

## **UC Merced**

### **Proceedings of the Annual Meeting of the Cognitive Science Society**

#### **Title**

More means more? Illusory causation between uncorrelated continuous events

#### **Permalink**

<https://escholarship.org/uc/item/67p779b9>

#### **Journal**

Proceedings of the Annual Meeting of the Cognitive Science Society, 45(45)

#### **Authors**

Liang, Lexin  
Chow, Julie YL  
Livesey, Evan

#### **Publication Date**

2023

Peer reviewed

# More means more? Illusory causation between uncorrelated continuous events

**Lexin Liang (llia6082@uni.sydney.edu.au)**  
School of Psychology, University of Sydney, Sydney  
NSW, Australia

**Julie Y.L. Chow (julie.chow@unsw.edu.au)**  
University of New South Wales, Sydney, NSW, Australia

**Evan J. Livesey (evan.livesey@sydney.edu.au)**  
School of Psychology, University of Sydney, Sydney  
NSW, Australia

## Abstract

Illusions of causality arise when people observe statistically unrelated events and yet form a belief that the events are causally linked. When participants observe a sequence of discrete binary events (e.g., a patient was either administered a treatment or no treatment, and subsequently recovers or does not recover from their illness), the frequency of the putative cause and outcome occurring inflates the illusion of causality. Recently, similar effects have been observed using outcomes of continuous magnitude. Participants are more likely to endorse the causal status of a (completely ineffective) cue if the target outcome (e.g., high magnitude outcomes) occur frequently. Here, we extended these findings by investigating how predictions and causal judgments for a cue of continuous magnitude were affected by the distribution of cue values presented. Participants observed cue values (dose of a fictitious medicine) sourced from either a continuous distribution or from two discrete values, and were followed by outcomes that were either continuous (Experiment 1) or binary in nature (Experiment 2). Our results show that participants were more likely to assume a linear relationship between drug dose and magnitude of recovery when cue dosage were predominantly high than when they were predominantly low.

**Keywords:** Causal learning; contingency learning; illusion of causality, cue density effect, illusory correlation

## The illusion of causality

Learning the relationship between potential causes and their putative effects is an important part of everyday life and an adaptive strategy for human survival (Blanco & Matute, 2018; Chow et al., 2021). We can use causal knowledge obtained from previous learning to guide our behaviors, in order to obtain preferred outcomes and avoid undesirable ones. However, a significant body of previous research has demonstrated that under certain conditions, we tend to perceive a causal relationship when none exists (Yarritu et al., 2014; Matute et al., 2015; Chow et al., 2019). This causal illusion can be innocuous or even adaptive in some situations; believing in one's ability to control events could make people feel hopeful and motivated in the face of uncertainty (Blanco, 2017). In many contexts, this kind of "false alarm" is much less costly than missing a relationship when it truly exists (Blanco, 2017; Blanco & Matute, 2018). However, the costs of believing in a non-existent relationship can be substantial,

for instance when individuals forgo effective medical treatment because of their belief in an ineffective alternative, such as homeopathy (Blanco et al., 2014).

Contingency learning is a key method that researchers use to study causal illusion. The most commonly used paradigm of contingency learning involves two events, a cue and an outcome. The two stimuli are binary in the sense that they are either present or absent, and different combinations of two stimuli can give rise to a 2×2 contingency matrix as summarized in Table 1. A widely-used normative index of contingency  $\Delta P$  is calculated by the following formula:  $\Delta P = P(O|C) - P(O|\sim C) = a/(a+b) - c/(c+d)$ , where a-d refer to the frequencies with which each of the 4 cell types occur. The  $\Delta P$  metric refers to the difference between the probability that the outcome occurs when the cue is present and the probability that the outcome occurs when the cue is absent. When the probability of outcome occurrence is the same regardless of cue occurrence (i.e.,  $\Delta P = 0$ ), the contingency between the two events is null.

Table 1: Contingency matrix containing four possible cue-outcome combinations.

	Outcome present	Outcome absent
Cue present	<i>a</i>	<i>b</i>
Cue absent	<i>c</i>	<i>d</i>

Previous literature has found that two manipulations consistently produce causal overestimation under a zero contingency. First, causal illusion tends to occur when the probability that the outcome occurs is high regardless of the cue's occurrence (i.e., increasing trial types *a* and *c*), a phenomenon called the outcome density (OD) effect (Matute et al., 2015; Chow et al., 2019). The OD effect potentially contributes to the popularity of pseudo-medicines which are usually claimed to treat mild diseases with high rates of spontaneous remission (Blanco et al., 2014). Second, causal illusion also tends to occur if the cue occurs more frequently regardless of outcome's occurrence (i.e., increasing trial types *a* and *b*), a phenomenon referred to as the cue density effect (CD effect) (Yarritu et al., 2014; Matute et al., 2015). Since many alternative and complementary medicines are

advertised as natural and free from side effects, they tend to be frequently used without concern, further contributing to the formation of causal illusion (Blanco et al., 2014).

Although these two effects have been replicated frequently in laboratory settings, it remains unknown to what extent they apply to everyday situations in which events are not neatly organized into binary categories. For instance, the possible effect of taking a different dose of a drug is difficult to capture by considering only the presence versus absence of treatment. Similarly, partial recoveries that commonly occur are not always well-described as either the presence or absence of the desired outcome. Therefore, how these ambiguous stimuli are perceived and used in forming beliefs is potentially informative for understanding real-world causal reasoning.

The problem of how people might use ambiguous stimuli in causal reasoning is often overlooked; dominant statistical models of contingency learning like  $\Delta p$ , assume that cue and outcome events can be neatly dichotomized as present or absent without consideration of how this is achieved in a variable environment. Several studies (e.g. Chow et al., 2019; Double et al., 2020) have found equivalent outcome density effects for binary and continuous outcomes, suggesting that it is not necessary for the effect to be neatly categorized as present or absent in order for causal illusions to flourish, at least when the putative cause is dichotomous and thus obviously present or absent.

Indeed, when only one of the events is continuous, there is some evidence that participants assimilate ambiguous stimuli in the direction of their causal hypothesis (Marsh & Ahn, 2009; Blanco, Moreno-Fernández & Matute, 2020). For example, if one believed a protein (outcome) was caused by tall bacteria (cue), and the protein was also shown to be consistently caused by intermediate-length bacteria (ambiguous cue), the intermediate-length bacteria tend to be spontaneously categorized as tall (Marsh & Ahn, 2009). This causal assimilation process inflates the perceived contingency between cue and outcome by increasing the number of a-cell trials. These findings are important as they suggest some flexibility in how ambiguous stimuli are interpreted, relative to the learner's causal model about the cue-outcome relationship. However, even in work on causal assimilation, the putative cause has not been completely continuous, but rather could be classified easily into three distinct categories (e.g., tall, short and intermediate bacteria). The question thus remains how do people learn about causal relationships when the cue—that is, the possible cause of the effect—varies in a completely continuous fashion.

There is evidence from the function learning literature that people tend to assume a linear monotonic relationship between events when cue and outcome events are continuous (Summers, Summers & Karkau, 1969). That is, when presented with some continuous cue dimension, participants form a causal hypothesis that increasing the cue value will lead to an increase in the outcome value. Applying this logic to the illusion of causality, if we hold constant the frequency of high magnitude outcomes, increasing the frequency of high magnitude cues (e.g., high doses of a drug) should lead

to greater illusions of causality, in line with the cue density effect. Although this would be consistent with demonstrations of the outcome density effect that use continuous outcomes (Chow et al., 2019; Double et al., 2020), it remains to be seen whether continuous cues support causal illusions in the same way. The primacy of the cue information, typically delivered before the learner makes a prediction about the likely effect, and its distinct role in mental models of cause and effect are reason enough to question whether people will process ambiguous cue values in the same way as ambiguous outcome values.

The current study aimed to investigate causal illusion and CD effect using continuous cues, presented as a possible cause of a continuous outcome (Experiment 1) or discrete and binary outcome (Experiment 2). We adopted a fictitious medical scenario similar to Chow et al.'s (2019) study, where participants were instructed to learn the efficacy of a fictitious (and ineffective) medicine, Calciucor, that did not correlate with a reported medical outcome, an increase in bone density in patients with Osteoporosis.

## Experiments 1 and 2

Participants in both experiments first learned 100 trials where a certain dosage was prescribed to a patient, and the patient recovered to a certain extent. A  $2$  (CD: high vs. low)  $\times 2$  (cue variability: discrete vs. continuous) between-subject design was adopted. The continuous cue and outcome distributions are illustrated in Figure 1. For outcomes, we used a high OD distribution since the CD effect had been shown to be more robust when outcome density is high (Blanco et al., 2013). In Experiment 1, this entailed a continuous negatively-skewed distribution ranging from 0-100, with larger values as the modal outcome. For Experiment 2, we used a discrete outcome that was either present or absent, but which was present on most trials.

To assess the results of causal learning, participants were asked to give two efficacy ratings. They were asked to rate the effect of the medicine (treatment rating) as well as the effect of increasing the dose (dosage rating) on treatment outcome. It should be noted that since we did not present cue-absent trials (i.e. in all instances, the patient received at least *some* medicine), answers for the treatment rating largely depend on participants' assumptions about spontaneous recovery rate of the disease, which the participant had no way of verifying during the experiment. Thus, we mainly focused on the dosage rating to analyse causal illusions. As the OD effect has been shown with both discrete and continuous outcomes (Chow et al., 2019; Double et al., 2020), we also expected to see a CD effect using discrete and continuous cues. Similarly, as outcome ambiguity did not produce systematic bias on the magnitude of the OD effect in previous work, we predicted there to be no interaction effect between cue variability and CD. In addition, participants made prediction ratings for recovery across 11 cue values (i.e., 0, 10, 20...100), and the linear trends of predicted outcomes as a function of the magnitude of the cue were used as complementary measurements to efficacy ratings. Consistent

with a CD effect, we hypothesized that the cue magnitude (i.e. treatment dose) would have a stronger positive effect on participants' predictions when they observed predominantly high doses of the medicine being administered (i.e. a high cue density).




Figure 1: Scatterplots showing all cue and outcome values presented to participants in the continuous low cue density condition (left panel) and continuous high cue density condition (right panel) of Experiment 1. Histograms display cumulative frequencies of cue and outcome values for all participants. The black line represents the best-fitting regression line, showing the zero correlation between cue and outcome values.

## Method

**Participants** A total of 162 undergraduate students in Experiment 1 and 163 undergraduate students in Experiment 2 completed the study online for partial course credit. In each experiment, participants were unsystematically allocated to one of four conditions. 11 participants in Experiment 1 and 12 participants in Experiment 2 were excluded from the data analyses because they took more than three times to pass the instruction understanding test. A further participant in Experiment 1 was excluded as they admitted writing down information while completing the study. The remaining 150 participants in Experiment 1 (110 female,  $M_{age}=20.68$  years) and 150 participants in Experiment 2 (97 female,  $M_{age}=19.68$  years) were distributed across the groups as follows. Experiment 1: 37 participants in the Low CD-Discrete Cue group, 35 in the Low CD-Continuous Cue group, 36 in the High CD-Discrete Cue group, and 42 in the High CD-Continuous Cue group. Experiment 2: 40 participants in the Low CD-Discrete Cue group, 37 in the Low CD-Continuous Cue group, 36 in the High CD-Discrete Cue group, and 37 in the High CD-Continuous Cue group.

**Design** The study used a 2 (CD: high vs. low)  $\times$  2 (cue variability: discrete vs. continuous) between-subject design. The distributions of cue and outcome values used in the continuous conditions of Experiment 1 are shown in Figure 1. In the continuous-cue conditions, participants observed cues sampled either from a positively skewed (low CD condition) or a negatively skewed (high CD condition) unimodal distribution. For the low CD condition, the sample of cues was created by an exponentially modified Gaussian distribution with a higher proportion of low cue values (distribution parameters:  $\mu=10$ ,  $\sigma=5$ ,  $\tau=25$ , and range=1-99,

yielding sample mean=32). For the high CD condition, a negatively skewed distribution was created by taking the complement of the low CD distribution (i.e.  $100 - C$ ), so that it contained a higher proportion of high cue values (sample mean = 68). Cue values were randomly generated from this distribution, with a further constraint that cue values in 80% of trials were below 50 in low CD condition, whereas cue values in 20% of trials were below 50 in high CD condition.

In the discrete cue conditions, participants were presented with the exact value of a lower cue and a higher cue. The two cue values were generated by taking the average of cue values below and above 50 in their corresponding CD distribution in continuous cue conditions. For the low CD condition, the lower-value cue of 24 occurred in 80% of trials and the higher-value cue of 66 occurred in 20% of trials. For the high CD condition, the low-value cue of 34 occurred in 20% of trials and the higher-value cue of 76 occurred in 80% of trials. In Experiment 1, all four conditions used an outcome distribution that was the same as the cue distribution in High CD-Continuous Cue condition. The cues and outcomes were paired randomly with the constraint that the correlation between cue and outcome values across all 100 training trials was zero. In Experiment 2, 80 outcome-present trials and 20 outcome-absent trials were shown. Again, assignment of outcome to cue value was randomized with the constraint that cue magnitude did not predict outcome occurrence.

**Stimuli and apparatus** Examples of the stimuli, ratings and prediction screens can be seen in Figure 2. The cue value was presented as the dose of medicine administered to a patient. In each training trial, a medicine bottle appeared in the center of the screen with a certain volume of medicine inside, and the cue value was shown below the bottle in milliliters (as shown in figure 2). In Experiment 1, trial-by-trial predictions about the magnitude of the outcome were made on a sliding scale from 0 to 100, and outcomes were subsequently presented as a red bar on a continuous scale from 0 to 100. In Experiment 2, participants rated the likelihood of the outcome occurring on a scale from 0 to 100 and the outcome presence versus absence was presented in text as either "Full recovery (normal bone density)" or "No change (low bone density)". In the testing stage, two efficacy ratings were made on a sliding scale from -100 (strong negative effect) to 100 (strong positive effect), with the midpoint 0 representing that the drug is completely ineffective (see Figure 2). The prediction ratings were made on a sliding scale from 0 to 100 the same as trial-by-trial predictions during training.

**Procedures** At the beginning of the experiment, participants were asked to imagine they were a medical researcher studying a new treatment for Osteoporosis. The new treatment Calciucor was designed to increase bone density in patients, but its efficacy and side effects remained unknown. Their objective in this study was to observe 100 patients who received a certain dose of Calciucor and recovered (to varying degrees in Experiment 1), and make judgements about the treatment's effectiveness.

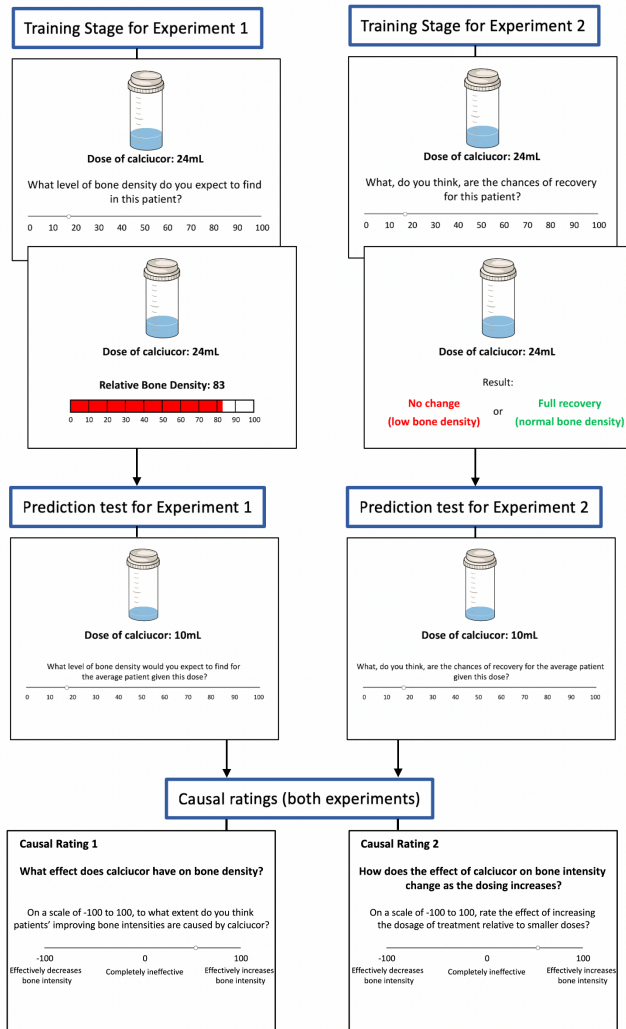


Figure 2: Schematic of the events observed by, and responses required from, participants in the training and test phases of Experiments 1 and 2.

In each training trial, participants were presented with a certain dose of medicine prescribed to a new patient. The dose of medicine was indicated by the volume of liquid inside the bottle and the number below. A sliding scale from 0 to 100 was presented below this cue stimulus, and participants were asked to predict the patient's extent of recovery by moving the slider to their desired location. Once a prediction was made, participants clicked continue and their prediction on the sliding scale was replaced by a horizontal bar showing the patient's actual recovery. This procedure was repeated so that participants learned 100 cue-outcome pairs in total.

During the testing stage, participants were first instructed to rate how effective Calciucor was in increasing bone density on a scale from -100 to 100 for the treatment rating. This was followed by the dosage rating asking what effect increasing the dose had on bone density, which was also answered on a scale from -100 to 100. Having made two efficacy ratings, participants were then presented with 11 prediction questions, where a certain dose of Calciucor

appeared on the screen followed by instructions to predict how much an average patient would recover if they were administered this dose. The doses in prediction questions range from 0 to 100 with an interval of 10 (i.e. 0, 10, 20...100), and the sequence of their presence was randomised. The study procedure is depicted in figure 2.

## Results

**Predictions during training** We first extracted linear functions of the effect of treatment dosage on prediction ratings made during training. A more positive slope indicates better predicted recovery for higher dose (i.e., stronger causal illusion). Figure 3 shows the slopes of training predictions as a function a treatment dosage in two experiments, separately for the discrete and continuous cue group.

These slopes were then analysed with a 2 (CD: low vs. high)  $\times$  2 (cue variability: discrete vs. continuous) between-subject ANOVA in two experiments. In Experiment 1, there was a significant effect of CD on slopes, with more positive slopes in the high-CD group ( $M = .22, SD = .33$ ) than the low-CD group ( $M = -.02, SD = .30, F(1,146) = 21.15, p < .001, \eta_p^2 = .13$ ), indicating stronger causal illusion in the high-CD group (i.e., the presence of CD effect). This effect of CD did not interact with cue variability ( $F < 1$ ), suggesting that CD effect was equivalent in discrete and continuous cue groups.




Figure 3: Slope estimates derived from participant ratings as a linear function of treatment dose during the training stages of Experiments 1 and 2.

Similarly, a CD effect was also found in Experiment 2, such that slope was significantly more positive in the high-CD group ( $M = .39, SD = .37$ ) than the low-CD group ( $M = .16, SD = .46, F(1,146) = 11.03, p = .001, \eta_p^2 = .07$ ). In addition, there was a significant effect of cue variability, such that the slopes were more positive in the continuous cue group ( $M = .42, SD = .41$ ) than the discrete cue group ( $M = .13, SD = .40, F(1,146) = 19.53, p < .001, \eta_p^2 = .11$ ). However, the CD effect in Experiment 2 did not interact with the effect of cue variability,  $F < 1$ . Thus, in both experiments, training predictions showed equivalent levels of CD effect in the discrete cue group and the continuous cue group.

**Prediction test ratings** Test predictions were analysed in two ways. First, predictions for 0mL dosage were separated



from predictions for non-zero dosages (10mL-100mL), and a mean prediction for non-zero dosages was calculated for each participant. A 2 (dosage: 0mL vs. >0mL) × 2 (CD: low vs. high) × 2 (cue variability: discrete vs. continuous) mixed measures ANOVA was conducted on test predictions.

In Experiment 1, there was a main effect of dosage on predictions, where predictions were significantly higher for >0mL dosages ( $M = 58.12, SD = 10.67$ ) than 0mL dosage ( $M = 44.13, SD = 33.10, F(1,146) = 23.95, p < .001, \eta^2_p = .14$ ). A main effect of cue variability was also found in Experiment 1, such that predictions were significantly higher in continuous cue group ( $M = 54.27, SD = 24.51$ ) than the discrete cue group ( $M = 47.80, SD = 26.24, F(1,146) = 5.33, p = .02, \eta^2_p = .04$ ). However, no 2-way or 3-way interactions were significant, including dosage by CD ( $F_s < 2.11$ ).

In Experiment 2, we again found significantly higher predictions for >0mL dosages ( $M = 65.65, SD = 14.17$ ) than 0mL dosage ( $M = 15.12, SD = 24.56, F(1,146) = 491.17, p < .001, \eta^2_p = .77$ ). This effect of dosage did not interact with CD ( $F(1,146) = 1.17, p = .28$ ) or cue variability ( $F < 1$ ), nor was there a three-way interaction between these variables,  $F < 1$ . Across both experiments, we found that participants gave higher prediction ratings when the cue value was >0ml, indicative of illusory causation. Thus, even without witnessing any patients given 0mL, participants judged the treatment to be effective at increasing bone density.

In our second analysis, predictions at 0mL were excluded to examine the perceived effect of increasing the dosage from 10mL to 100mL. Consistent with predictions during training, we extracted linear functions of the effect of treatment dosage on test predictions. The linear slopes were then analysed with a 2 (CD: low vs. high) × 2 (cue variability: discrete vs. continuous) between-subject ANOVA. As presented in Figure 4, the slope was significantly more positive in the high-CD group in both Experiment 1 ( $M_{diff} = .35, F(1,146) = 18.51, p < .001, \eta^2_p = .11$ ) and Experiment 2 ( $M_{diff} = .25, F(1,146) = 12.55, p < .001, \eta^2_p = .08$ ), indicating that CD effect was present in both experiments. In addition, there was a main effect of cue variability in Experiment 1, with significantly more positive slopes in the discrete cue group ( $M = .12, SD = .61$ ) than the continuous cue group ( $M = -.11, SD = .49, F(1,146) = 8.73, p = .004, \eta^2_p = .06$ ), but this effect was not significant in Experiment 2 ( $F < 1$ ). Critically, the significant CD effect in both experiments did not interact with cue variability ( $F < 1$ ), suggesting that CD effect was equivalent in the discrete and continuous cue groups.

**Efficacy test ratings** Since efficacy ratings were identical in Experiment 1 and 2, we combined efficacy ratings in two experiments and analysed each of them with a 2 (Experiment 1 vs. 2) × 2 (CD: low vs. high) × 2 (cue variability: discrete vs. continuous) between-subject ANOVA. Note that although direct comparisons across experiments bring their own limitations, in this instance we are only combining the two experiments to draw conclusions across them rather than interpreting differences between them.

Treatment rating asked participants about the effect of Calciucor on bone density (which cannot be directly verified from training because treatment was always given). As presented in Figure 5a, we found that treatment rating was significantly higher in Experiment 2 ( $M = 56.56, SD = 28.52$ ) than Experiment 1 ( $M = 29.21, SD = 29.08, F(1,292) = 64.64, p < .001, \eta^2_p = .19$ ). There was also a main effect of cue variability on treatment rating, with significantly higher ratings in the discrete cue group ( $M = 47.23, SD = 31.51$ ) than the continuous cue group ( $M = 38.60, SD = 31.71, F(1,292) = 5.80, p = .017, \eta^2_p = .02$ ). No other interactions were significant ( $F_s < 1.57$ ).

Dosage rating asked participants about the effect of treatment *dose* on bone density (where training should reveal that there was no relationship between dose and recovery). As presented in Figure 5b, there was a significant main effect of CD on dosage rating, where participants in the high-CD group ( $M = 19.60, SD = 30.96$ ) reported greater efficacy ratings than those in the low-CD group ( $M = 12.40, SD = 34.44, F(1,292) = 3.88, p = .049, \eta^2_p = .01$ ). Again, we found no significant interaction between CD and cue variability ( $F(1,292) = 3.20, p = .07$ ), suggesting that discrete and continuous cues produced equivalent CD effects, though this marginal non-significant interaction reflects the fact that the CD effect was at least numerically more distinct in the discrete conditions than continuous.

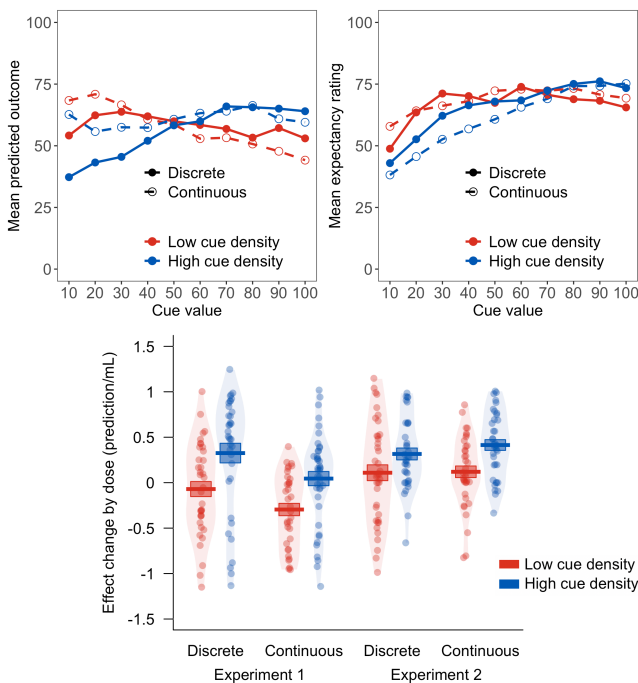


Figure 4: Test predictions as a function of dosage. Top panels: predictions made to doses in increments of 10 mL (10-100mL) for experiment 1 (left) and experiment 2 (right). Bottom panel: Slope estimates derived from test predictions as a linear function of treatment dose.

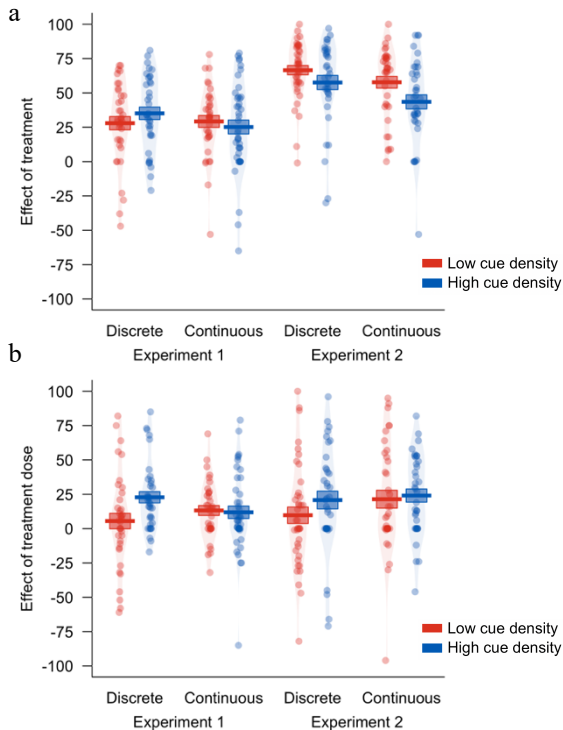


Figure 5: Test efficacy ratings in Experiments 1 & 2. (A) Treatment rating asked participants to judge the effect of Calciucor on bone density. (B) Dosage rating asked participants to judge the effect of treatment dose on bone density.

## General Discussion

The main purpose of this project was to investigate whether the CD effect could be replicated using continuous cue presentations in a medical context. In both experiments, we found support for the use of continuous cues in generating the CD effect. Particularly, we found that participants exposed to frequent high medicine dosages tended to predict more positive treatment outcomes, and gave higher causal judgement about the efficacy of increasing medication intake on increasing bone density. The evidence for this CD effect was consistent across predictions during training, test predictions, and efficacy ratings. While Chow et al. (2019) have demonstrated an OD effect using continuous outcomes, to our knowledge this is the first study to report the CD effect when cues were also presented in a continuous fashion.

Notably, and different from prior studies on causal illusion, participants in our experiments did not receive sufficient information to judge the efficacy of the treatment compared to no treatment, as they never witnessed patients taking 0mL during training. To assess participants' assumption about the medication efficacy, we compared participants' test predictions for 0mL vs. >0mL, and directly asked them how effective the drug was at increasing bone density in treatment rating. Both analyses revealed that participants in our experiments believed that the medication was effective

relative to no treatment. These results show a willingness for participants to infer the effect of no treatment on patient outcomes, despite not having direct experience with no treatment trials. One factor that appeared to influence treatment ratings was cue variability, with participants in the continuous-cue group giving lower treatment ratings than those in the discrete cue group. This suggests that witnessing richer events (i.e., continuous cues) could potentially act against participants' preconception about treatment efficacy. Likewise, increasing cue variability also seemed to reduce causal illusion, as evidenced by the less positive slopes of test predictions in the continuous-cue group. Although only observed in Experiment 1, this suggests that participants who have witnessed a larger sample of different medication intake were more resilient to the tendency to develop causal illusions about the effect of the dosage of the potential cause. These results provide tentative evidence to suggest that exposure to more diverse events may help to mitigate causal illusions.

Our results suggest that causal illusions are not restricted to scenarios where individuals overestimate the contingency between two discrete events. Rather, cue density also inflates causal illusions about the linear relationship between two continuous events, as in Experiment 1. This is most evident in participants' dosage ratings; participants who saw many patients given high dosage of Calciucor gave higher causal ratings than those in the Low CD condition. Our replication of CD effect using continuous cues signifies a need for future studies to further explore the boundary conditions of the CD effect. For example, is there a level of cue variability at which the CD effect no longer occurs? Exploring this question would help establish a causal learning theory that takes into account how continuous magnitude is parsed during learning (e.g. see Pacer & Griffiths, 2011). In addition, the extension of the CD effect to the use of continuous cues has important real-life implications. As most alternative and complementary medicines are advertised as harmless and free of side effects, people might have the tendency take higher dosages than recommended. Our results suggest this in itself may reinforce an incorrect belief that higher dosages result in better treatment outcomes. This belief could in turn increase medication intake, creating a vicious cycle in which patients take increasing doses over time. This behavior is costly at best, and could be detrimental to health if the medication is not as harmless as advertised.

## References

- Blanco, F. (2017). Positive and negative implications of the causal illusion. *Consciousness & Cognition*, 50, 56-68.
- Blanco, F., Barberia, I., and Matute, H. (2014). The lack of side effects of an ineffective treatment facilitates the development of a belief in its effectiveness. *PLoS ONE* 9(1).
- Blanco, F., & Matute, H. (2018). The Illusion of Causality: A Cognitive Bias Underlying Pseudoscience, in A.B.

- Kaufman, and J.C. Kaufman (eds), *Pseudoscience: The Conspiracy Against Science*, Cambridge, MA: MIT Press.
- Blanco, F., Matute, H. & Vadillo, M.A. (2013). Interactive effects of the probability of the cue and the probability of the outcome on the overestimation of null contingency. *Learning and Behavior*, *41*, 333–340.
- Chow, J. Y. L., Colagiuri, B., Rottman, B. M., Goldwater, M. B., & Livesey, E. J. (2021). Pseudoscientific health beliefs and the perceived frequency of causal relationships. *International Journal of Environmental Research and Public Health*, *18*, 11196.
- Chow J. Y. L., Colagiuri B, Livesey E. J. (2019). Bridging the divide between causal illusions in the laboratory and the real world: the effects of outcome density with a variable continuous outcome. *Cognitive Research: Principles and Implications*, *4*(1), 1-15.
- Double, K. S., Chow, J. Y. L., Livesey, E. J. & Hopfenbeck, T. N. (2020). Causal Illusions in the Classroom: How the Distribution of Academic Outcomes can Promote False Instructional Beliefs. *Cognitive Research: Principles and Implications*, *5*, 34.
- Freckelton, I. (2012). Death by homeopathy: issues for civil, criminal and coronial law and for health service policy. *Journal of Law & Medicine*, *19*, 454–478.
- Jenkins, H. & Ward, W. (1965). Judgement of contingency between responses and outcomes. *Psychological Monographs: General and Applied*, *79*, 1-17.
- Mandel, D. R. & Vartanian, O. (2009) Weighting of contingency information in causal judgement: Evidence of hypothesis dependence and use of a positive-test strategy. *The Quarterly Journal of Experimental Psychology*, *62*(12), 2388-2408.
- Marsh, J. K., & Ahn, W. K. (2009). Spontaneous assimilation of continuous values and temporal information in causal induction. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *35*(2), 334–352.
- Matute, H., Blanco, F., Yarritu, I., Díaz-Lago, M., Vadillo, M. A., & Barberia, I. (2015). Illusions of causality: how they bias our everyday thinking and how they could be reduced. *Frontiers in psychology*, *6*(888). <https://doi.org/10.3389/fpsyg.2015.00888>
- Matute, H., Yarritu, I., and Vadillo, M. A. (2011). Illusions of causality at the heart of pseudoscience. *British Journal of Psychology*, *102*(3), 392–405.
- Pacer, M. D., & Griffiths, T. L. (2011). A rational model of causal induction with continuous causes. *Advances in Neural Information Processing Systems*, *24*.
- Shou, Y., Smithson, M. (2015) Adapting to an Uncertain World: Cognitive Capacity and Causal Reasoning with Ambiguous Observations. *PLoS ONE* *10*(10).
- Soo, K. W., & Rottman, B. M. (2018). Causal strength induction from time series data. *Journal of Experimental Psychology: General*, *147*(4), 485–513.
- Vadillo, M. A., & Matute, H. (2007). Predictions and causal estimations are not supported by the same associative structure. *Quarterly Journal of Experimental Psychology*, *60*(3), 433–447.
- Vadillo, M. A., Musca, S.C., Blanco, F. et al.(2011). Contrasting cue-density effects in causal and prediction judgments. *Psychonomic Bulletin Review*, *18*, 110–115.
- Yarritu, I., Matute, H., and Vadillo, M. A. (2014). Illusion of control: the role of personal involvement. *Experimental Psychology*, *61*(1), 38–47.