Assessing Distributions of Causal Beliefs in the Illusory Causation Task

Jessica C. Lee (jessica.c.lee@sydney.edu.au)[12]

Julie Y.-L. Chow (julie.chow@unsw.edu.au)[2]

Jaimie E. Lee (jaimie.lee@unsw.edu.au)[2]

David W. Ng (david.ng2@unsw.edu.au)[2]

Peter F. Lovibond (p.lovibond@unsw.edu.au)[2]

1. School of Psychology, University of Sydney, NSW, 2006, Australia

2. School of Psychology, University of New South Wales, NSW, 2006, Australia

Abstract

The illusory causation effect describes the tendency to judge an unrelated cue and outcome to be causally related. The standard procedure for assessing the illusion is based on the implicit assumptions that participants start as naïve observers with no prior beliefs about the likely relationship between the cue and outcome, and that learning can be adequately captured as a point-estimate causal rating after null contingency training. Here, we use a novel distributional measure to assess participants' beliefs over a *range* of causal relationships *prior* to, as well as after, exposure to non-contingent cues and outcomes. Across two experiments with different causal scenarios and 50% cue and outcome density, we show that participants have an initial bias towards expecting a causal relationship between the cue and outcome, and that this bias is mostly corrected after exposure to the null contingency. We conclude that distributional measures of causal beliefs can offer novel insights in understanding the illusory causation effect.

Keywords: illusory causation, prior beliefs, distribution, contingency learning, null contingency

Introduction

A robust finding in the contingency learning literature is that when a cue and outcome are unrelated (i.e., a null contingency), participants mistakenly judge the cue and outcome to be causally related. This overestimation for a null contingency has been termed the *illusory causation* effect, and has been found across a range of causal scenarios (see Matute et al., 2019, for a review), with binary and continuous outcomes (Chow et al., 2019), and when the task requires passive observation as well as active responding from participants (Alloy & Abramson, 1979). The illusory causation effect has received a considerable amount of interest in the contingency learning literature due to its potential in modelling the formation of everyday false causal beliefs in the laboratory, and its implications for theoretical models of learning that generate normative predictions.

In a typical illusory causation experiment, participants are tasked with learning the relationship between a cue (e.g., a drug) and an outcome (e.g., recovery from a disease). They are then exposed to a series of training trials where the cue is either present or absent, and the outcome is either present or absent. These trial types can be represented as a 2x2 contingency table with 4 cell types: cell A (where the cue and outcome are present), cell B (where the cue is present but the outcome is absent), cell C (where the cue is absent but the outcome is present), and cell D (where the cue and outcome are both absent). The veridical contingency between the cue and outcome is captured by the ΔP statistic (Allan, 1980; Jenkins & Ward, 1965): p(O|C) – p(O|~C), which is the probability of the outcome occurring in the presence of the cue, minus the probability of the outcome occurring in the absence of the cue. When the difference between these conditional probabilities is 0, the cue and outcome are unrelated (i.e. there is a zero or "null" contingency). Although participants are generally capable of learning about cues and outcomes that have positive and negative contingencies (Shanks, 1987; Shanks & Dickinson, 1987), they reliably misjudge a non-contingent cue and outcome to be causally related. The magnitude of the illusory causation effect is exacerbated by selection of cell frequencies that implement high cue (cells A and B), and/or outcome density (cells A and C, e.g., Blanco et al., 2013).

Various accounts have been proposed to explain why the illusory causation effect occurs. Statistical models assume that participants implicitly compute ΔP, but give unequal weighting to different cell types. Preferential weighting to cell A (where the cue and outcome are both present) would explain why high cue and outcome density conditions generate stronger causal illusions, since both involve higher frequencies of cell A trials. Indeed, the order of importance of the cells has been shown empirically to be: A > B >= C > D (e.g., Wasserman, Dorner, & Kao, 1990), providing some support for these weighted statistical models.

Another class of explanations states that participants learn by updating the associative strength between the cue and outcome. Associative models such as the Rescorla-Wagner (RW) model (Rescorla & Wagner, 1972) specify learning as the updating of associative strength (V) via prediction error. The RW model can be used to model the illusory causation effect if it is assumed that an additional context cue is present for all 4 cells. This context cue provides a means for the model to account for learning on cell C trials, which otherwise would not occur since the RW model can only update associative strength for cues that are present on a given trial. The RW model can account for some key features

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

of contingency learning such as the observation that causal ratings increase with additional training for positive contingencies, and decrease over additional training for negative contingencies (Shanks, 1987).

Despite differences between the statistical and associative approaches, Chapman and Robbins (1990) showed that the asymptotic predictions of the Rescorla-Wagner model align with that of ΔP. Additionally, allowing for each cell type to have different alpha rates (i.e., salience) in the RW model leads to predictions in line with weighted ΔP (Wasserman, 1990; Wasserman, et al., 1993). In both cases however, the assumption that cell types should receive differential weighting is largely descriptive.

In this study, we investigate the illusory causation effect through a Bayesian lens by introducing two novel features that deviate from the standard procedure. Typically, participants provide a single, point estimate of the relationship between the cue and outcome on a causal rating scale (Matute et al., 2015). A rating that is significantly above 0 (the normative answer), is then interpreted as the existence of a causal illusion. While this method has proved to be reliable for capturing the illusory causation effect, it may mask complexity and uncertainty in what participants have learned. It is possible that participants entertain multiple hypotheses or beliefs about the way in which the cue and outcome might be related (see Lee et al., 2021).

Another feature of the standard illusory causation procedure is that participants complete the training phase first, and subsequently estimate the contingency between the cue and outcome. Implicit in this procedure is the assumption that whatever contingency judgement participants provide can be attributed to their learning during the training phase. Statistical and associative models of the illusory causation effect also implicitly assume that participants start as naïve or neutral observers. In other words, the metric used to model learning (ΔP or associative strength) is assumed to start at 0. Although previous studies have shown that participants do hold prior beliefs that can affect the magnitude of the causal illusion, these prior beliefs typically concern real-world contingencies that are manipulated to be more or less plausible (e.g. Béghin & Markovitz, 2023; Vicente et al., 2023; Fuselgang & Thompson, 2000, 2003). In contrast, we were interested in *measuring* prior beliefs in the standard version of the task where participants are given hypothetical and unknown cues and outcomes.

The aim of the current study was to test a novel distributional measure of causal beliefs in the illusory causation task. In addition to the conventional point-estimate causal rating measure, we used a distribution builder (Quentin, 2016) to assess participants' beliefs in 5 causal hypotheses (strong prevention, weak prevention, no effect, weak causation, strong causation). A further deviation to the standard procedure was that we assessed participants' beliefs prior to, as well as after, exposure to any cues and outcomes in a null contingency training phase. We used 50% cue and outcome density, such that any causal illusion displayed at the end of training would not be influenced by cue and outcome density effects (previous studies show a weak illusion under these conditions). Finally, we included a control group who was not assessed on their prior beliefs to examine whether the act of allocating belief to various causal possibilities affected learning (see Howe et al., 2021). We tested these aims in two experiments involving different causal scenarios.

Experiment 1

Experiment 1 used a cover story where participants were asked to assume they were a scientist investigating the effects of a fictitious chemical (chemical ORSPR1) on cell growth. This causal scenario was intended to be relatively neutral, since chemicals could have beneficial or detrimental effects on cell counts. Note however, that the presentation of the outcome was binary (cell growth either occurred or did not occur) on each training trial.

Method

Participants One hundred and one participants (56 female, 44 male, 1 other, *M* age = 42.0, *SD* age = 15.0) residing in Australia, USA, or the UK, were recruited from the Prolific.co platform. Participants were randomly allocated to the prior (n = 58) or control (n = 43) group. All participants passed the attention checks.

Materials The experiment was programmed with the jspsych library (de Leeuw, 2015) and hosted via JATOS (Lange et al., 2015). The distribution builder (Quentin, 2016) was adapted to assess distributional causal beliefs.

Procedure The experiment consisted of a training phase, test phase, and questionnaire phase. Participants were asked to imagine that they were a scientist working in a lab tasked with discovering what effect a hypothetical chemical (chemical ORSPR1, i.e., the cue) had on cell growth in live tissue samples (i.e., the outcome), which they would observe over a series of days (i.e., trials). Participants were explicitly told that they would observe a new tissue sample on each day.

During each training trial, participants were presented with text stating whether the tissue sample was treated or not treated with chemical ORSPR1. On cue present trials, a beaker icon was presented. On cue absent trials, the same image was presented greyed-out at 3% opacity. Once participants observed whether the cue was present or absent, they were prompted to make a prediction by pressing "A" for "no cell growth" or "L" for "cell growth". Outcome feedback was then presented for 2s as text ("cell growth did occur/did not occur"), accompanied by a petri dish icon if the outcome was present or greyed out (3% opacity) if the outcome was absent. In each block of training there were 6 repetitions of A, B, C, and D trials, presented in randomized order, and there were 2 blocks of training (48 trials total). The cue and outcome density were both 50%, and the inter-trial interval was 1s.

If participants were allocated to the prior group, they were asked to rate their initial beliefs prior to the training phase.

Participants were told that they would be asked for their beliefs in two different ways. For the first question, they were asked "what effect do you think chemical ORSPR1 has on cell growth" by allocating 20 balls in the distribution builder (see Figure 1) according to their strength of belief in each of the 5 options (I think chemical ORSPR1 "strongly prevents"/"weakly prevents"/"has no effect"/"weakly causes"/"strongly causes" cell growth). Participants allocated belief by clicking + and – buttons below each option, and instant visual feedback was provided as participants updated their beliefs. A running tally of the number of allocated balls was provided underneath the distribution builder, and participants could progress once all balls had been allocated.

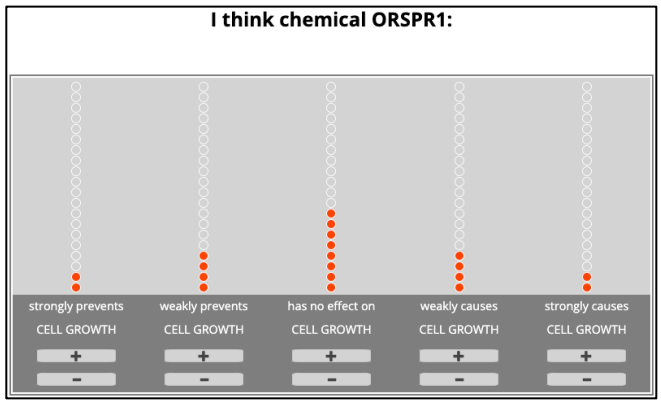

Figure 1: Screenshot of the distribution builder (Quentin, 2016) used to assess causal beliefs.

Participants then gave the standard causal rating on the next screen. Participants were asked to provide a single, best estimate to the same question by clicking on a visual analogue scale that ranged from "strongly PREVENTS" to "strongly CAUSES". The midpoint was marked with a tick and labelled "has NO effect on". A bidirectional rating scale was used to provide a conservative measure of the illusory causation effect (see Ng et al., 2023). Once participants were happy with their rating they could click Continue to progress.

After training, both groups were asked about their causal beliefs using both the distribution builder and visual analogue scale. Note that this was the second time that the prior group had to make these judgements, but the first time for the control group.

Results

Training For brevity, the training data will not be presented.

Test: Distribution Figure 2A shows the results from the distribution of participants' beliefs in a range of causal relationships at two time points (prior, and post-training).

To analyze the distribution of beliefs we computed two indices: a bias index obtained by multiplying the number of balls allocated to each option by coefficients [-2, -1, 0, +1, +2], and a shape index obtained by multiplying by coefficients [-2, +1, +2, +1, -2]. Note that the bias coefficients weight the "strong" options as double the "weak" options. Under this assumption, a bias index of 0 represents a lack of bias towards causal or prevention options, positive values indicate a causal bias, and negative values indicate a prevention bias. For the shape index, 0 represents a perfectly uniform distribution, positive values indicate an inverted-U distribution with maximal belief for "no effect", and negative values indicate a U-shape distribution with maximal belief for the two extreme options. Perfectly normative responding would be indicated by a bias index of 0 and a shape index of 40 (all 20 balls allocated to "no effect").

From Figure 2A, it is clear that prior to training, participants have the highest degree of belief in a potential causal relationship between the cue and outcome, and minimal belief in **"**no effect**"** or a preventative relationship. Unsurprisingly for the prior group, the bias index for the distribution of prior beliefs was significantly above 0, $t(57) = 5.86$, $p < .001$.

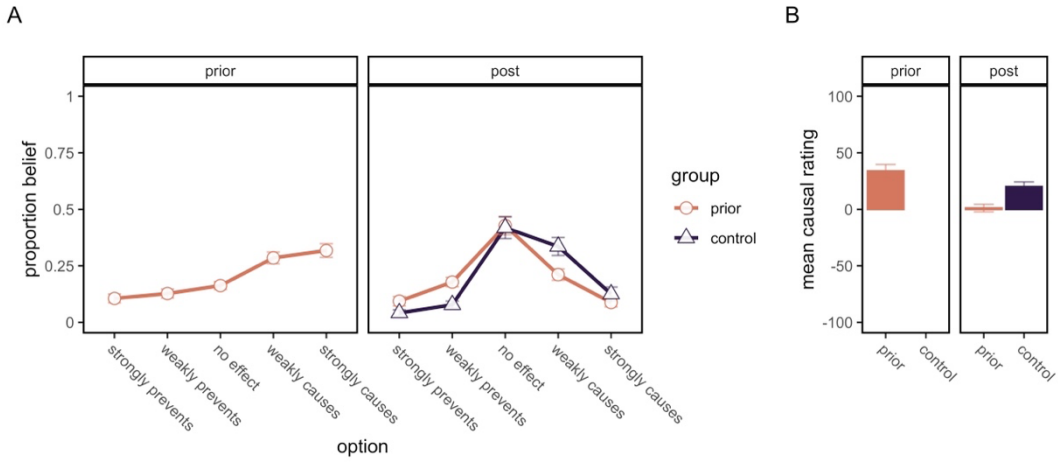

Figure 2. A) Mean proportion belief (i.e., allocated balls) for each option in the distribution builder, and B) mean causal ratings for each time point (prior to, and post-training) in Experiment 1. Error bars represent standard error of the mean.

However, after null contingency training, the distribution appears much more normative and symmetrical, with the highest amount of belief now at "no effect". In support of these observations, there was a significant decrease in the bias index, $t(57) = 5.18$, $p < .001$, and significant increase in the shape index, $t(57) = 7.66$, $p < .001$, as a result of training. The prior group's post-training bias index did not differ from 0, $t(57) = 0.348$, $p = .729$, while the control group's bias index showed a bias towards causal options, $t(42) = 5.83$, $p < .001$. There was a significant group difference in the post-training bias indices, $t(99) = 4.34$, $p < .001$, but no group difference in the shape indices, $t(99) = 0.17$, $p = .86$.

Test: Point estimate Figure 2B shows the results for the standard causal rating test of the causal illusion (0 is the normative answer) at two time points (prior to, and post-training). As can be seen in Figure 2B, the prior group's causal rating was significantly greater than 0 at the outset of training, $t(57) = 6.0$, $p < .001$, demonstrating an initial bias that the cue and outcome were causally related. After null contingency training, the prior group's causal rating reduced, and was no longer significantly above 0, $t(57) = 0.37$, $p = .714$. In contrast, the control group showed a significant causal illusion following the same null contingency training, $t(42) = 4.95$, $p < .001$. Causal ratings in the control group were significantly higher than the prior group post-training, $t(99) = 3.66$, $p < .001$.

Discussion

The main findings from Experiment 1 were that a) participants showed an initial bias towards expecting causal relationships prior to experiencing any training trials, b) this bias was fully corrected after exposure to a null contingency with 50% cue and outcome density (point-estimate causal rating did not differ from the normative answer and the distribution of belief was no longer causally-biased, with maximal belief for the normative option), and c) in contrast, the control group showed a significant illusion in their point-estimate ratings and a causal bias in their distribution of beliefs after training. Thus, the act of explicitly allocating belief to a range of causal options (prevention, no effect, causation) seemed to have inoculated against the illusion.

Experiment 2

Experiment 2 used a different, more commonly used causal scenario involving the relationship between a drug and recovery from a disease. This scenario was expected to produce stronger causal priors than in Experiment 1, since drugs are specifically developed to target particular diseases.

Method

Participants One hundred participants (66 female, 33 male, 1 other, *M* age = 37.3, *SD* age = 12.8) were recruited from the Prolific.co platform. Participants were randomly allocated to the prior (n = 49) or control (n = 51) group. One participant from the prior group was excluded from analysis due to failing the attention checks.

Procedure The procedure was identical to Experiment 1 except that in Experiment 2, participants were tasked with judging the relationship between a hypothetical drug (drug ORSPR1) and recovery from a newly discovered disease. The instructions and stimulus images were modified accordingly.

Results

Test: Distribution Figure 3A shows the results from the distribution builder. In contrast to our predictions, the drug-recovery scenario did not produce a stronger causal bias in participants' prior beliefs. In fact, belief in the drug being a weak cause of recovery was numerically stronger than belief in the drug being a strong cause, and there was a greater amount of belief allocated to "no effect" compared to Experiment 1.

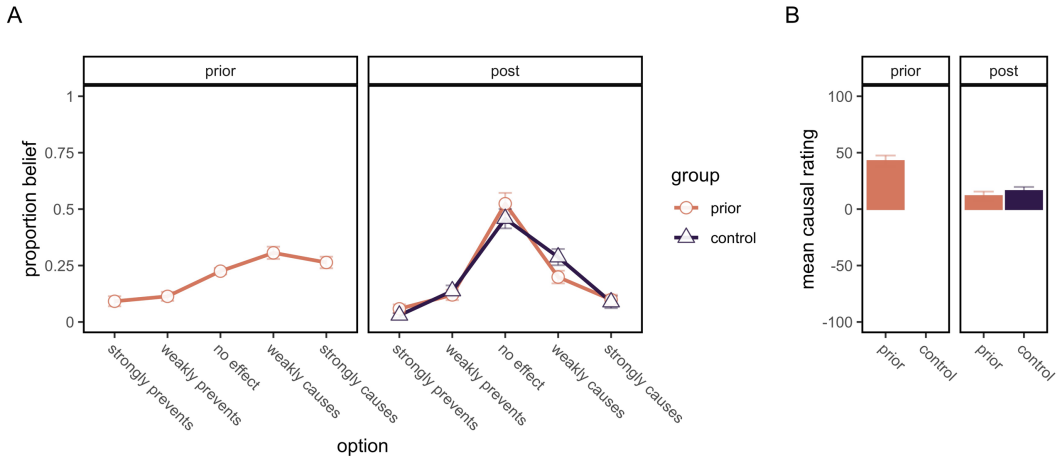

Figure 3. A) Mean proportion belief (i.e., allocated balls) for each option in the distribution builder, and B) mean causal ratings for each time point (prior to, and post-training) in Experiment 2. Error bars represent standard error of the mean.

Again unsurprisingly, the bias index for the prior distribution was significantly above 0, $t(47) = 5.20$, $p < .001$. Nevertheless, after null contingency training, the distribution for the prior group again became more peaked and symmetrical, with maximal belief for the "no effect" option. These observations were supported statistically, with a significant decrease in the bias index, $t(47) = 3.48$, $p = .001$, and significant increase in the shape index, $t(47) = 6.99$, $p < .001$, after null contingency training in the prior group. This time however, post-training group differences were not significant for the bias index, $t(97) = 1.07$, $p = .287$, nor the shape index, $t(97) = 0.292$, $p = .771$. Also in contrast to Experiment 1, the bias index was significantly different from 0 in both the prior, $t(47) = 2.42$, $p = .019$, and control groups, $t(50) = 3.60$, $p < .001$, after null contingency training.

Test: Point estimate Figure 3B shows the results for the typical point-estimate causal rating test. Similar to Experiment 1, the prior group's causal rating was significantly greater than 0 at the outset of training, $t(47) = 8.70$, $p < .001$, again demonstrating an initial bias towards expecting a causal relationship. However unlike Experiment 1, after null contingency training, both groups' ratings were significantly greater than 0, smallest $t(47) = 2.89$, $p = .006$, and there was no significant group difference, $t(97) = 0.84$, $p = .403$.

Discussion

The main results from Experiment 2 were that a) participants again showed an initial bias towards expecting causal relationships prior to training, b) unlike Experiment 1, this bias did not fully correct after null contingency training in either the causal rating or bias index, and c) also unlike Experiment 1, there were no group differences in the distribution indices nor the causal ratings after training, with both groups displaying a significant causal illusion. Thus, the inoculation effect observed in Experiment 1 was not observed in Experiment 2.

General Discussion

The first major finding in these experiments was that the distribution of prior beliefs was biased towards expecting causal relationships. In both experiments, the maximum amount of belief was allocated to the two causal options. Contrary to our predictions, the scenario involving the effect of a drug on recovery from a disease (Experiment 2) did not produce a stronger bias towards causal options than the scenario involving the effect of a chemical on cell growth (Experiment 1). The instructions for the two scenarios were devised to be as similar as possible (e.g., both involved the participant assuming the role of a scientist working in a lab), so it is possible that the scientific context overshadowed any effect of the specific cues and outcomes employed.

The second major finding was that reassuringly, the biased prior distribution of beliefs became much more normative after exposure to a null contingency with 50% cue and outcome density. In Experiment 1, the distributional prior became perfectly symmetrical, and the causal bias disappeared. This was consistent with participants' final causal rating, which did not differ significantly from the normative value of zero. In Experiment 2, although the shape of the distribution changed in a similar way to Experiment 1 after training, a residual causal illusion was seen in both the causal rating and in the bias index of the distribution.

It is important to highlight that in both experiments, both groups allocated maximal belief to the normative "no effect" option after null contingency training. Thus, our results show that on the whole, participants do in fact update their beliefs appropriately to arrive at the correct answer. This stands in contrast to interpretations of the classic illusory causation effect as a failure of acquisition, or a demonstration of non-normative learning. It should be emphasized that we do not believe our results speak against the idea of the illusory causation phenomenon being a genuine effect. However, our results do show that there is a high degree of uncertainty from participants in contingency learning tasks, and that participants entertain belief in a range of causal hypotheses, some of which are mutually exclusive with one another. The distributions we obtained were not the result of averaging over participants who allocated belief to a single option. Most participants spread their belief over multiple options, even after training. Thus, a single rating may not always be ideal since it fails to capture these complexities in learning.

Interestingly, while the drug-recovery scenario in Experiment 2 failed to generate more biased prior distributions, there was more residual bias in their post-training ratings. The results from the two cover stories are therefore somewhat contradictory, but may be explained if we distinguish between biases in the prior, and biases in updating. Since the drug-recovery scenario generated a similar distribution of prior beliefs to the chemical-cell growth scenario, the scenario differences must be due to differences in how participants *updated* their prior in response to trials. Thus, one advantage of the Bayesian approach is that differences between experimental groups or individuals can be attributed to differences in the prior (i.e., initial beliefs) and/or to the likelihood (i.e., updating, see Howe et al., 2021).

The final major finding was that we demonstrated that assessing participants' prior beliefs in a range of causal relationships can inoculate participants against showing the causal illusion. In Experiment 1, the prior group did not show the illusion in their final causal ratings, and the distribution of beliefs became perfectly symmetrical. In contrast, the control group who were not assessed on their prior beliefs showed a causal illusion in their causal rating and in the form of a non-symmetrical distribution. However, since this effect was not obtained in Experiment 2, this conclusion is tentative at present. The cover story seems to be critical, and it is unclear whether the inoculation effect would still be present when the cue or outcome occur more frequently, since these conditions are known to inflate the illusion. Unfortunately, it is not possible to discern whether the inoculation effect was due to exposure to the distribution builder, the causal rating,

or the combination of the two. It is also possible that merely alluding to the possibility of a preventative or null relationship in the instructions is sufficient. Certainly, the distributions that we obtained suggest that participants do not give much weight to these possibilities at the outset of training.

In any case, a simple change of causal scenario in Experiment 2 eliminated this inoculation effect. This result was particularly surprising because there were no apparent differences in prior beliefs between the two experiments, and the final causal rating observed in the control group was numerically higher for the chemical-cell growth scenario in Experiment 1 than the drug-recovery scenario in Experiment 2, which is the opposite pattern to what would be predicted if the drug-recovery scenario produced stronger priors and/or less updating. The current experiments were not set up to isolate the key determinant of any scenario differences, and thus we will withhold speculation on the reason for these differences. Nevertheless, our results highlight the need to consider the cover story when designing and reviewing studies on illusory causation.

Returning to explanations of the illusory causation effect, the biased prior distributions that we observed might offer one explanation for why participants weight the cells differently in contingency learning experiments. If participants have a prior belief that a cue and outcome are causally related, they might initially pay more attention to A and B trials to assess the validity of their primary hypothesis (i.e., a "positive test strategy", Mandel & Vartanian, 2009). C and D trials might only become relevant when participants have rejected the hypothesis that the cue and outcome are causally related. This conclusion is similar to the idea that beliefs about causal mechanisms change how participants use covariation information (Fugelsang & Thompson, 2001; 2003), as well as the idea that participants' encoding of covariation information is restricted to participants' focal sets (Cheng, 1997; Cheng & Novick, 1990). Eliciting participants' prior beliefs using the distribution builder may serve to expand participants' focal sets or consider other types of causal mechanisms, allowing them to consider the full range of contingencies.

An obvious and intriguing question that follows from our results is why participants should show an initial causal bias. The bias may exist because causal relationships are encountered more frequently in the real world. Another possibility is that prevention learning is more complex or abstract to represent. Learning about prevention necessarily requires 2 stages: learning that a cue or context causes an outcome, and then learning that the addition of another target cue prevents that outcome. In associative learning, Lee and Lovibond (2021) have shown that after this exact type of training, participants infer a variety of causal structures, with only a subset of participants learning that the target cue is a direct preventor of the outcome. A final but related possibility is that the biased prior was an artefact of the experimental context. Although we were careful to use neutral language (e.g., "Your job is to learn the relationship between…") in order to avoid hinting to participants that the underlying contingency was positive, it is possible that participants nevertheless assumed that this was the case. In other words, participants may assume that an experiment on learning is specifically an experiment on *causal* learning (see Szollosi & Newell, 2020). The laboratory setting employed in both cover stories may have had a similar effect if participants assumed that the scientists would only test chemicals or drugs that were good candidates to produce the effect in question. Some interesting questions for future research are whether a truly neutral causal scenario can be devised, to what extent participants' prior beliefs can be manipulated, and how cue and outcome density influence updating. For example, participants may fail to update their prior under high cue and/or outcome density as frequent cell A trials is consistent with a causal relationship.

One critique of our procedure is that because we used hypothetical stimuli that were relatively immune to real-world knowledge, the distributions of beliefs that we obtained were not genuine, or strongly held beliefs. To examine this possibility, we computed the correlation between prior, and post-training bias and shape indices in both experiments. In both experiments, there was a significant correlation between the prior and post-training shape indices, rs >= .40, ps <= .005, but not between the pre- and post-training bias indices, rs <= .24, ps >= .094. This pattern of results is probably due to the fact that many participants showed a bias index of 0 after training. This result might suggest that participants did not have strong prior beliefs. However, one could also argue that abandoning prior beliefs is a sensible thing to do in this context given the nature of the stimuli. It is also worth noting that participants who were uncertain could have allocated an equal amount of belief to all causal options (and some participants did). Thus, we believe that the distribution builder is a valid method of assessing causal beliefs.

In conclusion, we found that participants entertain a range of causal possibilities in illusory causation tasks. Participants showed an initial bias towards expecting a causal relationship, but this bias was largely corrected after exposure to a null contingency. We found consistent results from our novel distributional measure and the conventional causal rating, lending support to its reliability in assessing causal beliefs. We believe that including distributional measures of prior beliefs has the potential to offer novel insights in understanding the origins of false causal beliefs.

Acknowledgments

Jessica Lee was supported by a Discovery Early Career Researcher Award from the Australian Research Council [DE210100292]. Peter Lovibond was supported by an Australian Research Council Discovery Project Grant [DP190103738]. We would like to thank four anonymous reviewers for helpful comments on this paper.

References

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

Alloy, L. B., & Abramson, L. Y. (1979). Judgement of contingency in depressed and nondepressed students: Sadder but wiser? *Journal of Experimental Psychology: General, 108(4)*, 441-485.

Béghin, G., & Markovitz, H. (2023). Reasoning strategies and prior knowledge effects in contingency learning. *Memory & Cognition, 50*, 1269-1283.

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 & Behavior, 41*, 333-340.

Chapman G. B., & Robbins, S. J. (1990). Cue interaction in human contingency judgment. *Memory & Cognition, 18(5)*, 537-545.

Cheng, P. W. (1997). From covariation to causation: A causal power theory. *Psychological Review, 104*, 367-405.

Cheng, P. W., & Novick, L. R. (1990). A probabilistic contrast model of causal induction. *Journal of Personality and Social Psychology, 58*, 545-567.

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.

de Leeuw, J. R. (2015). jsPsych: A JavaScript library for creating behavioral experiments in a web browser. *Behavior Research Methods, 47(1),* 1-12.

Fuselgang J. A., & Thompson, V. A. (2000). Strategy selection in causal reasoning: When beliefs and covariation collide. *Canadian Journal of Experimental Psychology, 54(1)*, 15-32.

Fuselgang J. A., & Thompson, V. A. (2003). A dual-process model of belief and evidence interactions in causal reasoning. *Memory & Cognition, 31(5),* 800-815.

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

Lange K, Kühn S, & Filevich E. (2015). "Just Another Tool for Online Studies" (JATOS): An Easy Solution for Setup and Management of Web Servers Supporting Online Studies. *PLOS ONE 10(7)*: e0134073.

Lee, J. C., & Lovibond, P. F. (2021). Individual differences in causal structures inferred during feature negative learning. *Quarterly Journal of Experimental Psychology, 74(1)*, 150-165.

Lee, J. C., Lovibond, P. F., Hayes, B. K., & Lewandowsky, S. (2021). A mixture of experts in associative generalization. *Proceedings of the 43rd Meeting of the Cognitive Science Society.* (pp. 341-347). Cognitive Science Society.

Howe, P., Perfors, A., Walker, B., Kashima, Y., & Fay, N. (2021). Base rate neglect and conservatism in probabilistic reasoning: Insights from eliciting full distributions. *PsyArXiv*. https://doi.org/10.31234/osf.io/q7znk

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.

Matute, H., Blanco, F. & Díaz-Lago, M. (2019). Learning mechanisms underlying accurate and biased contingency judgments. *Journal of Experimental Psychology: Animal Learning & Cognition, 45*, 373-389.

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 can be reduced. *Frontiers in Psychology, 6*, 888.

Ng, D. W., Lee, J. C., & Lovibond, P. F. (2023). Unidirectional scales overestimate the illusory causation phenomenon. *Quarterly Journal of Experimental Psychology.*

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, 2*, 64-99.

Quentin, A. (2016). distBuilder v1.1.0 (1.1.0). Zenodo. https://doi.org/10.5281/zenodo.166736

Shanks, D. R. (1987) Acquisition functions in contingency judgement. *Learning and Motivation, 18*, 147-166.

Shanks, D. R. & Dickinson, A. (1987). Associative accounts of causality judgement. In G. H. Bower (Ed.), *The psychology of learning and motivation* (Vol. 21). New York: Academic Press.

Szollosi, A., & Newell, B. R. (2020). People as intuitive scientists: Reconsidering statistical explanations of decision making. *Trends in Cognitive Science, 24(12),* 1008-1018.

Vicente, L., Blanco, F., & Matute, H. (2023). I want to believe: Prior beliefs influence judgements about the effectiveness of both alternative and scientific medicine. *Judgment and Decision Making, 18*, e1.

Wasserman, E. A., Dorner, W. W., & Kao, S. F. (1990) Contributions of specific cell information to judgements of interevent contingency. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 16(3*), 509-521.

White, P. A. (1989). A theory of causal processing. *British Journal of Psychology, 80*, 431-454.