Title
Learning about Benefits and Side Effects of a Bogus Treatment are Similarly Influenced by the Frequency of the Outcome Occurring

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

Authors
Chow, Julie YL
Livesey, Evan

Publication Date
2023

Peer reviewed

Learning about Benefits and Side Effects of a Bogus Treatment are Similarly Influenced by the Frequency of the Outcome Occurring

Julie Y. L. Chow (julie.chow@unsw.edu.au)
School of Psychology, UNSW Sydney, Australia

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

Abstract

Detecting covariation in sequential events provides us with a powerful means of inferring the causal structure of our world. However, people often overestimate the causal relationship between unrelated events, a phenomenon referred to as *illusory causation*. This tendency is greatest when the putative effect occurs frequently; the widely replicated *outcome density (OD) effect.* Most laboratory research on illusory causation and the OD effect has focused on possible causes of a positive outcome, such as a drug that causes patient recovery. Despite its relevance, relatively few studies have examined illusory causation in cases where a cue is hypothesized to generate an unfavorable (negative) outcome, such as a drug that produces unwanted side effects. Here, we directly compared how people develop illusory beliefs about the generation of positive versus negative outcomes. We presented all participants with a drug treatment that was hypothesized to cause high readings of a fictitious cell count (but had no effect on cell count across a series of learning trials). We manipulated whether a high cell count occurred frequently or infrequently and whether a high cell count should be considered a beneficial medical outcome or an undesirable side effect. We found consistent evidence of an OD effect but no effect of the valence of the high cell count outcome. This suggests that illusory beliefs are not controlled by the desirability of the cause-effect relationship. We discuss implications for theories of applied causal reasoning.

Keywords: illusory causation; outcome density effect; causal learning; contingency learning; outcome valence

Introduction

The ability to accurately detect and interpret covariation information is important for the formation of accurate causal beliefs. For example, if a person develops a rash every time they eat peaches, they might learn that eating peaches causes an allergic reaction and thus may avoid peaches in the future. People are often accurate at assessing causal relationships through this type of contingency learning, that is, when observing sequentially occurring events and acquiring evidence about the likelihood of the outcome with and without the presence of a putative cause (Jenkins & Ward, 1965). However, people also tend to perceive a causal relationship between unrelated events, in which the outcome is not contingent on the putative cause (Alloy & Abramson, 1979). The illusory causation effect is most prevalent when the probability of the outcome occurring is high, irrespective of the putative cause (often referred to as the cue). This sensitivity to outcome frequency is commonly referred to as the *outcome density effect* (Alloy & Abramson, 1979). Illusory causation and the outcome density (OD) effect are thought to contribute to the development of pseudoscientific beliefs, including the use of pseudo-medicines that are shown to be ineffective (Matute, Yarritu & Vadillo, 2011); these treatments are also commonly used for illnesses that have a high rate of spontaneous remission such as the common cold (Blanco & Matute, 2020). The outcome density effect has previously been shown to be sensitive to causal instructions, such that it is the frequency of the *target* outcome occurring, as presented in the instructions about the putative cue-outcome relationship, results in inflated illusory beliefs (Blanco & Matute, 2015). These findings suggest that the effect may be influenced by the causal hypothesis the learner is entertaining, such that information consistent with that causal hypothesis (e.g., cue and target outcome co-occurring) is weighted more heavily, either in memory or during the information integration process, when making causal judgements (see Wasserman, Dorner & Kao, 1990).

Another class of explanations for the effect are associative models like the Rescorla-Wagner (RW) model (1972), which assume illusory causation is a product of pre-asymptotic learning. According to this account, learning about the cue-outcome relationship occurs more quickly than the context-outcome relationship due to increased salience of the cue relative to the context (no cue present). This is further inflated by the presence of many outcome-present trials when outcome density is high, or when the cue is present frequently (cue density effect). Thus, early in learning, the associative strength for the cue is much greater than the context, resulting in inflated causal beliefs about the cue-outcome relationship.

The vast majority of previous studies on illusory causation have focused on causal scenarios where the cue is thought to generate a desirable outcome. This desirability, stems either from the scenario presented in instructions, such as a patient who recovers from an illness after receiving a treatment cue (e.g., Matute et al., 2011) or based on the assumed goal of the task presented to the participant, such as a lightbulb turning on when a button is pressed (e.g., Alloy & Abramson, 1979). However, many real-world pseudoscientific beliefs involve negative illusions, where the potential consequence of an action or cue is undesirable. Examples of these include beliefs that vaccinations cause autism (Davidson, 2017) and eating genetically modified food causes cancer (Touyz, 2013). Note that these negative generative illusions are distinct from preventive illusions, where the cue is thought to

In M. Goldwater, F. K. Anggoro, B. K. Hayes, & D. C. Ong (Eds.), *Proceedings of the 45th Annual Conference of the Cognitive Science Society.*

prevent an undesirable outcome from occurring. In preventive illusions, the desired outcome is the *absence* of an undesired event, for instance, a herbal remedy preventing illness. Therefore, the hypothesized effect of the cue is still inherently positive. Indeed, while there is fairly good evidence for desirable preventive causal illusions of this nature (e.g., Aeschleman, Rosen & Williams, 2003; Alloy & Abramson, 1979; Hayashi & Modico, 2019), the development of false beliefs that are explicitly *undesirable*, such as a cue generating an unwanted effect, remains under-explored. This is important given the possibility that motivational factors (for instance those underpinning motivated reasoning) may play a role in how information is interpreted (e.g., Caddick & Rottman, 2021), and may be important at establishing illusory causation.

One study that has investigated the effect of positive and negative outcomes in a zero-contingency task is a study by Blanco, Barberia & Matute (2014). The primary goal of this experiment was simply to determine if the presence of side effects associated with a treatment decreases the tendency for people to over-administer the treatment cue, another procedure known to inflate illusory beliefs (i.e., cue density effect) by inflating cue-outcome coincidences. In this study, participants were either told that a novel drug would produce negative side effects or given only information about patient recovery and no information about potential side effects of the treatment. The authors found that participants told about the presence of side effects administered the treatment to patients on fewer trials, and subsequently reported lower efficacy ratings than participants who were not told about any side effects. The results of this study are important as they showed how the presence of negative outcomes might influence causal illusions. Other similar studies have also suggested that strategies effective at reducing positive illusions may in fact be counter-productive when the outcome is2egateve (Matute & Blanco, 2014). Thus, it is beneficial to consider how the intrinsic valence of the outcome might also influence biases that inflate false beliefs, like the OD effect.

To our knowledge, only one study has systematically investigated the probability of the target outcome on illusory beliefs with positive and negative outcomes. Rudski, Lischner & Albert (1999) manipulated outcome valence by using gains (positive) and losses (negative) in points on a prediction task. Participants were told that their goal was simply to make predictions about the outcome based on cues presented on the screen. Participants in the positive outcome condition either gained points or did not gain points on each trial, whereas participants in the negative outcome condition either lost or did not lose points on each trial. The contingency between cue and outcome was zero. The researchers also manipulated the probability of the outcome, such that the target outcome (gaining points in the positive condition, and losing points in the negative condition) occurred on either 75%, 50% or 25% of trials. The measure of illusory belief in this study was the superstitious rules generated by participants during the task, obtained through questionnaires administered across blocks of trials. The findings showed that although all participants were more likely to generate superstitious rules when the probability of the desired outcome was maximal (75% gain in the positive condition and 25% loss in the negative condition), participants in the positive outcome condition were significantly more likely to generate superstitious rules than participants in the negative outcome condition. The researchers concluded that the reduced rule generation in the negative outcome condition reflects increased sensitivity to the zero contingency between cues and outcomes. Thus, biases that inflate illusory beliefs are significantly attenuated when the outcome was presented as a negative event.

However, the results of Rudski et al (1999) could also be explained by a different cognitive bias: loss aversion. Generally, people perceive potential loss to be more salient than that of potential gain of the same magnitude, resulting in greater risk-aversion in a loss scenario than in a gain scenario (Kahneman & Tversky, 1979). Thus, the reduction in superstitious rule generation among participants in the loss condition may instead be a result of greater motivation to be accurate at learning the cue-outcome relationship in order to avoid loss. In medical decisions, loss aversion is important as it informs decisions about risky treatments and surgeries. However it is unclear how the valence of the outcome *alone* might influence how people learn about a cue-outcome relationship, and whether biases in learning like the OD effect might influence the illusion of causality to a similar degree when the target outcome is an undesirable event as when it is a desirable one. Thus the question that remains unanswered is whether people show equivalent illusory beliefs about a treatment and a health outcome when the target outcome is presented as a benefit of the treatment, and when it is presented as a side effect of treatment use.

The current study

The goal of the present study was to determine whether learning about the *same* cue-outcome relationship differs as a function of whether the target outcome is described as a positive or negative event. In particular, we were interested in whether participants would show differences in illusory causal beliefs and the OD effect when asked to evaluate a hypothesized cue-outcome relationship—that administration of a drug causes a high cell count—depending on whether the outcome (high cell count) was interpreted as a benefit or a side-effect. In the positive outcome condition, high cell counts were described as being indicative of recovery and a benefit of the treatment. In the negative outcome condition, high cell counts were described as being an unwanted side effect. Note here that valence refers to whether the outcome is framed in the causal instructions as being a positive or negative outcome; there was nothing inherently positive or negative about the physical depiction of the outcome itself. Thus, we were able to hold constant intrinsic salience of the outcome whilst manipulating its valence.

We measured participants' causal beliefs by asking them to rate how effective the drug was at causing a change in cell counts relative to no treatment. We also included a memory

question where participants were asked to report how many trials they remembered seeing where the patient had been given treatment or no treatment and subsequently recovered or not recovered. For a number of real-world health beliefs, variance in causal beliefs correlates with peoples' perception of event frequency such that their subjective judgements of cause-effect contingency are consistent with the strength of ratings endorsing the causal relationship (Chow et al., 2021). This contingency estimation question allowed us to investigate whether the valence of the outcome has an impact on memory, specifically for trials that confirmed the putative causal relationship (i.e., both cue and target outcome were present; "*a*-cell trials"). If outcome valence systematically influenced how we remember and process information, we should see differences in the impact of the OD manipulation as a function of outcome valence. In addition, participants might systematically differ in their memory of the number of trials where patients were given a novel treatment and subsequently reported high cell counts. Results from this experiment will allow us determine whether the desirability of the outcome is an important factor that influences the degree to which people overweigh cue-outcome coincidences in illusory causation, a strategy that may be appealing when the outcome is desirable but not when the outcome is undesirable.

Methods

Participants

Two-hundred and six participants were recruited from Amazon's Mechanical Turk and completed the entire study online (124 male, M_{age} = 30.9, SD_{age} = 9.68). Participant data was removed from analyses if they 1) failed a comprehension check presented immediately after the instructions more than two times, 2) failed the comprehension check at the end of the experiment (the two comprehension checks were identical), or 3) reported writing down information during the experiment. After exclusions, 173 datasets remained; Low OD-Positive, N = 44; High OD-Positive, N = 36; Low OD-Negative, N = 38; High OD-Negative, N = 55.

Figure 1: Screenshot example of a single training trial.

Stimuli & Apparatus

The experiment was programmed using the jsPsych library (de Leeuw, 2015). On cue-present trials, an image of a pill bottle was presented at the top of the screen together with the drug name "Cloveritol" below it in large bold font. On cue-absent trials, the image of the pill bottle was greyed out and the text "No Treatment" was presented instead. Both cue and no cue stimuli were presented on a blank background that was 300 x 440 pixels. The outcome "High cell count" or "Low cell count" was presented in bold text under the prompt "Actual observed JYL cell count:". An example screenshot of a single trial is shown in Figure 1.

Design

The study had a 2 (Outcome Valence: Positive vs Negative) × 2 (Outcome Density: Low vs High) between-subjects design. All participants were instructed that Cloveritol is linked to fictitious JYL cells in the blood, and the drug is thought to generate high JYL cell counts. In the positive valence condition, participants were told that high JYL cell counts were important for combating the effects of the virus that causes the illness, whereas in the negative valence condition, participants were told that high JYL cell counts may result in unwanted side effects.

Outcome density was manipulated by presenting participants with high JYL cell counts on either 80% of all trials (High OD), or on 20% of all trials (Low OD). Critically, the probability of high or low cell counts were the same for Cloveritol and No Treatment trials, i.e., zero contingency between treatment and high cell count. All participants experienced 50 cue-present and 50 cue-absent trials.

Procedure

On each trial, participants were presented with a new patient who was either administered Cloveritol or No Treatment. Participants were asked to make a prediction about whether the patient will have high or low cell counts by making a Z (Low cell count) or M (High cell count) key-press response. Having made a response, they were then shown the actual observed cell count for that patient. Note that unlike most loss aversion studies, where participants have a choice of whether to activate the cue (e.g., choosing whether or not to gamble for the chance of a reward), participants in our study were simply told whether Cloveritol or No Treatment was given to the patient on each trial.

At the end of the training phase, participants were required to provide a causal rating on the efficacy of Cloveritol (relative to no treatment) in causing a change in JYL cell counts. Causal ratings were made on a scale from -100 (Definitely lowers JYL cell counts) to 100 (Definitely increases JYL cell counts) with a mid-point of 0 (Completely ineffective). Participants in all conditions were presented with an identical rating scale. Finally, participants were asked to use their memory of the 100 training trials to report the number of patients who (a) were administered Cloveritol and had high JYL cell counts, (b) were administered Cloveritol and had low JYL cell counts, (c) were administered No

Treatment and had high JYL cell counts, and finally (d) were administered No Treatment and had low JYL cell counts. The four estimation questions were presented on a single screen. We were primarily interested in whether efficacy ratings for the treatment and recall estimates for the different trials differed as a function of outcome valence and outcome density. More positive efficacy ratings are indicative of greater causal illusions, where the veridical rating for a non-contingent relationship is zero. Similarly, contingency estimates for the probability of the outcome in the presence and absence of the cue should be zero if participants were accurate at identifying the null contingency between events.

Results

We report results from analyses using both Null Hypothesis Significance Testing (NHST) and Bayes Factor Analysis. Bayes Factors for main effects are reported as a likelihood ratio of the alternative model relative to the null model. Where there are more than two factors in the model, we reported the Bayes Factor Inclusion (BF_{incl}) across matched models (Rouder, Morey, Verhagen, & Wagenmakers, 2017), which provides an estimate for the evidence for the effect to equivalent models stripped of the effect.

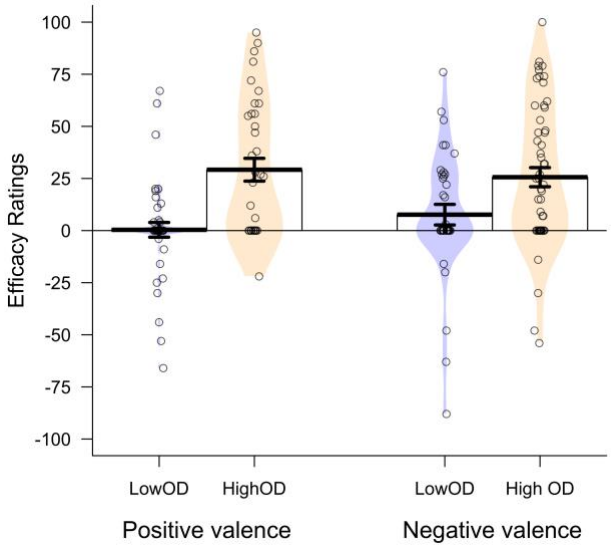

Figure 2. Efficacy ratings (+/-SEM) of Cloveritol at changing levels of JYL cell counts relative to no treatment as a function of outcome density (Low vs High OD) and outcome valence (positive vs negative).

Efficacy ratings

Results of the Analysis of Variance (ANOVA) comparing OD and outcome valence on participants' causal ratings are depicted in Figure 2. Results showed a main effect of OD, $F(1,169) = 25.6$, $p < .001$, $\eta_p^2 = .127$, $BF_{10} = 1.05e+4$, with greater causal ratings in the High OD ($M = 27.4$, $SD = 33.4$) compared to Low OD condition ($M = 7.81$, $SD = 27.0$). These results are consistent with the literature, where greater probability of the outcome occurring typically results in greater causal illusions.

There was no main effect of outcome valence, $F(1,169) = .153$, $p = .696$, $\eta_p^2 = 9.06e-4$, $BF_{10} = .260$, and no interaction between the two factors, $F(1,169) = 1.32$, $p = .252$, $\eta_p^2 = .008$, $BF_{incl} = .390$. Thus, we found no evidence that the illusion of causality and the outcome density effect differed as a function of whether high JYL cell counts were instructed to be a benefit or side effect of treatment use.

Contingency estimates from memory

We converted participants' estimates for each of the four cue-outcome trials to a proportion of the total number of trials reported. First, we analysed the proportion of *a*-cell trials (i.e. patient was given Cloveritol and had High cell counts) as a function of OD and outcome valence. These results are presented in Figure 3a. ANOVA revealed a main effect of OD, $F(1,169) = 182$, $p < .001$, $\eta_p^2 = .519$, $BF_{10} = 1.25e+26$. This result is not surprising given the greater number of *a*-cell trials in the High OD compared to the Low OD condition. There was no main effect of outcome valence, $F(1,169) = .228$, $p = .633$, $\eta_p^2 = .001$, $BF_{10} = .599$, and no interaction between OD and outcome valence, $F(1,169) = .153$, $p = .697$, $\eta_p^2 = 9.03e-4$, $BF_{incl} = .253$.

Visual inspection of Figure 3a (group mean compared to actual proportion of *a*-cell trials denoted by a dashed line; 0.1 for the Low OD condition and 0.4 for the High OD condition) indicate that there is greater overestimation of the proportion of cue-outcome coincidences in the Low OD condition compared to the High OD condition, where the average proportion of *a*-cell trials reported were closer to the veridical value (0.4). This overestimation in the Low OD condition compared to the High OD condition may simply be a result of the very low proportion of *a*-cell trials in the Low OD condition (0.1) which allowed for more room for overestimation to occur. Therefore, these results may simply reflect regression to the mean.

From all four cell estimates, we also computed a Δp score, which is an index of the contingency between Cloveritol use and High cell counts (Allan, 1980). Greater positive Δp scores are indicative of greater illusory causation where the likelihood of high cell counts is thought to be greater on treatment trials than no-treatment trials. Comparing participants' Δp score as a function of OD and outcome valence, we found only a main effect of OD, $F(1,169) = 4.17$, $p = .043$, $\eta_p^2 = .024$, $BF_{10} = 1.32$. There was no main effect of outcome valence, $F(1,169) = 1.28$, $p = .259$, $\eta_p^2 = .008$, $BF_{10} = .390$, and no interaction between the two factors $F(1,169) = 1.94$, $p = .166$, $\eta_p^2 = .011$, $BF_{incl} = .522$.

These results suggest that although participants were numerically more likely to misperceive a positive contingency between the cue and outcome when the overall probability of the outcome occurring is high, there was no difference in the bias for scenarios where the target outcome was presented as a benefit of the treatment, and when it was presented as an unwanted side effect.

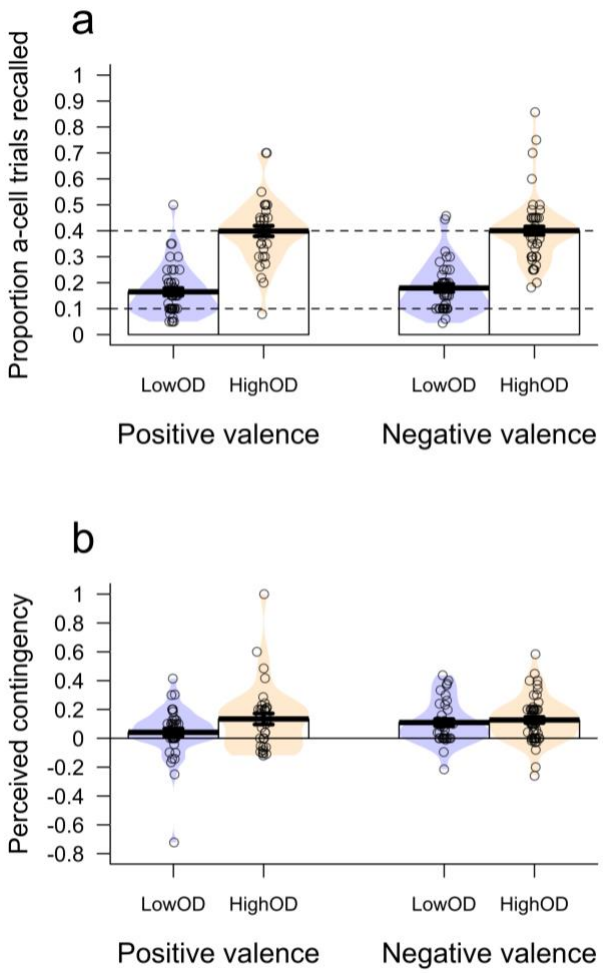

Figure 3. Average (a) proportion of cue-present and outcome-present trials reported, and (b) perceived contingency indexed by Δp, as a function of outcome density and valence. Dashed lines in Figure 3a indicate the actual proportion of cue-present/outcome-present trials in the Low OD (0.1) and High OD (0.4) condition.

Predictions across training

Average proportions of High cell count predictions made during training as a function of cue type (Cloveritol vs No treatment), OD and outcome valence is depicted in Figure 4.

Repeated-measures ANOVA revealed a significant main effect of cue type, $F(1,169) = 139.5$, $p < .001$, $\eta_p^2 = .452$, $BF_{10} = 3.72e+22$, with greater outcome predictions on cue trials than no cue trials. However, the bias to predict High cell count on cue trials compared to no cue trials did not differ as a function of OD, $F(1,169) = .049$, $p = .825$, $\eta_p^2 = 2.92e-4$, $BF_{incl} = .189$. There was no interaction between cue type and outcome valence, $F(1,169) = 1.97$, $p = .162$, $\eta_p^2 = .012$, $BF_{incl} = .523$, and no three-way interaction between cue-type, OD, and outcome valence, $F(1,169) = 1.38$, $p = .242$, $\eta_p^2 = .008$, $BF_{incl} = .464$. There was however a main effect of OD, $F(1,169) = 266.4$, $p < .001$, $\eta_p^2 = .612$, $BF_{10} = 9.33e+22$, with greater proportion of high cell count predictions when the target outcome occurred frequently ($M = .711$, $SD = .232$) than when it occurred infrequently ($M = .319$, $SD = .232$), indicative of some sensitivity to the base-rate of the outcome occurring. All other findings were not statistically significant, *largest* $F = 1.37$.

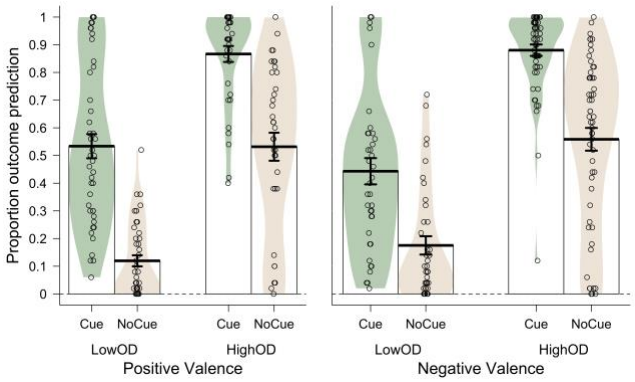

Figure 4. Average proportion of High cell count predictions made during the entire training phase as a function of cue type, OD and outcome valence.

Overall these results suggest that although participants show illusory causation, indexed by the difference in the proportion of High cell count predictions on treatment and no treatment trials, there was no evidence of an outcome density effect on this measure, and no difference in predictions as a function of outcome valence. Across all measures, we find consistent evidence of illusory causation. We also found evidence of the outcome density effect in test judgments of causality and trial frequency estimates. However there was no evidence of an effect of outcome valence on any of these measures.

General Discussion

The primary goal of this study was to determine whether the valence of the target outcome (positive vs negative event) had an impact on illusory causation and the outcome density effect. The findings show an OD effect, where causal ratings were higher when high cell counts occurred frequently than when they were infrequent. There was also some evidence that participants' contingency estimates were influenced by the frequency of the outcome occurring. Importantly, we found no evidence that the OD bias differed as a function of outcome valence across all measures. Together these results provide consistent evidence that there is no differential influence of outcome valence on the evaluation of a causal relationship between two events at least in this sort of task, which holds constant other factors like the probability of the cue, cognitive biases like loss aversion, and which manipulates outcome valence via instructions alone.

One important feature of the present study design is that all participants were presented with and asked to evaluate the *same* cue-outcome relationship, that is the relationship

between Cloveritol and High JYL cell counts. The only difference between the two valence conditions is whether a high cell count was framed as a benefit or side effect of the drug in the instructions. This is in contrast to previous studies (e.g., Aeschleman et al., 2003) where the causal structure of the cue-outcome relationship in the negative condition involves the prevention of an undesired outcome. Similarly, we avoided framing outcomes as losses versus gains, which allowed us to de-confound negative outcomes from losses. This is pertinent since events relating to a loss are perceived to be more salient, and may lead to other cognitive biases such as loss aversion. It should be noted however, that even with our instructions, it is plausible that some participants might still interpret unwanted side-effects of treatment as a loss (of health). However, given the generative framing of the causal instructions, i.e., Cloveritol causes high JYL cell counts, and the framing of the causal rating question in terms of changes to JYL cell counts rather than increase or decrease in health, it is more likely that participants would focus on the effect of the drug on the cell count rather than on patient health, thereby reducing the potential impact of loss aversion. Furthermore, even if the cell count outcome was perceived as a loss in the negative valence condition, participants were still constrained in this task to observe the consequences regardless (the outcome was unavoidable). Thus loss aversion is unlikely to drive a strong difference in the formation of causal beliefs about the effect of the treatment. This is supported by unpublished studies from our lab that have found instructional manipulations aimed at attenuating causal illusions to only have an indirect effect on causal beliefs, mediated by their impact on the participant's decision to choose to administer the cue or not, when cue administration on each trial was decided by the participant (Chow et al., 2023). There is little evidence that these interventions were able to directly impact causal beliefs when participants have no control over administration of the cue. Thus, if framing of the outcome as a side effect is perceived as a loss, it is more likely to impact illusory causation through the (reduced) frequency of cue administration available in other designs (e.g., Blanco et al., 2014), but not in the present "passive" version of the contingency learning task.

Nevertheless, our results showed that there was no evidence that framing high cell counts as a benefit of treatment or an unwanted side effect significantly influenced efficacy ratings or participants' estimation of the contingency between events. This finding is interesting since the causal relationship participants are falsely believing in is an appealing one in the positive valence condition, but not appealing in the negative condition. Thus, there is no evidence that biases that inflate illusory causation differ when the outcome is a desirable event that participants may be motivated to believe in, and when it is undesirable. This is in contrast to previous research that has found a differential effect of outcome valence on illusory beliefs, however as noted in the introduction, these results could be due to differences in causal structure (generating a desirable outcome vs preventing an undesirable one) or loss aversion, rather than outcome valence *per se*.

The greater positive causal belief ratings and overestimation of contingency between cue and outcome (Δp) in the High OD relative to Low OD condition provides further evidence that the number of cue-outcome coincidences strongly influence illusory beliefs. This is consistent with previous studies showing that greater administration of the treatment, which inflates the number of cue-outcome coincidences, lead to greater illusory belief that the treatment causes patient recovery (Blanco et al., 2014). The absence of a difference between the two valence conditions in the present experiment suggests that perhaps the effect of introducing potential side effects associated with a treatment on beliefs about treatment efficacy may be driven completely by the tendency for participants to (over)administer the treatment when there are no known side effects of treatment use (i.e., cue density bias). In our study, where participants had to directly evaluate the relationship between the treatment and negative outcome, infrequent side effects also resulted in a weak negative illusion. These findings have real-world import since most pseudo-medicines do not have noticeable side effects (or any effects), leading to both an over-administration of the treatment (i.e., cue density bias) *and* weak beliefs about the treatment-side effect relationship. Furthermore, these pseudo-medicines are typically used to treat minor illnesses where the patient is likely to recover (i.e., OD effect, Blanco & Matute, 2020).

In conclusion, we found that when asked to evaluate the same (null) cue-outcome relationship, the valence of the outcome did not matter in establishing illusory causation. What these findings suggest is that biases that contribute to illusory causation, like the tendency to attend to or overweigh cue-outcome coincidences, lends itself to confirming the causal hypothesis established by the task context and the instructions, regardless of whether the cue-outcome relationship is desirable or not. In other words, participants are biased towards confirming the hypothesis that Cloveritol leads to high JYL cell counts, independent of what high cell counts represent. A corollary of this is that what *does* seem to be important for establishing these false causal beliefs is the frequency with which the putatively linked events (in this case, Cloveritol and high cell counts) co-occur. Frequent cue-outcome coincidences increase opportunity for the causal hypothesis to be confirmed, and reduce the opportunity for learners to detect the null contingency. Manipulations of trial frequencies like cue and outcome density effects remain the strongest influence on illusory causation. Although illusory causation, and the outcome density effect in particular, are not reliant on these events being discrete or binary (see Chow et al., 2019; Double et al., 2020), this study suggests that the frequency with which the hypothesis—proposed in instructions or implied by the causal scenario—receives positive confirmation is the most crucial component in generating the illusory causation effect.

References

Aeschleman, S. R., Rosen, C. C., & Williams, M. R. (2003). The effect of non-contingent negative and positive reinforcement operations on the acquisition of superstitious behaviors. *Behavioural Processes, 61*(1-2), 37-45.

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

Alloy, L. B., & Abramson, L. Y. (1979). Judgment of contingency in depressed and nondepressed students: Sadder but wiser?. *Journal of experimental psychology: General, 108*(4), 441.

Barberia, I., Vadillo, M. A., & Rodríguez-Ferreiro, J. (2019). Persistence of causal illusions after extensive training. *Frontiers in Psychology, 10*, 24.

Blanco, F., & Matute, H. (2015). Exploring the factors that encourage the illusions of control: the case of preventive illusions. *Experimental psychology, 62*(2), 131–142. https://doi.org/10.1027/1618-3169/a000280

Blanco, F., & Matute, H. (2020). Diseases that resolve spontaneously can increase the belief that ineffective treatments work. *Social Science & Medicine, 255*, p 113012.

Blanco, F., Barberia, I., & 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), e84084.

Caddick, Z. A., & Rottman, B. M. (2021). Motivated Reasoning in an Explore-Exploit Task. *Cognitive Science, 45*(8), e13018.

Chow, J. Y. L., Goldwater, M. B., Colagiuri, B., & Livesey, E. J. (2023). *Instruction on the scientific method provides (some) protection against illusions of causality*. [Unpublished manuscript].

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.

Davidson, M. (2017). Vaccination as a cause of autism—myths and controversies. *Dialogues in clinical neuroscience, 19*(4), 403.

De Leeuw, J. R. (2015). jsPsych: A JavaScript library for creating behavioral experiments in a Web browser. *Behavior research methods, 47*(1), 1-12.

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.

Hayashi, Y., & Modico, J.G. (2019). Effect of response-independent delivery of positive and negative reinforcers on the development of superstitious behavior and belief in humans. *Behavior Analysis: Research and Practice, 19*(4), 327.

Jenkins, H. M., & Ward, W. C. (1965). Judgment of contingency between responses and outcomes. *Psychological monographs: General and applied, 79*(1), 1.

Kahneman, D., & Tversky, A. (1979). Prospect theory: An Analysis of Decision Under Risk. *Econometrica, 47*(2), 263-291.

Matute, H. (1996). Illusion of control: Detecting response-outcome independence in analytic but not in naturalistic conditions. *Psychological Science, 7*(5), 289-293.

Matute, H., & Blanco, F. (2014). Reducing the illusion of control when an action is followed by an undesired outcome. *Psychonomic bulletin & review, 21*(4), 1087-1093.

Matute, H., Yarritu, I., & Vadillo, M. A. (2011). Illusions of causality at the heart of pseudoscience. *British Journal of Psychology, 102*(3), 392-405.

Rescorla, R. A., & Wagner, A. R. (1972). A theory of Pavlovian conditioning: Variations in the effectiveness of reinforcement and nonreinforcement. *Classical Conditioning II: Current Research and Theory* (pp. 64-99). Appleton-Century Crofts.

Rouder, J. N., Morey, R. D., Verhagen, J., Swagman, A. R., & Wagenmakers, E. J. (2017). Bayesian analysis of factorial designs. *Psychological Methods, 22*(2), 304.

Rudski, J. M., Lischner, M. I., & Albert, L. M. (1999). Superstitious rule generation is affected by probability and type of outcome. *The Psychological Record, 49*(2), 245-260.

Touyz, L. Z. G. (2013). Genetically modified foods, cancer, and diet: myths and reality. *Current Oncology, 20*(2), e59.

Wasserman, E. A. , Dorner, W. W. & Kao, S. F. (1990). Contributions of Specific Cell Information to Judgments of Interevent Contingency. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 16*(3), 509-521.