

The Paradox of Time in Dynamic Causal Systems

Zachary J. Davis (zach.davis@nyu.edu)

Department of Psychology, NYU, 6 Washington Pl, New York, NY 10003 USA

Neil Bramley (Neil.Bramley@ed.ac.uk)

Psychology Department, The University of Edinburgh, 7 George Sq, Edinburgh, UK

Bob Rehder (bob.rehder@nyu.edu)

Department of Psychology, NYU, 6 Washington Pl, New York, NY 10003 USA

Abstract

Recent work has shown that people use temporal information including order, delay, and variability to infer causality between events. In this study we build on this work by investigating the role of time in dynamic systems, where causes take continuous values and also continually influence their effects. Recent studies of learning in these systems explored short interactions in a setting with comparatively rapidly evolving dynamics and modeled people as relying on simpler, resource-limited strategies to grapple with the stream of information (Davis et al., 2020). A natural question that arises from such an account is whether interacting with systems that unfold more slowly might reduce the systematic errors that result from these strategies. Paradoxically, we find that slowing the task indeed reduced the frequency of one type of error, but increased the error rate overall. To capture the differences between conditions, we introduce a novel *Causal Event Segmentation* model based on the notion that people compress the continuous scenes into events and use these to drive structure inference.

Keywords: causal learning; time; continuous; event cognition; interventions

Introduction

Learning about causal structure is central to higher level cognition, allowing people to predict the future, select beneficial actions, and make sense of the past. The study of how people learn causal structure has historically focused on simple scenarios involving the presence or absence of binary variables (e.g. did a patient take a drug, and did they get sick?). This has taught us much about how people use causal structure for a host of decisions (e.g. Ali et al., 2011; Fernbach & Erb, 2013; Hayes et al., 2014; Sloman, 2005). However, this focus on simple stimuli obscures other important questions, such as how we incorporate continuous covariation and temporal information into our causal judgments.

Time is central to our notions of causality (Hume, 1740), making it unsurprising that temporal contiguity is one of the strongest psychological cues to causality. Sophisticated expectations about delays between events shape causal judgments (Hagmayer & Waldmann, 2002; Pacer & Griffiths, 2012), intervention selections (Bramley et al., 2018), and goal directed actions (Buehner & May, 2003). People also judge that highly variable delays are less causal (Greville & Buehner, 2010) and use variability as a cue for structure in the absence of order or covariational cues (Bramley et al., 2018).

Prior work on the role of time in causality has focused on delay distributions, i.e. the time that it takes for one event to influence another, with events largely treated as punctate

rather than extended in time. In this project we instead study a fully continuous setting in which continuous valued causes continually affect rates of *change* of their effects, introducing a different set of representational challenges. Rather than reasoning directly about rates of occurrence of events or delay distributions between events, people must reason from unfolding timeseries data.

How might varying the speed at which a continuous system evolves affect what people learn about it? Extrapolating from the literature on events cited above, one might expect that a more slowly evolving system would make learners less likely to infer the presence of causal linkages between variables. Yet reduced speed may have advantages as well. In the setting originally explored by Davis et al. (2018), people were well described with a Local Computations (LC) model, which characterized them as focusing on establishing the relationship between pairs of variables independently, that is, rather than controlling for other variables—as one would by separately considering the full space of possible structural models wholesale. The key support for the LC model came from a particular characteristic error. Participants frequently inferred direct connections between variables that were indirectly related (e.g. in the network $X \rightarrow Y \rightarrow Z$ concluding incorrectly that additionally $X \rightarrow Z$). This is an error first observed in studies with binary variables (Fernbach & Sloman, 2009; Rottman & Keil, 2012). A natural question that arises is whether these errors in Davis et al. (2018) were due to participants failing to perceive the short lag in influence from X to Y and then to Z . Slowing the system will increase the difference in rates of change between direct and indirect effects, which we hypothesize will reduce these systematic errors.

We also aim to understand how people learn causal structure from a continuous flow of information by comparing different formal accounts of how people represent continuous information and use it to infer causal relationships. Firstly, we follow Davis et al. (2020) in describing people as computing likelihoods on the basis of the continuous dynamics directly—either considering all hypotheses in parallel (normative model), or focusing separately on individual edges (Local Computations variant). Secondly, we introduce a new preliminary account of how people might handle continuous information in time—the *Causal Event Segmentation* (CES) model—that characterizes people as segmenting the continuous stream into discrete events, and using those to infer causal

structure.

In summary, we ask two questions. Firstly, does slowing the dynamics of the system reduce the systematic errors that have been previously observed? We do find the expected reduction in those errors but at the cost of reducing overall accuracy. Secondly, how do people represent continuous information in dynamic systems? We find that a model describing people as segmenting continuous information into discrete events captures people’s behavior across conditions.

Ornstein–Uhlenbeck networks

The stimuli in our task were generated using a new approach for simulating continuous causal systems first proposed in Davis et al. (2018). See Davis et al. (2020) for a full explication of the generative process, but briefly Ornstein–Uhlenbeck (OU) networks represent causality with autoregressive processes that move towards a basin point as a function of time (Uhlenbeck & Ornstein, 1930). Importantly, however, when one variable is causally influenced by another variable (as defined by the causal structure of the OU network), this is modelled by making the effect’s basin point nonstationary, following some function of the state of its cause(s). We here further restrict these functions such that the effect either asymptotes to a value equal to the cause’s value (“regular” connections) or to the negative of the cause’s value (“inverted” connections). If a variable has more than one cause, we model it as attracted to the sum of the basin points defined by each of its causes. Formally, the change in a variable v_i following time t is given by

$$P(\Delta v_i^t | v^t, \omega, \sigma, \Theta) = \omega \left[\left[\sum_j \theta_{ji} \cdot v_j^t \right] - v_i^t \right] + N(0, \sigma) \quad (1)$$

where ω is the asymptote rate (“spring rigidity”), σ is the endogenous noise of each variable, and θ_{ji} is the causal impact of variable j on variable i (1 in the case of a direct, and -1 in the case of an inverse relationship). Simply put, the mean that variable i reverts to is the sum of the values of its causes, each first multiplied by their respective θ s. To accommodate interventions we use the “Do()” operator (Pearl, 2009), whereby an intervened on variable takes the assigned value with a probability of 1, ignoring all other endogenous or exogenous factors.

In this work we manipulate two aspects of the generative model. First, ω (rigidity) defines how tightly an effect tracks its cause. Higher values of ω result in an effect reaching its asymptote more quickly (instantaneously when $\omega = 1$). Second, the *update rate* determines the amount of real time (i.e. from the point of view of external observer) between each system update (i.e., the time from t to $t + 1$ when computing the change in variable v_i , Δv_i^t). Rigidity and update rate represent two alternative ways in which to manipulate system speed. Decreasing ω results in effects responding to their causes more slowly holding system noise (σ) constant. Decreasing the update rate also affects of the responsiveness of




Figure 1: Rate of change toward an asymptote of 80 in each condition. Stimuli were generated with a small amount of noise ($\sigma = 2$) to show the changing update rates by condition.

effects but in addition lowers the noise in the system from the point of view of an external observer. Fig. 1 illustrates these effects. We assess whether either of these manipulations result in a lower rate of local computations style errors.

Slowing the process does not come without drawbacks. Larger real time intervals between each application of Eq. 1 formally reduces the amount of evidence produced in a set period. To demonstrate this, we simulated 100 participants in each condition interacting with each of the 25 causal models tested in this experiment (Fig. 2B), and used the posterior entropy over graphs according to an optimal learner as a measure of the learnability of the systems.¹ There was a considerable increase in log posterior entropy from the fast-rigid condition (-224) to the fast-flexible (-55) and slow-rigid (-80) conditions, and even further increase in entropy for the slow-flexible condition (-13). These results demonstrate that while potentially reducing systematic, local computations style errors, slower dynamics come at the cost of reduced information with which to learn.

Method

Participants

205 participants (87 female, 1 no response; age $M = 37.2$, $SD = 11.8$) were recruited from Amazon Mechanical Turk using psiTurk (Gureckis et al., 2016). They were paid a base payment of \$3 plus performance related bonuses ($M = \$0.97$, $SD = \$0.46$) and the task took 32.6 minutes ($SD = 18.3$). Participants were randomly assigned to one of four conditions. Those who made a causal judgment before intervening on any slider on over 90% of trials were excluded (Table 1).

Materials

Participants interacted with a number of causal devices represented by three vertical sliders that moved on their own according to the hidden causal structure and OU process, but

¹The simulated participants intervened using an idealized version of the intervention strategies observed in Davis et al. (2018), intervening on each variable for 1/3 of the trial and evenly splitting their intervention choices between the extreme range of the variables (100 and -100).




Figure 2: Stimuli. (A) Task environment. Sliders turn blue when intervened on. (B) All tested causal graphs, presented in random order. Black arrowheads denote regular connections, white arrowheads denote inverse connections.

could also be intervened on, by clicking and dragging to set their levels, overriding their normal causes (see Fig. 2A).² The sliders were constrained to be between -100 and 100, and the buttons on the slider presented a rounded integer value in addition to moving up and down. A timer at the top of the page counted down from 45 seconds at 1s increments, and at the bottom of the page were six additional sliders (one for each potential causal relation). Responses could be one of three options: ‘Inverted’, ‘None’, or ‘Regular’, corresponding to $\theta < 0$, no relationship ($\theta = 0$), and $\theta > 0$, respectively. Participants were pretrained on these terms in the instructions.

Stimuli and Design

Participants were tested on 25 causal graphs (see Fig. 2B) that were roughly balanced across a number of factors, such as the number of inverted and regular links and the number of links between each variable. The graphs were presented in random order for a total of 25 trials. The OU parameters used during training and the test were $\sigma = 5$ and $\theta = [1, 0, -1]$ for regular, none, or inverse connections, respectively.

Participants were randomly assigned to one of four conditions that determined the other parameters, ω and update rate. We varied the rate at which the process unfolded, with the slider values updating at either 100 (“fast”) or 300ms (“slow”). Crossed with this, we manipulated the rigidity of the OU process, setting ω to be either 0.1 (“rigid”) or 0.05 (“flexible”). ω sets the rate at which the process asymptotes, when $\omega = 0.1$ the variables move 10% of the way toward their basin point in expectation, and when $\omega = .05$ the variables move 5% of the way toward their basin point in expectation (see Fig. 1). These two factors independently manipulate the rate of change of variables towards their asymptotic value, with ω changing the formal properties of the OU network itself, and update rate changing the experienced real-time pace of the process.

Table 1: Breakdown of conditions and number of participants. Counts prior to exclusions in parentheses.

Condition	Update rate	ω	n
fast-rigid	100ms	0.1	45 (54)
fast-flexible	100ms	0.05	42 (53)
slow-rigid	300ms	0.1	44 (51)
slow-flexible	300ms	0.05	38 (47)

Procedure

Participants first completed an interactive instruction section that used a sequence of videos to explain the nature and goals of the task, how to intervene, as well as the trial duration. They were instructed that, for a randomly selected trial, they would receive a bonus of \$0.25 for each correct causal link judgment (out of ‘no link’, ‘regular’ and ‘inverse’ for each of the 6 directed links). Importantly, this bonus scheme was demonstrated with a hypothetical participant who observed a chain network and correctly identified the two existing causal links but incorrectly added an additional direct link between the indirect effects. Participants were told that this participant received a reward of \$1.25 for the correct responses but missed out on an additional \$0.25 for marking the direct connection between indirect effects. Participants could not proceed to the task until they correctly answered five comprehension check questions probing if they knew the duration of each trial, the difference between a regular and inverted connection, that there can be more than one connection per network, and that they would have to provide a response for all six possible connections on each trial.

In the main task, participants completed 25 trials lasting 45 seconds each. A trial was initiated by pressing the “Start” button at the top of the page, whereupon the sliders began updating according to the OU process at either 100 or 300ms, depending on condition. Participants were free to click, drag, or hold any slider to any value for any amount of time overriding its normal causes. After releasing a slider, it continued

²See zach-davis.github.io for a demo.




Figure 3: Descriptive statistics. (A) Proportion of causal links correctly identified by condition. (B) Among variable pairs that were indirectly but not directly causally related, proportion that received a correct “none” response (indicating the absence of a direct connection). (C) Proportion of continuous interventions that are held at one value over time. As dynamics slow participants hold at the same value for longer. Plot points in panels (A) and (B) derived from maximum a posteriori predictions from three models. Horizontal dashed lines denote chance responding. Error bars denote standard errors of the mean.

to move according to the OU process.

Participants could make (and revise) their causal judgments at any point during the trial, but could not proceed to the next trial until they had entered a judgment for all six potential causal relations. No feedback was provided. After completing the 25 trials, participants were informed of their bonus and completed a brief post-test questionnaire.

Results

Descriptive Results

See Fig. 3 for descriptive results. Across all conditions, participants were above chance (.33) in identifying causal links ($M = .71$, $SD = .22$), $t(168) = 22.30$, $p < .001$. They were slightly more likely to recognize regular (.83) than inverse (.79) causal links, $t(168) = 3.89$, $p < .001$. Participants were also more likely to correctly classify causal links as the experiment progressed, as confirmed by a random effect regression with subject-level intercept and slope for trial number (Mean $\beta = .004$), $t(168) = 6.76$, $p < .001$.

Successful learning relies on effective interventions—i.e., ones that are extended in time and involve large swings of each variable’s value. In our task, a two-way ANOVA did not show differences in average duration of interventions ($M = 3.07s$) as a function of either update rate or spring rigidity (all F s < 1). However, the average range of interventions—defined as the minimum slider value subtracted from the maximum value during an intervention bout—varied substantially with condition: 138.5, 102.4, 89.0 and 69.1 in the fast-rigid, fast-flexible, slow-rigid, and slow-flexible conditions, respectively; an ANOVA revealed an effect of update rate, $F(1, 165) = 8.89$, $p = .003$, of rigidity, $F(1, 165) = 44.8$, $p < .001$, and no interaction.

That is, learners in slower systems generated less variable movement, a finding supported by the increase in propensity to hold variables at one value for extended periods of time as the dynamics slow (Fig. 3C).

Fig. 3A shows that, as expected, accuracy in identifying causal links decreased as the dynamics slowed. Although subjects were above chance in all conditions ($ps < .001$), a two-way ANOVA on accuracy showed decreased accuracy as a function of both update rate ($F(1, 165) = 9.9$, $p = .002$) and rigidity ($F(1, 165) = 7.4$, $p = .008$), but no interaction. Somewhat paradoxically, however, Fig. 3B shows that slowing the dynamics actually had the *opposite* influence on accuracy for indirect effects. As dynamics slowed, accuracy *improved*. A two-way ANOVA showed this was due to both update rate ($F(1, 165) = 15.1$, $p < .001$) and spring rigidity ($F(1, 165) = 5.1$, $p = .025$), with a small interaction between the two ($F(1, 165) = 4.5$, $p = .036$).

One reason for the reduced accuracy by condition may be that, as demonstrated above, slowing down the dynamics results in decreased information produced by the system. For example, a trial in the “slow” condition had 150 data-points, versus 450 in the “fast” condition. It was also shown that learners in slower systems perform less informative interventions, compounding the general decrease in information. To test whether accuracy differences between conditions was driven solely by overall system information, we computed the posterior entropy over the 729 possible graphs given the data that participants generated as they interacted with the sliders, according to a normative learner that knows the exact parameters of the system (but not causal structure). As expected, lower entropy was associated with higher ac-

accuracy, $r(167) = -.67, p < .001$, confirming that more system information aided learning. However, a two-way ANCOVA revealed that, even after controlling for differences in the conditions' posterior entropy, there remained a difference in accuracy between conditions, with main effects of speed ($F(1, 164) = 11.8, p < .001$) and spring rigidity ($F(1, 164) = 8.8, p = .004$), and an interaction between the two ($F(1, 164) = 7.6, p = .006$).

These results indicate independent effects of the two ways we slowed the dynamic system. Participants classified causal links less accurately both as the asymptote speed ω slowed from .1 to .05 and as the gap between t and $t+1$ increased from 100 to 300ms. However, as predicted, these manipulations *increased* accuracy on the key test of direct and indirect connections.

Modeling

To better understand participants' judgments we compared them to those of several causal structure learning models. In particular, we assessed two models that had previously been found to capture people's behavior, as well as a new model inspired by the event segmentation literature. For each participant and model, the model received as input the slider values and the participant's interventions and yielded a posterior distribution over the 729 causal graphs. We now briefly describe each model starting with our new addition.

Causal Event Segmentation model. The Causal Event Segmentation (CES) model describes people as segmenting continuous variables into events, and using these events as cues for causality. Roughly put, it conceives of people as drawing causal links only when their actions 'make something happen.' This is operationalized by taking all timepoints where a participant intervenes on a slider and drawing a causal link if another variable crosses some threshold value during that intervention bout. See Fig. 4 for visual demonstration of these principles.

The CES model assigns a central role to interventions, in that it only infers a causal link between an intervened upon variable as the "root" cause and a non-intervened upon "end" variable as the effect. More precisely, the CES model takes all timepoints where a variable is intervened on, and registers an event if any other variable crosses a threshold during that time. For example, the horizontal dashed line in Fig. 4 shows a threshold at which the CES model would register an event if another variable was being intervened upon.³ To determine the sign of the link, for all timepoints during the event (where the root variable is intervened upon and an end variable is above a threshold) the CES model compares the signs of the two variables. If on average the signs match, the model draws a "regular" causal link, if they do not match the model draws an "inverse" causal link.

If no end variables cross a threshold, no causal links are

³A threshold is an absolute value (e.g. 50), such that "crossing" a threshold could mean being greater than that value or less than its negative.




Figure 4: Illustration of CES model. During participant's intervention bout on X from approximately 1 to 3 seconds, Z crosses threshold (here shown as 50). Histograms on the right demonstrate distributions over fitted threshold values by condition, with colors corresponding to condition as in Fig. 1.

drawn. To only register events when a threshold is *crossed*, the CES model excludes all cases where a potential end variable is above threshold before the intervention begins. In addition, the CES model does not account for temporal delays between events. The only links that can be drawn originate from an intervened on node. Links are not drawn between two non intervened on nodes, even if they both cross the threshold. To account for uncertainty in participant judgments, per participant we fit a threshold parameter, as well as a guess parameter that corresponded to the probability of responding counter to the predictions of the events model.

Normative model. The normative model inverts the generative model to optimally infer the structure most likely to have produced the evidence (see Davis et al., 2020), starting each problem with a uniform prior, and assuming a static degree of parameter uncertainty. At each timepoint, the normative model evaluates the likelihood of each variable's value given the values of all other variables at the previous timepoint, as well as information about the participant's interventions, and updates the probability of each potential graph. Updating over all timepoints and multiplying by the prior probability of each graph (we assume uniformity) yields a posterior distribution over graphs.

Two choices were made to account for participant uncertainty. While participants see examples of OU networks in the instructions, they will of course not develop exact representations of key parameters ω , θ , and σ . We assume the same hyperpriors over these parameters as (Davis et al., 2020), representing a reasonable degree of uncertainty about their precise values while still maintaining qualitative similarity with the generating process. To account for uncertainty in participant judgments, a softmax function was applied to the posterior over graphs, with a separate temperature parameter fit for each participant to maximize the posterior probability of the graph chosen by participants.

Local Computations model. The local computations (LC) model is a broadly successful model of causal learning that




Figure 5: Evaluation measures for the primary models. (A) Proportion match between participant causal judgments and model’s estimate of the causal relationship. (B) Mean BIC per participant. The LC and normative models were fit with one per-participant softmax temperature parameter. The CES model was fit with two parameters per participant (threshold and guess parameters). (C) Number of participants best fit by each model, as measured by BIC.

describes people as focusing on pairs of variables, rather than evaluating the evidence with respect to the full space of possible structural models. At each timepoint, the LC model evaluates the individual contribution of each variable to each other variable’s value, ignoring the possible contribution of other links. In other respects the LC model is identical to the normative model. LC based models have been proposed as accounts of how people build causal models in a resource-efficient way (Bramley et al., 2017; Fernbach & Sloman, 2009). As mentioned, Davis et al. (2020) showed that the LC model best fit participants in a very similar task to this study’s fast-rigid condition. Here we test the extent to which these results generalize to different time characteristics.

Baseline model. The baseline assumes participants have an equal probability of responding for any graph. It has no fitted parameters.

Modeling Results

Fig. 5A shows the proportion of participants’ edge judgments that correspond with the most probable graph under each model. As the system slows down, participants are less likely to make the same judgments of causal relationships as the models. This suggests increasing noise in their judgments, a result that may cohere with the reduced information in the system as it slows.

Fig. 5B shows the relative performance of the models, as measured by mean Bayesian Information Criterion per participant. While all models outperform the baseline model, the distinction decreases as the system slows down. Overall, the CES model is the best-fitting model, although the normative and LC models narrowly outperform the CES model in one condition (slow-rigid).

This strong performance of the CES model is reflected in

the number of participants best fit by each model (Fig. 5C). Across all conditions, the CES model fits the majority of participants, although its separation from the other models is slightly ameliorated as the dynamics slow down.

Fig. 3 gives insight into why the CES outperforms the other models. While cell B shows that the LC model also predicts increased accuracy on indirect effects as the dynamics slow, only the CES model predicts the simultaneous *decrease* in overall accuracy.

Intervention Analyses

The above models can also shine light on *why* participants acted differently in the different conditions. Note that the informativeness of an intervention depends on the underlying learning model. For example, a good intervention for the CES model involves holding an intervened-on variable at a particular value for an extended period (providing the time needed for an event to be recorded). In contrast, for the normative model large swings in that variable are preferred. Under the assumption that learners choose interventions that are appropriate for how they learn, one can ask whether their interventions are more explicable under one model versus another. In particular, we ask which model provides the best account of how interventions varied over our experimental conditions. To this end, we compared, for each model, how each participant’s interventions *would have* fared when applied to each of the four experimental conditions. For example, if a participant in the fast-rigid condition held variable X at 100 for 10 seconds, we simulated that intervention in all four conditions. A model that learns as people do predicts that a participant’s interventions will lead to better learning when applied to his or her own condition as compared to one of the others—because those interventions were appro-

priate for that model in that condition. We fed this simulated participant data into our inference models. We found that the CES model was more accurate when participants' actions were matched with the condition in which they were actually performed ($t(168) = 5.15, p < .001$). However, this was not the case for the LC ($t(168) = .10, p = .92$) or normative ($t(168) = 1.02, p = .31$) models. This is consistent with the idea that participants acted to maximize learning under the CES model, roughly speaking this means they acted to produce many informative "events". Indeed, Fig. 3C confirms that participants in the slower conditions held variables at one end of the range longer than those in faster conditions—as appropriate for a CES learner.

Discussion

This paper investigated the impact of timing on causal learning in continuous dynamic systems. We predicted that reducing a system's speed might ameliorate a particular type of error captured by the local computations model—given $X \rightarrow Y \rightarrow Z$, incorrectly inferring a direct relationship between X and Z —on the grounds that learners would more readily perceive the lag in influence from X to Y and then to Z . Yet, we also noted that people are generally less likely to infer a causal relationship the greater the lag between cause and effect. In fact, we found just this paradoxical effect of time on learning. Whereas slowing the dynamics resulted in increased accuracy for indirect effects, it resulted in reduced accuracy overall. That is, rather than having a uniformly positive or negative effect, changes in system timing led to a trade-off between different types of errors.

We also introduced the CES model, finding that it fit the majority of participants across conditions. The success of this model fits nicely with work suggesting that people naturally segment continuous streams of information into discrete events (for review, see Zacks, 2020). That said, the CES model in its current form is highly exploratory with plenty of room for improvement and further testing. For one thing, it only infers a direct connection between an intervened-on root variable and end variable that registers an effect, whereas people have been shown to infer structure from linking sequences of events (Bramley et al., 2018). In addition, given the importance of interventions to produce events for the CES to learn from, a future direction would be modeling the CES's prescriptions for how one should intervene to maximize learning.

A limitation of this project is that it confounds the amount of information produced by the system with the rate of change, in that the 100ms update rate produced three times as many datapoints than a 300ms update rate in 45 seconds. We made this design choice because the alternative was to have 30 minute versus 90 minute experiments depending on condition. In the results section we showed that the overall amount of information produced by the system could not explain the difference between conditions, but a natural further experiment would be to test this directly by holding the total

number of observations constant.

Learning the relationships between continually shifting variables in real-time is as challenging as it is common. In this paper we identified factors that modulate performance in continuous dynamic environments, and proposed a new model for causal learning inspired by people's ability to abstract and discretize their experiences. We find support for the idea that, in these informationally rich settings, people use events triggered by their actions to infer causal structure.

References

- Ali, N., Chater, N., & Oaksford, M. (2011). The mental representation of causal conditional reasoning: Mental models or causal models. *Cognition*, 119(3), 403–418.
- Bramley, N. R., Dayan, P., Griffiths, T. L., & Lagnado, D. A. (2017). Formalizing neurath's ship: Approximate algorithms for online causal learning. *Psychological Review*, 124(3), 301.
- Bramley, N. R., Gerstenberg, T., Mayrhofer, R., & Lagnado, D. A. (2018). Time in causal structure learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 44(12), 1880.
- Buehner, M. J., & May, J. (2003). Rethinking temporal contiguity and the judgement of causality: Effects of prior knowledge, experience, and reinforcement procedure. *The Quarterly Journal of Experimental Psychology Section A*, 56(5), 865–890.
- Davis, Z. J., Bramley, N. R., & Rehder, B. (2018). Causal structure learning with continuous variables in continuous time. In *Proceedings of the 40th annual conference of the cognitive science society*. Madison, WI: Cognitive Science Society.
- Davis, Z. J., Bramley, N. R., & Rehder, B. (2020). Causal structure learning in continuous systems. *Frontiers in psychology*, 11, 244.
- Fernbach, P. M., & Erb, C. D. (2013). A quantitative causal model theory of conditional reasoning. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 39(5), 1327.
- Fernbach, P. M., & Sloman, S. A. (2009). Causal learning with local computations. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 35(3), 678.
- Greville, W. J., & Buehner, M. J. (2010). Temporal predictability facilitates causal learning. *Journal of Experimental Psychology: General*, 139(4), 756.
- Gureckis, T. M., Martin, J., McDonnell, J., Rich, A. S., Markant, D., Coenen, A., ... Chan, P. (2016). psiturk: An open-source framework for conducting replicable behavioral experiments online. *Behavior Research Methods*, 48(3), 829–842.
- Hagmayer, Y., & Waldmann, M. R. (2002). How temporal assumptions influence causal judgments. *Memory & Cognition*, 30(7), 1128–1137.
- Hayes, B. K., Hawkins, G. E., Newell, B. R., Pasqualino, M., & Rehder, B. (2014). The role of causal models in multiple judgments under uncertainty. *Cognition*, 133(3), 611–620.
- Hume, D. (1740). An enquiry concerning human understanding. In *Seven masterpieces of philosophy* (pp. 191–284). Routledge.
- Pacer, M. D., & Griffiths, T. L. (2012). Elements of a rational framework for continuous-time causal induction. In *Proceedings of the annual meeting of the cognitive science society* (Vol. 34).
- Pearl, J. (2009). *Causality*. Cambridge university press.
- Rottman, B. M., & Keil, F. C. (2012). Causal structure learning over time: observations and interventions. *Cognitive Psychology*, 64(1), 93–125.
- Sloman, S. A. (2005). *Causal models: How people think about the world and its alternatives*. Oxford University Press.
- Uhlenbeck, G. E., & Ornstein, L. S. (1930). On the theory of the brownian motion. *Physical Review*, 36(5), 823.
- Zacks, J. M. (2020). Event perception and memory. *Annual Review of Psychology*, 71, 165–191.