

UC Merced

Proceedings of the Annual Meeting of the Cognitive Science Society

Title

Causal inference and learned helplessness

Permalink

<https://escholarship.org/uc/item/8h3221hw>

Journal

Proceedings of the Annual Meeting of the Cognitive Science Society, 48(0)

Authors

Hew, Joshua Wen Zheng

Bramley, Neil R.

Publication Date

2026

Copyright Information

This work is made available under the terms of a Creative Commons Attribution License, available at <https://creativecommons.org/licenses/by/4.0/>

Peer reviewed

Causal inference and learned helplessness

Joshua W. Z. Hew (joshua.hew.25@ucl.ac.uk)¹ & Neil R. Bramley²

¹Department of Experimental Psychology, University College London

²Department of Psychology, University of Edinburgh

Abstract

Prolonged exposure to uncontrollable situations can cause individuals to become and remain dysfunctionally passive. This pattern, known as *learned helplessness*, is typically induced in lab settings using simple tasks, but real-world control involves complex, non-linear causal systems. In these environments, the ability to influence an outcome often diverges from the ease of achieving the specific result one *wants*. Moreover, ascribing agency to oneself is a non-trivial process that depends on prior mechanistic beliefs and counterfactual inference. To investigate these dynamics, we systematically manipulated structure, controllability, and reward prevalence while participants interacted with dynamic causal variables in real time. Whilst low levels of practical control reliably induced helpless behaviour, we found that this did not depend on reward prevalence or the accuracy of learners' causal beliefs.

Keywords: causal inference; learned helplessness; control; counterfactuals

Learned helplessness is a psychological phenomenon in which individuals, after experiencing uncontrollable adverse events, come to expect that their actions will have little or no effect on future outcomes. This expectation transfers to new situations, leading to pervasive and dysfunctional passivity. The phenomenon was first documented in animal work (Overmier & Seligman, 1967): dogs receiving unavoidable electrical shocks later made no attempt to escape even when escape was possible. This same pattern emerged in human adults in analogous tasks involving inescapable loud tones (Hiroto, 1974), and was shown to depend on how strongly participants attributed outcomes to external causes. Given these findings, learned helplessness is now a commonly used model of depression, since experimentally-induced helplessness replicates depressed individuals' diminished perception of control (Miller & Seligman, 1975).

Pathological control perception is hypothesised to play a central role in learned helplessness, where participants fail to recognise contingency between their actions and outcomes (Maier & Seligman, 1976). Subsequent findings, however, challenge the accuracy and utility of this account. The depressive realism hypothesis (Alloy & Abramson, 1979) suggests that depressed individuals actually make more accurate judgements about action-outcome contingencies than non-depressed individuals, who tend to overestimate their control. This fits a broader pattern: neurotypical individuals exhibit optimism bias (Sharot, 2011) and illusions of control, particularly when desired outcomes are common (Langer, 1975;

Langer & Rodin, 1976). Consistent with this, Teodorescu and Erev (2014) found that the prevalence of rewards in an experimental task played a larger role in preventing helpless behaviour than actual action-outcome contingencies.

Nevertheless, other work emphasises controllability as the key factor in learned helplessness. The opportunity to choose, even in arbitrary decisions, increases motivation and engagement (Cordova & Lepper, 1996). Choice is preferred over non-choice, and experienced as rewarding, as demonstrated across self-report, behavioural, and neuroimaging data (Leotti & Delgado, 2011). Rewards have little effect on performance in induction or problem-solving tasks (Osborn Popp et al., 2025). This suggests that where identifying action-outcome relations is key, participants may be motivated by an intrinsic desire for discovery and control. Consistent with this, curiosity and interest-based engagement may be stronger determinants of task performance than extrinsic rewards (Murayama, 2022).

Humans construct causal representations of their environment to facilitate reasoning, learning, and decision-making, particularly in contexts that require prediction, explanation or adaptive control (Sloman, 2005). Much distinctively human cognition can be construed as a “causal model-based control”, where success depends on learning and exploiting a causal model to achieve changing goals (Broadbent et al., 1986; Davis et al., 2018; Schulz et al., 2017). Learned helplessness scenarios fit this mould: participants must learn at test that a button press terminates an electric shock and use this knowledge to escape harm. Thus, this phenomenon can be understood as either a failure to construct a causal model that includes this affordance, or a failure to leverage that causal model to plan and perform an action that influences outcomes.

Several parallels make this framing appealing: causal strength judgements (e.g., Cheng, 1997) build on the same kinds of instrumental contingency utilised by Maier and Seligman (1976); research on attribution theory, which focuses on inferring the causes of social behaviour (e.g., Kelley, 1973), has informed later developments in learned helplessness theory, specifically those involving attributional styles (Abramson et al., 1978). While recent studies have explored causal model-based control in continuous time with continuous variables (Davis et al., 2018), the influence of individual differences and experimental manipulations on adaptive behaviour remains understudied. Furthermore, learned helplessness paradigms have largely relied on static, atemporal and well-

defined designs. Such designs fail to capture the complexity of real-world contexts that involve continuous functional relationships, delays, and feedback loops, where the timing and sequence of events are crucial for causal inference and effective control (Cartwright, 2004; Sloman & Lagnado, 2015).

The present study addresses a gap in the literature by investigating the role of causal inference in learned helplessness. We examine learned helplessness within a dynamic control task—an interactive online game where participants engage with continuous variables across discrete time-points. Adopting an exploratory approach that synthesises research on causal reasoning and dynamic decision-making, we manipulate features of the training phase to test two key predictions: that lower objective controllability will increase helpless behaviour, in line with prior work on perceived contingency, and that lower reward prevalence will similarly increase helplessness, consistent with findings on the illusion of control.

Methods

Participants were assigned to one of four experimental groups in a 2×2 between-subjects design, crossing reward region size (large vs. small) with the degree of control afforded by the dynamic causal system (high vs. low). Following classic learned helplessness paradigms, we divided the study into a training phase and a test phase. Control manipulations were restricted to the training phase; during the test phase, control was standardised such that participants in the previously low control groups gained the same level of influence as the high control groups. In contrast, the reward region size assigned to each participant remained constant across both phases. We note that this design choice was made to isolate the effect of causal control history and prevent potential confounds associated with shifting reward structures; we address the implications of this choice in the discussion.

Framework: Ornstein-Uhlenbeck process

Following Davis et al. (2020), we use an Ornstein-Uhlenbeck (OU) process to create causal environments where continuous variables interact across discrete time steps. In essence, this process captures how variables drift randomly over time while being pulled toward values determined by their causal parents. Formally, it characterises a variable’s random movement in time, biased towards a mean μ (Uhlenbeck & Ornstein, 1930). The change in variable x , Δx_t for a time step t is defined as:

$$\Delta x_t = \omega [\mu - x_t] + \mathcal{N}(0, \sigma) \quad (1)$$

where $\omega > 0$ is a value determining how strongly attracted x is to μ . The value of x is continuously perturbed with normally distributed noise of variance σ . The result is a “mean regressive random walk”, that is, a variable which evolves randomly over time but with a basin of attraction centred on μ .

This framework can be extended to capture causal dynamics involving multiple variables by making the attraction states of variables nonstationary, and defining them as functions of

their causal parents. To take the case of a three-variable causal system x , y , and z connected by a possibly cyclic causal graph G , the constant μ is replaced at time $t + 1$ with a function of the values of the causal parents at time t (Davis et al., 2020). The evolution of this system is then described by a system of equations of the form:

$$\Delta x_{t+1} = \omega \left[\left(\sum_{y \in a(x)} \theta_t^{yx} y_t \right) - x_t \right] + \mathcal{N}(0, \sigma) \quad (2)$$

where $a(x)$ are the parents of x in the causal graph. Here, x tends towards a new value determined by a linear combination of causal variables y and z , which are modified by θ representing the existence and strength of causal relationships $y \rightarrow x$ and $z \rightarrow x$. If $\theta_{yx} = 0$, y is not causally related to x ; if $\theta_{yx} > 0$, x tends towards a multiple of y ’s value (it is attracted); if $\theta_{yx} < 0$, x tends towards a negative multiple of y ’s value (it is repelled). Importantly, this relationship is independent for both directions, such that $\theta_{yx} = 0$ does not imply $\theta_{xy} = 0$. Cyclic relationships between two variables influencing each other such that $\theta_{yx} \neq 0$ and $\theta_{xy} \neq 0$ are also possible. If x has no parents, the system assumes it is attracted to 0.

Using this framework, we systematically vary the amount of control participants can exert over the system. For high control groups, we followed (Davis et al., 2020) in setting θ at $[1, -1]$ for “regular” and “inverse” causal relationships. This enabled participants to more easily assess causal strength and relationships. σ was set to 4 as a baseline level of noise. For low control groups, we drew $\theta_t \sim \text{Beta}(1, 3)$ independently on each step for regular connections and multiplied it with -1 for inverse connections. This meant that causes would often have little influence on their effects but occasionally have sizeable impact, making it difficult for participants to exert fine-grained prospective control since their actions had unpredictable effects. We also used $\sigma = 18$ to obscure the causal relationships by injecting a larger amount of exogenous noise into the system. All participants had stable $[1, -1]$ θ values and $\sigma = 4$ during the test phase. Figure 1a illustrates the different behaviour of each causal structure across the control manipulations.

We used six different causal structures across the experimental rounds, each consisting of three variables x , y , z , labelled “X”, “Y”, and “Z” (Figure 1b). The variables changed value in discrete time steps according to Equation 2. For each causal structure, participants were randomly and independently assigned to either counterbalance version A or B, with X and Y labels swapped between versions (Figure 1b). Participants were presented with a different structure in each round, requiring them to discover the underlying causal mechanisms while generalising task-level knowledge across rounds. All structures were rated as equally difficult by participants in a prior study. They were also designed so that influencing Z could be achieved through intervention on just one of X or Y.




Figure 1: (a) Behaviour of causal structures in counterbalance group A when perturbed with the same sequence of interventions (at timesteps 4, 8, 12, 13, 14, & 15), across both high and low control groups. The small reward region is shaded in green, and the large reward region is marked with dashed lines. (b) Causal structures counterbalanced between groups A and B. Black arrows represent regular relationships ($\theta = 1$), and white arrows represent inverse relationships ($\theta = -1$).

Experiment

The code, data, task demo, and preregistration link are available at: https://github.com/J2hwz/causal_helplessness. Ethical approval for this study was granted by the University of Edinburgh School of Philosophy, Psychology, and Language Sciences Ethics Committee (application reference: 141-2425/3). A power analysis suggested a minimum sample of 77. To ensure adequate power after potential exclusions, we recruited 100 participants via Prolific and compensated £2.00 each ($M_{\text{time}} = 19.7$; $SD_{\text{time}} = 5.3$). Following our preregistered criteria, seven participants were excluded for failing comprehension questions more than twice, or disengaging by switching browser tabs, yielding a final sample of 93 ($M_{\text{age}} = 43.3$; $SD_{\text{age}} = 13.0$; $N_{\text{female}} = 48$ [51.6%]). We implemented the experiment in JavaScript.

Experiment-level procedure Upon providing informed consent, participants viewed detailed instructions—including instructional videos and annotated diagrams. They were then required to pass a comprehension quiz to ensure their understanding of the task. The experiment comprised three training rounds and three test rounds. Between these two phases, par-

ticipants viewed their cumulative score and rated the task’s perceived controllability and difficulty on an 11-point Likert scale. Although the subsequent test phase was structurally identical, scores were reset at the beginning of the first test round. Upon completion, participants again viewed their score and provided a final set of controllability and difficulty ratings.

Throughout the task, we recorded several behavioural measures, including the frequency of interventions, the number of rewarded timesteps, and participants’ causal link selections. Following the experimental task, participants completed the Hope Scale (Snyder et al., 1991). We chose this over depression inventories such as the BDI (Beck et al., 2011) because our focus was on controllability rather than clinical symptoms. This was followed by demographic questions (e.g., age, gender), input method (mouse vs. trackpad), and an open-ended question on their control strategy.

Round-level procedure Each round, participants interacted with a task interface (Figure 2) featuring three variables (X, Y, and Z) displayed as separate sliders. To minimise the influence of prior knowledge that could inflate or deflate participants’ perceived control, variables were abstract rather than real-

world labels (Btesh et al., 2025; Yarritu & Matute, 2015). A line chart displayed the values of each variable over the most recent 15 timesteps, and a shaded reward region was centred on $y = 50$. This region spanned $[30, 70]$ in the large reward region groups and $[40, 60]$ in the small reward region groups (Figure 1a). All variables and UI elements were restricted to the range $[-100, 100]$; values pushed beyond these limits were truncated. Each round consisted of thirty 2.5-second timesteps (75 seconds total), with the current score and time displayed at the top.

By default, variables shifted in discrete time according to the transition process defined in Equation 2. At each timestep, participants could intervene on only one of X or Y by clicking, dragging, or holding the associated slider. Such interventions overrode the OU network dynamics and that variable's causal parents (Pearl, 2000). If a participant moved a slider, the variable took the most recent position set and remained there for the rest of the timestep; otherwise, the variable updated according to the OU network dynamics. Holding a slider across multiple timesteps was counted as a sequence of individual interventions (e.g., holding across timesteps 9-12 equalled 4 interventions).

To introduce a cost-benefit dynamic, each intervention incurred a 1-point cost, while maintaining the target variable Z within the reward region yielded 5 points per timestep. Participants could choose not to intervene, but this would rarely result in the target slider reaching or remaining in the reward region. Thus, inaction was generally non-optimal. This structure allows learned helplessness to manifest as dysfunctional disengagement when the expected immediate (or future) returns of acting are greater than zero. Participants in low control groups may stop interacting entirely after experiencing negative returns, resulting in no further rewards but also no loss of points.

After each round, participants identified the causal link they believed was most likely present in the previous round ($X \rightarrow Y$, etc.) from 12 possible edges. This targeted question probed which causal link participants leveraged the most in their strategy while keeping the experiment brief.

Results

Pre-registered analyses

Our primary finding was that low control induced helpless behaviour. The following analyses characterise this effect and explore several secondary factors. There were no significant differences in interventions across the three test rounds ($\beta = -0.19$, $SE = 0.32$, $t(185) = -0.60$, $p = .551$), so we fit all subsequent models on collapsed data. We evaluated two linear mixed effects models: A full model including all covariates and a reduced model containing only the variables of interest (control and reward region). Formal likelihood ratio testing confirmed that including additional predictors (age, trackpad use, self-report controllability, causal link identification, and Hope Scale) did not significantly improve model fit beyond the core variables ($\chi^2(5) = 1.02$, $p = .961$). We therefore




Figure 2: Task Interface screenshot from high control and high reward group. The reward region is shown in green. Participants could intervene on variable X or Y, or do nothing, and their goal was to control variable Z. In this case the true causal structure is graph 2A (see Figure 1). The participant was rewarded on steps 3-7 and 11-15. There are 15 steps remaining in the round.




Figure 3: Model predicted intervention counts (with 95% confidence intervals) based on control and reward region size, superimposed over raw data from these groups.

present the reduced model.

Model predictions are visualised in Figure 3. Consistent with our first prediction, participants who had low control during their training phase performed fewer interventions compared to their high control counterparts ($\beta = -4.57$, $SE = 1.73$, $t = -2.65$, $p = .010$). Contrary to our second prediction, the size of reward region had no significant association with intervention count ($\beta = -0.68$, $SE = 1.73$, $t = -0.39$, $p = .696$), suggesting that reward prevalence alone is insufficient to drive helplessness in this context. We also found no correlation between participant's reported Hope Scale and their number of interventions ($r(91) = -0.09$, $p = .404$).

Exploratory analyses

To better characterise the nature of this control effect, we conducted several exploratory analyses examining task per-

formance, perceived controllability, causal learning, and the types of interventions performed.

Task performance First, we used two-tailed Welch’s t -tests to assess if round-level scoring differed from no-control or random-control strategies. The no-control strategy involved making no interventions throughout the round, incurring no cost. The random-control strategy applied an intervention at each timestep to either X or Y, with the intervention value drawn uniformly from the range $[-100, 100]$. For each strategy, we simulated 1000 rounds for all 6 causal graphs within each of the 4 different experimental groups. During the test phase, all groups scored better than the random intervention strategy (all $p < .01$). However, performance relative to the no-control baseline varied by group. Whilst the low control–small reward region group did not perform significantly above this baseline ($p = .654$), all other groups did: low control–large reward region ($t(74.03) = 3.23, p = .002$), high control–small reward region ($t(74.00) = 2.61, p = .011$), and high control–large reward region ($t(65.01) = 4.18, p < .001$). This suggests that whilst participants generally learned to intervene effectively, the combination of low control and limited rewards left participants performing no better than had they not intervened at all.

Figure 4 shows within-round success by timestep with linear model fits. A smaller proportion of participants succeeded in the low control groups compared to the high control groups ($\beta = -0.04, SE = 0.02, t(114) = -2.22, p = .029$). Participants in the small reward region groups also saw decreases in scoring proportion ($\beta = -0.04, SE = 0.02, t(114) = -2.12, p = .036$). While the proportion of scoring participants increased by 1% each timestep ($SE = 0.009, t(114) = 10.69, p < .001$), there was a significant control $\times$ timestep interaction ($\beta = -0.002, SE = 0.0009, t(114) = -3.28, p < .001$), indicating that participants in low control groups had impaired improvement trajectories. Whilst both control and reward prevalence independently affected performance, previous experience with uncontrollability hampered performance.

Perceived controllability We used cumulative link mixed-models (CLMM) from the Ordinal package in R (Christensen, n.d.) to fit participants self-reported controllability at the 11 Likert levels (0 – 10). These results confirm that our manipulation influenced perceived as well as objective controllability: participants in the high control groups ($M = 4.53$) reported higher perceived controllability than low control groups ($M = 3.89, OR = 3.8, p = .008$). Consistent with the role of rewards in shaping illusions of control (Langer, 1975), those in the large reward region ($M = 4.57$) also reported higher perceived controllability than small reward region groups ($M = 3.85, OR = 2.8, p = .044$).

Identifying causal links Participants were above chance ($>25\%$) in identifying a causal link correctly in both training ($95\%CI[0.28, 0.40], p = .001$) and test ($95\%CI[0.27, 0.39], p = .003$) phases. This performance




Figure 4: Proportion of participants scoring per timestep during test rounds, with linear model fits and baseline strategy comparisons, across large and small reward regions.

was not uniform across all groups in the test phase. Only participants with small reward regions achieved above-chance accuracy in identifying correct causal links, for high ($n = 75, 95\%CI[0.31, 0.55], p < .001$) and low control ($n = 63, 95\%CI[0.25, 0.50], p = .04$) groups. Identification accuracy of participants with large reward regions did not differ significantly from chance. A mixed-effects logistic regression revealed that participants in small reward region groups were more likely to identify correct causal links ($OR = 1.92, 95\%CI[1.10, 3.37], p = .022$) compared to participants with large reward regions. This points to a more nuanced relationship between reward prevalence and causal learning than a simple effect on helpless behaviour alone.

Quality of interventions Here, we leverage the causal dynamics to retrospectively quantify the normative efficacy of each action taken, allowing us to examine how intervention quality during training related to test phase behaviour. We simulated what would have occurred had participants not intervened at a given timestep, for each intervention performed in the task. An example can be found in Figure 5. If a participant intervened on variable X at timestep t , we used the values of its parent variables at $t - 1$ within the causal structure to determine the pre-intervention value of X at t . Based on the current values at t , we then used Equation 2 to calculate the counterfactual value of Z at $t + 1$ and determine whether it was within the reward range. As in the experimental task, we used $\theta = [1, -1]$ for high control groups, whilst we used $\theta = [0.25, -0.25]$ for low control groups (expected value of the beta distribution).

By comparing the counterfactual value of Z to the actual value of Z, interventions could be classified as detrimental (reward lost due to intervention), inconsequential (no effect on reward outcome), necessary (reward only achieved with intervention), or redundant (reward achieved regardless). Counts of intervention types performed during the training phase are

shown in Table 1. The zero counts for detrimental and redundant interventions in the low control groups reflected the nature of the controllability manipulation itself. With weak causal influence and high noise, participants' actions rarely had sufficient impact to meaningfully help or hinder outcomes during training, and future designs might productively vary these manipulations more independently.




Figure 5: Visual explanation of intervention types. In each panel, X was intervened on at t (dark blue). The counterfactual value of Z at $t + 1$ (light green) is calculated based on the value X would have taken (light blue) without the intervention. Values of X (at $t + 1$) and Y (at all timesteps) have been omitted for clarity.

We found that some intervention types were significant predictors of test phase interventions. Specifically, a higher number of inconsequential ($\beta = 0.03, SE = 0.00, t(88) = 7.54, p < .001$), necessary ($\beta = 0.04, SE = 0.01, t(88) = 3.99, p < .001$), and redundant ($\beta = 0.06, SE = 0.03, t(88) = 2.18, p = .032$) interventions increased the number of test-phase interventions. Moreover, a greater number of necessary ($OR = 1.15, 95\%CI[1.09, 1.21], p < .001$) and redundant ($OR = 1.16, 95\%CI[1.04, 1.29], p = .008$) interventions was associated with higher perceived controllability. In contrast, inconsequential interventions reduced controllability ratings ($OR = 0.96, 95\%CI[0.93, 0.98], p < .001$).

Table 1: Count of intervention types in the training phase by group.

Group	Detri.	Incon.	Neces.	Redun.
High reward high control	90	810	248	94
Low reward high control	68	1140	50	196
High reward low control	0	1356	156	0
Low reward low control	0	1071	53	0

Discussion

This study provides a novel exploration of learned helplessness in a dynamic control setting. We found that participants who experienced uncontrollable training rounds interacted less in a test phase, and performed worse than those in high control groups. Whilst reward prevalence did not impact helplessness measures, participants in the low reward groups were more accurate in identifying causal links in their respective causal graphs. One interpretation is that lower rewards necessitated attention to causal relations, whereas in the high reward groups, participants could rely on chance or model-free strategies like win-stay-lose-sample (Robbins, 1952). This also supports the notion that when tasks require inductive reasoning, intrinsic motivations such as curiosity and perceived agency may outweigh extrinsic rewards in guiding behaviour.

The dynamic nature of our experimental framework offers a distinct methodological advantage over static paradigms: the ability to quantify the normative efficacy of actions given full knowledge of the causal dynamics, and thereby examine the factors that drive divergence from optimal behaviour in practice. This analysis revealed a dissociation between the feeling of control and the act of controlling. Whilst necessary and redundant interventions predicted greater test phase engagement and higher perceived controllability, inconsequential interventions predicted engagement but reduced perceived controllability. The brevity of our paradigm limits direct generalisation to clinical phenomena where helplessness develops over prolonged exposure, but the paradigm nonetheless offers an alternative perspective. The fine-grained analysis may complement longer-timescale studies by helping to pinpoint whether helpless behaviour stems from a lack of motivation or an inability to distinguish causal efficacy from noise.

Several limitations should be noted. First, our low control manipulation confounded stochastic causal strength and exogenous noise, making it difficult to isolate the precise component responsible for inducing helpless behaviour. Second, manipulating reward region size is relatively imprecise, as the impact of reward on helpless behaviour likely depends on nuanced properties such as reward frequency and probability, and may follow a non-linear relationship. Third, asking participants to identify only their most relied-upon causal link, rather than their full causal structure beliefs, limits the conclusions that can be drawn about participants' causal representations. Future work should address these limitations, potentially examining a smaller set of causal structures in greater depth to more precisely characterise these relationships.

Learned helplessness permeates various aspects of life, often extending beyond well-defined experimental trials. Our study demonstrates that helplessness can be induced within dynamic environments, where causal structure, controllability, and reward exert distinct influences across different stages. This framework provides a robust method for exploring helplessness and other goal-directed behaviour. By grounding these findings in causal representations, we offer a new perspective on how individuals navigate shifting environments.

References

- Abramson, L. Y., Seligman, M. E., & Teasdale, J. D. (1978). Learned helplessness in humans: Critique and reformulation. *Journal of Abnormal Psychology*, 87(1), 49–74. <https://doi.org/10.1037/0021-843X.87.1.49>
- 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>
- Beck, A. T., Ward, C. H., Mendelson, M., Mock, J., & Erbauch, J. (2011, September). Beck Depression Inventory. <https://doi.org/10.1037/t00741-000>
- Broadbent, D. E., FitzGerald, P., & Broadbent, M. H. P. (1986). Implicit and explicit knowledge in the control of complex systems. *British Journal of Psychology*, 77(1), 33–50. <https://doi.org/10.1111/j.2044-8295.1986.tb01979.x>
- Btsh, V., Bramley, N. R., Speekenbrink, M., & Lagnado, D. A. (2025). Less is more: Local focus in continuous time causal learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition*. <https://doi.org/10.1037/xlm0001451>
- Cartwright, N. (2004). Causation: One word, many things. *Philosophy of Science*, 71(5), 805–819. <https://doi.org/10.1086/426771>
- Cheng, P. W. (1997). From covariation to causation: A causal power theory. *Psychological Review*, 104(2), 367–405. <https://doi.org/10.1037/0033-295X.104.2.367>
- Christensen, R. H. B. (n.d.). Cumulative Link Models for Ordinal Regression with the R Package ordinal.
- Cordova, D. I., & Lepper, M. R. (1996). Intrinsic motivation and the process of learning: Beneficial effects of contextualization, personalization, and choice. *Journal of Educational Psychology*, 88(4), 715–730. <https://doi.org/10.1037/0022-0663.88.4.715>
- Davis, Z. J., Bramley, N. R., & Rehder, B. (2020). Causal structure learning in continuous systems. *Frontiers in Psychology*, 11. <https://doi.org/10.3389/fpsyg.2020.00244>
- Davis, Z. J., Bramley, N. R., Rehder, B., & Gureckis, T. (2018). A casual model approach to dynamic control. *Proceedings of the Annual Meeting of the Cognitive Science Society*, 40(0).
- Hiroto, D. S. (1974). Locus of control and learned helplessness. *Journal of Experimental Psychology*, 102(2), 187–193. <https://doi.org/10.1037/h0035910>
- Kelley, H. H. (1973). The processes of causal attribution. *American Psychologist*, 28(2), 107–128. <https://doi.org/10.1037/h0034225>
- Langer, E. J. (1975). The illusion of control. *Journal of Personality and Social Psychology*, 32(2), 311–328. <https://doi.org/10.1037/0022-3514.32.2.311>
- Langer, E. J., & Rodin, J. (1976). The effects of choice and enhanced personal responsibility for the aged: A field experiment in an institutional setting. *Journal of Personality and Social Psychology*, 34(2), 191–198. <https://doi.org/10.1037/0022-3514.34.2.191>
- Leotti, L. A., & Delgado, M. R. (2011). The inherent reward of choice. *Psychological Science*, 22(10), 1310–1318. <https://doi.org/10.1177/0956797611417005>
- Maier, S. F., & Seligman, M. E. (1976). Learned helplessness: Theory and evidence. *Journal of Experimental Psychology: General*, 105(1), 3–46. <https://doi.org/10.1037/0096-3445.105.1.3>
- Miller, W. R., & Seligman, M. E. (1975). Depression and learned helplessness in man. *Journal of Abnormal Psychology*, 84(3), 228–238. <https://doi.org/10.1037/h0076720>
- Murayama, K. (2022). A reward-learning framework of knowledge acquisition: An integrated account of curiosity, interest, and intrinsic–extrinsic rewards. *Psychological Review*, 129(1), 175–198. <https://doi.org/10.1037/rev0000349>
- Osborn Popp, P. J., Newell, B. R., Bartels, D. M., & Gureckis, T. M. (2025). Can cognitive discovery be incentivized with money? *Journal of Experimental Psychology: General*, 154(4), 1025–1037. <https://doi.org/10.1037/xge0001682>
- Overmier, J. B., & Seligman, M. E. (1967). Effects of inescapable shock upon subsequent escape and avoidance responding. *Journal of Comparative and Physiological Psychology*, 63(1), 28–33. <https://doi.org/10.1037/h0024166>
- Pearl, J. (2000). *Causality: Models, reasoning, and inference*. Cambridge University Press.
- Robbins, H. (1952). Some aspects of the sequential design of experiments. *Bulletin of the American Mathematical Society*, 58(5).
- Schulz, E., Klenke, E. D., Bramley, N. R., & Speekenbrink, M. (2017). Strategic exploration in human adaptive control. *Proceedings of the Annual Meeting of the Cognitive Science Society*, 39(0).
- Sharot, T. (2011). The optimism bias. *Current Biology*, 21(23), R941–R945. <https://doi.org/10.1016/j.cub.2011.10.030>
- Sloman, S. (2005, August). *Causal Models: How People Think about the World and Its Alternatives* (1st ed.). Oxford University Press New York. <https://doi.org/10.1093/acprof:oso/9780195183115.001.0001>
- Sloman, S., & Lagnado, D. (2015). Causality in thought. *Annual Review of Psychology*, 66(1), 223–247. <https://doi.org/10.1146/annurev-psych-010814-015135>
- Snyder, C. R., Harris, C., Anderson, J. R., Holleran, S. A., Irving, L. M., Sigmon, S. T., Yoshinobu, L., Gibb, J., Langelle, C., & Harney, P. (1991). The will and the ways: Development and validation of an individual-differences measure of hope. *Journal of Personality and Social Psychology*, 60(4), 570–585. <https://doi.org/10.1037/0022-3514.60.4.570>
- Teodorescu, K., & Erev, I. (2014). Learned helplessness and learned prevalence: Exploring the causal relations among perceived controllability, reward prevalence, and exploration. *Psychological Science*, 25(10), 1861–1869. <https://doi.org/10.1177/0956797614543022>

- Uhlenbeck, G. E., & Ornstein, L. S. (1930). On the theory of the Brownian Motion. *Physical Review*, 36(5), 823–841. <https://doi.org/10.1103/PhysRev.36.823>
- Yarritu, I., & Matute, H. (2015). Previous knowledge can induce an illusion of causality through actively biasing behavior. *Frontiers in Psychology*, 6. <https://doi.org/10.3389/fpsyg.2015.00389>