

This is an Open Access document downloaded from ORCA, Cardiff University's institutional repository: <https://orca.cardiff.ac.uk/id/eprint/127121/>

This is the author's version of a work that was submitted to / accepted for publication.

Citation for final published version:

Greville, W. James, Buehner, Marc J. and Johansen, Mark K. 2020. Causing time: Evaluating causal changes to the 'when' rather than the 'whether' of an outcome. *Memory and Cognition* 48 , pp. 200-211. 10.3758/s13421-019-01002-7

Publishers page: <http://dx.doi.org/10.3758/s13421-019-01002-7>

Please note:

Changes made as a result of publishing processes such as copy-editing, formatting and page numbers may not be reflected in this version. For the definitive version of this publication, please refer to the published source. You are advised to consult the publisher's version if you wish to cite this paper.

This version is being made available in accordance with publisher policies. See <http://orca.cf.ac.uk/policies.html> for usage policies. Copyright and moral rights for publications made available in ORCA are retained by the copyright holders.



Running head: Causing Time

Causing Time:

Evaluating causal changes to the *when* rather than the *whether* of
an outcome

W. James Greville^{1*}, Marc J. Buehner², Mark K. Johansen²

¹School of Psychology and Therapeutic Studies, University of South Wales, CF37 1DL.

²School of Psychology, Cardiff University, CF10 3AT.

*Author to whom correspondence should be addressed. Telephone: +44 1443 654 828. E-mail:

james.greville@southwales.ac.uk

Declarations of interest: none

Abstract

Studies of human causal learning typically conceptualize an effect as the presence or absence of an outcome or event in a given trial following a cause. However, causes may exert their influence in other ways; notably by advancing or postponing the time at which an outcome occurs. Prior research has not examined how humans evaluate causal changes where the change in timing itself is the effect of interest. This research took a first step in this direction by investigating whether participants can accurately judge cause-effect contingencies when the effect is a change in outcome timing, as distinct from outcome occurrence: A change to the *when* of the outcome rather than to the *whether*. Three experiments presented scenarios where a candidate cause could either advance or postpone an inevitable outcome by a given amount of time and with a given probability. Consistent with previous research on judgments about event occurrence, participants gave higher ratings to scenarios with greater contingency. These effects were generally consistent for actions that advanced or postponed the outcome. Overall, our findings demonstrate that people are sensitive to probabilistic contrasts involving causal changes in event timing.

Keywords: Causality, Temporal Contiguity, Contingency Judgment, Causal Learning

Causing Time: Evaluating causal changes to the *when* rather than the *whether* of an outcome

Introduction

Causal learning is the foundation of many cognitive abilities and enables us to predict and control our environment. More specifically, causal knowledge typically affords a reduction of uncertainty about events in our environment, even though it rarely eliminates it altogether: Not all smokers will suffer from cancer, but smoking increases the likelihood that cancer will develop. Consequently, most research has examined people's sensitivity to probability contrasts or contingencies, and evaluated how they are interpreted to form causal impressions (see Cheng & Buehner, 2012, and Shanks, Holyoak & Medin, 1996, for overviews). As a result, the literature is now rich with data on human causal judgment from statistical contingency information and with theories and models aimed at describing and explaining the cognitive processes underlying such judgments.

Conceptually, cause-effect contingencies are defined as the presence versus absence of a target event (usually referred to as the *effect* or *outcome*), given the presence versus absence of a candidate *cause* (Jenkins & Ward, 1965). Causal influence implies that the presence of the cause increases the likelihood that the effect occurs (as in 'smoking causes cancer'). Contingencies are specified with respect to a fixed temporal window, a trial, over which the cause and/or the effect are evaluated to have co-occurred or not co-occurred. In causal judgment experiments, this is implemented via a sequence of learning trials. In real life, such as in the smoking-cancer example, a specific time-frame is often implied, as in smoking will increase the likelihood of cancer over the course of an individual's lifetime.

Often though, causality changes not only *whether* something happens, but *when* it happens: Chemotherapy is linked to early menopause, and healthy lifestyle to later death. In both cases the cause does not change the likelihood of the event itself (both menopause and death will inevitably occur at some point), but instead changes aspects of its timing, by accelerating or delaying its occurrence. Indeed, one might argue that even in the smoking-cancer example, we might care more about *when* smoking might bring about cancer, rather than *whether*: Presumably, abstaining from smoking would be far less attractive if smoking-induced cancer were guaranteed to only set in after the age of 95. The real issue may be not so much about whether one gets cancer at all, but when one might get it, and whether smoking increases the chance of early cancer. In such situations, the emphasis is not on the occurrence of an outcome per se, but on its occurrence at or by a given time. In other words, the variable of interest is *when* the event occurs, and causality is based on whether the cause changes the *time* of its occurrence. Addressing these situations is the central aim of this research.

Prior research

Traditional approaches to causality have focused on probabilistic changes of event occurrence, but have not considered probabilistic changes to event time. Although there is a large body of literature on temporal contiguity and its role in causal induction (see Buehner, 2005), this work has been confined to scenarios where causes operate to change whether or not events occur; the focus on time in this literature is to determine to what extent delays between the cause and effect are detrimental to the discovery of causal regularities (Buehner & May, 2003; Greville, Cassar, Johansen, & Buehner, 2013; Shanks, Pearson, & Dickinson, 1989; Wasserman, Chatlosh, & Neunaber, 1983). Typically, all other things being equal, temporal contiguity of

cause and effect facilitates detection of a statistical relationship – that is, co-occurrence – between the two. More recent research has further demonstrated that temporal regularities – that is, when delays are either expected or predictable – further bolsters impressions of causality (Buehner & May, 2002, 2004; Greville & Buehner, 2010, 2016; Haering & Kiesel, 2012). However, these studies do not speak to situations where causes bring about changes in event timing rather than event occurrence. Although time is clearly closely tied to the evaluation of causation, the role of time in causation is still very unclear, particularly when time itself is actually part of the effect of interest.

A few studies have examined situations where a cause produces a change in event timing rather than affecting event occurrence. Wasserman and Neunaber (1986) for instance gave participants a task in which they scored points whenever a light illuminated, and participants could press a response key to influence the light. Unbeknownst to participants, the illumination of the light was inevitable (i.e. the outcome occurred on every trial), but pressing the response key could either advance or postpone its delivery. In one condition, the outcome by default occurred at the end of interval t (where t was either 3, 6 or 9s) but pressing the key would immediately present the outcome (and cancel its delivery at the end of t). Participants pressed the key more frequently during, and gave higher causal ratings to, this condition compared to other conditions where the timing of the outcome was either response-independent or where a response cancelled the outcome that would otherwise have occurred after t (thus effectively postponing the arrival of the next outcome). Participants thus demonstrated a degree of sensitivity to changes in event timing. Crucially though, participants were not actually asked if pressing the key affected the timing of the light, but rather if their responses affected the occurrence of the light; that is, they were instructed to focus on *whether* the light illuminated, and not *when*. Event

timing was thus seen only as an adjunct in the evaluation of event co-occurrence rather than itself being the feature of interest.

In other studies, Greville and Buehner (2007) found that the mere advancing (or postponing) of events in time was sufficient to elicit causal (or preventive) judgments from participants in situations where contingency was zero, thus suggesting that timing affects causal perception even when there is no statistical contingency. These experiments were however based on summary data presented in tabular format rather than real-time causal induction. Lagnado and Speekenbrink (2010) devised a real-time judgment problem to study the influence of timing on participants' judgment of whether a composite candidate cause (a bacterium in which three features – feelers, spots, and a tail – could each be present or absent) produced an outcome (stomach cramps). The presence of one of these features served as a 'hastener' which speeded up the arrival of the outcome (without affecting whether or not it occurred) – similar to the action of pressing the response key in Wasserman and Neunaber's (1986) study. However, in contrast to Wasserman and Neunaber's findings, no significant influence of adding a hastener on participants' causal judgments was found. Regardless, it is again important to note that both Greville and Buehner's (2007) and Lagnado and Speekenbrink's (2010) studies focused on how changes in event timing influenced perception of cause-outcome contingency, rather than on the detection and evaluation of the change in timing itself. In short, in all previous lines of inquiry, event timing has been the explanans rather than the explanandum. To our knowledge, no study thus far has defined an effect as being a change in outcome timing rather than a change to the probability of outcome occurrence and investigated the factors affecting causal learning in this case.

Conceptual Approach

The difficulty in examining how people make causal inferences regarding time is that time is a continuous dimension, while models of causality typically assume causes and effects to be binary; either present or absent (Allan, 1980; Cheng, 1997; Griffiths & Tenenbaum, 2005). From this assumption, a 2×2 table (Table 1) is used to represent co-occurrence patterns of cause and effect, and the learner is assumed to utilize information about the frequencies of each type of outcome (cells A, B, C, and D) to arrive at a causal conclusion. The influential ΔP model (Jenkins & Ward, 1965) for instance states that $\Delta P = P(e|c) - P(e|\neg c)$, where $P(e|c)$ is the likelihood of the effect in the presence of the cause, which can be calculated from the cells in the contingency table via $A/(A+B)$; $P(e|\neg c)$ is the likelihood of the effect in the absence of the cause and refers to $C/(C+D)$. If ΔP is positive, the cause is assumed to produce the effect, if it is negative, it prevents it, and if ΔP is zero, the cause is seen as not influencing the effect.

| Cause | Effect | |
|-----------------|-----------------|----------------------|
| | Present e | Absent $\neg e$ |
| Present c | A $e c$ | B $\neg e c$ |
| Absent $\neg c$ | C $e \neg c$ | D $\neg e \neg c$ |

Table 1: The cells of the 2x2 contingency matrix

However, time is by no means special in terms of concerning a continuous stimulus dimension. Many other relevant causal relations likewise concern changes on a continuum: Stock prices rise and fall to various extents in response to other economic trends or interventions, plants grow by small or large amounts (or not at all) following treatment with fertilizer, and allergic reactions in response to topical creams can range from mild to severe. Nevertheless, all these scenarios are regularly used in studies of probabilistic causation. Although in the natural world outside of the laboratory the stimulus dimension of interest is continuous, researchers have routinely simplified the task into one that considers a binary change on a continuous stimulus dimension, so that the causal intervention can simply lead to a rise or fall in stock prices, small or large plant growth, or allergic reactions (e.g. Ab Rashid & Buehner, 2013). We can likewise reduce causally-induced changes in event timing to binary outcomes – change or no change – which can be conceptualized as the presence or absence of the effect. So in its simplest case, time too can be described in a binary fashion, just as the presence or absence of the effect is a binary variable in conventional studies of contingency based causation, and as such can be represented in the same 2×2 contingency matrix. In essence, we can turn the question of *when* back into the question of *whether*, but instead basing this question on event *timing* rather than event *occurrence*; whether a change to the time of the outcome occurs rather than whether the outcome itself occurs.

Changes in event timing can and of course do occur probabilistically (taking aspirin will increase the likelihood of speedy relief from headache). Consequently, an effect e can be defined such that $P(e|c)$ refers to the probability with which an outcome happens at a target time (rather than simply whether it happens at all), given that a causal intervention has taken place (i.e. the headache goes away quickly after taking the aspirin). Similarly $P(e|-c)$ refers to the probability

that the event outcome happens at the target time even though the cause was absent (the headache quickly resolves without medication). Given that previous research has shown that people readily interpret probabilistic changes from one value on a continuous dimension to another (i.e. high/low price, small/large growth, mild/severe reaction, etc) largely in line with the deployed contingencies, we wanted to examine whether people would likewise evaluate probabilistic changes to the time of outcome occurrence in line with underlying contingencies.

The work presented here thus redefined what constitutes an effect: Rather than conceptualizing an effect as the occurrence of a specific outcome – such as a light flashing – we defined an effect as a change to the default time of outcome occurrence (i.e. whether performing a causal action makes the light flash earlier or later than it otherwise would have done). As in conventional studies on causal learning, we applied different contingencies with which performing a causal action altered the likelihood of the effect. Our interest was in whether participants are sensitive to changes in cause-effect contingency in situations when the effect is a change in outcome timing. From a computational perspective, our experiments thus asked whether human reasoners approach cause-effect contingencies pertaining to the *time of outcome occurrence* in the same way they approach contingency relations between cause and *probability of outcome occurrence*. In other words, are human reasoners in principle able to learn about causally induced changes to outcome timing in similar ways as they can learn about causally induced changes to outcome probability? Our hypothesis was that they are; however, it is by no means obvious that this will be the case. Because temporal contiguity is so intricately linked to causal inference (as outlined earlier), it is possible that causal actions that change the *time of outcome occurrence* automatically trigger causal inferences about the *probability of outcome occurrence*. Specifically, given that temporal contiguity is such a powerful cue to causality, it is

entirely possible that a cause that advances the time of outcome occurrence is perceived as a cause that generates the outcome, while a cause that postpones the time of outcome occurrence is deemed as one that is preventing the outcome – indeed this is exactly what Wasserman and Neunaber's (1986) results might be seen to suggest. Put differently again, our question can be rephrased as asking whether the continuous dimension of time of outcome occurrence (early/late) is equally amenable to probabilistic causal inference as other, continuous stimulus dimensions used in previous research such as size (small/big), strength (weak/strong), or value (low/high), if participants are given the opportunity to respond to it (as opposed to being simply asked about outcome occurrence).

It is probably useful at this point to clarify terminology: Most extant research on causal learning uses the terms *outcome* and *effect* interchangeably. This reflects the fact that in this research whether or not the outcome occurs is the matter of interest. In our case, however, the time of outcome occurrence is the matter of interest. Consequently, in our case, *effect* and *outcome* refer to two separate things: The effect is a change to the time of the (occurrence of the) outcome.

Methodological Approach

Our experimental paradigm comprised several distinct conditions, each consisting of multiple learning trials during which clicking a button could either advance or postpone a shape lighting up. This procedure is superficially similar to Wasserman and Neunaber's (1986) protocol, but with a number of important differences. Most obviously, Wasserman and Neunaber's participants were not asked to assess whether their responses produced changes in outcome timing, but whether the response produced the outcome. Our paradigm instead explicitly informed participants that the shape would always light up and instructed them to

focus on whether clicking the button changed the *time* at which the shape lit up. Therefore in our paradigm, the shape always lit up on every trial, irrespective of whether the participant clicked the button. Thus, the contingency between button press and the shape lighting up was always 0 – $P(e|c) = P(e|\neg c) = 1$; $\Delta P \rightarrow 0$. Pressing the button instead resulted in postponing or advancing the time of the shape lighting up according to a programmed contingency, similarly to Wasserman and Neunaber. Unlike in Wasserman and Neunaber, however, our dependent variable explicitly tracked participants' sensitivity to this change in outcome time.

Furthermore, in Wasserman and Neunaber's (1986) protocol only the first response in a given period produced the time change (and the first response always produced it), while any subsequent responses had no effect, rendering them irrelevant. Probabilistic relations between response and outcome timing thus were not directly manipulated so it was impossible to determine the influence of this on participants' judgments. In addition, if a response was made, the outcome was then immediately contiguous with the response (or was cancelled entirely until the next interval, effectively postponing the outcome). Thus the extent of advancement (or postponement) was likewise not systematically manipulated. In contrast, in our paradigm, participants were given a fixed window of opportunity in which to perform a single response in a given learning trial, after which no further responses were permitted. The outcome then followed, regardless of whether a response was made or not; however, if a response was performed, this had the effect of either advancing or postponing the outcome, with a given probability (see Figure 1 for an illustration). The probability and extent of any changes in outcome timing were systematically varied across multiple experimental conditions. Our paradigm therefore allowed our participants to specifically assess the effect of their responses – the cause – on the timing of the outcome (the effect of interest in this study), and so allowed us

to examine whether participants were accurate in judging contingency for causally-induced changes in event timing.

Experiment 1

We deployed a variant of a simple computer-based experimental task that has become an established paradigm for studies of contingency judgment since the work of Shanks, Pearson and Dickinson (1989). On the computer screen participants were shown the outline of a geometric shape and below this a button which they could press by clicking on it with the mouse. In traditional experiments, the crucial relationship that participants are required to judge is whether pressing the button causes the shape to light up. Here we instead asked participants to judge whether pressing the button changed the *time* at which the shape lit up. Our experiment consisted of several different conditions, each of which was in turn composed of a series of discrete trials. In each trial the shape always lit up exactly once (regardless of whether the button was pressed). If the button was pressed, this could advance or postpone the time at which the shape lit up. Participants were asked to judge the likelihood with which the time change occurred following a button press. Specifically, we asked participants to focus on *whether or not* the time change occurred, and not to take into account the amount of time by which the lighting up of the shape was advanced or postponed. Thus, our paradigm allowed us to discern whether participants' judgments displayed sensitivity to *contingencies of changes in outcome timing* in the same way that has been observed in studies examining changes to outcome occurrence.

Method

Participants:

Thirty-eight participants (32 females, 6 males, median age 18 years) from Cardiff University participated to receive either partial course credit or £4 payment. Prior research was used as a guide to sample size.

Design:

We manipulated three independent factors, each with two levels; *direction* of time change (advance vs postpone), *extent* of time change (2s vs 4s advancement or postponement) and *probability* of the time change occurring given the cause; $P(e|c)$ (0.6 vs 0.8). These factors combined in a $2 \times 2 \times 2$ fully within-subjects design to create eight conditions. The dependent variable was a causal rating (i.e. contingency judgment) provided by participants at the end of each condition. This was operationalized via a two-stage likelihood assessment of whether pressing the button produced a change in the timing of the outcome's occurrence (see below for details). In addition, as a consequence of engaging with the task, a number of secondary dependent measures were produced, such as participants' response rates; these are collectively termed "behavioral data" and are analysed along with the causal ratings in the results section.

Apparatus, Materials & Procedure:

Participants were tested in groups in a quiet 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.

Participants read on-screen instructions outlining the nature of the task (see Appendix), and then began the experiment. Each condition was composed of 20 trials, each trial lasting for

10 seconds. At the start of each condition, an outline of a shape was presented in the center of the screen, and below it was a red button which could be clicked with the mouse; following a mouse click, the button appeared to depress and became brighter for 250ms indicating that it had been pressed. Participants were able to press the button only when it was “active,” indicated by being colored in red (when the button was inactive, it became duller in appearance and no longer responded to mouse clicks). The button was active at the beginning of each trial for a period of 3s, before becoming inactive, thus providing a defined window of opportunity for participants to make a response (or not). The button also became inactive following a press and therefore could be pressed only once during a given trial. During each trial, the shape always lit up exactly once, for 500ms, at some point after the closure of the response window.

The crucial aspect then was the timing of the shape lighting up, and whether and how this was affected by pressing the button. The experiment was programmed such that within a given condition, the shape either lit up later on, towards the end of the trial, and pressing the button advanced the time of outcome; or the shape lit up earlier on in the trial, and pressing the button postponed the time of outcome. Table 2 summarizes the conditions in terms of the independent variables. Importantly, pressing the button resulted in probabilistic rather than deterministic changes to outcome timing. Specifically, if the button was not pressed then, in “advance” conditions, the shape lit up late with a probability of 0.8 and lit up early with a probability of 0.2, while in “postpone” conditions the shape lit up early with a probability of 0.8 and light up late with a probability of 0.2. In other words, there was a base rate of 0.2 of a change in the usual timing of the shape lighting up if no action was taken. The effect of pressing the button then was to alter these probabilities and increase the likelihood that the illumination of shape would be advanced/postponed following a button press, with a probability of either 0.8 or 0.6 (with

concomitant probability of 0.2 or 0.4 that there would be no time change), depending on the level of the independent variable operating in that condition. Thus, in an “advance 0.6” condition the shape would light up early with a probability of 0.6 if the button was pressed and with probability of 0.2 (i.e. the base rate) if the button was not pressed. The probability of a time change following a button press was constant within a given condition (but obviously varied across conditions).

The exact amount of time by which the lighting up of the shape was advanced or postponed could also change from one condition to the next, changing by either 2s or 4s. The shape could light up either 5s, 7s or 9s seconds after the start of the trial (2s, 4s, 6s after the closure of the 3s response window). In postpone conditions the default time was 5s and pressing the button could postpone the outcome (or rather, strictly speaking, increase the probability of the postponement) to either 7s (2s postponement) or 9s (4s postponement). Meanwhile in advance conditions the default time was 9s and pressing the button could advance the outcome to either 7s (2s advancement) or 5s (4s advancement). In other words, in the advance conditions, the default time of the outcome was the maximum experienced time between the trial start and the outcome, whereas in postpone conditions the default time was the minimum experienced time. Figure 1 provides a graphical representation of the trial structure. The combinations and timings of all conditions are detailed in Table 2.




Figure 1: Graphical representation of the trial structure and trial events. The example shown here is for the “advance, 4s, $p = 0.8$ ” condition where pressing the button increased the likelihood of the triangle flashing at the earlier time of 5s into the trial rather than the default of 9s into the trial. Panel A depicts a trial where no response is made and the probability of the “late” effect is 0.8 while the probability of an “early” effect is 0.2; Panel B depicts a trial where a response is made and the probability of the “late” effect is 0.2 while the probability of an “early” effect is 0.8.

Participants were informed that the shape would flash exactly once during each 10s trial, and that pressing the button would not affect whether or not the shape lit up; in other words, pressing the button was not the cause of *whether* or not the shape lit up. Instead, pressing the button could alter the *time* at which the shape lit up during the trial. Pressing the button could advance the time, postpone the time, or have no effect on when the shape lit up (relative to when the shape lit up if the button was not pressed).

Participants were advised that they should consider the overall effect across the set of 20 trials within each condition, but to evaluate the effect of the button separately for each condition. Participants were also advised that on some trials they should press the button, and on some trial not to press the button, in order to adequately assess how pressing the button altered the behavior of the shape, by comparing its behavior following their intervention compared to the baseline behavior. Any participants who failed to adhere to these requirements (either by always pressing or never pressing in a given condition) were removed from subsequent analyses.

At the end of each condition, the screen cleared and participants were asked to make a contingency judgment: to what extent pressing the button changed the *likelihood* that the shape lit up either early or late (participants were specifically instructed to ignore the *amount* of time). Participants first indicated whether pressing the button advanced or postponed the shape lighting up, and then provided a rating from 0-100 of the likelihood (if advancement was selected, the rating was recorded as a positive number and if postponement was selected the rating was recorded as a negative number; in either case, a rating of zero could be given to indicate no change). After confirming their rating participants progressed to the next condition. In total the experiment lasted around 25 minutes.

Results & Discussion

One participant failed to adhere to the experimental instructions (responding on all 20 trials for at least one condition). Their data was thus discarded from the analysis.

Manipulation check and behavioral data

An intrinsic difficulty with the free operant paradigm as utilized in the present experiment is that contingencies actually experienced can differ from the nominal programmed

contingencies (see Shanks & Dickinson, 1991; Wasserman & Neunaber, 1986). It was therefore necessary to check that our manipulation of probability was effective by examining the contingencies that were experienced by participants and confirming that they indeed differed significantly across conditions with different levels of probability (and further, that there were no unplanned systematic variations in probability experienced across conditions with different levels of direction or extent creating a confound). Participant response rate must also be examined as any systematic variations in the number of responses made across conditions could mean that the putative influence of the independent variables on causal ratings is mediated through levels of responding (see, e.g., Buehner & May 2003; Shanks & Dickinson, 1991, Vallée-Tourangeau, Murphy & Baker, 2005).

Table 3 shows means for actual $P(e|c)$, experienced ΔP and number of responses across all participants in each of the eight experimental conditions. Across all conditions, experienced probabilities and experienced ΔP were, on average, close approximations of their conceptual values. A $2 \times 2 \times 2$ (*direction* \times *probability* \times *extent*) ANOVA carried out on experienced $P(e|c)$ found a significant main effect of probability, $F(1,36) = 216.266$, $p < 0.0005$, $MSE = 0.017$, $\eta_p^2 = .857$, confirming that our manipulation was effective. Meanwhile no significant effects of either direction or extent were found (both $ps > 0.5$) Mean response rates by condition, Table 3 far right column, show that participants clicked the button on roughly half the trials, consistent with the instructions to explore both responding and nonresponding. There was a marginally significant effect of extent on number of responses emitted, $F(1,36) = 3.438$, $p = 0.072$, $MSE = 4.954$, $\eta_p^2 = .087$; specifically it appears conditions with the smaller time change of 2s elicited slightly higher levels of responding than those with the larger time change of 4s. There was no significant effect of either probability or direction on number of responses and no significant interactions (all $ps >$

0.2). Number of responses was however not correlated with causal ratings, $r = .001$. As such, number of responses was unlikely to be a mediating variable for the influence of the independent variables on causal ratings.

| Direction of time change | Probability of time change | Probability of "early" effect following button press | Probability of "late" effect following button press | Probability of "early" effect following no button press | Probability of "late" effect following no button press | Extent of time change | Time of "early" effect after start of trial | Time of "late" effect after start of trial |
|---|---|---|--|--|---|--------------------------------------|---|--|
| advance | 0.8 | 0.8 | 0.2 | 0.2 | 0.8 | 2s | 7s | 9s |
| | 0.6 | 0.6 | 0.4 | | | 4s | 5s | 9s |
| | 0.8 | 0.8 | 0.2 | | | | | |
| | 0.6 | 0.6 | 0.4 | | | | | |
| postpone | 0.8 | 0.2 | 0.8 | 0.8 | 0.2 | 2s | 5s | 7s |
| | 0.6 | 0.4 | 0.6 | | | 4s | 5s | 9s |
| | 0.8 | 0.2 | 0.8 | | | | | |
| | 0.6 | 0.4 | 0.6 | | | | | |

Table 2: The combination of the three independent variables to create eight different conditions, and the resulting probabilities and timings of the outcome within each condition. The key variables actively manipulated are emphasized in bold, while the remaining are corollaries of these.

| Condition | Conceptual $P(e c)$ | Experienced $P(e c)$ | Conceptual ΔP | Experienced ΔP | Mean responses (/20) |
|-------------------|------------------------|-------------------------|--------------------------|---------------------------|----------------------------|
| Advance, 0.8, 2s | 0.8 | 0.813 | 0.6 | .591 | 11.41 |
| Advance, 0.8, 4s | 0.8 | 0.816 | 0.6 | .623 | 11.30 |
| Advance, 0.6, 2s | 0.6 | 0.602 | 0.4 | .419 | 12.05 |
| Advance, 0.6, 4s | 0.6 | 0.587 | 0.4 | .406 | 11.03 |
| Postpone, 0.8, 2s | 0.8 | 0.807 | 0.6 | .600 | 11.24 |
| Postpone, 0.8, 4s | 0.8 | 0.827 | 0.6 | .638 | 10.95 |
| Postpone, 0.6, 2s | 0.6 | 0.606 | 0.4 | .417 | 11.35 |
| Postpone, 0.6, 4s | 0.6 | 0.565 | 0.4 | .339 | 10.86 |

Table 3: Conceptual and mean experienced $P(e|c)$, ΔP , and number of responses across all participants separately for each condition in Experiment 1.

Causal ratings

The dependent measure of primary interest was the causal ratings provided at the end of each condition. Firstly, participants correctly identified whether the outcome was being advanced or postponed, with consistently positive ratings given to advance conditions ($M = 41.00$) and negative ratings to postpone conditions ($M = -38.17$). In order to assess whether advancement and postponement differed in any meaningful way beyond this, and to simplify comparisons between experimental conditions, we subsequently recoded ratings from the postponement

conditions (i.e. converted negative ratings into positive, and positive into negative), effectively considering absolute rather than signed causal ratings. The following analyses are based on these recoded ratings.




Figure 2 – Mean contingency ratings across all participants in Experiment 1 as a function of change in outcome time. The left panel shows ratings for conditions where outcome time was advanced while the right panel shows ratings for conditions where outcome time was postponed. Solid lines and symbols represent probability of 0.8 of change in outcome timing while unfilled symbols and dashed lines represent probability of 0.6 of change in outcome timing. Error bars show standard error.

Figure 2 shows ratings broken down by probability and extent of time change, separately for advancement and postponement. A $2 \times 2 \times 2$ (*direction* \times *probability* \times *extent*) ANOVA calculated on absolute scores confirmed that the main effect of *direction* was not significant,

$F(1,36) = 0.142, p = .709, MSE = 3315.155, \eta_p^2 = .004$. This indicates that while participants clearly distinguished between advancement and postponement (as shown by the signed ratings), they did not respond in qualitatively different ways to these temporal changes (as shown by the absolute ratings). Of central interest however was participants' sensitivity to the *probability* of temporal change: Figure 2 clearly shows that participants gave higher ratings when $P(e|c)$ was 0.8 than when it was 0.6, $F(1,36) = 16.612, p < .0005, MSE = 2021.772, \eta_p^2 = .316$. The *extent* of temporal change also affected causal ratings, with larger changes eliciting higher ratings, $F(1,36) = 5.385, p = .026, MSE = 1932.393, \eta_p^2 = .130$. These patterns were generally consistent for both postponing and advancing the time of the effect. While there appears to be a more prominent difference between small and large temporal changes when there is also a high probability of advancement, none of the possible interactions were statistically significant (all $ps > 0.2$).

Experiment 1 thus demonstrated that participants were able to detect when their actions changed the temporal position of an outcome. Specifically, the causal efficacy of a participant's action lay in changing the time-point of an outcome, rather than influencing whether or not the outcome occurred. These time-changes were implemented probabilistically, and our data show that participants were clearly sensitive to the two probabilities we implemented. While participants clearly (and unsurprisingly) could distinguish between actions that advanced and postponed the outcome, our results also showed that beyond the sensitivity to the direction of temporal change, probabilistic changes to event time were evaluated similarly for both directions. That is, participants were equally sensitive to the likelihood with which their action advanced and postponed an outcome. In addition to this, participants also showed sensitivity to the extent with which an action caused a temporal change.

It is noteworthy that the sensitivity to the extent of temporal changes is not explained as a conflation of causal judgment with experienced temporal contiguity. Specifically, consider the research reviewed in the introduction, which showed that degrading cause-effect contiguity typically reduces judgments of causality (e.g. Shanks et al., 1989) – all else being equal, contiguous cause-effect relations are deemed stronger than delayed relations (see also Buehner, 2005). Applying this contiguity advantage to Experiment 1 could explain the sensitivity to the extent of temporal change in the advancement conditions: Greater advancement effected stronger action-outcome contiguity, and thus could lead to higher causal judgments. However, in the postpone conditions, a greater extent of postponing effected reduced action-outcome contiguity, yet was associated with stronger causal judgments. Thus, variations in the extent of temporal change impacted causal judgments in the same way both for actions that advanced as well as postponed the outcome. In other words, the sensitivity to extent does not reflect a conflation of experienced action-outcome contiguity with causal strength but instead seems to reflect greater confidence in or higher salience of larger temporal change.

Experiment 2A and 2B

Experiment 1 demonstrated that participants were sensitive to probabilistic changes in event timing, readily distinguishing between conditions where the probability of a change in outcome timing following a response was higher or lower. Experiment 1 also showed that causal judgments for actions that bring about a temporal change of an outcome are influenced by the extent of that temporal change, with larger changes eliciting higher judgments. Sensitivity to the salience of a causal change, or confidence in the observation of the change are also known to impact probabilistic causal judgment (for confidence see Griffiths & Tenenbaum, 2005, 2009;

for salience see Beckers, De Houwer, Pineño, and Miller, 2005). The following experiments aimed to replicate the main effect of interest – sensitivity to probabilistic changes in outcome timing – in light of the secondary effect of salience/confidence. To this end, we manipulated the extent of temporal change across a greater number of levels to assure ourselves that the main effect of contingency is robust across multiple levels of temporal extent. To streamline the experiment (and simplify interpretation of findings), we considered advancement and postponement separately across two experiments. This also allowed us to examine the robustness of the original finding when temporal change is restricted to only one direction within a given participant. Experiment 2A is concerned with causes that advance the time of an outcome, while Experiment 2B dealt with postponement. Each experiment deployed the same two levels of causal contingency as in Experiment 1, but featured three levels of temporal change: 2s, 4s, and 6s.

Method

Participants

44 participants (34 female, 10 males, median age 20 years) completed experiment 2A and 35 participants (21 female, 14 males, median age 20 years) completed experiment 2B.

Participants received either £4 payment, course credit, or the gratitude of the experimenters.

Design

Experiment 2A focused only on advancing the time of outcome, Experiment 2B on postponing it; direction of time change was thus no longer a within-subjects factor. Two factors, probability and extent of time change, remained. The levels of probability were 0.8 and 0.6 as for the previous experiment. An extra level was added to extent, such that the outcome could be

advanced (Experiment 2A) or postponed (Experiment 2B) by either 2s, 4s or 6s in time. The factors thus combined to create six conditions in two 2×3 fully within-subjects designs.

Apparatus, Materials & Procedure

The basic experimental setup was largely identical to Experiment 1, with only subtle changes in the instructions to participants to reflect the fewer conditions and that only advancing (Experiment 2A) or postponing (Experiment 2B) the time of outcome was possible. There was thus only one step involved in obtaining the dependent measure; participants were simply asked to provide a causal rating rather than first being asked whether pressing the button advanced or postponed the outcome. The study location and equipment were the same as for Experiment 1.

The ‘default’ outcome time was 9s into the trial for Experiment 2A and 3s into the trial for Experiment 2B. In Experiment 2A pressing the button probabilistically advanced the time of outcome by either 2, 4 or 6s (depending on the condition), to occur respectively 7, 5 or 3s into the 10s trial. In Experiment 2B pressing the button probabilistically postponed the time of the outcome by 2, 4, or 6s to occur respectively 5, 7, or 9s into the 10s trial. The response window, during which the button could be pressed, was shortened to 2s to accommodate these times. This prevented the outcome from coinciding with either the inactivation of the button or the end of the trial, which might have affected the detection of the outcome.

Results and Discussion

Two participants in Experiment 2A made no responses in one or more conditions and their data was thus discarded from the analysis.

Manipulation check and behavioral data

Mean experienced $P(e|c)$, ΔP and number of responses for both Experiment 2A and 2B are shown in Table 4. Experienced $P(e|c)$ once again varied significantly with probability both for Experiment 2A, $F(1,42) = 124.850$, $p < .0005$, $MSE = .014$, $\eta_p^2 = .748$, and Experiment 2B, $F(1,34) = 101.512$, $p < .0005$, $MSE = .022$, $\eta_p^2 = .749$, thus confirming the manipulation of probability was effective. There was no significant main effect of extent and no significant interaction in either experiment for experienced $P(e|c)$ (all $ps > 0.1$). Again, participants clicked the response button on slightly more than half the trials as per the response exploration instructions.

For Experiment 2A there was a marginally significant effect on number of responses of both probability, $F(1,42) = 3.755$, $p = .059$, $MSE = 7.772$, $\eta_p^2 = .082$, and extent, $F(2,84) = 3.092$, $p = .051$, $MSE = 5.123$, $\eta_p^2 = .069$; specifically, conditions with higher programmed values of $P(e|c)$ elicited slightly lower response levels and conditions with a higher extent of temporal change elicited slightly higher response levels). There was also a marginally significant positive correlation between responses and ratings, $r = .116$. For Experiment 2B there were no significant effects of either probability or extent on number of responses, nor was number of responses correlated with causal ratings (all $ps > 0.1$). Overall, results across the three experiments showed no consistent effects of the independent variables on number of responses made and no consistent relationship between responses and ratings.

Causal ratings

Mean causal ratings obtained in Experiments 2A and 2B are shown in Figure 3. Replicating the main effect of interest, both Experiment 2A and 2B found significant main effects of *probability*; for advancing (Experiment 2A) $F(1,42) = 5.158$, $p = .028$, $MSE =$

559.614, $\eta_p^2 = .109$ and for postponing (Experiment 2B), $F(1,34) = 23.199$, $p < .0005$, $MSE = 402.314$, $\eta_p^2 = .406$. Thus, participants were again able to differentiate between actions that changed the time of an outcome with high versus low probability. Importantly, this sensitivity was robust even when temporal changes were in one direction only for a given participant (either postponing or advancing).

| Experiment | Condition | Conceptual P(e c) | Actual P(e c) | Conceptual ΔP | Actual ΔP | Mean responses (/20) |
|------------------|--------------------------|----------------------|------------------|--------------------------|----------------------|----------------------------|
| Advance (2A) | P(e c) = 0.8, Extent =2s | 0.8 | .800 | 0.6 | .575 | 11.79 |
| | P(e c) = 0.8, Extent =4s | 0.8 | .769 | 0.6 | .597 | 11.53 |
| | P(e c) = 0.8, Extent =6s | 0.8 | .807 | 0.6 | .590 | 12.30 |
| | P(e c) = 0.6, Extent =2s | 0.6 | .586 | 0.4 | .427 | 12.44 |
| | P(e c) = 0.6, Extent =4s | 0.6 | .597 | 0.4 | .398 | 12.07 |
| | P(e c) = 0.6, Extent =6s | 0.6 | .584 | 0.4 | .425 | 13.00 |
| Postpone (2B) | P(e c) = 0.8, Extent =2s | 0.8 | .796 | 0.6 | .628 | 11.26 |
| | P(e c) = 0.8, Extent =4s | 0.8 | .803 | 0.6 | .613 | 11.86 |
| | P(e c) = 0.8, Extent =6s | 0.8 | .795 | 0.6 | .602 | 11.46 |
| | P(e c) = 0.6, Extent =2s | 0.6 | .636 | 0.4 | .450 | 11.26 |
| | P(e c) = 0.6, Extent =4s | 0.6 | .599 | 0.4 | .412 | 11.86 |
| | P(e c) = 0.6, Extent =6s | 0.6 | .548 | 0.4 | .378 | 11.20 |

Table 4: Mean actual $P(e|c)$ experienced, ΔP , and number of responses across all participants separately for each condition in Experiments 2A and 2B.

As before, the effect of *extent* of time change was also significant; for advancing (Experiment 2A), $F(2,84) = 26.165$, $p < .0005$, $MSE = 545.847$, $\eta_p^2 = .384$ and for postponing (Experiment 2B), $F(2,68) = 3.994$, $p = .023$, $MSE = 724.630$, $\eta_p^2 = .105$. The interaction between *probability* and *extent* was not significant in either case, $F(2,84) = 1.448$, $p = .241$, $MSE = 592.398$, $\eta_p^2 = .030$ and $F(2,68) = 0.932$, $p = .399$, $MSE = 737.024$, $\eta_p^2 = .027$ for advancing and postponing respectively.

The main effect of *probability* of temporal change was also found in a cross-experimental analysis combining both datasets from Experiments 2A and 2B, $F(1,76) = 23.718$, $p < .0005$, $MSE = 489.243$, $\eta_p^2 = .238$. Likewise, the main effect of *extent* of temporal change was also obtained, $F(2,152) = 22.7958$, $p < .0005$, $MSE = 625.829$, $\eta_p^2 = .231$. There was no significant effect of the between-subjects factor of *direction (experiment)*, $F(1,76) = 2.549$, $p = .115$, $MSE = 1964.977$, $\eta_p^2 = .032$. Comparing the two data sets in this way did however find a marginally significant *direction* \times *extent* interaction, $F(2,152) = 2.784$, $p = .065$, $MSE = 625.829$, $\eta_p^2 = .035$, suggesting that participants are slightly more sensitive to the extent of the temporal change when the intervention advances an outcome than when it postpones it. Finally, Bonferroni-corrected pairwise comparisons further exploring the main effect of extent across both studies found significant differences between temporal changes of both 2s and 6s ($MD = 18.821$, $p < .0005$), and 2s and 4s ($MD = 12.8143$, $p < .0005$), but not between 4s and 6s ($MD = 6.007$, $p = .130$), suggesting a differential sensitivity to the extent of temporal change consistent with Weber's law.




Figure 3 – Mean contingency ratings across all participants in Experiments 2A and 2B as a function of change in outcome advancement (A)/postponement (B) in time. Solid lines and symbols represent probability of 0.8 of outcome advancement/postponement while unfilled symbols and dashed lines represent probability of 0.6. Error bars show standard error.

General Discussion

Where causal relations unfold over time, the ultimate occurrence of an event may be of far less importance than its timing. Many fields often adopt a time-based approach to assessing causality; in geriatric medicine for instance, anticholinergic drugs aim to postpone loss of function in patients with dementia despite being unable to prevent it altogether. Here, we evaluated whether humans make causal judgments in the same way when the effect of interest is a change in outcome timing as when it is the occurrence (or nonoccurrence) of an outcome. Results across three experiments indicate that this is indeed the case and that people are able to

track statistical changes in the occurrence or non-occurrence of a change in outcome timing much as they track changes in the occurrence or non-occurrence of an outcome itself. As such, our results are an existence proof that people can represent causal changes to event time just as they can represent changes to other (non-temporal) stimulus dimensions such as size or price. This is despite time, unlike these other dimensions, having a seemingly direct influence on causal cognition whereby temporal contiguity typically promotes causal inference. Our finding that time can be encoded as a stimulus dimension fits well with earlier observations that people learn to associate specific times with particular outcomes (e.g. Haering & Kiesel, 2012; Young, Rogers & Beckmann, 2005).

Our results also showed that participants' causal judgments were additionally influenced by the extent of temporal change. This aspect of our findings may potentially be due to perceptual salience and judgment confidence effects: More extreme changes from a default event time may be more noticeable, which in turn may result in greater confidence that a change has occurred. Salience and confidence effects have already been demonstrated to influence causal judgments of probabilistic event occurrence (see Beckers et al., 2005; Griffiths & Tenenbaum, 2007, 2009). Although the extent of temporal change unquestionably biases judgements of whether a temporal change has been brought about, nonetheless the key finding of sensitivity to probabilistic changes of event time is robust and is not explained by this bias.

In summary, our results conceptually replicate well-established patterns from conventional causal learning studies and transfer them to situations where the effect of interest concerns outcome timing rather than ultimate occurrence. Specifically we have demonstrated that participants show sensitivity to $P(e|c)$ (and by extension ΔP) when assessing the probabilistic relation between a candidate cause and a change in the timing of a specific

outcome. We have shown that sensitivity to temporal change is tempered by judgment biases associated with greater salience of or confidence in more extreme temporal changes. Further research might examine whether other well-established judgment biases from probabilistic causal induction such as sensitivity to outcome-density and base-rate (Vallée-Tourangeau, Murphy, & Baker, 2005; Vallée-Tourangeau, Murphy, Drew, & Baker, 1998; Wasserman, Elek, Chatlosh, & Baker, 1993) still occur when causal manipulations concern outcome timing. Indeed, recent work by Chow, Livesey, Colagiuri and Don (2018) has shown that outcome density biases do occur in causal inference problems dealing with continuous magnitudes (although Chow et al did not consider time as a stimulus dimension).

Previous research, both empirical and theoretical, has implicated contingency and temporal proximity as influences on causal induction. Specifically, stronger cause-outcome contingency leads to higher causal judgment (all things being equal), while reduced temporal proximity between cause and outcome degrades causal judgment (again all things being equal). However, little research has evaluated the combination of contingency and temporal proximity in terms of a cause delaying or advancing the occurrence of an outcome. The findings presented in this article suggest that whether reduced temporal proximity between cause and outcome is detrimental for causal judgment depends on whether the focus of the causal judgment is on the *when* versus the *whether* of the outcome.

Open Practices Statement

The raw data for all experiments is available at (<http://osf.io/7fbvh>). The experiments were not pre-registered.

References

- Ab Rashid, A. A., & Buehner, M. J. (2013). *Causal reasoning with continuous outcomes*. Paper presented at the 35th annual meeting of the Cognitive Science Society.
- 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.
- Beckers, T., De Houwer, J., Pineño, O., & Miller, R. R. (2005). Outcome additivity and outcome maximality influence cue competition in human causal learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 31(2), 238.
- Buehner, M. J. (2005). Contiguity and covariation in human causal inference. *Learning and Behavior*, 33(2), 230-238.
- 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. doi:10.1080/02724980244000675
- 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.
- Burns, P., & McCormack, T. (2009). Temporal information and children's and adults' causal inferences. *Thinking & Reasoning*, 15(2), 167-196.

- Cheng, P. W. (1997). From covariation to causation: A causal power theory. *Psychological Review*, 104(2), 367-405.
- Cheng, P. W., & Buehner, M. J. (2012). Causal Learning. In K. J. Holyoak & R. G. Morrison (Eds.), *The Oxford Handbook of Thinking and Reasoning* (pp. 210-233). Oxford, England: Oxford University Press.
- Chow, J., Don, H., Colagiuri, B., Livesey, E. (2018). Illusory causation and outcome density effects with a continuous and variable outcome. In T.T. Rogers, M. Rau, X. Zhu, & C. W. Kalish (Eds.), *Proceedings of the 40th Annual Conference of the Cognitive Science Society* (pp. 1491-1496). Austin, TX: Cognitive Science Society.
- Greville, W. J., & Buehner, M. J. (2007). The influence of temporal distributions on causal induction from tabular data. *Memory & Cognition*, 35(3), 444-453.
- Greville, W. J., & Buehner, M. J. (2010). Temporal predictability facilitates causal learning. *Journal of Experimental Psychology: General*, 139(4), 756–771.
- Greville, W. J., & Buehner, M. J. (2016). Temporal predictability enhances judgements of causality in elemental causal induction from both observation and intervention. *The Quarterly Journal of Experimental Psychology*, 69(4), 678-697.
doi:10.1080/17470218.2015.1041535
- 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.
- 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. (2009). Theory-based causal induction. *Psychological Review*, 116(4), 661.
- Haering, C., & Kiesel, A. (2012). Time in action contexts: Learning when an action effect occurs. *Psychological Research*, 76(3), 336-344.
- Jenkins, H., & Ward, W. (1965). Judgment of contingencies between responses and outcomes. *Psychological Monographs*, 7, 1-17.
- Lagnado, D. A., & Sloman, S. (2004). The advantage of timely intervention. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 30(4), 856.
- Lagnado, D. A., & Sloman, S. A. (2006). Time as a guide to cause. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 32(3), 451.
- 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*, 41B(2), 139-159.
- Vallée-Tourangeau, F., Murphy, R. A., & Baker, A. G. (2005). Contiguity and the outcome density bias in action–outcome contingency judgements. *Quarterly Journal of Experimental Psychology*, B(2), 177–192.
- Vallée-Tourangeau, F., Murphy, R. A., Drew, S., & Baker, A. G. (1998). Judging the importance of constant and variable candidate causes: A test of the power PC theory. *Quarterly Journal of Experimental Psychology: Human Experimental Psychology*, 1, 65-84.

- 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.
- Wasserman, E. A., & Neunaber, D. J. (1986). College students' responding to and rating of contingency relations: The role of temporal contiguity. *Journal of the Experimental Analysis of Behavior*, 46(1), 15-35.
- Young, M. E., Rogers, E. T., & Beckmann, J. S. (2005). Causal impressions: Predicting *when*, not just *whether*. *Memory & Cognition*, 33(2), 320-331.

Appendix

Instructions given to participants were as follows:

In this task you have to judge to what extent you can change what is happening on the computer screen in front of you. Specifically, there will be a series of 8 judgment problems, each lasting just over 3 minutes. During each judgment problem you will see the outline of a shape on the screen. Each judgment problem is divided into 20 periods of 10 seconds each. The shape will light up exactly once in each 10s period, and it may light up relatively early or late within each 10s period. The 10s periods are delineated by a brief blanking of the screen.

There is also a button beneath the shape, which you can press by clicking on it with the mouse. The button will be active for some time at the beginning of each 10s period and it will be glowing red when it is active. The button will become inactive (and no longer glowing red) once it is pressed, or if it is not pressed after a certain amount of time. You may press the button at any time while it is active; clicking on it while it is inactive will have no effect.

Pressing the button when it is active may change the time when the shape lights up. On some problems the button may postpone the shape lighting up, while on others it may advance it. Your task then is to assess whether and to what extent pressing the button changes the TIME of when the shape lights up.

To reiterate, the shape will ALWAYS light up once per 10s interval, and your task is not to judge WHETHER pressing the button makes the shape light up. Instead, you are to judge whether pressing the button changes WHEN the shape lights up.

In order to make an informed judgment it is beneficial for you to press the button during some 10s periods, and refrain from pressing it during others.

After completing the trial, participants first indicated whether they thought pressing the button advanced or postponed the outcome, and were then presented with the following question:

You indicated that pressing the button [advanced/postponed] WHEN the shape lit up. Setting aside HOW MANY SECONDS pressing the button [advanced/postponed] the time when the shape lit up, to what extent did the button increase the LIKELIHOOD that the shape lit up [early/late]?

Please provide a number between 0 and 100 where:

0 means that pressing the button NEVER [advanced/postponed] the light-up time.

50 means that pressing the button [advanced/postponed] the light-up time SOME OF THE TIME.

100 means that pressing the button ALWAYS [advanced/postponed] the light-up time.

Use intermediate numbers to express intermediate beliefs.