



King's Research Portal

Document Version
Peer reviewed version

[Link to publication record in King's Research Portal](#)

Citation for published version (APA):

Shankar, S., Byrom, N., van Tilburg, W., & Rakow, T. (in press). On the interchangeability of presentation order for cause and effect: Experimental tests of cue and outcome density effects. *Quarterly Journal of Experimental Psychology*.

Citing this paper

Please note that where the full-text provided on King's Research Portal is the Author Accepted Manuscript or Post-Print version this may differ from the final Published version. If citing, it is advised that you check and use the publisher's definitive version for pagination, volume/issue, and date of publication details. And where the final published version is provided on the Research Portal, if citing you are again advised to check the publisher's website for any subsequent corrections.

General rights

Copyright and moral rights for the publications made accessible in the Research Portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognize and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the Research Portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the Research Portal

Take down policy

If you believe that this document breaches copyright please contact librarypure@kcl.ac.uk providing details, and we will remove access to the work immediately and investigate your claim.

**On the interchangeability of presentation order for cause and effect: Experimental tests
of cue and outcome density effects.**

Sahana Shankar

King's College London & Royal

Holloway University of London

Nicola Byrom

King's College London

Wijnand A. P. van Tilburg

University of Essex

Tim Rakow

King's College London

**Author Accepted Version for the Quarterly Journal of Experimental Psychology,
accepted for publication 01 October 2024**

Word count: [13100]

Sahana Shankar Psychology Department; Institute of Psychiatry, Psychology and Neuroscience (IoPPN); King's College London; London, United Kingdom and University of Royal Holloway, Egham, United Kingdom. Tim Rakow; Psychology Department; Institute of Psychiatry, Psychology and Neuroscience (IoPPN); King's College London; London, United Kingdom. Wijnand A. P. van Tilburg; Department of Psychology, University of Essex; Colchester, United Kingdom. Nicola Byrom; Psychology Department; Institute of Psychiatry, Psychology and Neuroscience (IoPPN); King's College London; London, United Kingdom.

All procedures performed involving human participants were in accordance with the ethical standards of the institutional research committee. The authors declare that there are no potential conflicts of interest with respect to the research, authorship, and/or publication of this article. All authors consented to the submission of this manuscript. Corresponding

author: Sahana Shankar; Psychology Department; Royal Holloway University of London,
Egham; sahana.shankar@rhul.ac.uk

LISS-DTP funded this research [ES/P0000703/1]. For the purpose of Open Access,
the author has applied a CC BY public copyright license to any Author Accepted Manuscript
(AAM) version arising from this submission.

Abstract

Studies of cue-outcome contingency learning demonstrate outcome-density effects: participants typically overestimate contingencies when the outcome event is relatively frequent. Equivalent cue-density effects occur, though these have been examined less often. Few studies have simultaneously examined both those event density effects or have manipulated the presentation order of the events, limiting knowledge of whether these phenomena share underlying principles. We report three well-powered experiments to address those gaps. Participants judged the effectiveness of a medical treatment after viewing a series of pairings for two events, a cause (treatment given vs. not) and an effect (patient recovered vs. not). Experiment 1 manipulated both event densities independently. We then manipulated the presentation order for the cause and the effect, alongside a manipulation of effect density (Experiment 2a) or cause density (Experiment 2b). Experiment 1 found a large main effect of event-density ($\eta_p^2 = .55$), which was qualified by a significant interaction between event type and density level ($\eta_p^2 = .10$) whereby effect density had greater impact than cause density. Experiments 2a and 2b found effects for effect-density ($\eta_p^2 = .60$) and cause-density ($\eta_p^2 = .31$). The effects of cause-effect presentation order were always small and non-significant. We conclude that effect-density manipulations had substantial impact on contingency judgments, and cause-density manipulations less so. Moreover, it matters little which event (cause or effect) is seen first. These findings have implications for contingency, associative, probabilistic, and causal models of contingency judgment; primarily, that people may be more sensitive to the causal status of events than to their temporal order of presentation.

On the interchangeability of presentation order for cause and effect: Experimental tests of cue and outcome density effects

Successfully identifying associations between events is fundamental to cognition (Shanks, 2010). This form of pattern detection is a basic building block of the learning required for adaptive functioning, ranging from learning which words represent which objects or ideas, to learning which outcome(s) reliably follow a specific type of action (Batterink & Neville, 2011; Shanks, 1993; Shelton & Martin, 1992; Werker et al., 1998). Learning theorists have therefore sought to understand whether and how associations and causes can be accurately identified—and to identify any conditions that hinder people’s ability to accurately gauge the strength of an association or causal path between a cause and the subsequent effect. One such hindrance relates to the relative frequency, or ‘density’, of events. For example, when an allergic reaction occurs, individuals typically will know the effect (allergic reaction) and may also identify a putative cause (allergen). Changes to cause-density (allergen prevalence) or effect-density (allergy prevalence) can alter judgments of strength of association, even when the actual relationship between events remains constant.

The most prominent of such density related biases occur when, due to high frequency of the cause or effect, a cause-effect association is perceived when none exists (Allan et al., 2005; Matute, 1995; Musca et al., 2010). This can lead people to attribute a causal relationship between two independent events, called the illusion of causality (Matute et al., 2015). This phenomenon is important because it represents a common error in human learning that is thought to contribute to the development and maintenance of superstitious beliefs and pseudoscientific thinking (Matute et al., 2015).

In this paper, we examine whether biases arise in judgments of contingency as cause (cue) density and effect (outcome) density are manipulated. These manipulations are made while varying the presentation order of putative causes and effects to see how these variations

influence judgments. Within much of the associative and contingency learning literature, causes are referred to as cues whilst effects are referred to as outcomes. We will use the terminology of cue and outcome across this manuscript.

Alloy and Abramson (1979) conducted one of the first studies of the outcome density related bias, using a modified paradigm previously introduced by Jenkins and Ward (1965). In a series of contingency problem trials, participants were invited to choose whether or not to press a button to assess the relationship between button pressing and the light turning on. These researchers manipulated both the contingency between button pressing and the light turning on, as well as the overall frequency of the light turning on—referred to as the outcome density. In some conditions the probability that the light would turn on after pressing the button was higher than for not pressing the button, i.e., there was a positive contingency between button pressing and light turning on. In other conditions, the probabilities of the light turning on were equal for button presses and non-presses, equivalent to zero contingency. In some conditions the overall probability of the light turning on was higher than 50%, whereas in others it was not, representing high and comparatively low outcome density, respectively. After completing 40 trials, participants' judgments of the relationship between cue (button pressing) and effect (light illumination) were measured. An outcome density related bias was found, whereby the relationship between cue and outcome was judged to be stronger when outcome density was high rather than low, even when there was no contingency between the events (Alloy & Abramson, 1979).¹ Further studies have found outcome density related biases to be robust across variations of the original paradigm (Blanco et al., 2013; Crump et al., 2007; Gillan et al., 2014; Matute et al., 2007; Msetfi et al., 2005; Vallée-Tourangeau et al., 2005). These variations include an allergy paradigm

¹ This original study shows that the tendency to ascribe cause when none existed was diminished in a sub-group of participants with depression—a finding which some describe as a demonstration of 'depressive realism.'

(assessing which foods cause allergy; Blanco et al., 2013) and a medical treatment paradigm (judging the effectiveness of a treatment for a disease; Blanco et al., 2011).

Though less frequently studied than outcome density, bias in judgment related to cue density have also been found for contingency learning. Using an allergy paradigm, Vadillo et al. (2011) manipulated cue density to be either low or high, and found a cue density bias whereby judgments of association increased with cue density (Vadillo et al., 2011). Specifically, participants believed more strongly that a substance was an allergen when that substance featured in more than 50% of learning trials, despite the contingency between substance and allergic reaction remaining constant. Perales et al (2005) extended this work by manipulating cue density, and simultaneously varying whether or not there was a positive contingency. An effect of cue density was found, with contingency judgments being higher in the high density conditions compared to the low density conditions. However, while outcome density was held constant within each contingency, it varied between the positive and zero contingencies—thereby confounding contingency and outcome density and making it difficult to assess whether contingency level moderates the size of the outcome density bias.

Few studies have simultaneously (and systematically) examined cue density and outcome density biases. In one such rare investigation, Blanco et al (2013) used an allergy paradigm with zero cue-outcome contingency, and employed a 2×2 design with cue density (high vs low) and outcome density (high vs low) to examine the interactive effects of cue and outcome densities. They found an effect of outcome density in all conditions, and a small effect of cue density, but only when outcome density was high. They also ran a similar experiment with a neutral scenario where the stimuli consisted of geometrical figures and found similar effects. However, Blanco et al did not manipulate the order of cue and outcome presentation. So, while their findings may be interpreted in line with outcome and cue density biases, it is conceivable that these biases actually relate to the order of event occurrence; that

is, manipulating the density of the ‘second event’ to be presented may have a stronger effect than manipulating the density of the ‘first event’ presented, regardless of whether these events are potential causes or effects. Here, an important distinction in terminology is that ‘cues’ are not always ‘causes’ and ‘effects’ are not always ‘outcomes’. This is, because ‘cue’ refers to the event presented *first*, whilst ‘outcome’ refers to the *event to be predicted* from the cue. In most cases, this means cues are causes and outcomes are biases (i.e., following the natural order to temporal precedence). However, the order in which causes and effects are *presented* (i.e., made known) can be switched. If so, because the presentation order of cue (first) and outcome (last) remains unchanged, then the effect is a ‘cue’ while the cause is an ‘outcome’. As such, for clarity from this point onwards, ‘cue’ refers to the first presented event, and ‘outcome’ is the second presented event; and rather than referring to cue and outcome density biases, we will use the terms *cause-density* and *effect-density* biases, respectively.

Overview of Research Questions and Experiments

We systematically examined the relative size of cause-density and effect-density in contingency learning, through a three-experiment investigation that incorporated and improved upon methods used in previous studies. These experiments provided tests of contingency learning theory, probabilistic contrast models, causal theories, and associative models. In Experiment 1, we compared cause-density and effect-density biases directly, by manipulating these event densities in a contingency learning paradigm. By investigating cause-density and effect-density biases systematically, we investigated whether the same cognitive principles underlie both effects. If the effect sizes differ and/or one effect is conditional on the level of the other variable (with the other variable always having an effect) then this suggests that the mechanisms operating on these events are (at least partially) distinct, such as the importance of causal status assigned to causes and effects. This first

experiment thus closely replicated the study by Blanco et al (2013), who found evidence for both effect- and cause-density biases, with the former exceeding the latter in size.

After confirming the presence and relative sizes of effect and cause-density biases, we then extended our investigation in Experiments 2a and 2b by additionally manipulating the presentation order for the cause and the effect. Critically, this allowed us to examine the impact of presentation order independently of cause versus effect density. This extension to the contingency literature reflects that while, logically, a cause must precede an effect, sometimes people observe or learn of an effect *before* they see or hear of its possible causes. For example, a doctor may observe signs of infection in patients before learning about the possible causes these patients were exposed to. A manager might see that their salesforce has made a number of successful sales before they find out what steps their workers took in an attempt to achieve those successes. Some indication that the presentation order of the cause and the effect can be important for processing information comes from the causal theory literature on semantic research. Fenker et al. (2005) presented participants with pairs of word which described events as causes (e.g., rain) or effects (e.g., flood). With a temporal manipulation of these words, participants' reactions were quicker when cause preceded effect compared to when effects preceded the cause. However, it is not clear how such an effect would translate to contingency learning; for example, whether or how the biases that are associated with increasing the event density are affected by event presentation order. Therefore, our experiments address a key aspect of theory that may have important applications beyond the lab.

Manipulating the event presentation order afforded the exploration of interaction effects involving *both* density level and event presentation order. This allowed us to test whether the influence of cause-density or effect-density on contingency judgment is altered by reversing the order in which event information is seen. Thus, to contrast causes and effects

on the basis of their roles in defining contingencies (as distinct from their order of observation) the temporal confound must be removed. Manipulating the presentation order of causes and effects allows one to do so, and the impact of this temporal ordering can be directly assessed. For example, it may be that manipulating cause-density and effect-density have similar impact on contingency judgments, provided that they occupy the same temporal position (e.g., both presented last as the effect), but not when occupying different temporal positions (e.g., cause presented as cue, effect presented as outcome). This latter ordering is the default in past studies (Blanco et al., 2013; Perales et al., 2005; Vadillo et al., 2011). We could therefore assess whether previous observations that outcome-density biases are larger than cue-density biases instead reflect that density biases are larger for whichever event is the cause rather than the effect. As such, we wished to explore the effect of this temporal ordering and did not know how this would affect the relative size of effects for cause density and effect density.

General Method

We examined event-density effects in three contingency learning experiments. Manipulations of cause-density and effect-density were examined in the context of medical judgments about fictitious illnesses. We employed a standard contingency learning paradigm throughout these studies with important incremental variations from one experiment to the next. We outline the common design and experiment idiosyncrasies below.

Design

Each experiment had a $2 \times 2 \times 2$ mixed design, with judgment of effectiveness of a medication as the dependent variable. Three factors were varied in each experiment: event manipulated, session order, and density level.

Each experiment required participants to complete two sessions, with cause-density manipulated in one session and effect-density manipulated in the other. Thus, the density event manipulated (e.g., cause-density vs. effect-density) was varied within subjects. The order of these sessions was randomly assigned and formed our second independent variable (between-subjects, e.g., cause condition then effect condition vs. effect condition then cause condition). The sessions were a minimum of three days apart to minimize carry-over effects. Within each session, we manipulated density level to be high [0.8] vs. low [0.2] for the corresponding event (see Table 1). The density levels for effect refer to the probability of effect given cause and probability of the effect given no cause (for which probabilities were identical because the cause-effect contingency was always zero). The higher the effect-density the more frequently the effect occurs with and without the cause. The density levels for cause refer to the probability of the cause given the effect and probability of the cause given no effect. The higher the cause-density the more frequently the cause occurs with and without the effect. Manipulating densities in this way created a low and a high-density block of trials, each with 40 trials. When cause-density was manipulated, effect-density was held at

Table 1
Density levels across conditions

Table 1a
Low cause density at Zero Contingency

Medication	Recovery	
	Present	Absent
Present	4	4
Absent	16	16
	P(O/R)=.2	P(O~/R)=.2
	$\Delta p = .2 - .2 = 0$	

Table 1b
High cause density at Zero Contingency

Medication	Recovery	
	Present	Absent
Present	16	16
Absent	4	4
	P(O/R)=.8	P(O~/R)=.8
	$\Delta p = .8 - .8 = 0$	

Table 1c
Low effect density at Zero Contingency

Medication	Recovery		
	Present	Absent	
Present	4	16	P(O/R)=.2
Absent	4	16	P(O~/R)=.2
	$\Delta p = .2 - .2 = 0$		

Table 1d
High effect density at Zero Contingency

Medication	Recovery		
	Present	Absent	
Present	16	4	P(O/R)=.8
Absent	16	4	P(O~/R)=.8
	$\Delta p = .8 - .8 = 0$		

a medium density of 0.5, for both high and low cause-density; and vice versa when effect-density was manipulated. Previous studies have often used only high and low densities (Blanco et al., 2013), for this study we chose the non-manipulated density to be held at a medium level whilst the manipulated density was either high or low. We chose this set up to reduce the possible interaction effects found previously between effect and cause densities (Blanco et al., 2013). The block shown first was chosen randomly. Thus, each participant provided a measurement for each of the four density-event by density-level combinations and did so in one of two session orders. Different medication and illness names were used for each of these four combinations of density event by density level, assigned at random (see Table S1). The blocks were counterbalanced.

Participants

Participation was via Prolific.co (Prolific, 2020), an online research participant platform with approximately 100,000 registered volunteers at the time of data collection. Only those aged 18 or above and resident in the UK could participate. All participants declared they were fluent in English (Table 2 for demographics).

Task and Procedure

The task was based on Blanco et al (2014; images used with permission) and was implemented via the online Gorilla experiment generator (Gorilla Experiment Builder, 2020). After signing up, participants were prompted to complete the first session of the experiment.

Table 2

Demographics of participants who completed both sessions for each experiment

	Experiment 1 (<i>N</i> =85)	Experiment 2a (<i>N</i> =96)	Experiment 2b (<i>N</i> =97)
Age, modal group, (<i>range</i>)	21-25, (18-70)	26-30 (18-71+)	31-35 (18-70)
Female gender, <i>N</i> (%) <i>of sample</i>)	50 (58.8%)	62 (64.6%)	60 (61.9%)

After 3 days, participants were invited to complete the second session of the experiment, which followed the same procedure as the first session. In each session, participants read an information sheet, provided consent, answered demographic questions (first session only), and then read task instructions. In the task, participants imagined that they were a doctor specialising in new illnesses and their potential treatments. The participants' instructions explained that their patients were suffering from a (named) syndrome and may be treated with a specific medication. They then proceeded to the first trial.

Within each trial, the following images were shown (Figure 1), in a standard presentation order: fixation cross for 750ms, unhealthy patient image for 2,000ms, treatment cause image (medication image vs. no medication image) for 3,000ms, treatment effect image (recovered patient vs. not recovered patient) for 3,000ms. These images automatically followed each other, without input from the participant. Between trials, participants clicked 'next' to continue. Each block contained 40 trials presented in a random order; each session had 2 blocks (low vs. high density), with block order randomised.

At the end of each block, participants judged medication effectiveness in response to the prompt: "To what extent was the medication effective in treating the illness the patients were suffering from?" To do so they used a slider, ranging from 0 ("it was not effective at all") to 100 ("it was perfectly effective"). There was no slider handle displayed until a response was given to avoid judgment anchoring (Mussweiler et al., 2012).

Each session lasted approximately 20 minutes, and at the end of the entire study participants read a short debriefing. Participants received UK £5.00 for completing the experiment. The studies received ethical approval from the Research Ethics Committee of King's College London, MRSP-19/20-19281.

Experiment 1: Comparing Cause-density and Effect-density Biases

Method

This experiment investigated the effect on contingency judgment of equivalent manipulations of cause-density and effect-density, to see if, and by how much, effect-density biases and cause-density biases differ in size. Previous research found cause-density biases to be smaller than effect-density biases (Blanco et al., 2013). As such, this study adds to the literature by conducting a highly powered, within-subjects comparison of cause-density and effect-density manipulations.


Participants

Ninety-one volunteers participated, with 85 participants completing both sessions; see Table 2 for participant demographics. For a 2×2 within-subjects interaction of assuming (a) an attenuation interaction pattern, (b) any mean differences being 0.5 (all $SD = 1$), (c) $\alpha = .05$, (d) 40 people per group, (e) 0.5 common correlation with no sphericity, a power of 87% is achieved. The study thus provides an adequate sample size to test the hypothesis.

Design and Analysis

The design and procedure were as described in the General Method above. Three factors were varied in the experimental design: manipulation of event (within subjects: cause-density vs. effect-density) by density level (within subjects: high density vs. low density) by session order (between subjects: cause-density manipulation first vs. effect-density manipulation first). Half of the sample completed the cause-density condition during the first

Figure 1
Stimuli images



session, and the effect-density condition in the second session, with the other half of the sample completing the effect-density condition first and the cause-density second. This was included as a between-subjects factor in the analysis. By accounting for whether participants saw the cause-density or event-density first in the first session, the analysis automatically accounts for the condition participants were in for the second session.

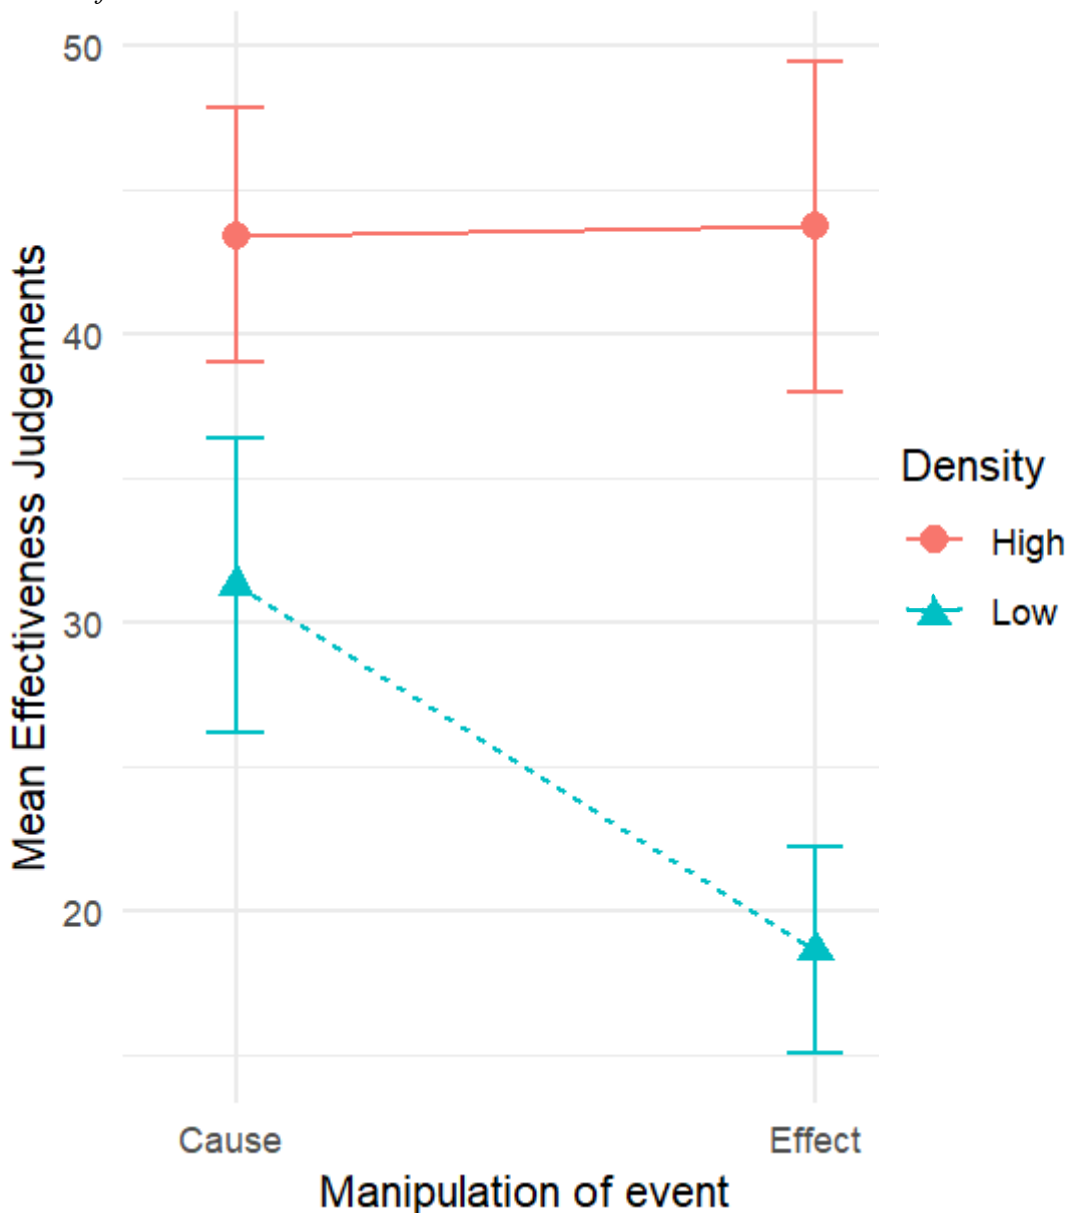
We ran a $2 \times 2 \times 2$ mixed analyses of variance (ANOVA), with the dependent variable of judgments of the effectiveness of the medication. In this factorial analysis, it is the effects involving density level that are of prime interest for our hypotheses because these relate to the impact that cause/effect-density has on contingency judgments.

Results and discussion

As predicted, there was a significant main effect of density level, $F(1, 83) = 101.91, p < .001, \eta_p^2 = .55$ (very large effect). Participants judged the medication to be less effective in the low-density conditions ($M = 24.88, SD = 17.22$) than the high density conditions ($M = 43.58, SD = 19.06$). There was a medium-sized main effect of event on the participants judgments of effectiveness, $F(1, 83) = 11.16, p = .001, \eta_p^2 = .12$. Participants judged the medication to be less effective in the effect-density condition ($M = 31.09, SD = 18.06$) than the cause-density condition ($M = 37.37, SD = 18.38$). That is, overall participants provided lower judgments of the effectiveness of the medication in conditions where the effect-density was being manipulated, than conditions where the cause-density was manipulated. There was an interaction effect of density level by event, $F(1, 83) = 8.93, p = .004, \eta_p^2 = .10$. As shown in Figure 2 and Table 3, the density effect, that is, the difference between judgments made in a high density condition and judgments made in a low density condition, was nearly three times larger in the effect-density condition compared to the cause-density condition. This comparison places the point estimate of the density effect as nearly three times larger in the effect-density condition compared to the cause-density condition.

To further understand the two-way interaction, both pairs of simple main effects were examined. Two paired-samples *t*-tests verified that the density effect was present in both the cause and the effect-density conditions. Participants judged the effectiveness of the medication to be significantly lower in the low density condition than the high density condition, both in the cause-density condition, $t(84) = 4.54, p < .001$, and in the effect-density condition $t(84) = 8.74, p < .001$.

Figure 2
Mean effectiveness judgments by manipulation of event for high and low densities, with 95% confidence intervals



Paired-samples *t*-tests were also conducted to verify whether there was a significant effect of event in both the high density and low density condition. Participants’ judgments of effectiveness of the medication did not differ significantly between manipulations of cause and of effect in the high density condition, $t(84) = -.092, p = .927$. However, judgments did differ significantly in the low density condition, $t(84) = 5.10, p < .001$, with participants giving lower judgments for the effect-density condition than for the cause-density condition (Table 3).

We found a significant interaction between session order and density level, $F(1, 83) = 8.93, p = .004, \eta_p^2 = .10$ (medium-sized effect). This suggests that the size of the density effect varies according to the order in which events are manipulated; cause-density manipulation first or effect-density manipulation first. This two-way interaction is further qualified by the significant three-way interaction for event manipulated by density level by

Table 3
Mean (SD) response by group (Cause then Effect and Effect then Cause), split by session and manipulation of cause or effect

	Judgments split by group				Overall judgments	
	Cause then Effect		Effect then Cause		Type of Density Manipulation	
	1 st session: Cause	2 nd session: Effect	1 st session: Effect	2 nd session: Cause	Cause	Effect
High density condition	46.95 (21.49)	39.61 (27.33)	47.59 (26.19)	40.18 (19.81)	43.45 (20.80)	43.74 (16.89)
Low density condition	27.83 (20.34)	17.78 (17.19)	19.39 (16.69)	34.52 (26.72)	31.29 (23.95)	18.61 (16.85)
95% CI for mean difference	11.69, 26.56	13.94, 29.72	19.76, 36.65	-2.50, 13.81	6.54, 17.77	19.41, 30.85
Density effect size <i>d</i>	0.91	0.96	1.28	0.24	0.54	1.49

session order, $F(1, 83) = 5.51, p = .021, \eta_p^2 = .06$. This suggests that the differential impact on judgment of the cause-density and effect-density manipulations varied as a function of which session took place first (viz. cause condition vs. effect condition). Therefore, we compared the size of the density effect across each combination of session order, and manipulation of event (Table 3 and Figure 3). This demonstrates that the density effect was present for both the cause-density and the effect-density manipulation when this manipulation took place in the first session, with both being large effects (paired samples t-test for difference between high and low for cause, $t(41) = 5.20, p < .001$, and for effect, $t(44) = 6.74, p < .001$ in the first session). In the second session, the effect-density effect was present (paired samples t-test for effect of density for the effect condition in second session, $t(41) = 5.59, p < .001$) but there was no significant effect of cause-density manipulation (paired samples t-test for effect of density in the cause condition in second session, $t(44) = 1.40, p = .169$). In sum, from the two-way and three-way interactions involving density level and event presentation order we conclude that the effect of density level is less consistent for manipulations of cause-density than it is for manipulations of effect-density of an equivalent scale, such that a cause-density effect on contingency judgment is only detected in some of the conditions of this experiment. Context or contrast effects can counteract the effect of cause-density as indicated by the moderating effect of session order. This is in line with previous research that examined the interactive effects of cause densities and effect densities on contingency judgments, which found that cause effects were only found with high effect densities (Blanco et al., 2013).


All other main effects and two-way interactions were small and non-significant, there being no main effect of session order, $F(1, 83) = .47, p = .467, \eta_p^2 = .006$; no significant interaction between density and session order, $F(1, 83) = .92, p = .342, \eta_p^2 = .01$; and no

significant interaction between event or session order, $F(1, 83) = 1.65, p = .202, \eta_p^2 = .02$ (see Supplementary Materials for additional analyses within each session).

In sum, we examined and compared the influence of cause-density and effect-density on judgments of medication effectiveness in an online contingency-learning paradigm. The findings suggest that both cause-density and effect-density can impact contingency judgments, but that the impact is smaller and less stable for manipulations of cause-density as compared to equivalently-sized manipulations of effect-density.

Figure 3

Mean effectiveness judgments by group (Cause then Effect, Effect then Cause), by session, and manipulation of event for high and low densities, with 95% confidence intervals



Experiment 2a and 2b: Density Effects across Event Presentation Orders

Method

Experiments 2a and 2b investigated whether the impact of cause-density and effect-density on judgments of effectiveness of medication depends on the order in which the cause and the effect are presented. Experiment 2a investigated whether this manipulation of cause-effect presentation order moderated the effect of *effect-density* on contingency judgment. In a symmetrical fashion, Experiment 2b investigated whether the same manipulation of cause-effect presentation order moderated the effect of *cause-density* on contingency judgment.

A similar design to the one outlined in the General Methods was employed. However, instead of manipulating the density event, the order of event presentation was manipulated. In the trials of one session, the cause was shown prior to the effect, as per Experiment 1. This thus followed the presentation of fixation cross, unhealthy patient image, treatment cause image then treatment effect image. In the trials of the other session, the effect was shown first, followed by the cause; the image order being fixation cross, unhealthy patient image, treatment effect image then treatment cause image. Participants across all conditions were given the same task instructions, there were no mention of the reverse order of events in the instructions.

Experiment 2a: Effect-density and Event Presentation Order

Participants

Participants were 110 Prolific members; 107 participants completed the first session, and 96 participants completed both sessions (see Table 2).

Results and Discussion

We ran a $2 \times 2 \times 2$ mixed ANOVA. The dependent variable was judgment of medication effectiveness. The three independent variables were event presentation order (within subjects: cause then effect vs. effect then cause), density level (within subjects: high

vs. low density), and session order (between subjects: cause-then-effect in first session vs. effect-then-cause in first session). Effect-density biases were expected in both event presentation order conditions.

As predicted, and replicating the findings of Experiment 1, there was a large and statistically significant main effect of effect-density, $F(1, 94) = 141.34, p < .001, \eta_p^2 = .60$. Participants judged the medication to be less effective in the low density condition ($M = 19.16, SD = 11.51$) than the high density condition ($M = 43.83, SD = 22.16$). There was no main effect of event presentation order, $F(1, 94) = 2.64, p = .107, \eta_p^2 = .03$ (small effect); judgments of medication effectiveness did not differ significantly between the cause first ($M = 30.13, SD = 16.82, 95\% CI [19.77, 31.05]$) and effect first ($M = 32.86, SD = 16.42, 95\% CI [19.03, 28.70]$) conditions. As this is a well-powered experiment with reasonably narrow *CI*s for means and mean differences, these findings suggest that event presentation order makes little difference to the judgment of medication effectiveness.

Our analysis showed no interaction effect of event presentation order by density level, $F(1, 94) = 0.21, p = .647, \eta_p^2 = .002$; the size of the density effect did not vary significantly by event presentation order (Table 4). Thus, neither here, nor for the main effect of event presentation order, do we find statistically significant evidence that event presentation order

Table 4

Experiment 2a and 2b mean (SD) judgment of medication effectiveness response by event presentation order and density

	Experiment 2a: Outcome density		Experiment 2b: Cue density	
	Cause First	Effect First	Cause First	Effect First
High density response	42.93 (26.96)	44.70 (23.99)	46.55 (19.10)	46.87 (21.31)
Low density response	17.52 (14.94)	20.83 (15.65)	35.70 (20.33)	31.58 (23.23)
95% CI for mean difference	19.77, 31.05	19.03, 28.70	6.07, 15.61	9.63, 20.95
Density effect size <i>d</i>	1.17	1.18	0.55	0.69

made a difference. Additionally, there was no significant three-way interaction for event presentation order by density level by session order, $F(1, 94) = 0.00, p = .972, \eta_p^2 = .000$ (see Supplementary Materials for analysis on session order).

None of the other effects in this ANOVA had any bearing on the hypotheses of this experiment, and all were small in size. There was no significant main effect of session order, $F(1, 94) = 0.00, p = .997, \eta_p^2 = .000$; and no significant interaction effect for density level by session order, $F(1, 94) = .07, p = .791, \eta_p^2 = .001$, or for event presentation order by session order, $F(1, 94) = 3.29, p = .073, \eta_p^2 = .03$.

Our primary interest in this experiment was whether event presentation order moderated the effect-density bias. We found no significant evidence for such moderation. However, null-hypothesis significance testing is not designed to evaluate evidence for the absence of an effect. We undertook two analyses to further evaluate how well the data aligned with the null hypothesis that the effect-density bias is not affected by the order in which events are presented.

First, we used confidence intervals to identify the plausible size of the (population) effect that event presentation order has on the effect-density bias (Greenwald, 1975). To do so, we computed the effect-density bias for each participant (i.e., the difference in judgment, for high effect-density minus low effect-density condition). The effect size for the difference in effect (cause first minus effect first) was $d = .047, 95\% \text{ CI } [-0.153, +0.247]$. From this, we conclude that if event presentation order moderates the size of the effect-density bias, that moderating effect is likely very small (Cohen, 2013).

Second, we computed a Bayes Factor to assess the evidence for the point null hypothesis of zero moderating effect of presentation order ($d = 0$) against an alternative prior (Dienes, 2014). Following the principles outlined by Dienes (2019), we used an alternative prior that was informed by our previous data for the size of the effect-density bias and by

theoretical considerations that specified a reasonable range of predicted effect sizes. We therefore used a normally distributed alternative prior with a mean of $d = 0$ ($SD = 0.75$). As suggested by Dienes (2019), this SD was chosen to be half of the effect size for the effect-density manipulation that had been observed in Experiment 1.² We found a Bayes factor of $BF_{01} = 6.66$, demonstrating that the odds of the data are around 6 to 7 times greater under the null hypothesis than under the alternative prior that we tested. This is conventionally described as ‘substantial’ evidence for H_0 over H_1 .

Experiment 2b: Cause-density and Event Presentation Order

Participants

One-hundred volunteers participated, with 97 participants completing both sessions (see Table 2).

Results and Discussion

We performed a $2 \times 2 \times 2$ mixed ANOVA. The dependent variable was judgment of medication effectiveness. The three independent variables were: event presentation order (within subjects: cause then effect vs. effect then cause); cause-density (within subjects: high vs. low density); and session order (between subjects: cause-then-effect in first session vs. effect-then-cause in first session). As per Experiment 2a, it is the interaction effects involving *both* event presentation order and cause-density that are important for the main contribution of this experiment. In line with Experiment 1, a cause-density bias was expected in both

² With a normal alternative prior centred on $d = 0$, and because most of the density in a normal distribution lies within 2 SDs of its mean, this generates a reasonable range from $d = -1.5$ to $d = +1.5$ for the moderating effect of presentation order on effect-density bias. The lower bound of $d = -1.5$ represents the case that all of the effect-density bias relies on the cause being presented first rather than second. We regard that lower bound as reasonable because we know of no model that predicts a reversal of the effect density bias under any conditions. The upper bound of $d = +1.5$ reflects the case whereby the effect-density bias is substantially larger (i.e., it doubles in size) when the cause is presented after the effect as compared to when the cause is presented first as in most studies (including our Experiment 1). Between these bounds, having a normal distribution with most of its density close to zero reflects that small- and medium-sized effects are more common than large ones in psychological research.

event presentation order conditions, but these biases were expected to be smaller, on average, than the effect-density bias in Experiments 1 and 2a.

Replicating the findings of Experiment 1, there was a main effect of cause-density $F(1, 95) = 42.09, p < .001, \eta_p^2 = .31$ (large effect). Participants judged the medication to be less effective in the low density condition ($M = 33.63, SD = 17.34$) than the high density condition ($M = 46.68, SD = 15.34$). There was no main effect of event presentation order on judgments of effectiveness $F(1, 95) = 0.89, p = .887, \eta_p^2 = .009$. Participant judgment of medication effectiveness did not differ significantly between the cause first ($M = 41.12, SD = 15.88, 95\% CI [37.92, 44.32]$) or effect first ($M = 39.19, SD = 17.02, 95\% CI [35.75, 42.62]$) conditions.

There was no statistically significant event presentation order by density level interaction, $F(1, 95) = 1.68, p = .198, \eta_p^2 = .02$ (small effect). The size of the density effect did not vary significantly by event presentation order (Table 4). Together with the absence of a main effect of event presentation order, this suggests that there is little evidence that event presentation order made a difference. There was no significant three-way interaction of event presentation order by density level by session order, $F(1, 95) = 0.12, p = .731, \eta_p^2 = .001$ (see Supplementary Materials for further details on session order).

None of the other effects in this ANOVA analyses have any bearing on the hypotheses of this experiment, and all were small in size. There was no significant main effect of session order, $F(1, 95) = 2.09, p = .152, \eta_p^2 = .02$; and no significant interactions for density by session order, $F(1, 95) = 1.26, p = .265, \eta_p^2 = .01$, or event presentation order by session order, $F(1, 95) = 2.66, p = .106, \eta_p^2 = .03$.

As for Experiment 2a, we re-examined the (non-significant) moderating effect of event presentation order on the cause-density effect. This moderating effect of event presentation order was small, $d = -0.132$, with a 95% CI of $[-0.332, 0.068]$ suggesting that

the moderating effect is unlikely to be greater than small-to-medium in size. Our Bayesian analysis again pitted a point null hypothesis of $d = 0$ against a normal alternative prior with $M = 0$ and SD of half the effect size found in Experiment 1 ($d = 0.27$). The data were again more probable under the null than under the alternative prior, with the Bayes Factor of, $BF_{01} = 1.35$. However, because this Bayes Factor is close to 1, the data are largely inconclusive for the comparison between H_0 and H_1 . Overall, the findings of Experiments 2a and 2b suggest that event presentation order makes little or no difference to the judgment of medication effectiveness or to the size of the effect-density and cause-density biases.

General Discussion

Considering all three experiments together, they replicated the previous findings for effect-density and cause-density biases in contingency learning (Blanco et al., 2013), showing that effect-density biases are larger than cause-density biases, on average. The relative difference in size appears to depend on other experimental manipulations. In addition to this, it seems event presentation order has little effect on effectiveness judgments (Table 5). Across the experiments, the size of effects for cause and effect-density were relatively consistent.

The first experiment compared the effect-density bias to the cause-density bias to test if they differed in magnitude. Consistent with the findings from Blanco et al. (2013), who conducted a similar study but with a between-subjects design with only 27 participants per

Table 5

Comparison of means (SD), CIs and effect size for cue and outcome density across experiments 1, 2a, and 2b

	Effect Density		Cause density	
	Experiment 1	Experiment 2a	Experiment 1	Experiment 2b
High density response	43.74 (16.89)	43.81 (25.46)	43.45 (20.80)	46.71 (20.19)
Low density response	18.61 (16.85)	19.18 (15.35)	31.29 (23.95)	33.64 (21.87)
95% CI for mean difference	[19.41, 30.85]	[20.95, 28.32]	[6.54, 17.77]	[9.39, 16.75]
Density effect size d	1.49	1.17	0.54	0.62

condition, the effect-density bias is a robust effect. In contrast, the cause-density bias seems more dependent on context, as evidenced by the fact that the cause-density bias was present in the first session but not in the second session. The reduced cause-density bias in session two could result from what was seen, or the judgment made, in session one. Similar cross-session context effects have been observed in other studies of judgment and preference, including when—as we did here—sessions are separated by at least one day (Rakow et al., 2020). We suggest that such context effects deserve further attention in contingency learning to better understand the role of relative or comparative judgments play when people assess contingency from observed or remembered samples of information (Stewart et al., 2006). Whatever the mechanism responsible, it does seem that, compared to the effect-density bias, the cause-density bias relies on a larger density manipulation to show a clear effect.

Experiments 2a and 2b examined whether the event density effect differed when the presentation order of the cause and effect was manipulated. We find it intriguing that this manipulation made no difference to the presence of density effects, and little difference to their size. For example, the finding that presentation order seems not to matter allows us to cautiously discount the possibility that the cause-density bias is smaller than the effect-density bias as a result of the cause being typically presented before the effect (i.e., designating the cause as the ‘cue’ while the effect is the to-be-predicted ‘outcome’). After all, little changed when we reversed event presentation order. Moreover, the finding suggests that contingency judgments do *not* rely on events being encountered in their temporal order of occurrence with the putative cause preceding the effect.

Our experiments therefore clarify two key results that any theory of contingency learning must be able to explain. First, it clarifies that density effects in contingency learning are larger for effect-density manipulations than for cause-density manipulations. Second, it

clarifies that contingency judgments and the size of density effects are mostly unaffected by the presentation order of the contingent events (cause and effect).

Theoretical contributions

Contingency learning theory

The contingency learning account is a normative account of causal belief, stating that the belief about events should correspond with the difference in the probability that two events co-occur, $P(\text{Event 1}|\text{Event 2})$, and the probability that one of these events occurs in isolation, $P(\text{Event 1}|\text{No Event 2})$. This difference is quantified as ΔP —the degree of contingency (Allan, 1980). Notably the contingency account makes no assumptions about the nature of events; all events, even causes and their consequent effects, are expected to be treated equally.

A contingency is positive if an event is more likely to result in another event, which results in a positive ΔP , likewise a negative contingency or negative ΔP occurs when the presence of an event reduces the likelihood of another event from occurring. A zero contingency occurs when ΔP is zero, hence there is no difference in the probability of two events co-occurring or the one event occurring in isolation. For both positive and negative ΔP values, there is some relationship between the two events. Our experiments all used zero contingencies (i.e., $\Delta P = 0$). As such, according to contingency learning theory, participants should judge there to be no relationship between the two events (i.e., between the cause and effect). This was not what we observed; participants consistently identified a positive contingency between events.

There have been various attempts to use the contingency learning account to explain density biases in judgments (e.g., Msetfi et al., 2007). These consider the influence that factors such as the inter-trial interval or response rates might have on judgments. These factors unfortunately provide us with little assistance to explain our pattern of results. We

believe that the differential impact of manipulating cause density vs. effect density that we observed is problematic for the contingency learning account as a descriptive account. Put simply, it cannot explain *differences* in judgment between cause density and effect density conditions as we have found in this paper. For example, if the two events in a contingency are represented as event 1 = cause and event 2 = effect, there is no reason, within the contingency model, why these two events should be treated differently when calculating the probability or density of these events—they are just events. Only when contingency (ΔP) is calculated are the conditional probability of outcome given the presence of the cue, and the conditional probability of the outcome given the absence of the cue are these events treated differently. As seen in previous research (Blanco et al., 2013) when calculating event probabilities both contingency events are treated the same. To put this another way, according to the contingency table (Table 6), trials of cells B and C are identical as they are simply the occurrences of one event in the absence of the other event. Therefore, no systematic differences in contingency judgments between conditions should arise.

Notably, and contrary to the prescriptions of contingency learning theory, Wasserman (1990) found that people tend to weigh the cells of a (summarized) contingency table according to $A > B > C > D$ when judging contingencies for data presented in contingency tables. Thus, people tend to weigh cell A trials the most, whilst cell D trials are given the least weighting. For example, this implies in the medication treatment scenarios used in our experiments, participants will give most weight to occurrences of patients given medication and recovering, followed by patients given medication not recovering, then patients no given

Table 6
Contingency table with cells A, B, C and D

Cue	Outcome	
	Present	Absent
Present	A	B
Absent	C	D

medication and recovery, and the least weight to patients who do not receive medication and do not recover. In addition to this, Wasserman also found that the trials of cells A and D lead to increased positive judgments whilst cell C and B counts decrease contingency judgments. To note, Wasserman is not examining event 1 and event 2 but rather the experiments are focusing on a cause and effect, similar to the work in our experiments. Thus, applying Wasserman's work to our medical treatment scenario, frequent occurrences of patients given medication and recovering (i.e., cell A observations) should result in more positive judgments. Notably, increasing effect density *or* increasing cause density increases such occurrences. However, frequent observations of patients given medication not recovering (i.e., cell B observations) decrease judgments of contingency. Such occurrences become relatively more frequent with increasing cause-density *but not* with increasing effect-density. Therefore, Wasserman's analysis predicts a smaller effect on contingency judgment from manipulating cause-density as compared to manipulating effect-density. This is because the frequent cell B observations that occur when cause-density is high partially counteract the influence of the frequent cell A observations that occur when either cause-density or effect density is high.

By considering the differential weighting of cells (Wasserman, 1990), the differences we found in the level of density effect between cause- and effect-density manipulations are accounted for. Specifically, when examining the low contingency conditions for both cause-density (Table 1a) and effect-density (Table 1c), cells A and D have the same frequency counts, whilst cells B and C are reversed. Therefore, considering the cell weights people apply, there are more counts of cell B for the low effect-density condition compared to the low cause-density condition, which participants would weigh more. This higher weighting would decrease the positive contingency judgments, because cells B and C decrease positive contingency judgments. This matches the findings of our experiments, whereby for the low

density conditions' participants have given more positive judgments for the low cause-density condition than the low effect-density condition.

When examining the high contingency conditions for both cause-density (Table 1b) and effect-density (Table 1d), again cells A and D have the same frequency counts, whilst cells B and C are reversed. For the high cause-density condition there are more counts of cell B, than there are for the high effect-density condition. As such, lower judgments would be expected for the high cause-density condition, than the high effect-density condition, which also reflects the results of the current study. Here the Wasserman extension to contingency learning theory is able to predict the findings of Experiment 1, as well as the differences between Experiments 2a and 2b for the difference in the size of the density effect.

An assumption underlying the contingency account, and other human contingency learning and association accounts, that the cause occurs first, and then the effect occurs (De Houwer & Beckers, 2002). However, this does not fit with the data of our experiments. The data suggests that there is something fundamentally different between causes and effects, as a first event ('cue') vs second event ('outcome') effect was not found.

The Contingency account, with the Wasserman extension may be able to account for our data. However, it must be assumed that both the cause-values and effect-values are encoded in memory, the cause and effect could be observed in either order (as was the case in Experiments 2a and 2b). Therefore, if we assume the account does not need information about a cause to precede the information about effect, it can account for the findings that presentation order of cause and effect did not significantly impact the size of event-density effects.

Associative perspective

The associative perspective considers the associative strength of two events, the conditioned stimulus (CS/ cue) and unconditioned stimulus (US/ outcome). When these two

events are paired it leads to an internal connection between the events. The strength of connection is determined by how often the events are paired together and how often events occur in isolation. If context is also factored into the association between cue and US, such that we consider how context and cue might compete for associative strength, the associative account can also be influenced by how often both the cue and outcome are absent (De Houwer & Beckers, 2002). These models assume learning occurs on a conditioning trial only to the extent that the outcome is surprising. On the first conditioning trial when a cue is first paired with a US, as nothing signals the outcome, the outcome is surprising and therefore learning occurs. As conditioning trials proceed, cue will come to predict the outcome and the outcome will become less surprising. At some point the cue will predict the outcome perfectly and at this point no further learning will occur.

Overall, the associative perspective can predict both cause-density and effect-density biases (Musca et al., 2010). However, these effects would be expected to decrease in size the more trials that occur (Musca et al., 2010; Vadillo & Barberia, 2018). This is because the overall strength of an association between a cause and an effect is determined by the predictive value of the cause, rather than the overall number of coincidences between both events. In the short term, before associative strength reaches its asymptote, the ‘coincidence’ can have a temporary biasing effect. The effect-density bias is understood as a transient pre-asymptotic bias arising from accidental cause-effect pairings before the context becomes an effective competitor (Barberia et al., 2019). If there is unequal salience for the events of the contingency, the associative models can predict effect-density bias. Although if the occurrence of the effect is less salient than its non-occurrence, then associative models can predict positive effect-density bias. Similarly the account can predict cause-density bias. Vadillo & Barberia (2018) demonstrated that the Rescorla-Wagner model predicts both cause-density and effect-density biases in judgment at extreme probabilities (0.2 and 0.8) and

also demonstrated differences in the sizes of these effects, with cause-density bias being much smaller than effect-density bias at both high and low probabilities. While in our studies, when the probability of cause was 0.2 or 0.8, the probability of outcome was 0.5, and vice versa, which does not match the exact conditions of the simulations by Vadillo & Barberia (2018), their findings reflect the differences in cause- and effect-density biases found in our data.

In regard to event presentation order, the associative account assumes that the cue is followed by the outcome, as otherwise there would be no surprise element and hence no learning. Therefore, switching the presentation order of the cue and the outcome should disrupt learning. As such, the setup of associative models would suggest that they cannot naturally account for the finding that contingency judgments were largely unaffected by reversing event presentation order to place the effect before the cause. At minimum, this suggests that for an associative account to explain our data, the cue and outcome must be arbitrarily defined—the first presented event is being associated with the second presented event regardless of the assigned labels of cause and effect. Thus, in line with our data it would make sense to speak of surprise at cause-value for a given effect-value in the same way that (in the framework of the associative account) one speaks of surprise at effect-value for given cause-value. However, that approach to (arbitrary) definition of the cue and outcome cannot easily account for the larger size of the effect-density bias as compared to the cause-density bias. This is because, given the equal size of our cause-density and effect-density manipulations, it would not be clear why surprise for effect values would be larger than surprise for cause values. Consequently, we believe that the standard associative account does not fully account the findings of our paper. However, we do note that some extended associative accounts such as the HeiDI, which considers reciprocal associations between cues and outcomes, may suggest that the order of events within a trial may not be critical for

learning to take place (Honey et al., 2020), which when combined with the earlier points in our discussion suggests that the associative learning account can be adapted to the findings of this paper.

Probabilistic contrast models

The probabilistic contrast model is another rule-based account similar to the contingency account. According to the model, individuals' reason about cause-effect relationships by recalling all the events of the contingency and calculate a statistic similar to ΔP (Cheng, 1997; Cheng & Novick, 1990). Thus, individuals compare situations that differ by whether the cause is present or absent. By comparing these situations, inferences about the relationship between cause and effect are learnt. The following equation is then used by the model,

$$P_i = \frac{\Delta P}{(1 - P(e|\bar{i}))}$$

In this equation for the power for event i (P_i), e , is the occurrence of the effect and, i , is the occurrence of the event. When the cause is preventative, the denominator is $P(e|\bar{i})$. So, the power for an event to cause an effect is a joint function of the overall contingency (P) and the probability of the effect in the absence of the event.

Probabilistic models can predict higher judgments of control in positive and negative contingency conditions when effect-density is high. However, they cannot predict effect-density bias when the contingency is zero. However, it is argued that zero contingency is misperceived by people as being slightly positive or negative. With this assumption, the model can predict increased judgments of contingency when effect-density is high (Cheng & Novick, 1990; Kutzner et al., 2008). This may, for example, occur in within-subjects designs where participants judge multiple contingencies. If this was only an issue in within-subjects designs, then effect-density bias should not be present in between-subjects designs—but it is. For our experiments, these probabilistic contrast models would predict that effect-density and

cause-density biases should be small or absent in the first session. However, we found large effects. Moreover, in Experiment 1, we found a larger cause-density bias in the first session than in the second session. Although if the model were to assume that in general zero contingencies are always perceived as being either slightly negative or positive the model would provide support for the findings. Although if that is the case, there would need to be further details on when zero contingency situations would be thought of as being negative or positive.

Similar to other models, the base assumption of probabilistic contrast models is that causes are followed by effects. Taking the probabilistic contrast model calculation as an example, power is driven by the idea that for any cause-and-effect pairing, there are always potential alternative causes. For example, if one has a fever (effect), the typically known cause may be having the flu, however alternative causes could be an infection or heat exhaustion. Accordingly, knowledge of the base rate of the effect (i.e., the relative frequency of the effect) will let one to make a contrast between the candidate cause and alternative causes such as experimental context, and causal power is defined by contingency of an effect contrast. Therefore, this model neither explicitly states anything regarding presentation order, nor about cause-density biases. However, the findings of this paper are consistent with the more general idea of the model that individuals estimate causal relationships by using information beyond the contingency information given.

Causal model theory

Causal model theories explain how people make statistical inferences in predictive reasoning (from observed causes to effects) and diagnostic reasoning (from observed effects to causes) (Fenker et al., 2005; Waldmann, 2000). Causal model theory proposes that individuals hold abstract knowledge about the properties of causal relations that interacts with their learning of statistical information (Waldmann & Holyoak, 1992). The theory has three

key assumptions. Firstly individuals have a predisposition to learn relationships between events as directed links from cause to effect (predictive), even in situations where information regarding the effect is presented first (diagnostic). Consequently, people make quicker and stronger predictive reasoning judgments compared to diagnostic reasoning judgments (Fenker et al., 2005). Secondly the strength of the causal relationship perceived between cause and effect is dependent on the contingency between the cause and effect. The stronger the contingency between the cause and effect, the stronger that relationship is perceived to be. Finally, even though relationships between events are directional, directed from cause to effect, individuals can still make both predictive inferences and diagnostic inferences. That is, even though people are better at making inferences from causes leading to effect (predictive inference), they are still able to infer links from effects to the likely causes (diagnostic inferences).

In our experiments, causal model theory can explain positive judgment in zero contingency conditions, due to the fact that the model assumes that people hold abstract knowledge about the world. The stimuli used in these experiments, are related to medication and illness, and in general people tend to hold the belief that medications will make one feel better and cure the illness they have. As such, the general knowledge people have about medications may have led participants to provide positive contingency judgments in the zero contingency conditions. When considering the effect-density bias, it could be argued that participants come into the experiment assuming a positive relationship between the medication (cause) and getting better (effect) and use the information provided in the experimental conditions to update their statistical inference by each trial. This could then lead to participants giving positive but lower judgments in the low probability condition, and positive and higher judgments in the high probability judgments. Therefore, the model would be expected to predict effect-density biases. The model would also be able to predict cause-

density biases, because a similar logic could be applied here. However, it is not clear to us that causal model theory would predict that cause-density biases should be smaller than effect-density biases (as was seen in our data). For example, there seems no reason to assume that a general expectation that medications are efficacious would, by itself, result in a stronger perceived causal link when effect density is high as compared to when cause density is similarly high. Therefore, if causal model theory is to explain the asymmetry in density biases, that asymmetry must presumably be a feature or bug in the process by which participants learned statistical information rather than a product of their general knowledge about medication and cure.

In regard to event presentation order, causal model theory does state that people can make inferences between causes and effects regardless of which order they are presented in. However the model also assumes that people are predisposed to making predictive inferences from causes to effects. As such, the model implies that people should be able to make contingency judgments in both presentation-order conditions, but that the judgments might be more accurate in the predictive conditions compared to the diagnostic condition. We did find that participants made contingency judgments in both the predictive and diagnostic conditions – but with no evidence that manipulating presentation order altered their judgment. Therefore, the possibility that contingency judgments might differ between the predictive condition (cause seen first) and the diagnostic condition (effect seen first) was not supported by our findings. However, if it is instead argued that people are able to extract the information from diagnostic conditions to make the links predictive in nature, then it would be expected that the density effects would be similar across both conditions.

Theoretical summary

Many of the learning models presented here are not able to fully account for the findings of this paper. Whilst the reasoning as to why has been outlined here, the more

important consideration to take away is the fact that these models need to incorporate the assumptions that people are more sensitive to the causal status of events (i.e. whether something is a cause or effect) than to their temporal order of presentation or discovery.

Limitations and Future Directions

As is common for this research on density effects in contingency learning (Blanco et al., 2013; Msetfi et al., 2007; Vallée-Tourangeau et al., 2005), we focused exclusively on zero-contingency conditions. In the interest of understanding the robustness of cause and effect-density biases, it would – of course – be valuable to examine density effects for non-zero contingencies, both for positive and negative contingencies. Examining both direction of contingency would require adopting a bi-directional judgment scale (e.g., -100 to +100) in place of the standard 0-100 scale that we used in the current experiments, which allowed participants to consider only one direction of association. Even when negative contingencies are not examined, there are good reasons to use such a bi-directional scale. A scale bounded at zero does not permit participants to give a negative judgment, which is what the probabilistic contrast model predicts that some participants will wish to do, and which any model that allows for judgmental ‘noise’ might expect of some participants. As such, even when there is no reason to suppose that the *average* participant will perceive a negative contingency when no contingency exists, providing scales that allow individual participants to make judgments in line with whatever they perceive permits better tests of theory.

Our data are in line with Wasserman’s (1990) extension to the contingency learning model, which proposes that some cause-effect combinations receive additional (or reduced) weight when contingencies are judged. While Wasserman’s proposal predicts our data successfully, it provides a summary *description* (*as if* model) rather than a process-level *explanation* for our data. Thus, we do not know whether the apparent differential weighting for the A-B-C-D combinations is due to attention, memory, expectation, belief-updating, or

something else. It is a challenge to understand which of these (or other) candidate processes are responsible for the density effects we have observed. However, given the large size of these effects, we think the study of density effects offers good opportunities to better understand the cognitive processes that underpin learning via experimental manipulations that affect these processes. For instance, one possibility is that participants judge contingency by trying to recall the joint distribution of cause-values and effect-values that they observed. Therefore, it could be valuable to test what impacts the memory for that distribution; and to see whether any biases in memory predict (or are predicted by) density effects in contingency judgment.

The session order effect that we found in Experiment 1 whereby the cause density bias was not present in the second session, also points to the value of further research into the processes that underpin contingency judgment. Across all experiments the sessions were separated by three days and the same density levels were used; therefore, it is not clear why this effect occurred in Experiment 1 but not Experiment 2a or 2b. One difference between Experiment 1, and 2a and 2b, is that in Experiment 1 participants saw both cause and effect manipulated, whilst in the latter experiments only one of these event manipulations was seen. Consequently, the session order effect seen in Experiment 1 may be due to participants seeing effect-density manipulations first, which may make participants insensitive to cause-density manipulations. To examine this further, a first step would be to establish the replicability of the order effect that we found for cause-density. If found to be reliable, researchers could then examine whether – as we have suggested – judgments made in a second session are influenced by observations or judgments made in an earlier session.

The framing of the contingency rating questions in our experiments were always framed in terms of the cause leading to the effect (i.e. a ‘cause-framed’ question, how effective was the *medication* in treating the illness?). Recent research suggests that effect-

framed vs. cause-framed questions can lead to differences in causal ratings (Camara & Livesey, 2023) with effect-framed questions leading to a smaller effect-density bias and larger cause-density bias compared to cause-framed questions. The current studies used cause-framed questions, and found cause-density biases were still smaller than the effect-density biases. Therefore, the research could be expanded to examine whether using effect-framed questions reduces the size of the cause-density bias, or even eliminates it..

Conclusion

Our research has demonstrated the robustness of effect-density manipulations on contingency judgments, and their greater impact in comparison to equivalent manipulations of cause-density. In addition to this, event presentation order was found to have negligible effect on judgments. Seemingly, cause-density biases in contingency judgment are reliably smaller than effect-density biases. Furthermore, contingency judgments do not seem to rely on events being encountered in their temporal order of occurrence with the cause preceding the effect. Together, these findings suggest that learning models should incorporate the assumption that people are more sensitive to the causal status of events than to the temporal order of those events.

References

- Allan, L. G., Siegel, S., & Tangen, J. M. (2005). A signal detection analysis of contingency data. *Learning & Behavior*, *33*(2), 250–263. <https://doi.org/10.3758/BF03196067>
- 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), 441–485. <https://doi.org/10.1037/0096-3445.108.4.441>
- Barberia, I., Vadillo, M. A., & Rodríguez-Ferreiro, J. (2019). Persistence of Causal Illusions After Extensive Training. *Frontiers in Psychology*, *10*.
<https://doi.org/10.3389/fpsyg.2019.00024>
- Batterink, L., & Neville, H. (2011). Implicit and Explicit Mechanisms of Word Learning in a Narrative Context: An Event-related Potential Study. *Journal of Cognitive Neuroscience*, *23*(11), 3181–3196. https://doi.org/10.1162/jocn_a_00013
- Blanco, F., Matute, H., & Vadillo, M. A. (2011). Making the Uncontrollable Seem Controllable: The Role of Action in the Illusion of Control. *Quarterly Journal of Experimental Psychology*, *64*(7), 1290–1304.
<https://doi.org/10.1080/17470218.2011.552727>
- 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*(4), 333–340. <https://doi.org/10.3758/s13420-013-0108-8>
- Camara, A. M., & Livesey, E. (2023). Question Framing Modulates the Cause Density and Effect Density Biases in Causal Illusions. *Proceedings of the Annual Meeting of the Cognitive Science Society*, *45*(45). <https://escholarship.org/uc/item/48t5b1g8>
- Cheng, P. W. (1997). From covariation to causation: A causal power theory. *Psychological Review*, *104*, 367–405. <https://doi.org/10.1037/0033-295X.104.2.367>

- Cheng, P. W., & Novick, L. R. (1990). A probabilistic contrast model of causal induction. *Journal of Personality and Social Psychology, 58*, 545–567.
<https://doi.org/10.1037/0022-3514.58.4.545>
- Cohen, J. (2013). *Statistical Power Analysis for the Behavioral Sciences*. Academic Press.
- Crump, M. J. C., Hannah, S. D., Allan, L. G., & Hord, L. K. (2007). Contingency judgments on the fly. *Quarterly Journal of Experimental Psychology, 60*(6), 753–761.
<https://doi.org/10.1080/17470210701257685>
- De Houwer, J., & Beckers, T. (2002). A Review of Recent Developments in Research and Theories on Human Contingency Learning. *The Quarterly Journal of Experimental Psychology Section B, 55*(4b), 289–310. <https://doi.org/10.1080/02724990244000034>
- Dienes, Z. (2014). Using Bayes to get the most out of non-significant results. *Frontiers in Psychology, 5*. <https://www.frontiersin.org/articles/10.3389/fpsyg.2014.00781>
- Dienes, Z. (2019). How do I know what my theory predicts?. *Advances in Methods and Practices in Psychological Science, 2*(4), 364–377.
<https://doi.org/10.1177/2515245919876960>
- Fenker, D. B., Waldmann, M. R., & Holyoak, K. J. (2005). *Assessing causal relations in semantic memory. 33*(6), 1036–1046.
- Gillan, C. M., Morein-Zamir, S., Durieux, A., Fineberg, N. A., Sahakian, B. J., & Robbins, T. W. (2014). Obsessive–compulsive disorder patients have a reduced sense of control on the illusion of control task. *Frontiers in Psychology, 5*.
<https://doi.org/10.3389/fpsyg.2014.00204>
- Greenwald, A. G. (1975). Consequences of prejudice against the null hypothesis. *Psychological Bulletin, 82*, 1–20. <https://doi.org/10.1037/h0076157>

- Honey, R. C., Dwyer, D. M., & Iliescu, A. F. (2020). HeiDI: A model for Pavlovian learning and performance with reciprocal associations. *Psychological Review*, *127*(5), 829–852. <https://doi.org/10.1037/rev0000196>
- Jenkins, H. M., & Ward, W. C. (1965). Judgment of contingency between responses and outcomes. *Psychological Monographs: General and Applied*, *79*(1), 1–17. <https://doi.org/10.1037/h0093874>
- Kutzner, F., Freytag, P., Vogel, T., & Fiedler, K. (2008). Base-Rate Neglect as a Function of Base Rates in Probabilistic Contingency Learning. *Journal of the Experimental Analysis of Behavior*, *90*(1), 23–32. <https://doi.org/10.1901/jeab.2008.90-23>
- Matute, H. (1995). Human Reactions to Uncontrollable Outcomes: Further Evidence for Superstitions Rather Than Helplessness. *The Quarterly Journal of Experimental Psychology Section B*, *48*(2), 142–157. <https://doi.org/10.1080/14640749508401444>
- 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*. <https://doi.org/10.3389/fpsyg.2015.00888>
- Matute, H., Vadillo, M. A., Vegas, S., & Blanco, F. (2007). Illusion of Control in Internet Users and College Students. *CyberPsychology & Behavior*, *10*(2), 176–181. <https://doi.org/10.1089/cpb.2006.9971>
- Msetfi, R. M., Murphy, R. A., & Simpson, J. (2007). Depressive realism and the effect of intertrial interval on judgments of zero, positive, and negative contingencies. *Quarterly Journal of Experimental Psychology*, *60*(3), 461–481. <https://doi.org/10.1080/17470210601002595>
- Msetfi, R. M., Murphy, R. A., Simpson, J., & Kornbrot, D. E. (2005). Depressive realism and outcome density bias in contingency judgments: The effect of the context and

- intertrial interval. *Journal of Experimental Psychology. General*, 134(1), 10–22.
<https://doi.org/10.1037/0096-3445.134.1.10>
- Musca, S. C., Vadillo, M. A., Blanco, F., & Matute, H. (2010). The role of cue information in the outcome-density effect: Evidence from neural network simulations and a causal learning experiment. *Connection Science*, 22(2), 177–192.
<https://doi.org/10.1080/09540091003623797>
- Mussweiler, T., Englich, B., & Strack, F. (2012). Anchoring Effect. In *Cognitive illusions* (pp. 195–212). Psychology Press.
- Perales, J. C., Catena, A., Shanks, D. R., & González, J. A. (2005). Dissociation between judgments and outcome-expectancy measures in covariation learning: A signal detection theory approach. *Journal of Experimental Psychology. Learning, Memory, and Cognition*, 31(5), 1105–1120. <https://doi.org/10.1037/0278-7393.31.5.1105>
- Prolific*. (2020). [Computer software]. Prolific. <https://www.prolific.co>
- Rakow, T., Cheung, N. Y., & Restelli, C. (2020). Losing my loss aversion: The effects of current and past environment on the relative sensitivity to losses and gains. *Psychonomic Bulletin & Review*, 27(6), 1333–1340. <https://doi.org/10.3758/s13423-020-01775-y>
- Shanks, D. R. (1993). Human instrumental learning: A critical review of data and theory. *British Journal of Psychology*, 84(3), 319–354. <https://doi.org/10.1111/j.2044-8295.1993.tb02486.x>
- Shanks, D. R. (2010). Learning: From Association to Cognition. *Annual Review of Psychology*, 61(1), 273–301. <https://doi.org/10.1146/annurev.psych.093008.100519>
- Shelton, J. R., & Martin, R. C. (1992). How semantic is automatic semantic priming? *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 18, 1191–1210.
<https://doi.org/10.1037/0278-7393.18.6.1191>

- Stewart, N., Chater, N., & Brown, G. D. A. (2006). Decision by sampling. *Cognitive Psychology*, 53(1), 1–26. <https://doi.org/10.1016/j.cogpsych.2005.10.003>
- Vadillo, M. A., & Barberia, I. (2018). A comparator-hypothesis account of biased contingency detection. *Behavioural Processes*, 154, 45–51. <https://doi.org/10.1016/j.beproc.2018.02.009>
- Vadillo, M. A., Musca, S. C., Blanco, F., & Matute, H. (2011). Contrasting cue-density effects in causal and prediction judgments. *Psychonomic Bulletin & Review*, 18(1), 110–115. <https://doi.org/10.3758/s13423-010-0032-2>
- Vallée-Tourangeau, F., Murphy, R. A., & Baker, A. G. (2005). Contiguity and the outcome density bias in action–outcome contingency judgments. *Quarterly Journal of Experimental Psychology: Section B*, 58(2), 177–192. <https://doi.org/10.1080/02724990444000104>
- Waldmann, M. R. (2000). Competition among causes but not effects in predictive and diagnostic learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 26(1), 53–76. <https://doi.org/10.1037/0278-7393.26.1.53>
- Waldmann, M. R., & Holyoak, K. J. (1992). Predictive and Diagnostic Learning Within Causal Models: Asymmetries in Cue Competition. *Journal of Experimental Psychology: General*, 121(2), 222–236.
- Werker, J. F., Cohen, L. B., Lloyd, V. L., Casasola, M., & Stager, C. L. (1998). Acquisition of word–object associations by 14-month-old infants. *Developmental Psychology*, 34(6), 1289–1309. <https://doi.org/10.1037/0012-1649.34.6.1289>

Supplementary Materials

Instructions given to all participants across conditions of the experiments.

The instructions were also accompanied by sample images of the experimental trials images.

“Imagine that you are a doctor who works at the hospital. You specialise in new diseases and their potential treatments. The crisis induced by the new diseases may be reduced by using new treatments. However, these treatments are still in their testing phase, therefore the reliability of the treatments have not been proven yet. ### When you click the ‘Start’ button, you will be presented with a series of images which will show unhealthy patients suffering from a disease. You will also be shown information on patient recovery and any treatment the patient may have experienced. See below for images that correspond to this. After you have seen several patients, their treatments and recovery, you will be asked to assess the effectiveness of the new treatment in treating the new disease. You will be presented with 2 new diseases and their respective drug treatments. You will be given time for a break between each disease block.”

Density order effects within each session

Experiment 1

We ran a $2 \times 2 \times 2$ mixed ANOVA with manipulation of event and density as within-subjects factors and density order during first session (low first or high first) as independent variables, and effectiveness judgments as dependent variable. There was no significant main effect of density order during first session $F(1, 83)=2.01, p=.152, \eta_p^2=.025$. There were no interaction effect of manipulation of event \times density order during first session, $F(1, 83)=.34, p=.56, \eta_p^2=.004$, density \times density order during first session, $F(1, 83)=.55, p=.462, \eta_p^2$

=.007, manipulation of event \times density \times density order during first session, $F(1, 83)=.02$, p
=.887, η_p^2 =.00.

We ran a $2 \times 2 \times 2$ mixed ANOVA with manipulation of event and density as within-subjects factors and density order during second session (low first or high first) as independent variables, and effectiveness judgments as dependent variable. There was no significant main effect of density order during second session $F(1, 83)=1.92$, p =.170, η_p^2 =.02. There was a significant interaction effect of manipulation of event \times density order during second session, $F(1, 83)=5.63$, p =.020, η_p^2 =.06. There was no significant interaction effect of density \times density order during second session, $F(1, 83)=3.72$, p =.057, η_p^2 =.04, manipulation of event \times density \times density order during second session, $F(1, 83)=1.38$, p
=.243, η_p^2 =.02.

Experiment 2a: Effect density

We ran a $2 \times 2 \times 2$ mixed ANOVA with event presentation order and density as within-subjects factors and density order during first session (low first or high first) as independent variables, and effectiveness judgments as dependent variable, to investigate whether density order during first session has an order effect. There was a significant main effect of first session density order, $F(1, 94)=4.22$, p =.043, η_p^2 =.043, participants gave lower judgments if they were in the low density first condition ($M=28.35$, $SD=20.60$), compared to the high density first condition ($M=34.27$, $SD=19.34$), suggesting there may have been an anchoring effect present that is based on the first judgment they made. There was a significant 3 way interaction of event presentation order \times density \times density order during first session, $F(1, 94)=9.30$, p =.003, η_p^2 =.09 and presentation order \times density order during first session, $F(1, 94)=12.05$, p <.001, η_p^2 =.11. There was no interaction effect of density \times density order during first session, $F(1, 94)=1.14$, p =.290, η_p^2 =.01.

We also ran a $2 \times 2 \times 2$ mixed ANOVA with event presentation order and density as within-subjects factors and density order during second session (low first or high first) as independent variables, and effectiveness judgments as dependent variable. There was no significant main effect of density order during second session $F(1, 94)=1.18, p =.280, \eta_p^2 =.01$. There were no significant interaction effects of event presentation order \times density order during second session, $F(1, 94)=.80, p =.374, \eta_p^2 =.01$, or density \times density order during second session, $F(1, 94)=.043, p =.836, \eta_p^2 =.00$, manipulation of order of event \times density \times density order during second session, $F(1, 94)=.07, p =.789, \eta_p^2 =.00$.

Experiment 2b: Cause density

We ran a $2 \times 2 \times 2$ mixed ANOVA with event presentation order and density as within-subjects factors and density order during first session (low first or high first) as independent variables, and effectiveness judgments as dependent variable, to investigate whether density order during first session has an order effect. There was no significant main effect of first session density order, $F(1, 95)=0.19, p =.665, \eta_p^2 =.002$, judgment judgment. There was a significant 3 way interaction of event presentation order \times density \times density order during first session, $F(1, 95)=6.13, p =.015, \eta_p^2 =.06$. There was no interaction effect of density \times density order during first session, $F(1, 95)=1.98, p =.163, \eta_p^2 =.02$. , or no

interaction effect of presentation order \times density order during first session $F(1, 95)=2.82, p = .097, \eta_p^2 = .03$.

We also ran a $2 \times 2 \times 2$ mixed ANOVA with event presentation order and density as within- subjects factors and density order during second session (low first or high first) as independent variables, and effectiveness judgments as dependent variable. There was no significant main effect of density order during second session $F(1, 95)=.14, p = .712, \eta_p^2 = .00$. There were no significant interaction effects of event presentation order \times density order during second session, $F(1, 95)=.114, p = .736, \eta_p^2 = .00$, or density \times density order during second session, $F(1, 95)=.010, p = .923, \eta_p^2 = .00$, manipulation of order of event \times density \times density order during second session, $F(1, 95)=.028, p = .895, \eta_p^2 = .00$.