Stimulus control and delayed outcomes in a human causality judgement task

Phil Reed

Swansea University

Correspondence address: Phil Reed,
Department of Psychology,
Swansea University,
Singleton Park,
Swansea, SA2 8PP, UK.
e-mail: p.reed@swansea.ac.uk

Short title: causality judgement.

Published in: *Journal of Experimental Psychology: Animal Learning and Cognition*.

Abstract

Three experiments examined the impact of delayed outcomes on stimulus control of causal judgements using an interdimensional generalisation procedure. Human participants rated the causal effectiveness of responses on multiple schedules, and then underwent a generalisation test. In Experiment 1, a 3s unsignalled outcome delay reduced ratings of causal effectiveness, relative to an immediate outcome, but had higher ratings compared to a component lacking outcomes. In a generalisation test, incremental generalisation gradients, indicating inhibitory control, were found for the stimulus associated with delayed outcomes when comparison was with immediate outcomes; but decremental gradients, indicating excitatory control, were found when the comparator lacked outcomes. In Experiment 2, signalled 3s outcome delays produced higher causal ratings than unsignalled delays; with unsignalled delays producing incremental (inhibitory), and signalled delays producing decremental (excitatory), generalisation gradients when compared against each other. In Experiment 3, relative to immediate outcomes, unsignalled delays produced incremental (inhibitory) gradients, and signalled delays produced no gradient. These findings suggest similar factors may control judgements of causality as control conditioned responding.

Keywords: human causal judgements; delayed outcomes; signalled delays; generalisation gradients; inhibitory; excitatory.

The impact of temporal contiguity on judgements of causal effectiveness has been examined over a wide range of paradigms: perceptual causality (Michotte, 1946; Schlottman, 1999), judgements of the relationship between events (Lee, Hayes, & Lovibond, 2018); judgements of the effectiveness of actions (Greville & Buehner, 2010; Reed, 1999), and in dynamic video-game contexts (Young & Nguyen, 2009). In most paradigms, temporal contiguity facilitates judgements of causal relationships between events, such that a delay inserted between actions and outcomes reduces judgements of causality, relative to when the outcome is presented immediately following an action or event (Greville & Buehner, 2010; Michotte, 1946; Reed, 1999; van Elk, Salomon, Kannape, & Blanke, 2014). Such delayed relations may not be detected (Greville & Buehner, 2010; Michotte, 1946), may be underestimated (Reed, 1999), or may require additional resources to produce an impression of causality or efficacy (Mendelson & Shultz, 1976; Schlottmann, 1999; Waldmann, 1996). In contrast, temporal contiguity between actions or events and outcomes can produce higher judgements causality, even when the action action-outcome contingency is poor (Anderson & Sheu, 1995; Mendelson & Shultz, 1976; Siegler & Liebert, 1974; Wasserman & Neunaber, 1986).

Initial investigations of temporal contiguity and causal perception were conducted by Michotte (1946); participants watched one stimulus (A) move towards and collide with another stimulus (B), after which A remained stationary, and B started to move. Strong impressions of causality were reported if B moved immediately, but a delay in the onset of B's motion attenuated or abolished this impression. In similar studies, participants have reported a preference for contiguous but poor predictors of outcomes (Mendelson & Shultz, 1976; Siegler & Liebert, 1974). In studies using an analogue of instrumental conditioning, participants judge whether a response (e.g., a response to a keyboard) causes an outcome (e.g., an event on the monitor) to occur (Greville & Buehner, 2010; Reed, 1992; Wasserman

& Neunaber, 1986). Participants report greater judgements of causal effectiveness when the outcome occurred immediately after the response, compared to when it is delayed (Greville & Buehner, 2010; Reed, 1992).

Despite the ubiquity of these effects, and their potential importance for understanding the nature of learning about action-outcome relationships, relatively little work has focused on the effects of temporal contiguity, as opposed to contingency, in the field of human causal judgement (see Buehner, 2005, for a review). Several theoretical suggestions have been made as to the mechanisms that may underlie delayed-outcome effects, such as those based on associative learning (Allan, 1993; Dickinson & Balleine, 2000), that have suggested contiguity is important in the formation of response-outcome associations. Contiguity effects have been accommodated by associative (Allan, 1993), as well as by knowledge-based and probabilistic views (Einhorn & Hogarth, 1986; Waldmann, 1996), and rate-based theories (Gallistel & Gibbon, 2000; Gallistel, Craig, & Shahan, 2019; Miller & Barnet, 1993). In these views, weakening of action-outcome contiguity alters how that relationship is encoded (Buehner, 2005). Waldmann (1996) suggests that contiguity alters placement of events in a 2x2 contingency table (i.e. presence/absence of response/outcome; where the ratio outcomes given a response to other relationships determines causal judgements). Contiguous response-outcome pairings are interpreted as outcomes given a response, but non-contiguous instances are interpreted as both response without outcomes, and outcomes without a response. Although there is some degree of interpretive difficulty in using rate-based theories to explain contiguity data (see Gallistel & Balsam, 2014), one possible account given by Gallistel and Gibbon (2000; Baum, 1983) suggests that the degree of contiguity alters the assessment of the degree to which outcome rates alter as a consequence of changes in response rates.

It may also be that suggestions invoked from conditioning with delayed reinforcement may be useful in helping to explain the effects on causal judgments (Reed, 1999). Delayed

reinforcement in instrumental conditioning has received several thorough reviews (Boakes & Costa, 2014; Lattal, 1984; Lattal, 2010; Renner, 1964; Tarpy & Sawabini, 1974), and has been noted to have numerous impacts on instrumental responding beyond its effect on response rate (Lattal, 2010; Reed & Reilly, 1990; Richards, 1974). One effect of delayed instrumental reinforcement is on stimulus control over behaviour, and levels of generalisation between stimuli (Lattal, 1984; Richards & Marcattilio, 1978). Although these effects have been widely investigated in instrumental tasks using nonhumans (Richard, 1974; Richards & Marcattilio, 1978; Guttman, 1959), there is a relative paucity of evidence relating to human causal judgement studies (but see Lee et al., 2018; Lee, Lovibond, & Hayes, 2019; Livesey & McLaren, 2009). Theoretical accounts of causal judgement derived from animal learning studies should predict similar impacts on stimulus control with human causal judgements (Allan, 1993; Reed, 1999), although other theoretical views have also been used to explain such data (see Buehner, 2005; Lee et al., 2018; Wong & Lovibond, 2017). There are also some instances where similar manipulations conducted with humans produce different outcomes to those noted with nonhumans (Lee et al., 2019; Livesey & McLaren, 2009), so it is far from clear that every manipulation performed on human participants will mirror the effects noted from the study of nonhumans.

Stimulus control and generalisation play important roles in evaluative attributions (Dack, McHugh, & Reed., 2009; Lee et al., 2018), and exploring generalisation of causal judgements would provide important data for understanding everyday behaviours (Shepard, 1987). Moreover, there are clinical implications of generalisation abilities, as maladaptive generalisation is noted with many conditions, such as panic disorder (Lissek et al., 2010), posttraumatic stress disorder (Grillon & Morgan, 1999; Morey et al., 2015), and generalised anxiety disorder (Lissek, Kaczkurkin, Rabin, Geraci, Pine, & Grillon, 2014). These

considerations, in addition to potential theoretical implications, suggest this may be a useful area to investigate in terms of causal judgements.

Following training with a single visual wavelength, Guttman and Kalish (1956) varied stimuli along that visual dimension and noted approximately symmetrical gradients of generalised responding across the spectrum with a peak at the training value. Lee et al. (2018) have noted a similar effect for human judgements of causation, in that judgements of causal efficacy were higher in the presence of stimuli closest to an initial training hue (but see Lee et al., 2019; Livesey & McLaren, 2009, for instances where generalisation does not follow the same nonhuman pattern). However, evidence drawn from interdimensional generalisation protocols suggests that the nature of the generalisation gradient is determined by alternative sources of outcome, which is important in determining whether stimuli associated with response-reinforcer relationships become excitatory or inhibitory (Guttman, 1959; Richards & Hittesdorf, 1978; Richards & Marcattilio, 1978). In such studies, participants are trained on a multiple schedule, often with two components, whose stimuli are drawn from two different dimensions. Participants are then given a generalisation test for stimuli on the dimension associated with one of the components. A generalisation gradient with the highest point at the trained stimulus, and lower levels of responding as the distance from that point increase (a 'decremental gradient'), is taken to indicate excitatory control by the trained stimulus. A generalisation gradient with the lowest point at the trained stimulus, and increasing responding as the distance from that stimulus increases (an 'incremental gradient'), is taken to imply inhibitory stimulus control (Guttman, 1959; Jenkins, 1965; Rilling, 1977). Guttman (1959) noted that stimuli associated with the weaker of two components on a multiple schedule (e.g., a schedule providing a lower frequency of reinforcement) served as an inhibitory stimulus.

In relation to delayed reinforcers, incremental generalisation gradients (evidencing inhibitory control) are obtained around a stimulus associated with a component offering delayed reinforcement on a multiple VI 1-min VI 1-min schedule when the comparison is immediate reinforcement (Richards, 1974, 1975). Richards and Marcattilio (1978) found such an incremental (inhibitory) generalisation gradient using pigeon subjects when the target (tested) component offered delayed reinforcement and the other comparator component was associated with immediate reinforcement. However, they noted decremental generalisation gradients, indicating excitatory conditioning, for the stimulus associated with a delayed reinforcement component, when alternative had been associated with extinction (see also Richards & Hittesdorf, 1978). Thus, the function of the stimulus associated with delayed reinforcement is dependent on the status of the comparison stimulus. If human causal judgements in tasks in which participants are asked to rate the effectiveness of their responses follow the same rules as these conditioning studies, then this effect should be noted for causal judgements. The initial aim of the current study was to replicate this effect in a causal judgement paradigm, based on that employed by Lee et al. (2018), but presenting delayed outcomes, to determine if this is an instance where human and nonhuman patterns of generalisation can cohere.

While it is widely known causal impressions decline with a temporal delay between action and outcome, events occurring during the delay can bridge the temporal gap between action and outcome (Buehner & May, 2003; Gruber, Fink, & Damm, 1957; Reed, 1992; 1999; Young & Falmier, 2008a; 2008b). For example, if a tone is presented during an action-outcome delay, levels of causal judgements are partly restored (Reed, 1999; Rescorla, 1982; Young & Falmier, 2008b; Weller, Schwarz, Kunde, & Pfister, 2017). It could be suggested that the addition of the delay signal removes the delay in these conditions (Buehner, 2005), as Reed (1999) found the effectiveness of the delay signal depended on its contiguity with the

response. However, delay signals serve multiple functions (Lattal, 1984), often altering the impact of the conditioning procedure stimulus control (Reilly, Schachtman, & Reed, 1996; Richards & Hittesdorf, 1978). In terms of interdimensional generalisation studies, Richards and Hittesdorf (1978) found that pigeons, trained on a multiple schedule containing schedules with signalled and unsignalled reinforcement delay, showed incremental (inhibitory) generalisation gradients for stimuli associated with an unsignalled, but not signalled, delays of reinforcement schedule. A second aim of the current studies was to examine whether signalling the outcome delay would change the nature of the stimulus control in human causality judgements.

The current studies employed an interdimensional generalisation protocol to test stimulus functions associated with unsignalled and signalled outcome delays in causal judgement tasks. The generalisation gradients around stimuli associated with unsignalled and signalled outcome delays were examined in comparison to immediate outcomes, conditions lacking outcomes, and with respect to each other. Should similar results to conditioning studies be noted, it would extend the range of factors that are functionally similar, and provide further data which theoretical accounts of causality judgement would need to accommodate.

Experiment 1

Experiment 1 examined the degree to which an action-outcome delay would reduce judgements of causal effectiveness in a multiple schedule (Reed, 1999). As causal relations with fixed temporal action-outcome delays have been found to be rated more strongly than those with variable delays (Greville & Buehner, 2010), a comparison between two delay conditions (fixed vs. variable) was included. Further, the experiment examined whether

results reported by Richards and Marcattilio (1978) for instrumental conditioning with pigeons would occur in the human causality judgement procedure. Richards and Marcattilio (1978) tested generalisation along the stimulus dimension associated with a delayed outcome (reinforcer), after training on a two-component multiple schedule. Their results showed an incremental (inhibitory) gradient when the alternative comparison component was associated with immediate reinforcement, but a decremental (excitatory) gradient when the comparison was associated with the component lacking outcomes.

In the current study, a procedure similar to that reported by Lee et al. (2018) to test intradimensional generalisation in causality judgements was employed for an interdimensional procedure. Participants responded on a two-component schedule. All responses in both components had a .5 probability of being followed by an outcome (i.e. a multiple random ratio (RR) 2, RR-2 schedule of outcome presentations was in operation). One component was signalled by a black cue, and the other component was signalled by a stimulus of blue-green hue. Three groups of participants were tested. One group received immediate outcomes (black) or fixed 3-s delayed outcomes (blue-green); a second group received immediate outcomes (black) or variable delay outcomes (blue green); and a final group received no outcomes (black) or fixed 3-s delayed outcomes (blue-green). Following training, participants were tested in conditions lacking outcomes across a range of stimuli from the blue-green dimension.

Method

Transparency and Openness.

We report how we determined sample size, data exclusions, and manipulations. Data are available on request. Data were analysed using SPSS version 26.

Participants.

A sample of 150 volunteers (70 males and 80 females) was recruited thorough an online platform (Survey Circle). Participants were aged between 18 to 47 years (mean = 22 $\pm$ 5.7 SD). Exclusion criteria were that subjects could not by below 18 years of age, and had to have English as a first language. Participants received no financial remuneration for their participation, but earned credits allowing them to utilise the recruitment websites. G-Power calculation implied that for 95% power, with a $p < .05$ criteria, and a medium effect size ($f' = .25$), that 117 participants would be required for a 3x2 mixed-model ANOVA, with minimal correlations between measures. As a result, it was decided to recruit 50 participants per group, in line with numbers recruited by Lee et al. (2018). In total 267 individuals started the experiment, but 117 did not complete, and their data were discarded. Ethical approval was obtained through the University psychology Ethics Committee.

Apparatus.

The experiment was programmed using Visual Basic 6.0, and was presented through the Gorilla online system, designed to run on desktop and laptop computers. The volunteers were randomly assigned to groups until 50 had been assigned to any particular group, after which participants were assigned to the remaining groups. Participants made responses using a standard PC keyboard and mouse. The screen was white. For one component, a black circle stimulus (200 pixels in diameter) was presented in the centre of the screen near the top. For the other component, a blue-green circle of the same size was presented at the same location. Following Lee et al. (2018) the stimuli were created using the HSB (hue, saturation, brightness) scale. Responses to the spacebar of the keyboard would sometimes produce a 250ms triangle flash below the coloured cue towards the bottom of the screen. During the test phase, 7 blue-green stimuli were presented, created by varying hue, while keeping

constant saturation (100%) and brightness (75%). The minimum and maximum hue values were .403 (145^{o} = green) and .555 (200^{o} = blue), with equal spacing between stimuli along the dimension giving a midpoint (S4) of .479 (172^{o}). A schematic showing the creation of the endpoint and S4 stimuli is shown in Appendix 1, based on the website: HSB Color Picker (codepen.io).

Procedure.

The design for this study, and all subsequent studies reported here, can be seen in the schematic Table 1.

Table 1 about here

Pre-training: During training participants were instructed to determine the causal effectiveness of a spacebar response in making a triangle appear on the monitor (see Reed, 1999). They were presented with the following instructions:

"*Your task is to judge the extent to which responses to the space bar cause a triangle to appear on the computer screen. You will be presented with different problems (signalled by coloured circles on the screen), each lasting a short period. In each, responses may or may not cause the triangle to appear. You can press the space bar as often or little as you wish, but it may be helpful to press it some of the time, and not to press it some of the time. The relation between responding and whether the triangle appears will be constant within each problem, but they could differ between problems.*"

Following these instructions, a black circle was presented for 45s. During this period, outcomes were programmed to occur with a .5 probability after each response (i.e. an RR-2 schedule), and were presented immediately after a response when they were presented. There

was no visual feedback (other than the programmed outcomes) following a response in this, or any subsequent phase. Following this period (45s), the circle disappeared, and the participants were asked to move a cursor on the screen along a line, using the mouse, to indicate how effective the responses had been. The line was 10cm long on the screen. It had no markings apart from at each of the ends. The ends of the line were labelled: 'Not effective at all' (left end), and 'Completely effective' (right end). Above the line appeared the instructions: "*Overall, to what extent do you feel pressing the button controlled the triangle lighting up in this condition?*" (Greville & Buehler, 2010). Participants moved the cursor to the point along the line that they felt indicated the degree of relationship between a response and an outcome. If the cursor was placed at the extreme left, this was scored 0; if it was at the extreme right, this scored 100; and 5cm along (in the middle) was scored 50, etc. After positioning the cursor at the appropriate point, participants were instructed to press Return to start the next 45s exposure to these contingencies. There were 4 such exposures.

Training: During the initial training, all outcomes were presented with a .5 probability after each response. Participants received exposure to two components, each associated with a different stimulus (i.e. a multiple RR-2, RR-2 schedule). For one group (.5, .5FD), one component, in which outcomes were presented immediately, was signalled by the presentation of the black stimulus, as above. The other was signalled by the presence of a blue-green circle, which was the hue at the midpoint of the continuum. Outcomes were delayed by 3s in this component. If an outcome was scheduled, then it was presented 3s after the response. Responses made during this period could also produce outcomes, according to the same schedule. For a second group (.5, .5VD), these contingencies were operative, with the exception that the delay varied between 1s and 5s, with a mean of 3s, with each delay being randomly selected. The final group (Ext, .5FD), received no outcomes for responses when the black stimulus was sown, and outcomes with a .5 probability and a 3s fixed delay,

when the blue-green stimulus was presented. Each component lasted for 45s, and was followed by a rating of causal effectiveness. There was no counterbalancing adopted to make use of the previous training with immediate outcomes for the black component. There were 8 exposures to the black, and 8 exposures to the blue-green stimulus, randomly presented the participants, with the limitation that only two components of any type could be presented successively.

Generalisation test: Prior to test, participants were presented with instructions that they would now be asked to rate the likely causal effectiveness of a response in the presence of different stimuli. The test phase consisted of two presentations of each of the 7 blue-green stimuli, and 4 presentations of the black stimulus. The stimulus was presented for 1000ms, then disappeared from the screen, and the instructions and rating line appeared. The order of presentation of the coloured stimuli was randomised.

Results and Discussion

Figure 1 about here

Figure 1 shows the group-mean ratings of causal effectiveness for the immediate (or lacking outcomes) component, and the delay components, for the three groups. These ratings were the average of those made in the last three presentations of each component where performance was taken to have become more stable (all analyses on these training data are conducted only on these terminal performance data). Inspection of these data shows that ratings for the immediate components were higher, and ratings for the component lacking outcomes was lower, than for either the fixed or variable delay components. This result replicates previous findings (Greville & Buehner, 2010; Reed, 1999). There was little

difference between the fixed or variable delay conditions. Examination of the ratings for the unsignalled delay condition revealed slightly higher ratings for the group with the comparison lacking outcomes than for the groups with the immediate outcome comparisons. This small sized difference is similar to a behavioural contrast effect sometimes reported in interdimensional generalisation studies, but not by Richards and Marcattilio (1978).

These impressions were confirmed by a two-factor mixed-model analysis of variance (ANOVA) with group as a between-subject factor, and component (comparison versus target) as a within-subject factor. These analyses were supplemented by Bayesian analysis to establish the probability of the outcome suggested by the ANOVA. The results of the ANOVA, along with the effect size (and 95% confidence intervals), and the appropriate Bayes statistic (using a uniform non-informative prior), are shown. These analysis revealed significant main effects of group, $F(2,147) = 169.49$, $p < .001$, $\eta^2_p = .535[.427:.616]$, $p(H_1/D) = .999$, and component, $F(1,147) = 53.93$, $p < .001$, $\eta^2_p < .268[.154:.376]$, $p(H_1/D) = .999$, and a significant interaction, $F(2,147) = 325.57$, $p < .001$, $\eta^2_p = .689[.608:.745]$, $p(H_1/D) = .999$. Simple effect analysis revealed a significant effect of group for the comparison component, $F(2,147) = 485.81$, $p < .001$, $\eta^2_p = .868[.830:.892]$, $p(H_1/D) = .999$. Tukey's Honestly Significant Difference (HSD) tests revealed significantly higher ratings by each of the groups with immediate conditions compared to those of the group with the condition lacking outcomes, ps $< .05$. Simple effect analysis revealed a significant effect of group for the target component, $F(2,147) = 13.86$, $p < .001$, $\eta^2_p = .159[.060:.258]$, $p(H_1/D) = .999$. Tukey's HSDs tests revealed higher ratings for the unsignalled delay group with the comparison lacking outcomes than either of the two unsigalled groups with the immediate outcome comparisons, ps $< .05$.

Table 2 about here

Table 2 shows the group-means for the response and outcomes per min, the probability of an outcome given a response, and the experienced action-outcome delay, averaged over the last three presentations to each component. There was little difference in the response rates emitted by the three groups across the two components. This lack of difference is often reported in studies of causality judgment that use outcomes as opposed to reinforcers, and focus the task of judgements of causality (cf. Reed, 1993; 2001). A two-factor mixed-model ANOVA (group x component) conducted on the response rates revealed no main effects of group, $F < 1$, $\eta^2_p = .008[.000:.049]$, $p(H_0/D) = .989$, or component, $F(1,147) = 1.66$, $p = .199$, $\eta^2_p = .011[.000:.066]$, $p(H_1/D) = .841$, or interaction, $F < 1$, $\eta^2_p = .008[.000:.051]$, $p(H_0/D) = .987$.

A two-factor mixed-model ANOVA (group x component) conducted on the outcome rates unsurprisingly revealed significant main effects of group, $F(2,147) = 56.63$, $p < .001$, $\eta^2_p = .435[.313:.526]$, $p(H_1/D) = .999$, and component, $F(1,147) = 40.27$, $p < .001$, $\eta^2_p = .215[.108:.323]$, $p(H_1/D) = .999$, and a significant interaction, $F(2,147) = 75.11$, $p < .001$, $\eta^2_p = .505[.389:.587]$ $p(H_1/D) = .999$. Simple effect analyses revealed a significant effect of group for the comparison component, $F(2,147) = 132.53$, $p < .001$, $\eta^2_p = .643[.549:.705]$, $p(H_1/D) = .999$, with Tukey's HSD tests revealing higher outcomes/min for the two groups with immediate outcomes as opposed to lacking outcomes, ps $< .05$. There was no simple effect of group for the target component, $F < 1$, $\eta^2_p = .006[.000:.044]$, $p(H_0/D) = .983$.

Following from these differences in outcome rates due to the component lacking outcomes, a two-factor mixed-model ANOVA (group x component) revealed significant main effects of group, $F(2,147) = 484.18$, $p < .001$, $\eta^2_p = .868[.829:891]$, $p(H_1/D) = .999$, and

component, $F(1,147) = 574.51$, $p < .001$, $\eta^2_p = .799[.743:.836]$, $p(H_1/D) = .999$, and a significant interaction, $F(2,147) = 623.82$, $p < .001$, $\eta^2_p = .894[.836:.913]$, $p(H_1/D) = .999$, for the probability of an outcome given a response. Simple effect analysis revealed a significant effect of group for the comparison component, $F(2,147) = 1057.32$, $p < .001$, $\eta^2_p = .935[.916:.947]$, $p(H_1/D) = .999$. Tukey's HSD tests revealed greater pO/R for both groups with an immediate outcome than the group lacking outcomes, ps $< .05$. There was no simple effect of group for the target component, $F < 1$, $\eta^2_p = .003[.000:.031]$, $p(H_0/D) = .983$.

A one way ANOVA conducted on the action-outcome delay revealed no difference between the groups, $F(2,147) = 1.38$, $p = .254$, $\eta^2_p = .018[.000:.072]$, $p(H_0/D) = .896$. A one-way ANOVA conducted on the response rates during the delay period revealed no difference between the groups, $F < 1$, $\eta^2_p = .006[.000:.012]$, $p(H_0/D) = .999$.

Figure 2 about here

The group-mean mean judgements during were taken for the black stimulus, previously associated with immediate or no outcome conditions, and were highly similar to those obtained during training. These values were: 58.8 ($\pm$ 11.8) for the .5, .5FD group; 59.5 ($\pm$ 12.7) for the .5, .5VD group; and 9.1 ($\pm$ 5.5) for the Ext, .5FD group. Figure 2 shows the group-mean judgments for each of the stimuli used in the generalisation test (S4 was the original training stimulus). Inspection of these data for the stimuli related to the delay conditions from training (blue-green) shows incremental gradients, suggesting inhibitory control, for both the fixed and variable delay groups. This mirrors the response rate findings reported by Richards and Marcattilio (1978; Richards & Hittesdorf, 1978). There was little difference between these delayed outcome groups, with the variable delay group having slightly steeper gradients; a finding again similar to that seen in Richards and Marcattilio

(1978). In contrast, the group with an unsignalled delay and a comparison lacking outcomes demonstrated decremental gradients, representing excitatory control. It should be noted that the gradient was asymmetrical. This may have been the result of the lack of counterbalancing of the colours for the stimuli in the immediate and comparison components, with one colour potentially having either intrinsic positive or negative valence.

This description was corroborated by a two-factor mixed-model ANOVA (group x stimulus) which revealed significant main effects of group, $F(2,147) = 11.45$, $p < .001$, $\eta^2_p = .135[.043:.231]$, $p(H_1/D) = .997$, and stimulus, $F(6,882) = 16.97$, $p < .001$, $\eta^2_p = .103[.064:.136]$, $p(H_1/D) = .998$, and a significant interaction, $F(12,882) = 63.08$, $p < .001$, $\eta^2_p = .462[.410:.495]$, $p(H_1/D) = .997$. Trend tests conducted on each group revealed significant 'incremental' quadratic trends for groups with fixed, $F(1,882) = 170.22$, $p < .001$, $\eta^2_p = .162[.120:.205]$, $p(H_1/D) = .999$, and variable, $F(1,882) = 210.50$, $p < .001$, $\eta^2_p = .192[.148:.237]$, $p(H_1/D) = .999$, delays and an immediate outcome comparison; and a significant decremental quadratic trend for the unsignalled delay group with the comparison lacking outcomes, $F(1,882) = 267.82$, $p < .001$, $\eta^2_p = .232[.187:.278]$, $p(H_1/D) = .997$, groups. There were no linear trends for any group, all Fs < 1. To test whether there was a difference between the fixed and variable delay groups in terms of generalisation, a separate two-factor ANOVA (group x stimulus) was conducted just on these two groups. This revealed no significant interaction between group and stimulus, $F(6,588) = 1.46$, $p < .188$, $\eta^2_p = .014[.000:.029]$, $p(H_0/D) = .999$.

These results replicate many previous demonstrations that an unsignalled delayed outcome reduces judgments of causal effectiveness (Greville & Buehner, 2010; Reed, 1999; van Elk et al., 2014). These ratings were reduced in the absence of differences in the rate of outcomes, or the probability of an outcome given a response. The stimulus associated with the unsignalled delay controlled a generalisation gradient of causal judgements (see also Lee

et al., 2018). Again, that there can be transfer of judgment from a trained to a novel stimulus through generalisation has been shown previously (Dack et al., 2009' Lee et al., 2018). A novel finding is that the nature of the gradient depends on the nature of the comparison stimuli during training, an effect noted for response rates in instrumental discrimination learning (Guttman, 1959; Richards & Marcattilio, 1978). When trained against an immediate outcome, the gradient around the stimulus associated with the delayed outcome suggested inhibitory control. In contrast, when trained against a component lacking outcomes, this gradient suggested excitatory control. This was the result expected if the factors controlling causal judgements are similar to those controlling instrumental learning (Allan, 1993; Dickinson & Balleine, 2000; Reed, 1999).

Experiment 2

Richards and Hittesdorf (1978; Experiment 2) explored generalisation gradients for cues associated with unsignalled and signalled delays of reinforcement. Pigeons responded on a multiple VI 60s, VI 60s schedule: one component always associated with unsigalled delays of reinforcement, the other with signalled delayed reinforcement. Signalled delays maintained higher response rates than unsignalled delays (see also Lattal, 1984; Reed & Reilly, 1990). Incremental (inhibitory) generalisation was obtained for unsignalled components, implying relative to signalled components that an unsignalled delay exerted inhibitory control (see also Reilly et al., 1996, for a similar discussion). No test of the alternative comparison for signalled delays compared to unsiganlled delays was made. The second study examined whether such an effect would also been seen with signalled and unsignalled delays of outcome in a human causality judgement task. As in Experiments 1, outcomes were programmed to occur with a .5 probability following each response. There

were two components presented (i.e. participants responded on a multiple RR-2, RR-2 schedule), with components signalled either by black or a blue-green hue. Two groups were tested, with both groups receiving a component with signalled delayed outcomes and a component with unsignalled delayed outcomes. Following training, participants were tested in the absence of outcomes across a range of stimuli from the blue-green dimension; for one group, this dimension had been associated with an unsignalled delay, and for the other this had been associated with the signalled delay.

Method

Participants and apparatus.

A sample of 100 volunteers (19 males and 81 females), with a mean age of 21 $\pm$ 3.2; range = 18 - 31) was recruited as described in Experiment 1. G-Power calculation implied that for 95% power, with a $p < .05$ criteria, and a medium effect size ($f' = .25$), that 96 participants would be required for a 2x2 mixed-model ANOVA, with minimal correlations between measures. In total 175 individuals started the experiment, but 75 did not complete, and their data were discarded. The apparatus was as described in Experiment 1

Procedure.

Training: No pretraining was given, as it was hoped not to establish any prior expectation of outcome immediacy. Participants were presented with the instructions as described in Experiment 1, and then were exposed to two-components, each signalled by a different coloured stimulus (i.e. a multiple RR-2, RR-2 schedule). For one group (.5S, .5U), in one component (S), outcomes were presented with a .5 probability, following a 3s signalled delay after the response. The signal was a row of 6 'X's (Reed, 1999), presented under the black stimulus. In the other component, associated with the presence of a blue-

green circle, as described in Experiment 1, outcomes were delayed by 3s with no signal. Responses during outcome delays were recorded, but had no programmed consequences in either component. The second group (.5U, .5S) received the components associated with the alternative cuing arrangement (black was unsignalled delayed outcomes, and blue-green was a signalled delayed outcome). There was no visual feedback, other than programmed outcomes or delay stimuli, following a response. Each component lasted for 45s, followed by a rating of causal effectiveness as described in Experiment 1. There were 8 exposures to the black, and 8 exposures to the blue-green stimulus, randomly presented the participants, with the limitation that only two components of any type could be presented successively.

Generalisation test: Prior to test, participants were presented with instructions that they would now be asked to rate the likely causal effectiveness of a response in the presence of different stimuli. The test phase consisted of two presentations of each of the 7 blue-green stimuli, and 4 presentations of the black stimulus. The coloured stimulus was presented for 1000ms, then disappeared from the screen, and the instructions and rating line appeared. The order of presentation of the coloured stimuli was randomised.

Results and Discussion

Figure 3 about here

Figure 3 shows group-mean ratings of causal effectiveness for the two groups in both components. These ratings were the average of those made in the last three presentations of each component where performance was taken to have become more stable (all analyses on these training data are conducted only on these terminal performance data). Inspection of these data shows that ratings for the signalled outcome delays were higher than for

unsignalled delay components. This result replicates previous findings (Buehner & May, 2003; Reed, 1999; Young & Falmer, 2008a).

A two-factor mixed-model ANOVA (group x component) conducted on the judgements revealed no significant main effects of group, $F(1,98) = 1.38$, $p = .242$, $\eta^2_p = .013[.000:.089]$, $p(H_0/D) = .832$, or component, $F < 1$, $\eta^2_p = .004[.000:.064]$, $p(H_0/D) = .888$, but there was a significant interaction, $F(1,98) = 264.27$, $p < .001$, $\eta^2_p = .715[.617:.775]$, $p(H_1/D) = .999$. Simple effect analyses revealed judgements for the signalled delay were higher in comparison component, $F(1,98) = 113.65$, $p < .001$, $\eta^2_p = .537[.400:.632]$, $p(H_1/D) = .999$. and in the target component, $F(1,98) = 107.08$, $p < .001$, $\eta^2_p = .522[.3783:.619]$. $p(H_1/D) = .999$.

Table 3 about here

Table 3 shows the group-means for the response and outcomes per min, the probability of an outcome given a response, and the experienced action-outcome delay, averaged over the last three exposures to the components, for each component. For the Delay group, the comparison was a signalled delay, and the target was an unsignalled delay. For the signalled delay group, the comparison was a delay, and the target was a signalled delay. A two-factor mixed-model ANOVA (group x component) was conducted on the response rate data revealed no main effects of group, $F < 1$, $\eta^2_p < .001[.000:.050]$, $p(H_0/D) = .999$, or component, $F < 1$, $\eta^2_p < .008[.000:.077]$, $p(H_0/D) = .999$, or interaction, $F < 1$, $\eta^2_p < .008[.000:.076]$, $p(H_0/D) = .999$. As in Experiment 1, and previous studies (Reed, 1999), response rates are not often impacted in these sorts of causality judgement studies.

A two-factor mixed-model ANOVA (group x component) conducted on the outcome rates revealed greater outcome rates in the comparison component, $F(1,98) = 4.48$, $p = .037$,

η^2_p = .043[.000:.143] , $p(H_1/D)$ = .481, but there was no main effect of group, $F < 1$, η^2_p = .002[.000:.054], $p(H_9/D)$ = .999, nor was there an interaction, $F < 1$, η^2_p = .004[.000:.064], $p(H_9/D)$ = .999. A two-factor mixed-model ANOVA (group x component) conducted on the probability of an outcome given a response revealed no significant main effects of group, $F < 1$, η^2_p = .001[.000:.004], $p(H_0/D)$ = .981, or component, $F(1,98) = 1.34$, $p = .249$, η^2_p = .013[.000:.088], $p(H_0/D)$ = .833, nor was there an interaction, $F < 1$, η^2_p = .008[.000:.077], $p(H_0/D)$ = .909. A two-factor mixed-model ANOVA (group x component) conducted on the actual action-outcome delay revealed no significant main effects of group, $F(1,98) = 1.84$, $p = .177$, η^2_p = .018[.000:.099], $p(H_0/D)$ = .785, or component, $F(1,98) = 1.44$, $p = .233$, η^2_p = .014[.000:.091], $p(H_0/D)$ = .822, nor was there an interaction, $F < 1$, η^2_p = .001[.000:.004]. $p(H_0/D)$ = .999. A two-factor mixed-model ANOVA (group x component) conducted on the delay response rates revealed no significant main effects of group, $F < 1$, η^2_p = .001[.000:.004], $p(H_0/D)$ = .999, or component, $F(1,98) = 1.56$, $p = .215$, η^2_p = .015[.000:.097], $p(H_0/D)$ = .856, nor was there an interaction, $F < 1$, η^2_p = .010[.000:.023]. $p(H_0/D)$ = .999.

Figure 4 about here

The group-mean mean judgements during were taken for the black stimulus, and were highly similar to those obtained during training. These values were: 18.4 (± 8.4) for the .5U component of the .5U, .5s group; and 42.0 (± 12.7) for the .5S component for the .5S, .5U group. Figure 4 shows the group-mean judgments for each of the stimuli used in the generalisation test (S4 was the original training stimulus). Inspection of these data for the target condition (Blue-green) shows incremental gradients, suggesting inhibitory control, for

the unsignalled delay group, but an excitatory decremental gradient for the signalled group. A two-factor mixed-model ANOVA (group x stimulus) conducted on these data revealed a significant main effect of group, $F(1,98) = 7.40$, $p = .008$, $\eta^2_p = .070[.005:.186]$, $p(H_1/D) = .792$, but not stimulus, $F(6,588) = 1.51$, $p = .173$, $\eta^2_p < .016[.000:.031]$, $p(H_0/D) = .999$, there was a significant interaction, $F(6,588) = 79.92$, $p < .001$, $\eta^2_p = .462[.400:.508]$, $p(H_1/D) = .999$. Trend tests conducted on each group revealed significant quadratic trends for the delay, $F(1,588) = 196.92$, $p < .001$, $\eta^2_p = .280[.210:.317]$, $p(H_1/D) = .999$; and signalled, $F(1,588) = 159.12$, $p < .001$, $\eta^2_p = .221[.164:.278]$, $p(H_1/D) = .999$. There was no linear trends for the delay group, $F(1,588) = 1.48$, $p = .229$, $\eta^2_p < .003[.000:.017]$, $p(H_0/D) = .864$, but there was for the signalled group, $F(1,588) = 9.42$, $p = .003$, $\eta^2_p = .016[.002:.043]$ $p(H_1/D) = .929$.

Overall, these results replicated the effect of signalling a delayed outcome on elevating judgements of causal effectiveness relative to an unsignalled outcome (Buehner & May, 2003; Reed, 1999). They also extended the findings regarding stimulus control, to show that, as with instrumental conditioning, stimuli associated with an unsignalled delay show apparent inhibitory control, when trained against a signalled delay (Richards & Hittesdorf, 1978). The results for the signalled delay showed that, relative to an unsignalled delay comparison, there were excitatory generalisation gradients. This is the first time that such a comparison has been directly examined, but would be predicted from the findings of higher ratings of causal effectiveness in the signalled, relative to unsignalled. delay conditions. As with Experiment 1, the data add to the literature showing parallels between the factors controlling instrumental learning and those controlling causal judgements (Allan, 1993; Dickinson & Balleine, 2000).

Experiment 3

The final study extended the range of conditions studied by comparing the generalisation gradients observed when comparing unsignalled and signalled outcome delays to an immediate comparison component. This would allow replication of the results noted in Experiment 1 for unsinalled delays, but also explore whether cues associated with a signalled delay would have the same inhibitory properties as those associated with an unsignalled delay. In the study reported by Richards and Hittesdorf (1978; Experiment 1), pigeons responded on a multiple VI 60s, VI 60s schedule: one component associated with immediate reinforcement, the other with delayed reinforcement. Signalled delays maintained higher response rates than unsignalled delays (see also Lattal, 2010; Reed & Reilly, 1990). Incremental (inhibitory) generalisation gradients were obtained for both types of delayed reinforcement schedules, although gradients were shallower for the signalled components than the unsignalled delay components. However, when using a similar procedure, where the availability of reinforcement is signalled, cues associated with such signalled-reinforcement components have been shown to produce flat generalisation gradients when trained in contrast to immediate reinforcement components (Lander, 1971; McCoy, Parker, & McGee, 1976), so predictions about the effects of the signalled component are not as clear as for the unsignalled component.

In the current experiment, participants responded on a multiple RR-2, RR-2 schedule of outcome presentations; one component signalled by a black cue, and the other component signalled by a blue-green cue. Three groups of participants were tested, with all groups receiving immediate outcomes in one component (black). In the other component (blue-green), one group received unsignalled 3-s delayed outcomes; a second group received 3s

signalled outcome delays; and a final group received immediate outcomes. Participants were then tested without outcomes across a range of stimuli from the blue-green dimension.

Method

Participants and apparatus.

A sample of 150 volunteers (58 males and 92 females), with a mean age of 20 $\pm$ 2.9; range = 18 - 29) was recruited as described in Experiment 1. G-Power calculation implied that for 95% power, with a $p < .05$ criteria, and a medium effect size ($f' = .25$), that 117 participants would be required for a 2x2 mixed-model ANOVA, with minimal correlations between measures. In total 214 individuals started the experiment, but 64 did not complete, and their data were discarded. The apparatus was as described in Experiment 1

Procedure.

Pre-training: Participants were presented with the instructions as described in Experiment 1. Participants then received training with immediate outcomes programmed with a .5 probability after each response (i.e. an RR-2 schedule), cued by a black circle, as described in Experiment 1.

Training: During training, participants were exposed to two components, each signalled by a different coloured stimulus (i.e. a multiple RR-2, RR-2 schedule). All groups received immediate outcomes in one component cued by a black circle. There was no counterbalancing to capitalise on any expectancies of immediate outcomes created during pertaining. For one group (.5, .5U), in the other component (blue-green), outcomes were presented after a 3s delay. For a second group (.5, .5S), these contingencies were operative, with the exception that the delay was signalled with a row of 6 X, as described in Experiment 2. The final group (.5, .5) received immediate outcomes in the blue-green component. There

was no visual feedback, other than programmed outcomes or delay signals, following a response in this or any subsequent phase. Each component lasted for 45s, and was followed by a rating of causal effectiveness. There were 8 exposures to the black, and 8 exposures to the blue-green stimulus, randomly presented the participants, with the limitation that only two components of any type could be presented successively.

Generalisation test: Prior to test, participants were presented with instructions that they would now be asked to rate the likely causal effectiveness of a response in the presence of different stimuli. The test phase consisted of two presentations of each of the 7 blue-green stimuli, and 4 presentations of the black stimulus. The coloured stimulus was presented for 1000ms, then disappeared from the screen, and the instructions and rating line appeared. The order of presentation randomised.

Results and Discussion

Figure 5 about here

Figure 5 shows the group-mean ratings of causal effectiveness for the immediate outcome comparison (black) components, and for the target delay or immediate outcome components (blue-green), for the three groups. These ratings were the average of those made in the last three presentations of each component where performance was taken to have become more stable (all analyses on these training data are conducted only on these terminal performance data). Inspection of these data shows that ratings for the immediate outcome components were higher than ratings for the delayed outcome components, and ratings for the signalled delay outcome were higher than those for the unsignalled delayed outcome. This result corroborates the data from Experiments 1 and 2, and from previous studies (Buehner &

May, 2003; Reed, 1999; Young & Falmier, 2008a). These impression were confirmed by a two-factor mixed-model ANOVA (group x component), which revealed significant main effects of group, $F(2,147) = 134.50$, $p < .001$, $\eta^2_p = .584[.479:.655]$, $p(H_1/D) = .999$, and component, $F(1,147) = 345.07$, $p < .001$, $\eta^2_p = .701[.622:.755]$, $p(H_1/D) = .999$, and a significant interaction, $F(2,147) = 93.93$, $p < .001$, $\eta^2_p = .561[.452:.635]$, $p(H_1/D) = .999$. Simple effect analyses revealed no difference in judgements for the immediate condition, $F(1,147) = 1.19$, $p = .203$, $\eta^2_p = .008[.000:.058]$, $p(H_0/D) = .989$, but a group difference for the target condition, $F(1,98) = 226.28$, $p < .001$, $\eta^2_p = .606[.508:.676]$, $p(H_1/D) = .792$. Tukey's HSDs tests revealed significant pairwise differences between all groups, ps $< .05$.

Table 4 about here

Table 4 shows the group-means for the response and outcomes per min, probabilities of an outcome given a response, and the experienced action-outcome delay, averaged over the last three exposures to the components, for each component. A two-factor mixed-model ANOVA (group x component) revealed the immediate group had a higher response rate than the other two groups, $F(2,147) = 5.74$, $p = .004$, $\eta^2_p = .072[.008:.155]$, $p(H_1/D) = .651$, but there was no main effect of component, $F(1,147) = 2.11$, p .149, $\eta^2_p = .014[.000:.072]$, $p(H_1/D) = .809$, or interaction, $F < 1$, $\eta^2_p = .001[.000:.038]$, $p(H_1/D) = .991$. The same pattern of data was shown for the outcome rates, with a two-factor mixed-model ANOVA (group x component) revealing a significant main effect of group, $F(2,147) = 6.69$, $p = .002$, $\eta^2_p = .083[.013:.170]$, $p(H_1/D) = .812$, but not component, $F(1,147) = 1.65$, $p = .201$, $\eta^2_p = .011[.000:.066]$, $p(H_0/D) = .843$, nor interaction, $F < 1$, $\eta^2_p = .001[.000:.037]$, $p(H_0/D) = .991$. A two-factor mixed-model ANOVA (group x component) conducted on the probability of an outcome given a response revealed no significant main effects of group, $F(2,147) = 1.34$, $p =$

.264, η^2_p = .017[.000:.071], $p(H_0/D)$ = .974, or component, $F < 1$, η^2_p = .001[.000:.001], $p(H_0/D)$ = .999, nor was there an interaction, $F(2,147) = 2.93$, $p = .056$, η^2_p = .038[.000:.107], $p(H_0/D)$ = .969. Unsurprisingly, a between-subject ANOVA conducted on the actual action-outcome delay in the target component revealed a significant main effect of group, $F(2,147) = 2843.17$, $p < .001$, η^2_p = .975[.967:.979], $p(H_1/D)$ = .999. Tukey's HSDs tests revealed only significant pairwise differences the immediate outcome condition and each of the two delay groups, ps < .05. Likewise, a between-subject ANOVA conducted on the delay response rate revealed a significant main effect of group, $F(2,147) = 1116.09$, $p < .001$, η^2_p = .919[.923:.983, $p(H_1/D)$ = .999. Tukey's HSDs tests revealed only significant pairwise differences the immediate outcome condition and each of the two delay groups, ps < .05.

Figure 6 about here

The group-mean mean judgements during were taken for the black stimulus, associated with the immediate condition during training, and were highly similar to those obtained during training. These values were: 65.3 (± 14.2) for group .5, .5U; 65.9 (± 13.4) for group .5, .5S; and 65.5 (± 11.5) for group .5, .5. Figure 6 shows the group-mean judgments for each of the stimuli used in the generalisation test (S4 was the original training stimulus). Inspection of these data shows an incremental gradient, suggesting inhibitory control, for both the unsignalled delayed outcome group. This replicates the findings from the present Experiment 1, and is consistent with those reported for pigeons response rates by Richards and Hittsdorf (1978). There was also a shallower incremental gradient for the signalled delayed outcome group, again similar to that reported by Richards and Hittsdorf (1978). There was little evidence of a gradient for the immediate outcome group. A two-factor mixed-model ANOVA (group x stimulus) conducted on these data revealed significant

main effects of group, $F(2,147) = 726.61$, $p < .001$, $\eta^2_p = .908[.881:.924]$, $p(H_1/D) = .999$, and stimulus, $F(6,882) = 36.61$, $p < .001$, $\eta^2_p = .199[.150:.239]$, $p(H_1/D) = .999$, and a significant interaction, $F(12,882) = 15.99$, $p < .001$, $\eta^2_p = .178[.125:.212]$, $p(H_1/D) = .999$. Trend tests conducted on each group revealed significant quadratic trends for the delay, $F(1,882) = 72.33$, $p < .001$, $\eta^2_p = .075[.045:.111]$; and signalled, $F(1,882) = 6.82$, $p < .001$, $\eta^2_p = .008[.000:.023]$; but not for the immediate group, $F < 1$, $\eta^2_p = .001[.000:.001]$. There were no linear trends for any group, all Fs < 1. To test whether there was a difference between the unsignalled and signalled delay groups in terms of generalisation, a separate two-factor ANOVA (group x stimulus) was conducted just on these two groups. This revealed significant main effects of group, $F(1,98) = 532.48$, $p < .001$, $\eta^2_p = .844[.788:.8786]$, $p(H_1/D) = .999$, and stimulus, $F(6,588) = 58.58$, $p < .001$, $\eta^2_p = .374[.310:.422]$, $p(H_1/D) = .999$, and a significant interaction between group and stimulus, $F(6,588) = 20.61$, $p < .001$, $\eta^2_p = .173[.115:.220]$, $p(H_1/D) = .999$.

These results corroborate suggestions that cues associated with unsignalled and signalled delayed outcomes in human causality judgements conform to similar principles as those associated with delayed reinforcers in instrumental learning tasks. In this regard, they offer support to views suggesting that these mechanisms control such causal judgements (Allan, 1993). They also replicate findings reported by Richards and Hittesdorf (1978) regarding inhibitory properties accruing to cues associated with signalled reinforcers, when initially trained in contrast with immediately presented outcomes/reinforcers. These results stand in contrast to a signaled reinforcement procedure, in which a cue is presented prior to the delivery of reinforcement when that reinforcer becomes available (Lander, 1971; McCoy et al., 1976). The signal can be as long lasting in the two procedures, but the associated cues only accrue inhibitory properties when a delay is associated. As in the current study, these inhibitory cues were not as pronounced as those noted for unsignalled delay conditions.

General Discussion

The current series of experiments explored the effects of delayed outcomes on stimulus control of causal judgements using an interdimensional generalisation procedure (Guttman, 1959). The studies replicated previous findings regarding the impact of an outcome delay on causal ratings (Greville & Buehner, 2010; Reed, 1999, Young & Falmier, 2008a). They additionally noted that whether the cue appeared to become inhibitory or excitatory depended on the comparison component in training, replicating previous work from nonhuman instrumental conditioning studies (Richards & Hittesdorf, 1978; Richards & Marcattilio, 1978). The findings have some relevance to theoretical suggestions concerning the mechanisms underlying causal judgement (Allan, 1993; Gallistel & Gibbon, 2000; Miller & Barnet, 1993; Waldmann, 1996), especially those based on the view that similar mechanisms are at play in causal judgment paradigms as operate in instrumental learning procedures (Allan, 1993; Dickinson & Balleine, 2000).

Relative to an immediate outcome component, unsignalled outcome delays reduced ratings of causal effectiveness (Experiments 1 and 3). A signalled outcome delay similarly reduced ratings of causal effectiveness relative to an immediate outcome (Experiment 3), but improved ratings relative to the unsignalled delay (Experiments 2 and 3). Generalisation tests of the stimulus dimension cuing the unsignalled delayed outcome components demonstrated incremental (U shaped) generalisation gradients, when initial training had been in comparison to a component with immediate outcomes (Experiments 1 and 3), and with a signalled outcome delay (Experiment 2). However, relative to a component lacking outcomes, a decremental (inverted U shaped) generalisation curve (Experiment 1). For signalled outcome delays, incremental generalisation gradients were likewise found in comparison to immediate

outcomes (Experiment 3), but decremental gradients were found relative to unsignalled outcome delay components.

Thus, the current studies corroborate previous findings showing that when an outcome follows an action after a delay the relationship may not be detected (Greville & Buehner, 2010; Michotte, 1946), may be underestimated (Reed, 1999; van Elk et al., 2014), or may require additional information (instructions) to generate an impression of causality (Einhorn & Hogarth, 1986; Mendelson & Shultz, 1976; Schlottmann, 1999). Such effects of an unsignalled delay on causal judgements are similar to those noted in instrumental conditioning (Lattal, 2010; Reed & Reilly, 1990; Richards & Hittesdorf, 1978; Tarpy & Sawabini, 1974), which has prompted some to suggest they support a view that the mechanisms responsible for instrumental learning, notably the assumption of temporal contiguity being critical for learning in most instances (Boakes & Costa, 2014; Lattal, 2010), are also operational in procedures investigating human causal judgements (Allan, 1993).

However, such effects can also be explained by knowledge-based and probabilistic views (Einhorn & Hogarth, 1986; Waldmann, 1996), and rate-based theories (Gallistel & Gibbon, 2000), and have been taken by some to show the rules determining how an action-outcome relationship is encoded (see Buehner, 2005, for a discussion). For example, Waldmann (1996) suggests that contiguity alters the interpretation of events in a 2x2 contingency table governing causal judgments, involving the presence or absence of the response and outcome. A contiguous response-outcome pairing counts as an outcome given a response, but a delay places the instance as both a response without an outcome, and an outcome without a response. Thus, a delay shifts attribution from the response to other potential causes. Alternatively, Gallistel and Gibbon (2000) propose participants track rates of outcomes, and how this rate changes after responses are made, responding more when outcome rate increase after response rates increase (see also Baum, 1983, for a similar

analysis form an operant conditioning viewpoint). In this view, contiguity determines whether an outcome contributes to the rate conditional on the presence, or absence, of the cause. However, it should be noted that Gallistel and Balsam (2014) have questioned whether temporal contiguity should be interpreted in this manner. They suggest that, even with long cue-outcome delay, there is good learning about the relationship, but responding is displayed differently. Of course, this suggestion then leaves open the nature of the performance rules which may govern the expression of such learning, which could be viewed as vague.

Thus, both instrumental and cognitive accounts can accommodate the effects of delayed outcomes. There have been some attempts to untangle these predictions (e.g., Buehner & May, 2002; 2003, 2004). For example, Buehner and May (2003) presented participants with explicit instructions about potential delays, and noted that knowledge of potential delays improved assessments of delayed relations. Buehner (2005) suggests that this pattern of data cannot be accommodated by instrumental views, since judgements of the same event-outcome relationship altered according to higher level knowledge, which instrumental accounts fail to represent. Although it is not clear that this latter assumption is correct, it is clear that all theories can offer plausible accounts of the action of a response-outcome delay.

The current effect of inserting a stimulus between the action and the delayed outcome in ameliorating the detrimental impacts of the delay, and restoring levels of causal judgement (see also Reed, 1999; Young & Falmier, 2008b; Weller et al., 2017), also mirror those noted in instrumental learning (see Lattal, 1984; Reilly et al., 1996). These mechanisms may include the development of second-order conditioning/conditioned reinforcement, whereby the signal acquires the properties of the outcome with which it is associated (Williams, 1991). Other possibilities include the signal marking the preceding response, by making it more

discriminable or salient than other potential causes, and enhancing subsequent association with the outcome (Reed, 1999); or bridging, whereby the stimulus perceptually bridges the response and outcome across time, facilitating their direct association (Rescorla, 1982).

It is possible that the effects could be explained by cognitive views by making somewhat assumptions about the impact of delay signals on the perception and rules formed. For example, Waldman (1996) suggests a delayed response-outcome pairing might be evaluated as an instance of an outcome given a response, if there is some reason to assume a delayed causal relation. Presumably a delay signal serves this purpose. Eihorn and Hogarth (1986) suggest that relations between non-contiguous events can be judged causal in the presence of a chain (such as a delay signal) that bridges the temporal gap. How this differs from bridging hypotheses (see Rescorla, 1982), derived from conditioning studies, is unclear. This suggestion also entails an assumption that, when participants assume a causal mechanism necessitates a delay, immediate pairings will not be credited with causality. Thus, the role of non-contiguity in cognitive theories is more complex than in instrumental theories.

As the view that instrumental and causal learning are similar seems to provide a parsimonious account, in the sense of only requiring one set of mechanisms, it could be felt that this approach may be best suited to explaining such results. Further evidence regarding the similarities between the effects seen in causal judgement and instrumental learning came from a study reported by Lee et al. (2018), who noted that, in generalisation tests, causal efficacy ratings were higher in the presence of stimuli closest to an initial training cue. It should be noted that, although Lee et al. (2018) found results similar to those seen for nonhumans, their explanation of their results was quite different, and involved explaining peak shift as the result of combining multiple rule-based generalisations about the relationship between stimuli. In fact, many cognitive models have been used to explain

generalisation, such as Bayesian-based accounts (Lee, Mills, Hayes, & Livesey, 2021), and rule-based theories (Wong & Lovibond, 2017). The difficulty that seems apparent in these accounts is in specifying which rules are combined under which circumstance, as this appears to differ between participants, and makes the explain appear somewhat post hoc (Lee et al., 2018).

Evidence from interdimensional generalisation studies of instrumental learning suggests that the stimulus function and generalisation gradient are determined in comparison by alternative outcomes during training (Guttman, 1959; Richards & Marcattilio, 1978). Decremental generalisation gradients, with greatest responding at the target stimulus value and less responding with greater distance from that target, are taken to indicate excitatory control, and incremental generalisation gradients, with lowest responding at the target value) are taken to imply inhibitory stimulus control (Guttman, 1959; Jenkins, 1965). Guttman (1959) noted that stimuli associated with the weaker of two components on a multiple schedule (e.g., a schedule providing a lower frequency of reinforcement) served as an inhibitory stimulus. The current results with a human causal judgement procedure are similar to those noted previously for pigeons, extending he parallels, and being difficult for existent views of causal judgement not based on instrumental conditioning to explain in any straightforward manner.

The present experiments and previous studies with nonhumans have shown the generalisation gradient around a cue associated with a delayed outcome component depends on the training context (Richards & Hittesdorf, 1978; Richards & Marcattilio, 1978). If compared with a component lacking outcomes, a decremental gradient was obtained; if compared with immediate outcomes, an incremental gradient is noted. To the extent that these gradients reflect inhibitory (incremental gradients) and excitatory (decremental gradients) stimulus control (Jenkins, 1965; Rilling, 1977), a relativity view is supported

(Guttman, 1959; see also Miller & Matzel, 1988). When the outcome component is weaker schedule (i.e. compared to immediate outcomes or signalled delays) the cue associated with the component is inhibitory; in contrast, when it is stronger (i.e. relative to a component lacking outcomes), its correlated-stimulus is excitatory.

These results are similar to previous findings in other conditioning paradigms (Auge, 1974; Miller & Matzel, 1988). For instance, Auge (1974) noted a stimulus associated with an FI 30-s schedule could be a conditioned reinforcer in comparison to a leaner schedule (e.g., FR-200), but not in comparison to a richer schedule (FR-10). These results show the importance of the context in conditioning effects, which has also been noted for studies of delayed reinforcement (Reed & Reilly, 1990; Reilly et al., 1996); unsignalled delays contributing to the strength of the context and interfering with response learning (Reed & Reilly, 1990), while signalled delays remove this effect (Reilly et al., 1996). In this case, the signal would increase the associative strength of the response, and make its generalisation gradients, in comparison to weaker responses, decremental implying excitatory strength.

These effects may allow some views of causal judgement to explain effects of unsignalled and signalled delays by assigning the events to different cells of a probabilistic judgement matrix (see Buehner, 2005), or may be interpretable in terms of rate dependence (Gallistel & Gibbon, 2000; Miller & Barnet, 1993). The current effects on generalisation gradients, which parallel those already explored an explained for instrumental conditioning, would require additional assumption by these views, regarding the action of a stimulus on the rule or knowledge-base that produce judgements. These could, of course, be valid, but it would need further experimentation to determine if they could add anything to the already existent mechanisms for instrumental learning. It should be noted, however, that similarity of the empirical effects between instrumental conditioning and causal judgement studies, in this instance, does not provide incontrovertible evidence for the operation of parallel mechanisms.

There are instances where humans show different patterns of generalisation to nonhumans (e.g., Lee, Lovibond, & Hayes, 2019; Livesey & McLaren, 2009).

In addition to these concerns, a number of issues could be further explored based on these studies. Firstly, there may be a relationship between training performance and the subsequent shape of the generalisation gradients obtained at test. It may be that the target gradient at test (flat, incremental, or decremental) could be predicted based on the relative judgements in training. Unfortunately, due to the length of training, most participants in the current study had strongly discriminated performance at the end of training, with little variation, making the range of training performances unlikely to be suitable for teasing out any interesting effects in this regard. Future studies could examine this possibility by reducing the level of training provided. Secondly, the response scale employed in these experiments gave options ranging from "not effective at all" to "completely effective". This response scale was modelled on those previously used in studies of causal judgements involving delayed outcomes (e.g., Greville & Buehler, 2010). However, it might be that "completely effective" could imply either a positive (i.e. generating) relationship, or a negative (i.e. inhibitory) relationship. From the consistency of results obtained in the current studies across participants, it appears that this was not likely to be the case, or, at least, that all participants were using the scale in a similar way to each other with respect to these alternative interpretations. Nevertheless, there are other approaches to assessing such judgements; for example, Wasserman and Neunaber (1986) used a bipolar scale, which may be a better choice given the possibility of inhibitory relationships. Future studies could assess the impact of response scales on judgements.

In summary, the current results extend the range of circumstances in which instrumental and causal judgement procedures have been shown to produce the same types of effects. Although several theoretical suggestions have been made as to the mechanisms that

may underlie judgement effects, such as those based on associative learning, knowledge-based and probabilistic views, and rate-based theories, those based on the assumption that the two procedures reflect the operation of the same mechanisms seem to offer the most parsimonious account of these data.

References

Allan, L.G. (1993). Human contingency judgments: Rule based or associative? *Psychological Bulletin*, *114*(3), 435.

Anderson, J.R., & Sheu, C.F. (1995). Causal inferences as perceptual judgments. *Memory & Cognition*, *23*, 510-524

Auge, R.J. (1974). Context, observing behavior, and conditioned reinforcement. *Journal of the Experimental Analysis of Behavior*, *22*, 525-533.

Baum, W.M. (1983). Matching, statistics, and common sense. *Journal of the Experimental Analysis of Behavior*, *39*(3), 499.

Boakes, R.A., & Costa, D.S. (2014). Temporal contiguity in associative learning: Interference and decay from an historical perspective. *Journal of Experimental Psychology: Animal Learning and Cognition*, *40*(4), 381.

Buehner, M.J. (2005). Contiguity and covariation in human causal inference. *Learning & Behavior*, *33*(2), 230-238.

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*, 865–890.

Dack, C., McHugh, L., & Reed, P. (2009). Generalization of causal efficacy judgments after evaluative learning. *Learning & Behavior*, *37*(4), 336-348.

Dickinson, A., & Balleine, B.W. (2000). Causal cognition and goal-directed action.

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

Gallistel, C.R., & Balsam, P.D. (2014). Time to rethink the neural mechanisms of learning and memory. *Neurobiology of Learning and Memory*, *108*, 136-144.

Gallistel, C.R., & Gibbon, J. (2000). Time, rate, and conditioning. *Psychological Review*,

107, 289-344

Gallistel, C.R., Craig, A.R., & Shahan, T.A. (2019). Contingency, contiguity, and causality in conditioning: Applying information theory and Weber's Law to the assignment of credit problem. *Psychological Review*, *126*(5), 761.

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

Grillon, C., & Morgan, C.A., III. (1999). Fear-potentiated startle conditioning to explicit and contextual cues in Gulf War veterans with posttraumatic stress disorder. *Journal of Abnormal Psychology*, *108*, 134 –142

Gruber, H.E., Fink, C.D., & Damm, V. (1957). Effects of experience on perception of causality. *Journal of Experimental Psychology*, *53*, 89–93

Guttman, N. (1959). Generalization gradients around stimuli associated with different reinforcement schedules. *Journal of Experimental Psychology*, *58*, 335-340.

Guttman, N., & Kalish, H.I. (1956). Discriminability and stimulus generalization. *Journal of Experimental Psychology*, *51*, 79 – 88.

Jenkins, H.M. (1965). Generalization gradients and the concept of inhibition. In D. Mostofsky (Ed.), *Stimulus generalization*. Pp. 55-61. Palo Alto: Stanford University Press.

Lander, D.G. (1971). Stimulus control following response reduction with signalled reinforcement. *Psychonomic Science*, *23*(5), 365-366.

Lattal, K.A. (1984). Signal functions in delayed reinforcement. *Journal of the Experimental Analysis of Behavior*, *42*(2), 239-253.

Lattal, K.A. (2010). Delayed reinforcement of operant behavior. *Journal of the Experimental Analysis of Behavior*, *93*(1), 129-139.

Lee, J.C., Hayes, B.K., & Lovibond, P.F. (2018). Peak shift and rules in human

generalization. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *44*(12), 1955.

Lee, J.C., Lovibond, P.F., & Hayes, B.K. (2019). Evidential diversity increases generalisation in predictive learning. *Quarterly Journal of Experimental Psychology*, *72*(11), 2647-2657.

Lee, J.C., Mills, L., Hayes, B.K., & Livesey, E.J. (2021). Modelling generalisation gradients as augmented Gaussian functions. *Quarterly Journal of Experimental Psychology*, *74*(1), 106-121.

Lissek, S., Rabin, S., Heller, R.E., Lukenbaugh, D., Geraci, M., Pine, D.S., & Grillon, C. (2010). Overgeneralization of conditioned fear as a pathogenic marker of panic disorder. *The American Journal of Psychiatry*, *167*, 47–55.

Lissek, S., Kaczkurkin, A. N., Rabin, S., Geraci, M., Pine, D.S., & Grillon, C. (2014). Generalized anxiety disorder is associated with overgeneralization of classically conditioned fear. *Biological Psychiatry*, *75*(11), 909-915.

Livesey, E.J., & McLaren, I.P. (2009). Discrimination and generalization along a simple dimension: peak shift and rule-governed responding. *Journal of Experimental Psychology: Animal Behavior Processes*, *35*(4), 554.

McCoy, J.F., Parker, F.C., & McGee, G. (1976). Signalled reinforcement and dimensional stimulus control. *The Journal of Psychology*, *93*(1), 121-128.

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

Michotte, A.E. (1963/1946). *The perception of causality* (T.R. Miles, trans.). London: Methuen.

Miller, R.R., & Barnet, R.C. (1993). The role of time in elementary associations. *Current*

Directions in Psychological Science, *2*, 106- 111.

Miller, R.R., & Matzel, L.D. (1988). The comparator hypothesis: A response rule for the expression of associations. In *Psychology of learning and motivation* (Vol. 22, pp. 51-92). Academic Press.

Morey, R.A., Dunsmoor, J.E., Haswell, C.C., Brown, V.M., Vora, A., Weiner, J., . . . the VA Mid-Atlantic MIRECC Workgroup. (2015). Fear learning circuitry is biased toward generalization of fear associations in posttraumatic stress disorder. *Translational Psychiatry*, *5*, e700.

Reed, P. (1992). Effect of a signalled delay between an action and outcome on human judgement of causality. *Quarterly Journal of Experimental Psychology*, *44B*, 81–100

Reed, P. (1993). Influence of the schedule of outcome presentation on causality judgements. *The Quarterly Journal of Experimental Psychology Section A*, *46*(2), 327-345.

Reed, P. (1999). Role of a stimulus filling an action-outcome delay in human judgments of causal effectiveness. *Journal of Experimental Psychology: Animal Behavior Processes*, *25(1)*, 92–102.

Reed, P. (2001). Schedules of reinforcement as determinants of human causality judgments and response rates. *Journal of Experimental Psychology: Animal Behavior Processes*, *27*(3), 187.

Reed, P., & Reilly, S. (1990). Context extinction following conditioning with delayed reward enhances subsequent instrumental responding. *Journal of Experimental Psychology: Animal Behavior Processes*, *16*(1), 48.

Reilly, S., Schachtman, T.R., & Reed, P. (1996). Signaled delay of reinforcement: Effects of postconditioning manipulation of context associative strength on instrumental performance. *Learning and Motivation*, *27*(4), 451-463.

Renner, K.E. (1964). Delay of reinforcement: A historical review. *Psychological Bulletin*,

61(5), 341.

Rescorla, R.A. (1982). Effect of a stimulus intervening between CS and US in autoshaping. *Journal of Experimental Psychology: Animal Behavior Processes*, *8(2)*, 131–141.

Richards, R.W. (1974). Inhibitory stimulus control and the magnitude of delayed reinforcement. *Journal of the Experimental Analysis of Behavior*, *21*, 501-509.

Richards, R.W. (1975). Reinforcement delay, behavioral contrast, and inhibition. *Psychological Record*, *25*, 281-291.

Richards, R.W., & Hittesdorf, W.M. (1978). Inhibitory stimulus control under conditions of signalled and unsignalled delay of reinforcement. *The Psychological Record*, *28*(4), 615-625.

Richards, R.W., & Marcattilio, A.J. (1978). Stimulus control and delayed reinforcement. *Learning and Motivation*, *9*(1), 54-68.

Rilling, M. (1977). Stimulus control and inhibitory processes. In W.K. Honig & J.E.R. Staddon (Eds.), *Handbook of operant behavior*. pp. 432-480. Englewood Cliffs, N. J.: Prentice Hall.

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

Shepard, R.N. (1987). Toward a universal law of generalization for psychological science. *Science*, *237*, 1317–1323.

Siegler, R.S., & Liebert, R.M. (1974). Effects of contiguity, regularity, and age on children's causal inferences. *Developmental Psychology*, *10*, 574-579

Tarpy, R.M., & Sawabini, F.L. (1974). Reinforcement delay: A selective review of the last decade. *Psychological Bulletin*, *81*(12), 984.

van Elk, M., Salomon, R., Kannape, O., & Blanke, O. (2014). Suppression of the N1

auditory evoked potential for sounds generated by the upper and lower limbs. *Biological Psychology*, *102*, 108–117.

Waldmann, M.R. (1996). Knowledge-based causal induction. *Psychology of Learning and Motivation*, *34*, 47-88.

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*, 15-35

Weller, L., Schwarz, K.A., Kunde, W., & Pfister, R. (2017). Was it me? – Filling the interval between action and effects increases agency but not sensory attenuation. *Biological Psychology*, *123*, 241-249.

Williams, B.A. (1991). Marking and bridging versus conditioned reinforcement. *Animal Learning & Behavior*, *19*(3), 264-269.

Wong, A.H., & Lovibond, P.F. (2017). Rule-based generalisation in single-cue and differential fear conditioning in humans. *Biological Psychology*, *129*, 111-120.

Young, M.E., & Falmier, O. (2008a). Color change as a causal agent: Revisited. *American Journal of Psychology*, *121*, 129–157

Young, M.E., & Falmier, O. (2008b). Launching at a distance: The effect of spatial markers. *Quarterly Journal of Experimental Psychology*, *61*, 1356–1370.

Young, M.E., & Nguyen, N. (2009). The problem of delayed causation in a video game: Constant, varied, and filled delays. *Learning and Motivation*, *40*(3), 298-312.

Table 1: Designs for the pretraining and training phases of the three experiments. .5 = probability of an outcome given a response

	Pretraining	Training
Experiment 1	.5 immediate outcome	.5 immediate outcome; .5 3s fixed delay .5 immediate outcome; .5 3s variable delay No outcome; .5 3s fixed delay
	Pretraining	**Training**
Experiment 2	-	.5 3s signalled delay; .5 3s unsignalled delay .5 3s unsignalled delay; .5 3s signalled delay
	Pretraining	**Training**
Experiment 3	.5 immediate outcome	.5 immediate outcome; .5 3s unsignalled delay .5 immediate outcome, .5 3s signalled delay .5 immediate outcome, .5 immediate outcome

Table 2: Experiment 1. Group-mean (standard deviation) for response and outcomes per min, experienced action-outcome delay, and experienced delays and response rate during the delay period for the delayed components, averaged over the last three exposures to the components, for each component. .5 = pO/R; FD = 3s fixed delay; VD = 3s variable delay; Ext = lacking outcomes; Comparison = comparison component; Target = target component.

	.5, .5FD		**.5, .5VD**		**Ext, .5FD**	
	Comparison	Target	Comparison	Target	Comparison	Target
Resp rate (resp/min)	20.8 (7.7)	19.8 (9.3)	22.4 (8.6)	19.9 (7.6)	21.3 (7.0)	21.3 (7.8)
Out rate (out/min)	10.4 (3.8)	9.9 (4.7)	11.2 (4.3)	9.9 (3.8)	0	10.6 (3.9)
pO/R	.5 (.1)	.5 (.1)	.5 (.1)	.5 (.1)	0	.5 (.1)
Delay (s)		2.6 (.6)		2.7 (.5)		2.5 (.6)
Delay rate (resp/min)		19.8 (9.5)		19.8 (7.6)		21.1 (.79)

Table 3: Experiment 2. Group-mean (standard deviation) for response and outcomes per min, experienced action-outcome delay, and experienced delays and response rate during the delay period, averaged over the last three exposures to the components, for each component. .5 = pO/R; D = 3s fixed delay; S = 3s signalled delay; Comparison = comparison component; Target = target component.

	.5S, .5U		**.5U, .5S**	
	Comparison	Target	Comparison	Target
Resp rate (resp/min)	19.0 (8.5)	19.0 (6.5)	20.2 (5.8)	18.4 (5.8)
Out rate (out/min)	7.1 (3.7)	6.4 (3.0)	7.2 (2.4)	5.9 (2.7)
pO/R	.3 (.2)	.3 (.1)	.4 (.1)	.3 (.1)
Delay (s)	2.7 (.3)	2.7 (.3)	2.8 (.4)	2.8 (.3)
Delay rate (resp/min)	20.2 (5.8)	18.7 (6.8)	18.9 (8.5)	17.9 (6.3)

Table 4: Experiment 3. Group-mean (standard deviation) for response and outcomes per min, experienced action-outcome delay, experienced action-outcome delay, and experienced delays and response rate during the delay period for the delay components, averaged over the last three exposures to the components, for each component. .5 = pO/R; FD = 3s fixed delay; VD = 3s variable delay; Ext = lacking outcomes; Comparison = comparison component; Target = target component

	.5, .5U		**.5, .5S**		**.5, .5**	
	Comparison	Target	Comparison	Target	Comparison	Target
Resp rate (resp/min)	20.1 (6.5)	18.4 (5.0)	19.1 (5.6)	18.5 (4.4)	22.0 (6.0)	21.2 (8.2)
Out rate (out/min)	10.0 (3.5)	9.4 (2.7)	9.3 (2.8)	9.3 (2.6)	11.3 (2.8)	10.6 (4.1)
PO/R	.5 (.1)	.5 (.1)	.5 (.1)	.5 (.1)	.5 (.1)	.5 (.03)
Delay (s)		2.8 (.3)		2.8 (.3)		0 (0)
Delay rate (resp/min)		18.5 (5.9)		18.8 (5.2)		0

Figure 1: Experiment 1. Mean judgements of causality after training. .5 = pO/R; FD = 3s fixed delay; VD = 3s variable delay; Ext = lacking outcomes; Comparison = comparison component; Target = target component. Error bars = 95% confidence limits.

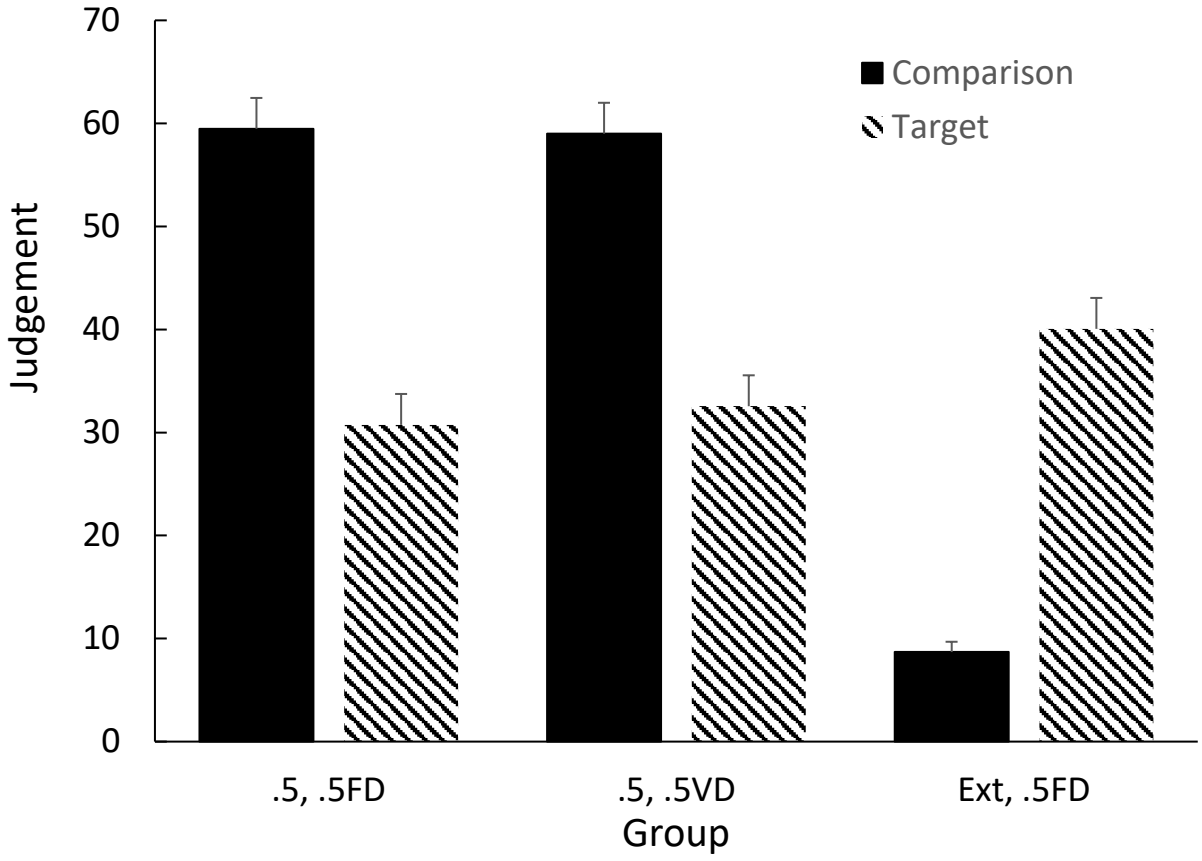

Figure 2: Experiment 1. Group-mean judgments for each of the stimuli used in the generalisation test (S4 was the original training stimulus). .5 = pO/R; FD = 3s fixed delay; VD = 3s variable delay. Error bars = 95% confidence limits.

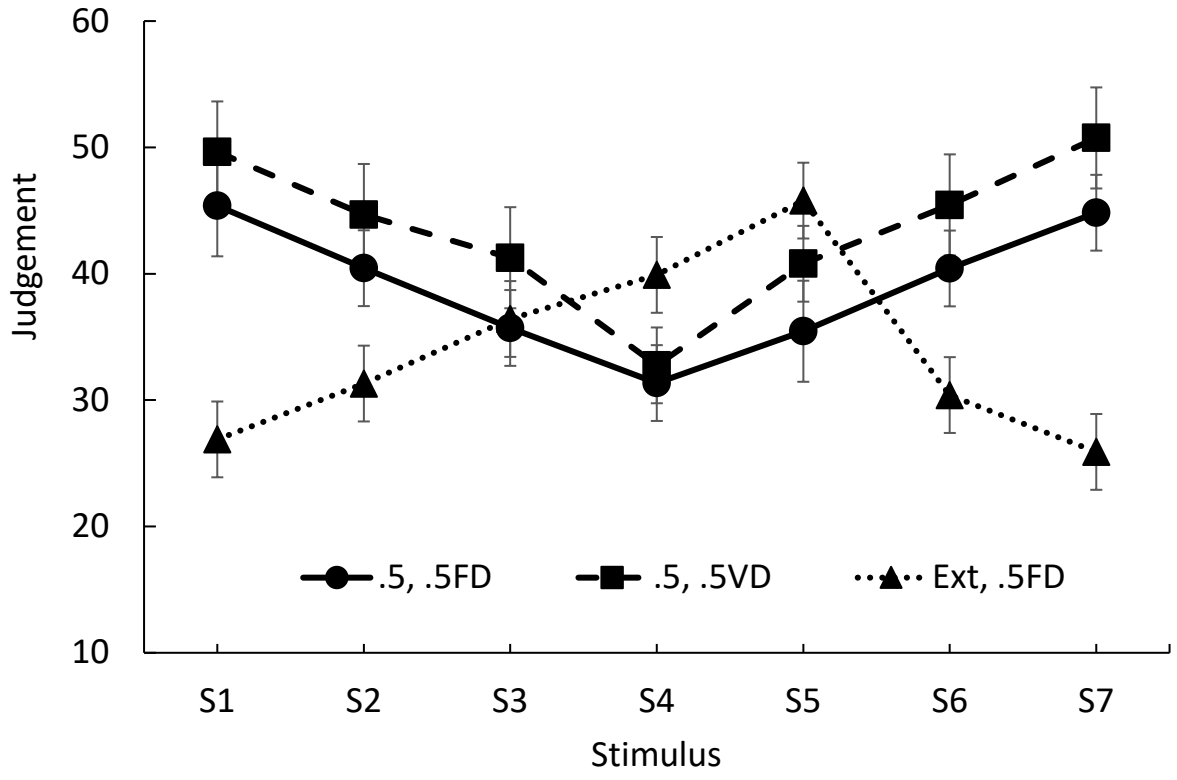

Figure 3: Experiment 2. Mean judgements of causality after training. .5 = pO/R; U = 3s unsignalled delay; S = 3s signalled delay; Comparison = comparison component; Target = target component. Error bars = 95% confidence limits.

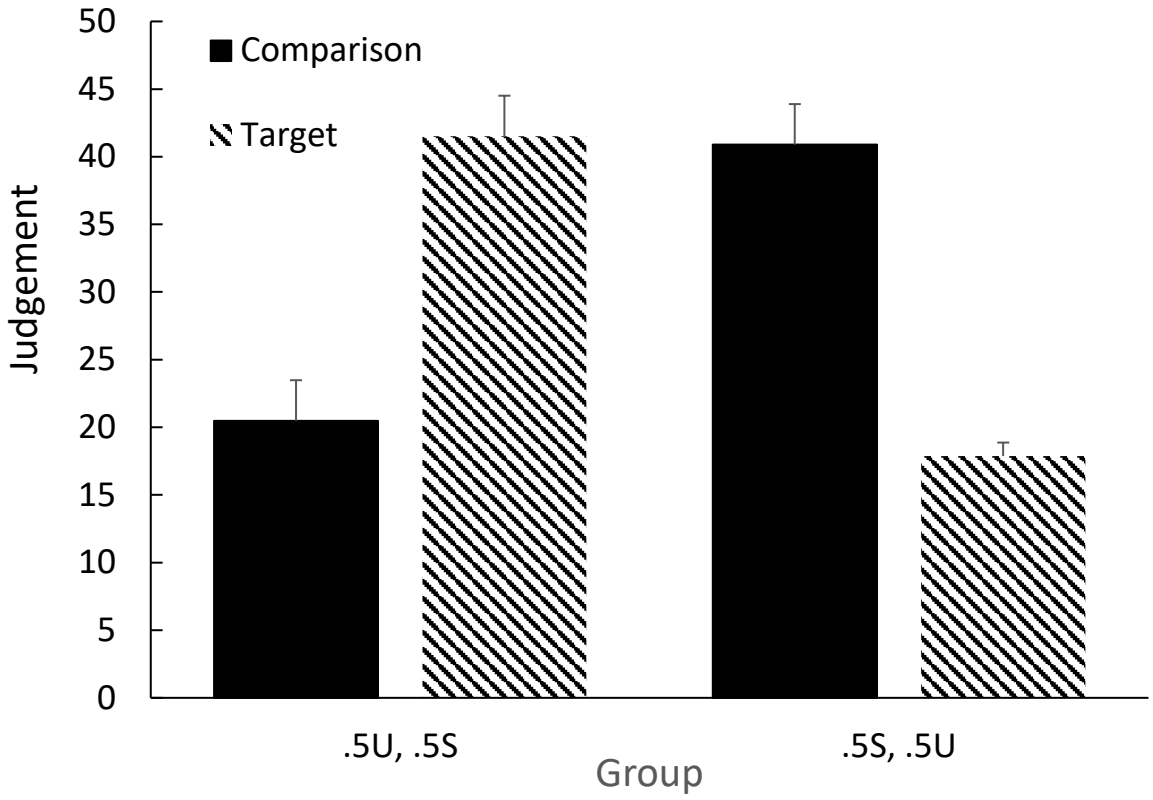

Figure 4: Experiment 2. Group-mean judgments for each of the stimuli used in the generalisation test (S4 was the original training stimulus). .5 = pO/R; U = 3s unsignalled delay; S = 3s signalled delay. Error bars = 95% confidence limits.

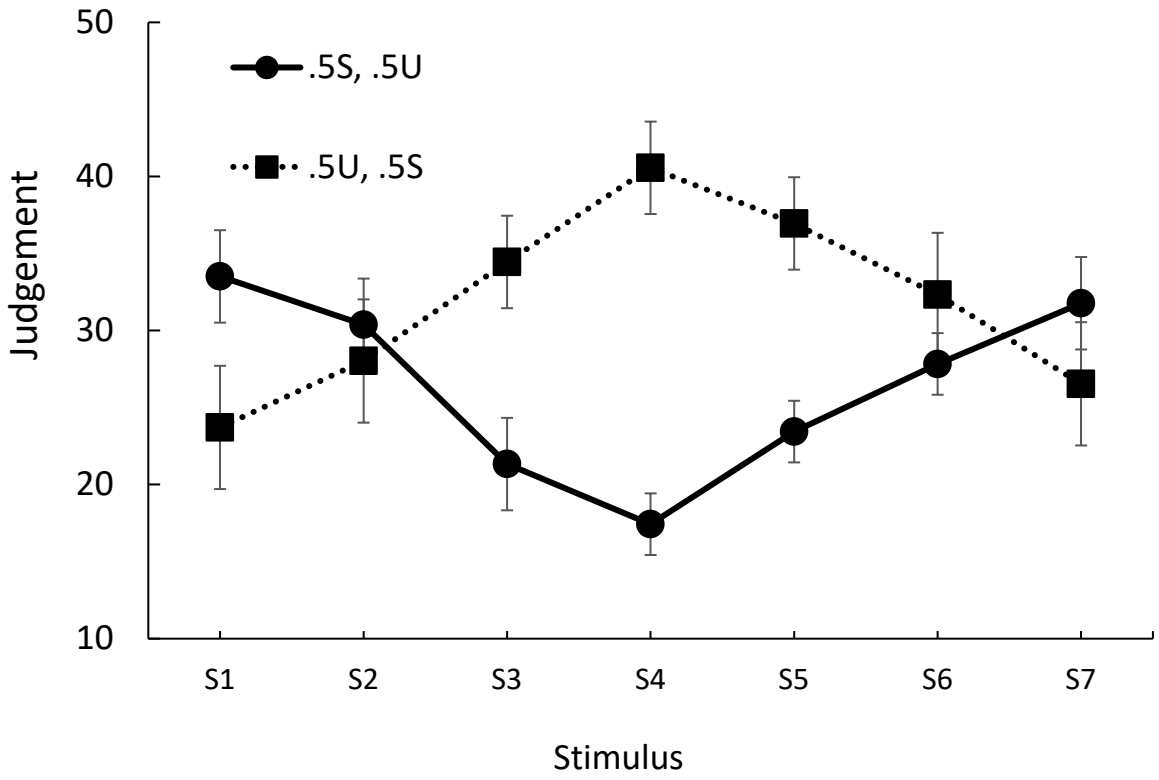

Figure 5: Experiment 3. Mean judgements of causality after training. .5 = pO/R; U = 3s unsignalled delay; S = 3s signalled delay; Comparison = comparison component; Target = target component. Error bars = 95% confidence limits.

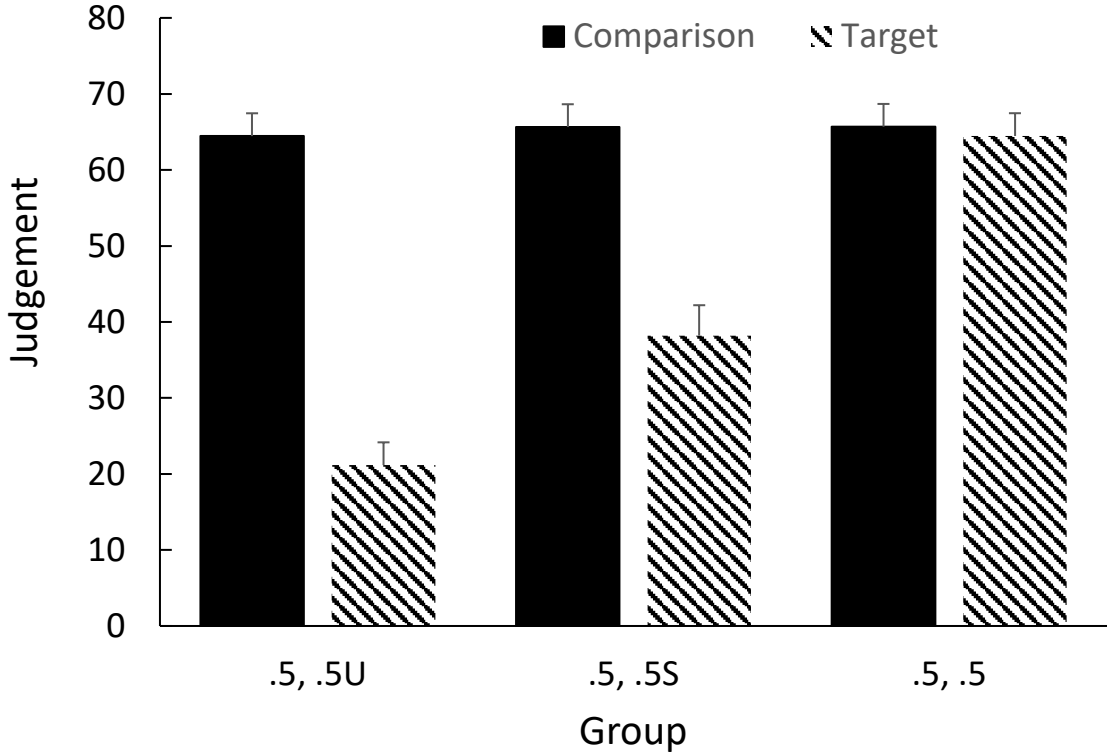

Figure 6: Experiment 3. Group-mean judgments for each of the stimuli used in the generalisation test (S4 was the original training stimulus). .5 = pO/R; U = 3s unsignalled delay; S = 3s signalled delay. Error bars = 95% confidence limits.

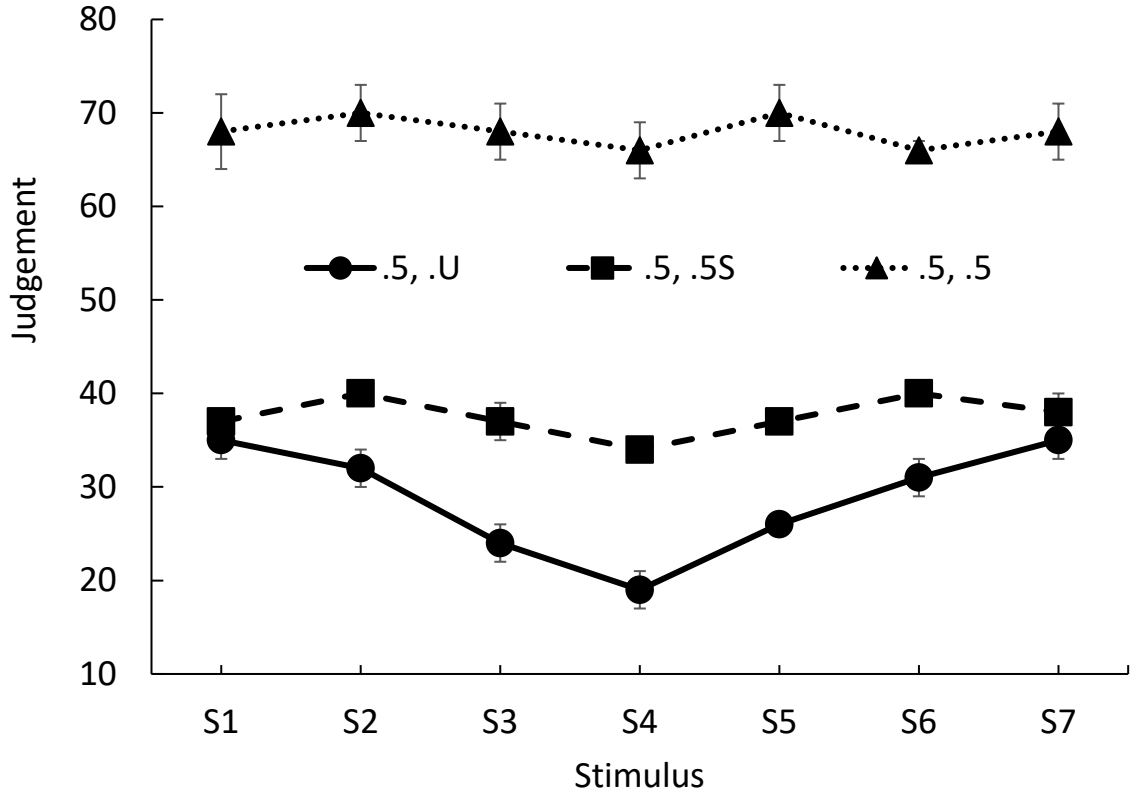

Appendix 1

Screenshot of the stimuli variables used to create the coloured stimuli using the HSB system from the website: HSB Color Picker (codepen.io).