# Continuous-time Bayesian causal structure learning explains dopamine and behavioral anomalies in associative learning experiments

**John J. Vastola (john_vastola@hms.harvard.edu)**
Department of Neurobiology, Harvard Medical School, Boston, MA, USA

## Abstract

**Cause-effect learning is a core competency of sufficiently intelligent animals, and has been invoked to try to explain anomalous results (in both behavior and phasic dopamine activity) in certain associative learning experiments. But it is unclear how to mathematically formalize the problem of cause-effect learning, especially in a way that accommodates conditional independence structure, priors, and temporal structure (i.e., event order and proximity in time). We propose a novel Bayesian framework for modeling cause-effect learning which incorporates each of those aspects, yet remains relatively simple and has few free parameters. We study salient mathematical properties of our framework, including how inference is affected by topological structure in the assumed causal graph. Finally, we apply our framework to explain associative learning experiments, and find that it parsimoniously accounts for many otherwise puzzling observations. For example, our model explains the observation that cue-reward associations can be weakened by providing free reward at other times (contingency degradation), but only if the free reward is uncued. It also explains the observation that associations can be learned in fewer trials if each trial is longer. Our results suggest a new way to think about cause-effect learning, and support the idea that animals exploit nontrivial (causal) state representations even in simple associative learning settings.**

## Introduction

Cause-effect learning is a fundamental feature of intelligence (Hume, 1748; Gopnik & Schulz, 2007; Penn & Povinelli, 2007). It generically involves observing temporal sequences of events, and then making inferences about causal relationships based on the extent to which one event reliably predicts another, prior beliefs about which kinds of events are likely to be causally related, and the temporal ordering and proximity of events. Existing mathematical models of cause-effect learning are largely based on structural causal models (SCMs) (Pearl, 2009), and naturally capture conditional independence structure and the influence of prior beliefs (Griffiths & Tenenbaum, 2009), but generally fail to model temporal information, which is known to play a dominant role in human cause-effect learning (Lagnado & Sloman, 2006; Bramley, Gerstenberg, Mayrhofer, & Lagnado, 2018). Our first goal is to construct a theory of (Bayesian) cause-effect learning that addresses this issue. Our second goal is to leverage this account to study puzzling associative learning experiments in which cause-effect learning has been implicated (Jeong et al., 2022), with *contingency degradation* experiments—in which a cue-reward association is weakened by providing a free reward at another time, but only if that reward is uncued—being a paradigmatic example (Garr et al., 2024; Qian et al., 2024).

## Framework for Bayesian causal inference

We introduce a framework for modeling causal relationships in event sequences (Fig. 1a), motivated by three insights. First, if $A$ causes $B$, the continued presence of $A$ may not be necessary for $B$ to happen; cues can 'cause' reward even with cue-reward delays in trace conditioning experiments. Second, one cause often produces only one effect (e.g., one reward). Third, effects sometimes fail to follow their causes (e.g., probabilistic rewards). These considerations motivate modeling causality via latent (unobservable) 'causal power' variables, which are created and destroyed in discrete amounts. For example, one can imagine a cue producing one unit of causal power, which is then 'used up' to produce reward (Fig. 1b). Our full generative model allows events to occur spontaneously and/or as effects of some previous cause, and for effects to sometimes fail to follow their causes (Fig. 1c). We view learners as performing Bayesian inference over possible causal structures (Fig. 1d); a cue-reward association is considered learned if a 'cue causes reward' model is likelier than an 'independent' model.

## Explaining associative learning experiments

The number of trials required to make an association in simple cue-reward conditioning experiments depends mainly on the ratio of the cue-reward interval and intertrial interval (Gallistel & Gibbon, 2000), which is not consistent with naive temporal difference learning accounts (Gershman, 2024), but is a generic prediction of our theory (Fig. 2a-b). Our theory is consistent with the related observation (Burke et al., 2024) that animals learn in fewer trials if each trial is longer (Fig. 2c).

As usual, we identify dopamine with prediction errors, which the causal model affects since it defines 'state'. Our theory predicts that possible causes 'compete' to explain their effects. For example, if cue 1 and cue 2 are both present and cause reward, the likelihood accounts for both possibilities (Fig. 2d). This fact explains contingency degradation: the animal's best hypothesis is that reward can both spontaneously occur and be caused by the cue, which devalues the cue, since it may not be responsible for every reward that directly follows it (Fig. 2e). Cuing the second reward removes this ambiguity, and hence the effect (Fig. 2f). Our model shares features with previous attempts to explain contingency degradation, like invoking cause-effect learning (Jeong et al., 2022), and assuming animals exploit a nontrivial model of state transitions (Qian et al., 2024), but substantially generalizes these attempts and grounds them in normative (Bayesian) principles.

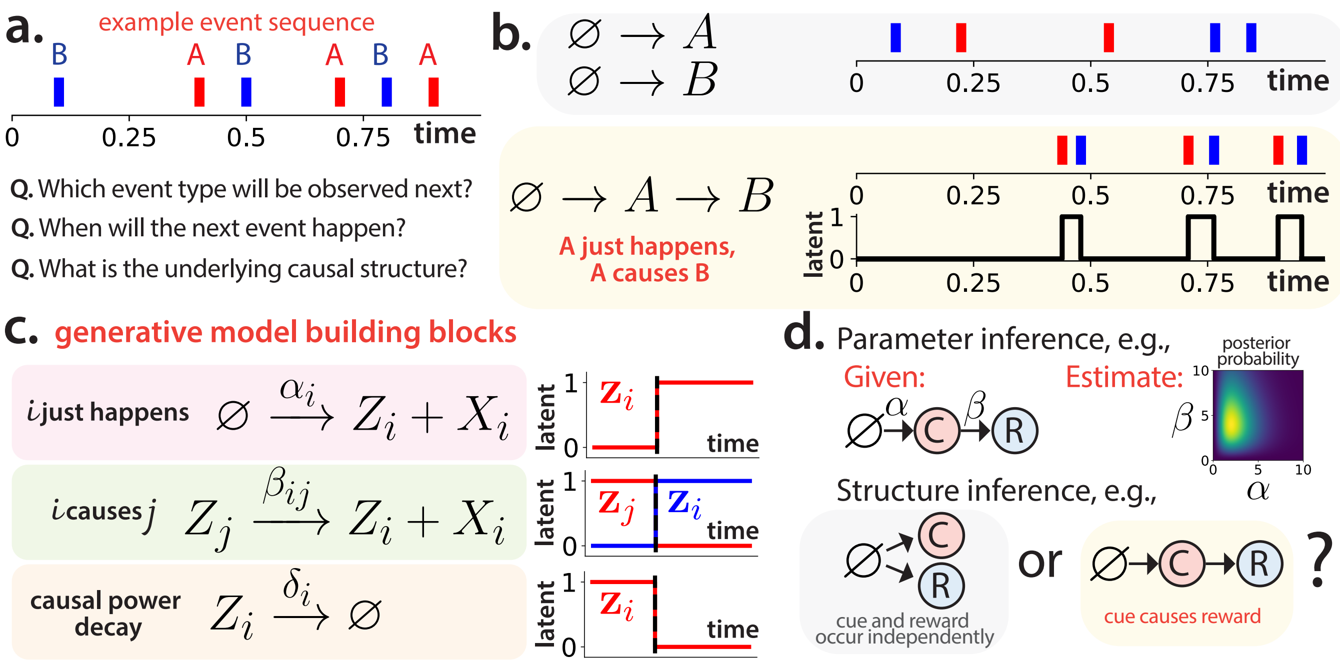


Figure 1: **Novel generative model for causality. a.** Example event sequence and relevant questions for learner. Events are characterized by a type (here: A or B) and time stamp. **b.** Samples from two generative models: one where A and B are independent, and one where A causes B. **c.** $X_i$ denotes event $i$ and $Z_i$ denotes the associated causal power variable. Our generative model allows three types of phenomena: events can occur spontaneously, events can be caused by other events, and causal power can decay. **d.** Two types of inference: parameter and structure inference. The former is required for the latter.

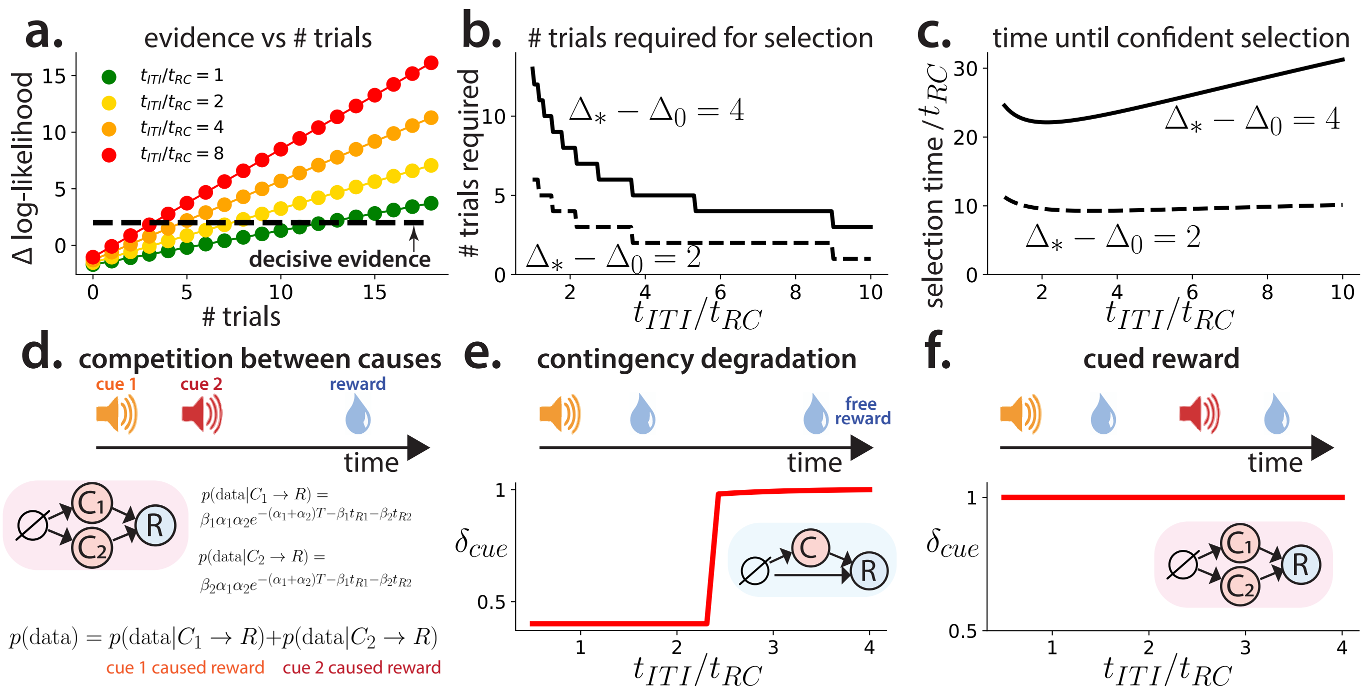


Figure 2: **Explaining anomalous associative learning experiments.** $t_{RC}$: cue-reward delay. $t_{ITI}$: intertrial interval. **a.** Evidence (causal vs independent model log-likelihood difference $\Delta$) in a cue-reward conditioning experiment as a function of # trials. **b.** Number of trials required for model selection given two different priors $\Delta_0$. **c.** Elapsed time until selection. **d.** When two cues can cause reward and both are present, likelihood involves sum of terms representing each possibility. **e.** If $t_{ITI}$ is sufficiently short, free rewards greatly devalue a previously learned cue. **f.** Without causal ambiguity, there is no contingency degradation effect.

## References


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), 1880–1910. (Place: US Publisher: American Psychological Association) doi: 10.1037/xlm0000548

Burke, D. A., Taylor, A., Jeong, H., Lee, S., Wu, B., Floeder, J. R., & K Namboodiri, V. M. (2024). Reward timescale controls the rate of behavioral and dopaminergic learning. *bioRxiv*. doi: 10.1101/2023.03.31.535173

Gallistel, C. R., & Gibbon, J. (2000). Time, rate, and conditioning. *Psychological Review*, *107*(2), 289–344. (Place: US Publisher: American Psychological Association) doi: 10.1037/0033-295X.107.2.289

Garr, E., Cheng, Y., Jeong, H., Brooke, S., Castell, L., Bal, A., ... Janak, P. H. (2024). Mesostriatal dopamine is sensitive to changes in specific cue-reward contingencies. *Science Advances*, *10*(22), eadn4203. doi: 10.1126/sciadv.adn4203

Gershman, S. J. (2024, Jul). *Rate estimation revisited.* PsyArXiv. doi: 10.31234/osf.io/47q5y

Gopnik, A., & Schulz, L. (2007). *Causal Learning: Psychology, Philosophy, and Computation*. Oxford University Press. doi: 10.1093/acprof:oso/9780195176803.001.0001

Griffiths, T. L., & Tenenbaum, J. B. (2009). Theory-based causal induction. *Psychological Review*, *116*(4), 661–716. (Place: US Publisher: American Psychological Association) doi: 10.1037/a0017201

Hume, D. (1748). *An enquiry concerning human understanding.* London, England: Andrew Millar.

Jeong, H., Taylor, A., Floeder, J. R., Lohmann, M., Mihalas, S., Wu, B., ... Namboodiri, V. M. K. (2022). Mesolimbic dopamine release conveys causal associations. *Science*, *378*(6626), eabq6740. doi: 10.1126/science.abq6740

Lagnado, D. A., & Sloman, S. A. (2006). Time as a guide to cause. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *32*(3), 451–460. (Place: US Publisher: American Psychological Association) doi: 10.1037/0278-7393.32.3.451

Pearl, J. (2009). *Causality*. Cambridge University Press.

Penn, D. C., & Povinelli, D. J. (2007). Causal cognition in human and nonhuman animals: A comparative, critical review [Journal Article]. *Annual Review of Psychology*, *58*(Volume 58, 2007), 97-118. doi: https://doi.org/10.1146/annurev.psych.58.110405.085555

Qian, L., Burrell, M., Hennig, J. A., Matias, S., Murthy, V. N., Gershman, S. J., & Uchida, N. (2024). The role of prospective contingency in the control of behavior and dopamine signals during associative learning. *bioRxiv*. doi: 10.1101/2024.02.05.578961