in the range.

If you have no further questions, click "OK" to start.

Coverstory for Grenade-Delay Scenario in Experiment 3

You are a military officer on an army training site. Today's training objective is to test some weapons used in your unit. In particular, your task is to find out whether triggering a "FIRE!" button produces explosions in the training range.

In order to find out whether the equipment works correctly, you and a colleague have positioned yourselves in two individual command posts, which allow a perfect view of the training range a few miles away. You each have a grenade launcher; when you click "FIRE!", it launches shells into the training range. The shells have to travel from your post to the site, so a few seconds should pass between your clicking of the "FIRE!" button and the explosion.

Remember that there is also another officer, who is also testing his explosives at the same time. Thus, if you observe an explosion, it may be because you clicked the "FIRE!" button or because the other officer caused an explosion.

You can choose at any time whether or not to click on "FIRE!". You can click it as often or as little as you like. However, because of the nature of the task it is to your advantage to click it some of the time and not to click it some of the time.

You will work on the task for 2 minutes.

The relationship between your clicking "Fire!" and explosions in the range will be constant throughout the task.

At the end of the problem you will be asked how strongly clicking your "FIRE!" button causes explosions in the range.

If you have no further questions, click "OK" to start.

Author Note

Marc J Buehner and Jon May, Department of Psychology, University of Sheffield.

This work has been funded by the EU T&MR network TACIT and ESRC grant R000223692. This work is part of a PhD thesis carried out by the first author under supervision of the second author. We thank Pat Cheng, Sabine Pahl, and Mat White for their helpful comments on an earlier draft. Correspondence concerning this article should be addressed to Marc J. Buehner, Department of Psychology, University of Sheffield, Western Bank, Sheffield S10 2TP, England, UK. Electronic mail may be sent to marc.buehner@sheffield.ac.uk.

Table 1 Results from Experiment 1. Control conditions included 24 background events.

Standard deviations in brackets. N=18

Condition	P(etc)	Delay	Scenario	
			Lightbulb	Grenade
Control	0.00	n/a	3.44 (11.78)	11.39 (20.13)
	0.75	0s	71.67 (25.38)	69.17 (28.56)
Experimental		0s	80.00 (9.08)	81.56 (11.83)
	0.75	2s	59.17 (29.57)	56.11 (35.67)
		5s	50.17 (36.61)	52.39 (43.61)

Figure Captions

Figure 1. Experiment 1: Mean ratings of causal strength for experimental conditions.

Error bars indicate standard errors.

Figure 2. Experiment 1: Distribution patterns of causal ratings.

Figure 3. Experiment 2: Mean ratings of causal strength. Error bars indicate standard errors.

Figure 4. Experiment 2: Distribution patterns of causal ratings.

Figure 5. Experiment 3: Mean ratings of causal strength. Error bars indicate standard errors.