Less is more: Adaptive strategies in continuous time causal learning

Victor J. Btesh, Neil R. Bramley, Maarten Speekenbrink and David A. Lagnado

Presenter: Hanqi Zhou
Date: 2024.10.29

Background

Causal learning from observational data

Association/correlation

- Hume (1777/1975) argued that causal relations are not directly observable and therefore must be inferred on the basis of observable cues.
  - E.g., contiguity, temporal order -> *co-occurrence* of cause and effect

Causal learning from observational data
Association/correlation

- Hume (1777/1975) argued that causal relations are not directly observable and therefore must be inferred on the basis of observable cues.
  - E.g., contiguity, temporal order -> *co-occurrence* of cause and effect
- ΔP model (Jenkins & Ward, 1965)
  - Contingency: $\Delta P = P(E \mid C) - P(E \mid \neg C)$

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

Causal learning from observational data

Association/correlation

- Hume (1777/1975) argued that causal relations are not directly observable and therefore must be inferred on the basis of observable cues.
    - E.g., contiguity, temporal order -> *co-occurrence* of cause and effect
- ΔP model (Jenkins & Ward, 1965)
    - Contingency: $\Delta P = P(E \mid C) - P(E \mid \neg C)$

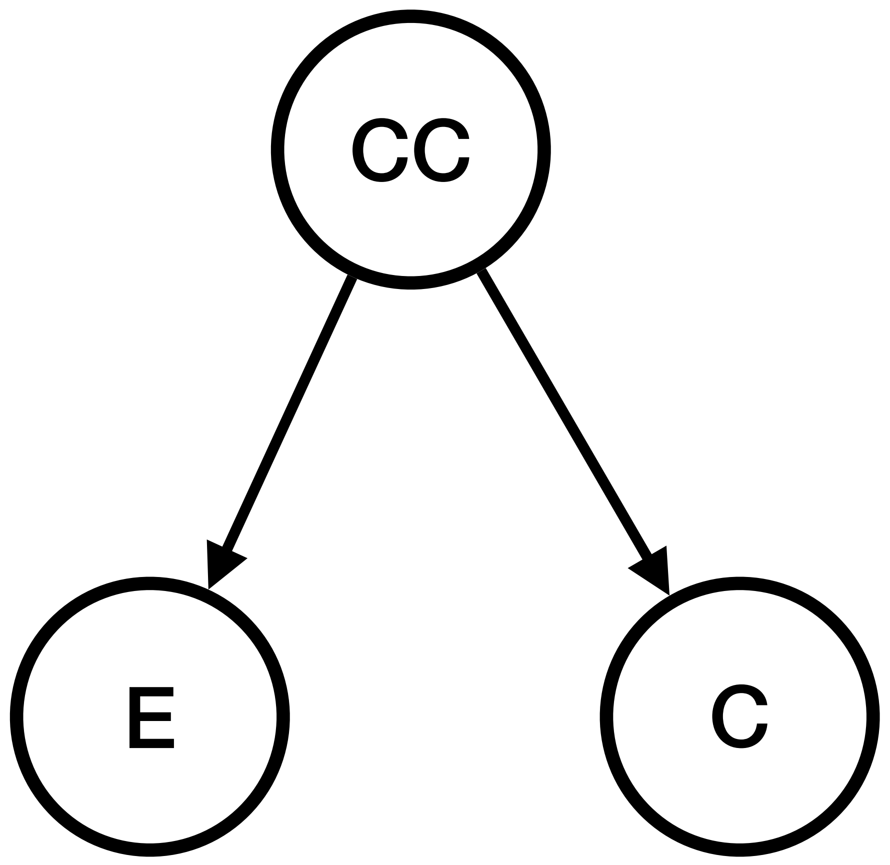

Causal learning from observational data

Causal strength

- Propose causal representations as opposed to associative ones.
- Power PC theory (short for a causal power theory of the probabilistic contrast model; Cheng, 1997)
- Consider background cause: $P = \dfrac{P(E \mid C) - P(E \mid \neg C)}{1 - P(E \mid \neg C)}$

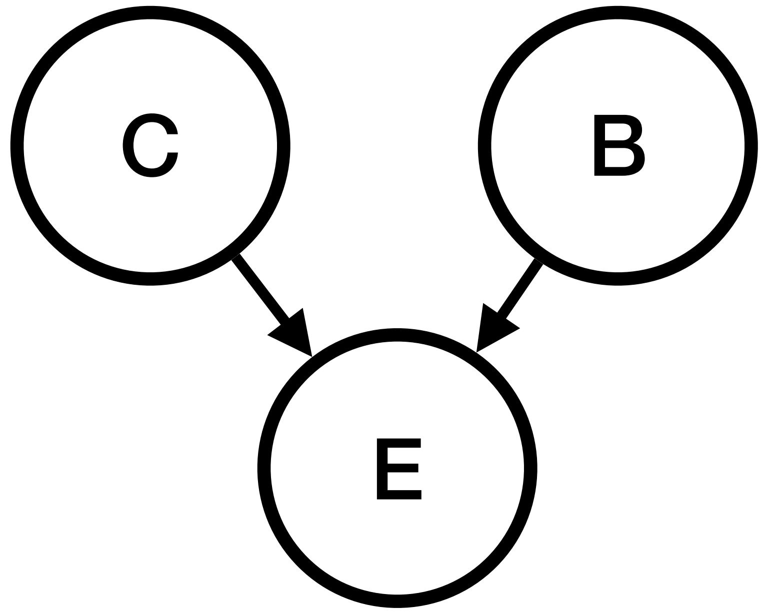

Cheng, P. W. (1997). From covariation to causation: A causal power theory. *Psychological review*, *104*(2), 367.

Causal learning from observational data

Causal strength & structure

- Representation: causal Bayesian network (Pearl, 2000)
  - Probabilistic graphical model that represents a set of variables and their conditional dependencies via a directed acyclic graph (DAG)
  - A causal network is a Bayesian network with the requirement that the relationships be causal.
- Goal:
  - Structure learning & parameter estimation

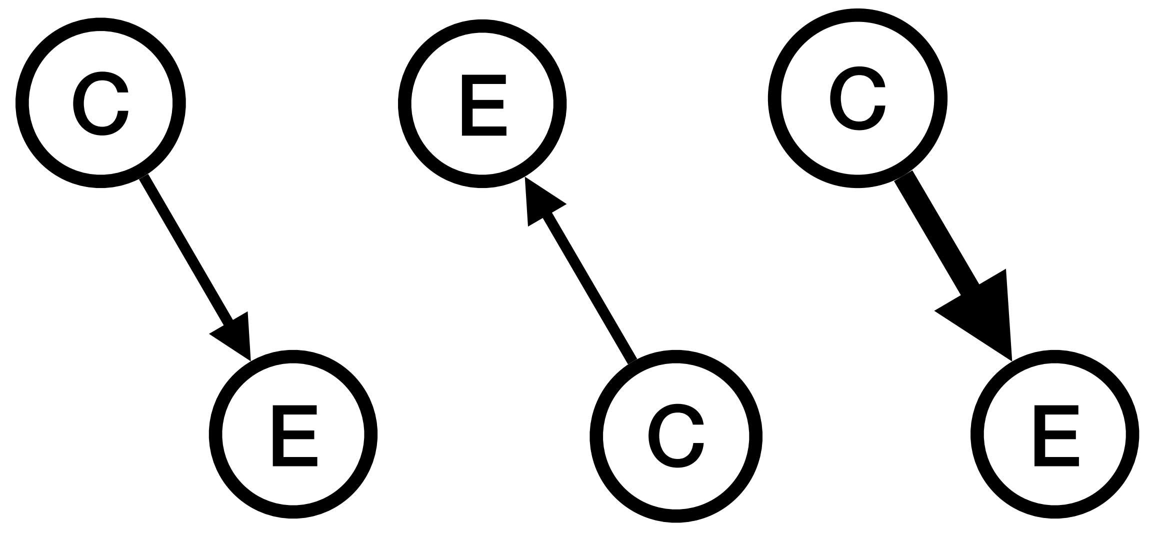

Causal learning from observational data

Causal strength & structure

- Causal support (Griffiths & Tenenbaum, 2005)
  - Evidence the data provide for the existence of a causal structure
  - Support $= \log \dfrac{P(D \mid G_1)}{P(D \mid G_0)}$

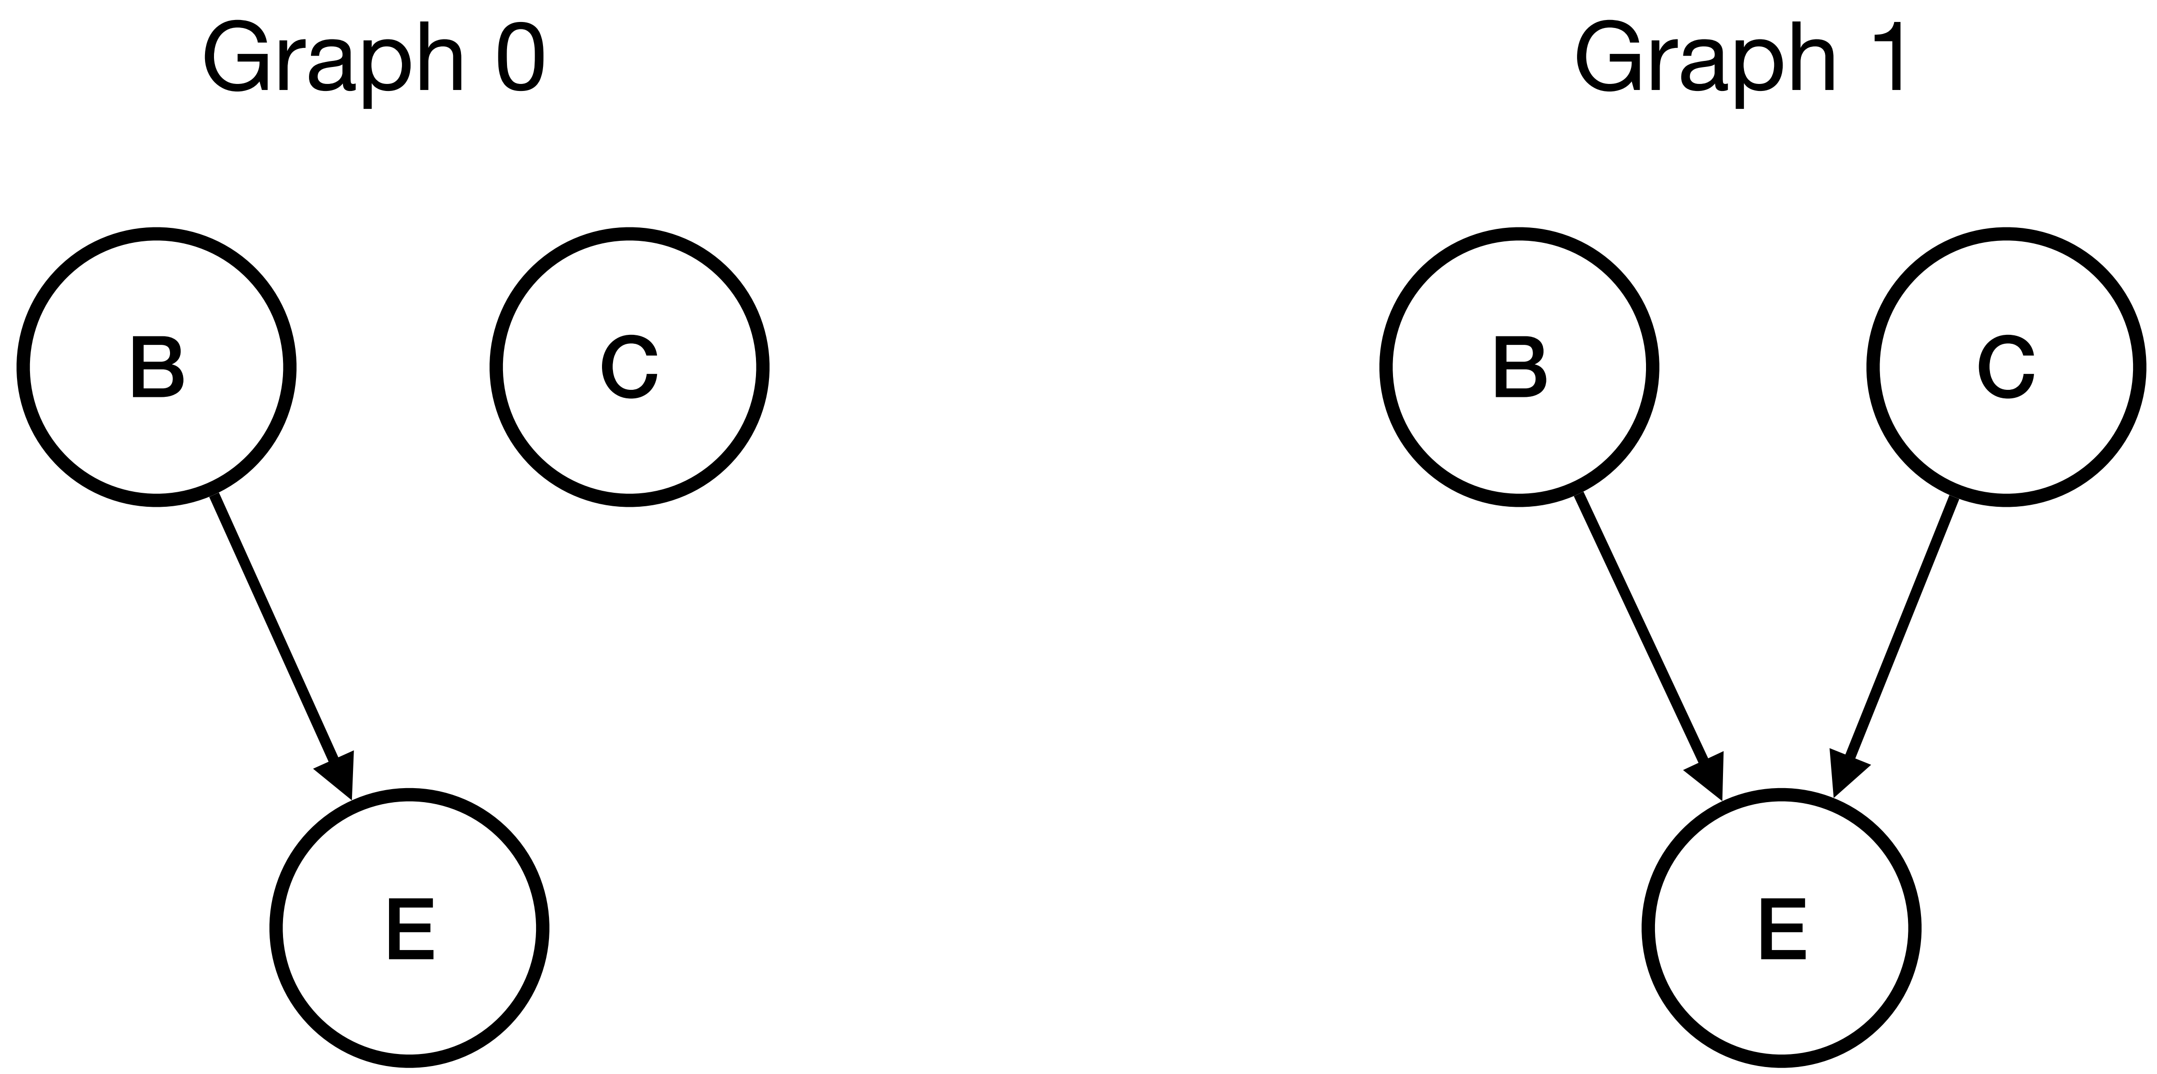

Griffiths, T. L., & Tenenbaum, J. B. (2005). Structure and strength in causal induction. *Cognitive psychology*, *51*(4), 334-384.

Causal learning from observational data

Causal strength & structure

- Complexity of structure learning
  - Two hypotheses: Support $= \log \frac{P(D \mid G_1)}{P(D \mid G_0)}$
  - Full Bayesian model: $P\left(G_i \mid D\right) = \frac{P\left(D \mid G_i\right) P\left(G_i\right)}{\sum_{G \in \mathcal{G}} P(D \mid G) P(G)}$

Causal learning from observational data

Causal strength & structure

- Complexity of structure learning
- Full Bayesian model $P\left(G_i \mid D\right) = \dfrac{P\left(D \mid G_i\right) P\left(G_i\right)}{\sum_{G \in \mathcal{G}} P(D \mid G) P(G)}$

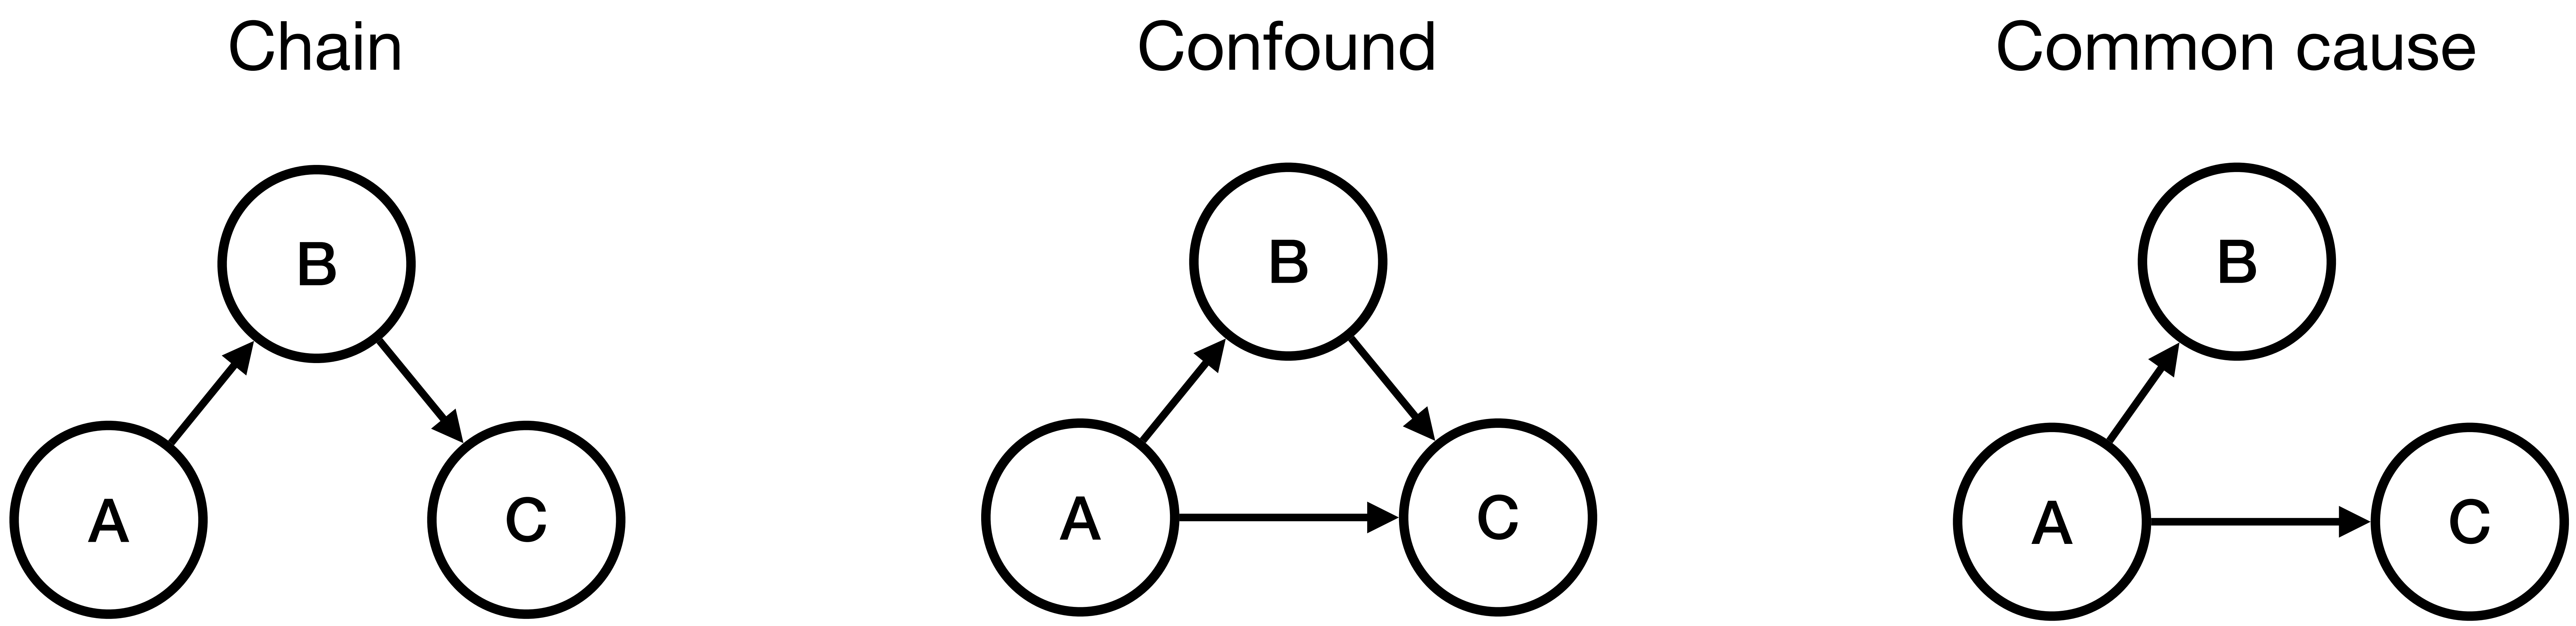

Causal learning from observational data

Causal strength & structure

- Complexity of structure learning
- Full Bayesian model $P\left(G_i \mid D\right) = \dfrac{P\left(D \mid G_i\right) P\left(G_i\right)}{\sum_{G \in \mathcal{G}} P(D \mid G)P(G)}$

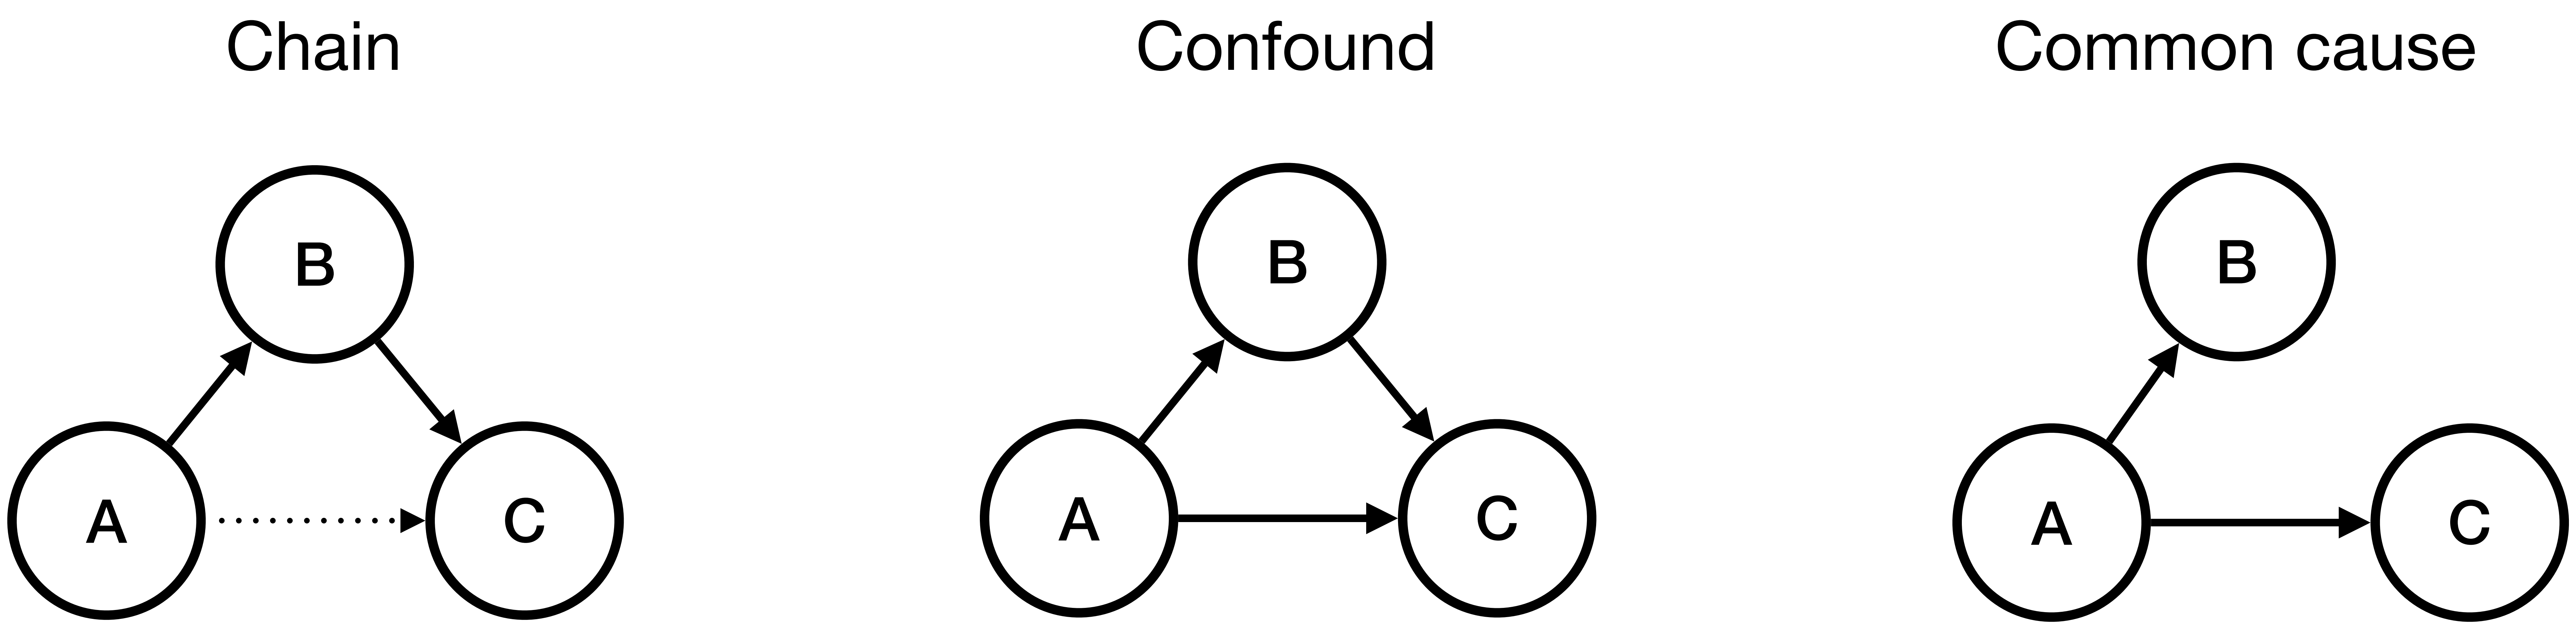

Causal learning from heuristics

Descriptive model of model updating

- Single-effect (Waldmann, 2008) & local computations (Fernbach & Sloman, 2009)
  - $D^t \rightarrow G^t_{\text{update}}$
  - $G^{t+1} = G^t \cup G^t_{\text{update}}$
- Local updating (Neurath ship; Bramley et al., 2017)
  - Local modification: $\{G_i^{t+1}\} \in \text{edit}(G^t)$
  - Inference: $p(G_i^{t+1} \mid D^t) \propto \text{dist}(G_i^{t+1}, G^t) + p(D^t \mid G_i^{t+1})$

Waldmann, M. R., Cheng, P. W., Hagmayer, Y., & Blaisdell, A. P. (2008). Causal learning in rats and humans: A minimal rational model. *The probabilistic mind. Prospects for Bayesian cognitive science*, 453-484.

Fernbach, P. M., & Sloman, S. A. (2009). Causal learning with local computations. *Journal of experimental psychology: Learning, memory, and cognition*, 35(3), 678.

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.

Causal learning from intervention

Motivation

- Theory
    - Intervention has been deemed the hallmark of causality (e.g., Pearl, 2000; Woodward, 2003) in the sense that A causes B if and only if a sufficient intervention to change the state of A would also change B.

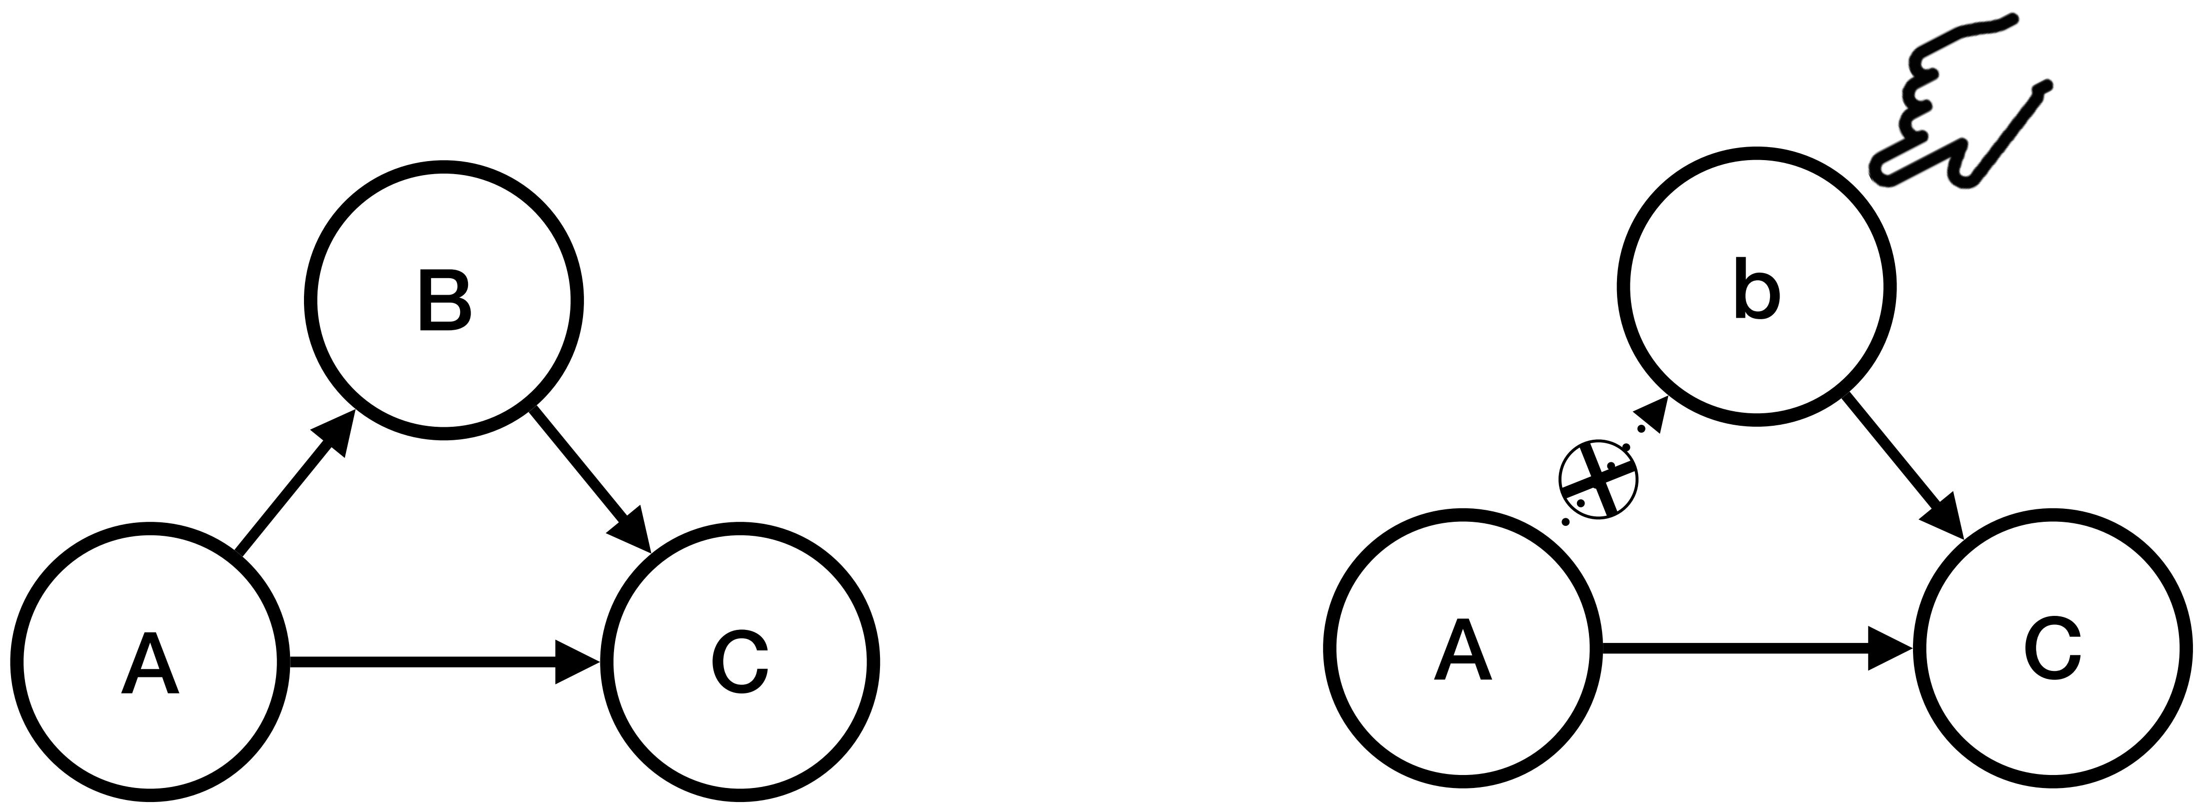

Causal learning from intervention
Motivation

- Empirical evidence
  - A common finding is that given just observational data, participants’ inferences are strikingly suboptimal.
    - Lagnado and Sloman’s (2004) Experiment 1, three variables, observing 50 trials of how events covaried -> 5 out of 36 participants identified the correct model
  - Accuracy improvement before and after intervention (Steyvers et al., 2010)

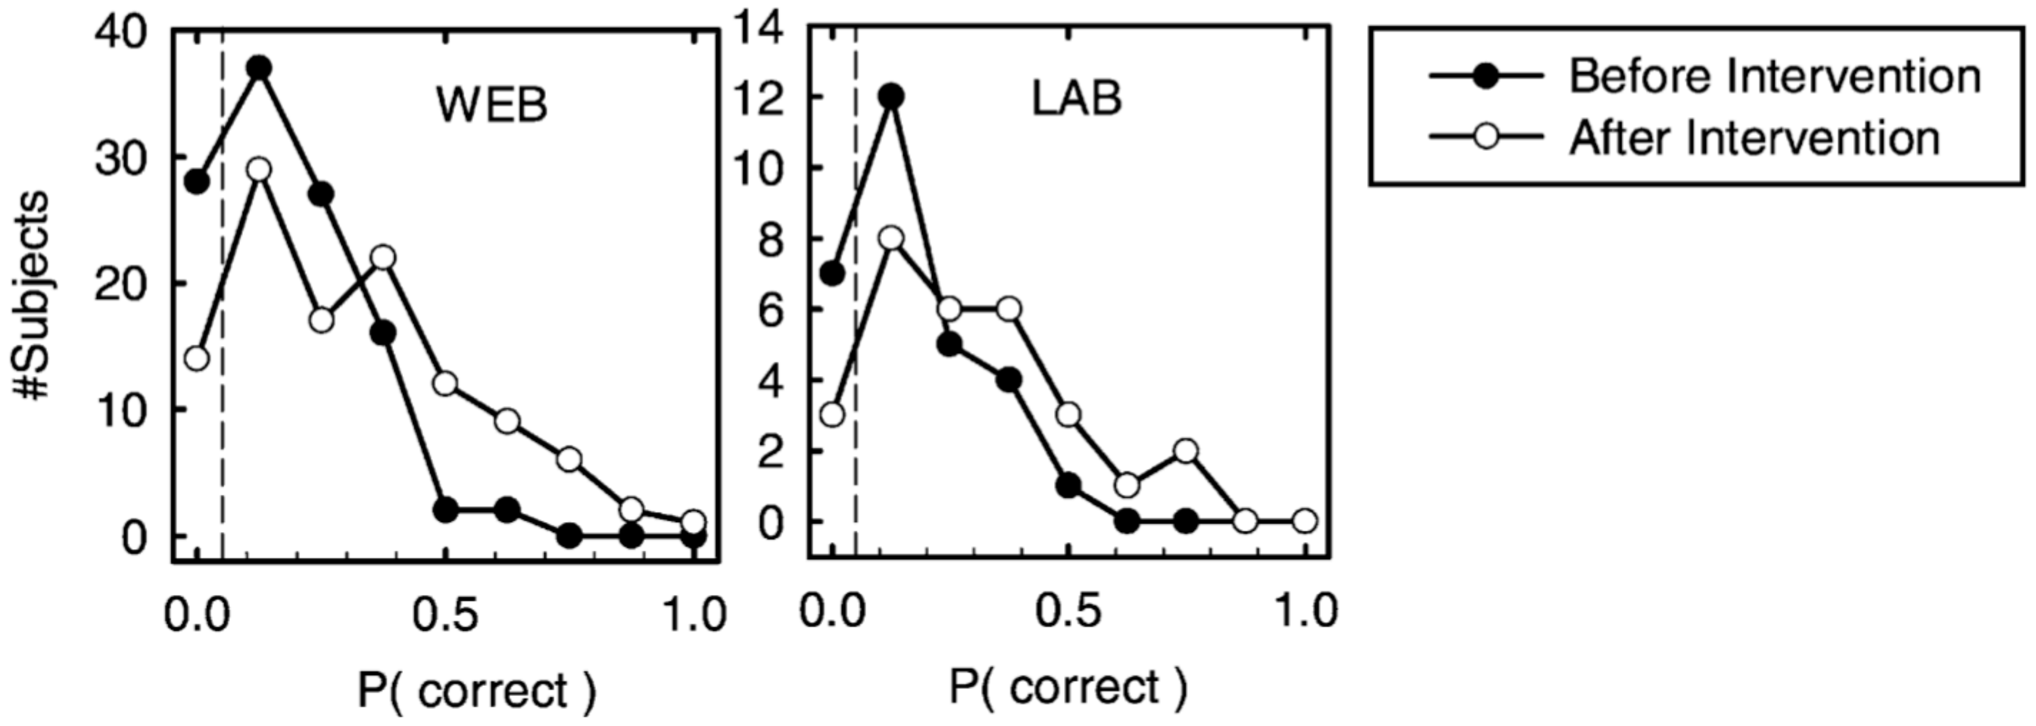

Lagnado, D. A., & Sloman, S. (2004). The advantage of timely intervention. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *30*(4), 856.
Steyvers, M., Tenenbaum, J. B., Wagenmakers, E. J., & Blum, B. (2003). Inferring causal networks from observations and interventions. *Cognitive science*, *27*(3), 453-489.

Causal learning from *active* intervention

Descriptive model

- Complexity of the structure
  - Heuristics of updating the model
- Complexity of the space of intervention choices (active causal learning)
  - Heuristics of choosing interventions

Causal learning from *active* intervention

Descriptive model of intervention strategies

- Conservative & Forgetful (Bramley et al., 2014)
  - Prioritize interventions that maximize *information gain*, showing a preference for reducing overall uncertainty - $H^t(G) - H^t(G \mid D = \{\text{obs, interv}\})$
  - Forgetful but conservative
- Adaptive (Coenen et al., 2014)
  - Information gain + confirmation

Bramley, N. R., Lagnado, D. A., & Speekenbrink, M. (2015). Conservative forgetful scholars: How people learn causal structure through sequences of interventions. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *41*(3), 708.

Coenen, A., Rehder, B., & Gureckis, T. M. (2015). Strategies to intervene on causal systems are adaptively selected. *Cognitive psychology*, *79*, 102-133.

Causal learning from active intervention in continuous time

Motivation

- Choosing *where* to intervene -> choosing *when* and *where* to intervene

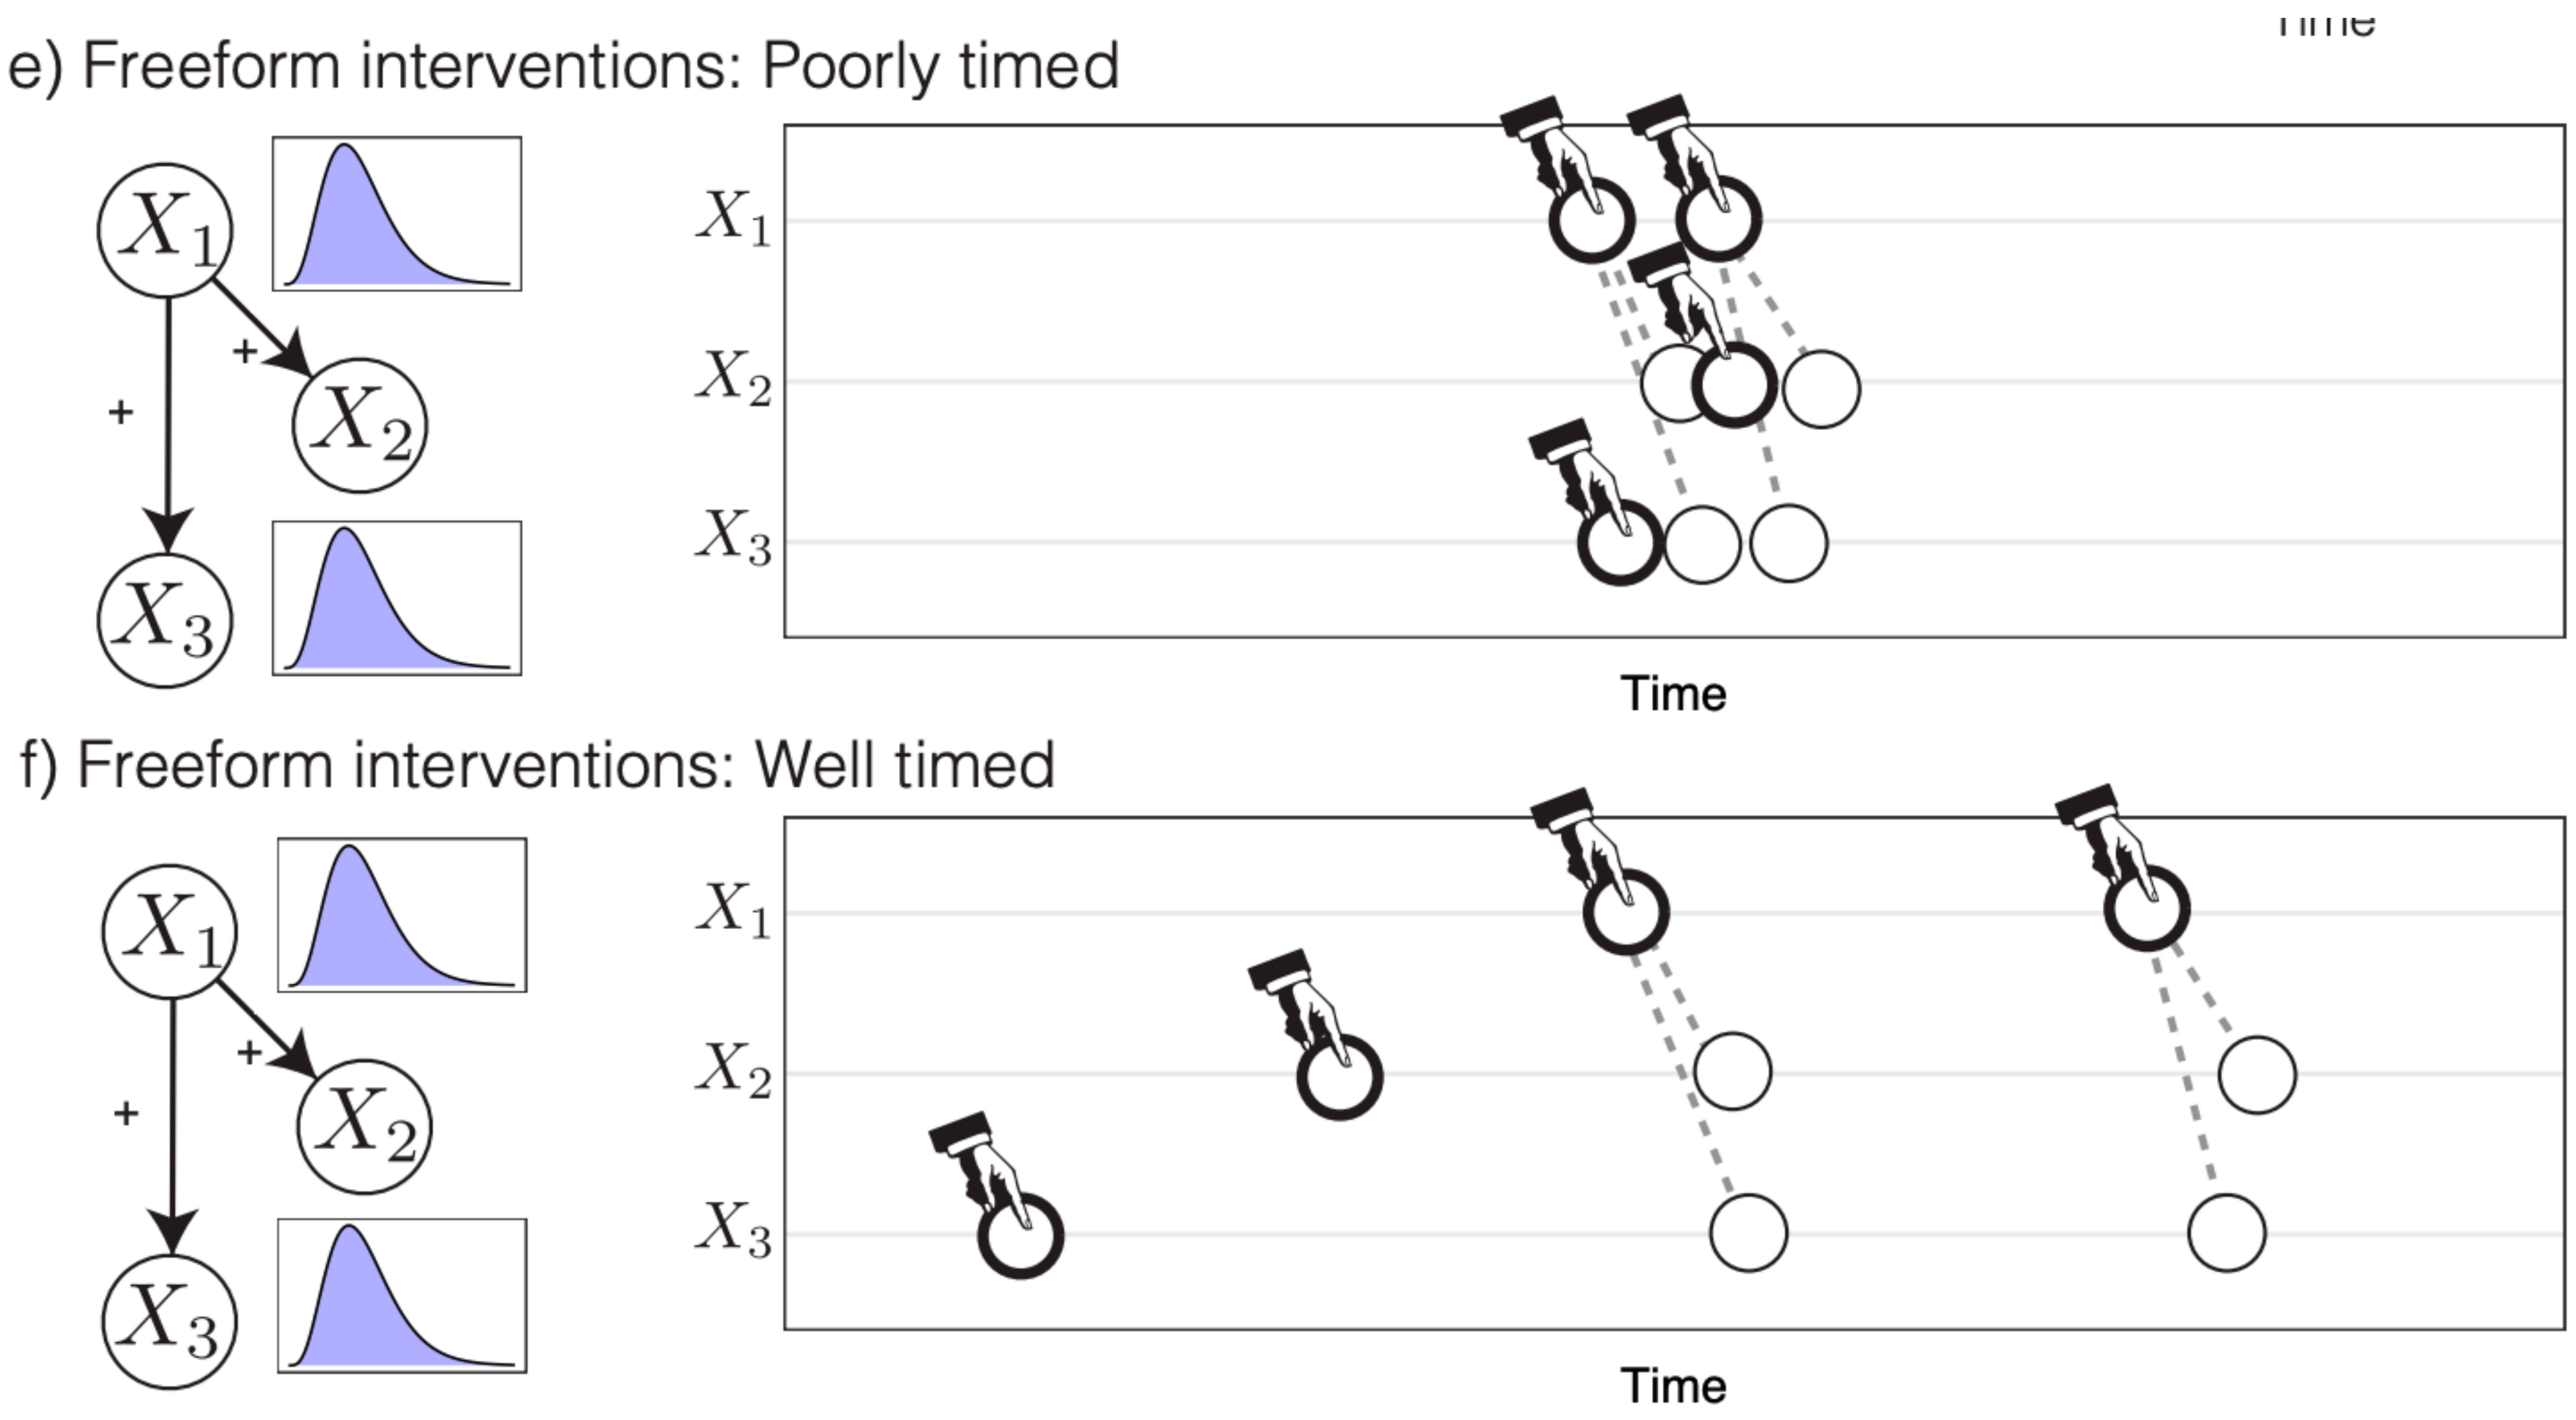

Bramley, N. R., Gerstenberg, T., Mayrhofer, R., & Lagnado, D. A. (2019). Intervening in time. *Time and causality across the sciences*, 86-115.

Motivation

Position of this work

- Descriptive model of causal learning (heuristics)
    - model updating and intervention selection
- Active, continuous time

Motivation of this work

- Existing works: algorithmic complexity of learning
    - It focused on identifying the key parameters of inference models which best match participants rather than on its interaction with prior knowledge
- This work: the integration of new data given pre-existing beliefs

Main

- In the present study, we investigate the inference mechanisms that people use to learn complex causal structures in continuous time, and how this learning is influenced by prior beliefs.

Hypothesis

- We expect participants' judgements to be best captured by computational models that selectively consider data during or *at least around intervening times* and which focus *on the direct links from the intervened upon node*.

- We expect participants to be less accurate in domain-specific contexts when the *ground truth graph diverges from their prior judgements*. We propose two alternative but not exclusive mechanisms by which such a result may occur:

    1. Direct single dependence: participants' posterior beliefs are influenced by their prior beliefs solely by biasing the *inference process*

    2. Direct and action-mediated dependence: participants' posterior beliefs are influenced by their prior beliefs by affecting the entire *active learning process (inference and interventions)*.

Experiment and behavior results

Paradigm

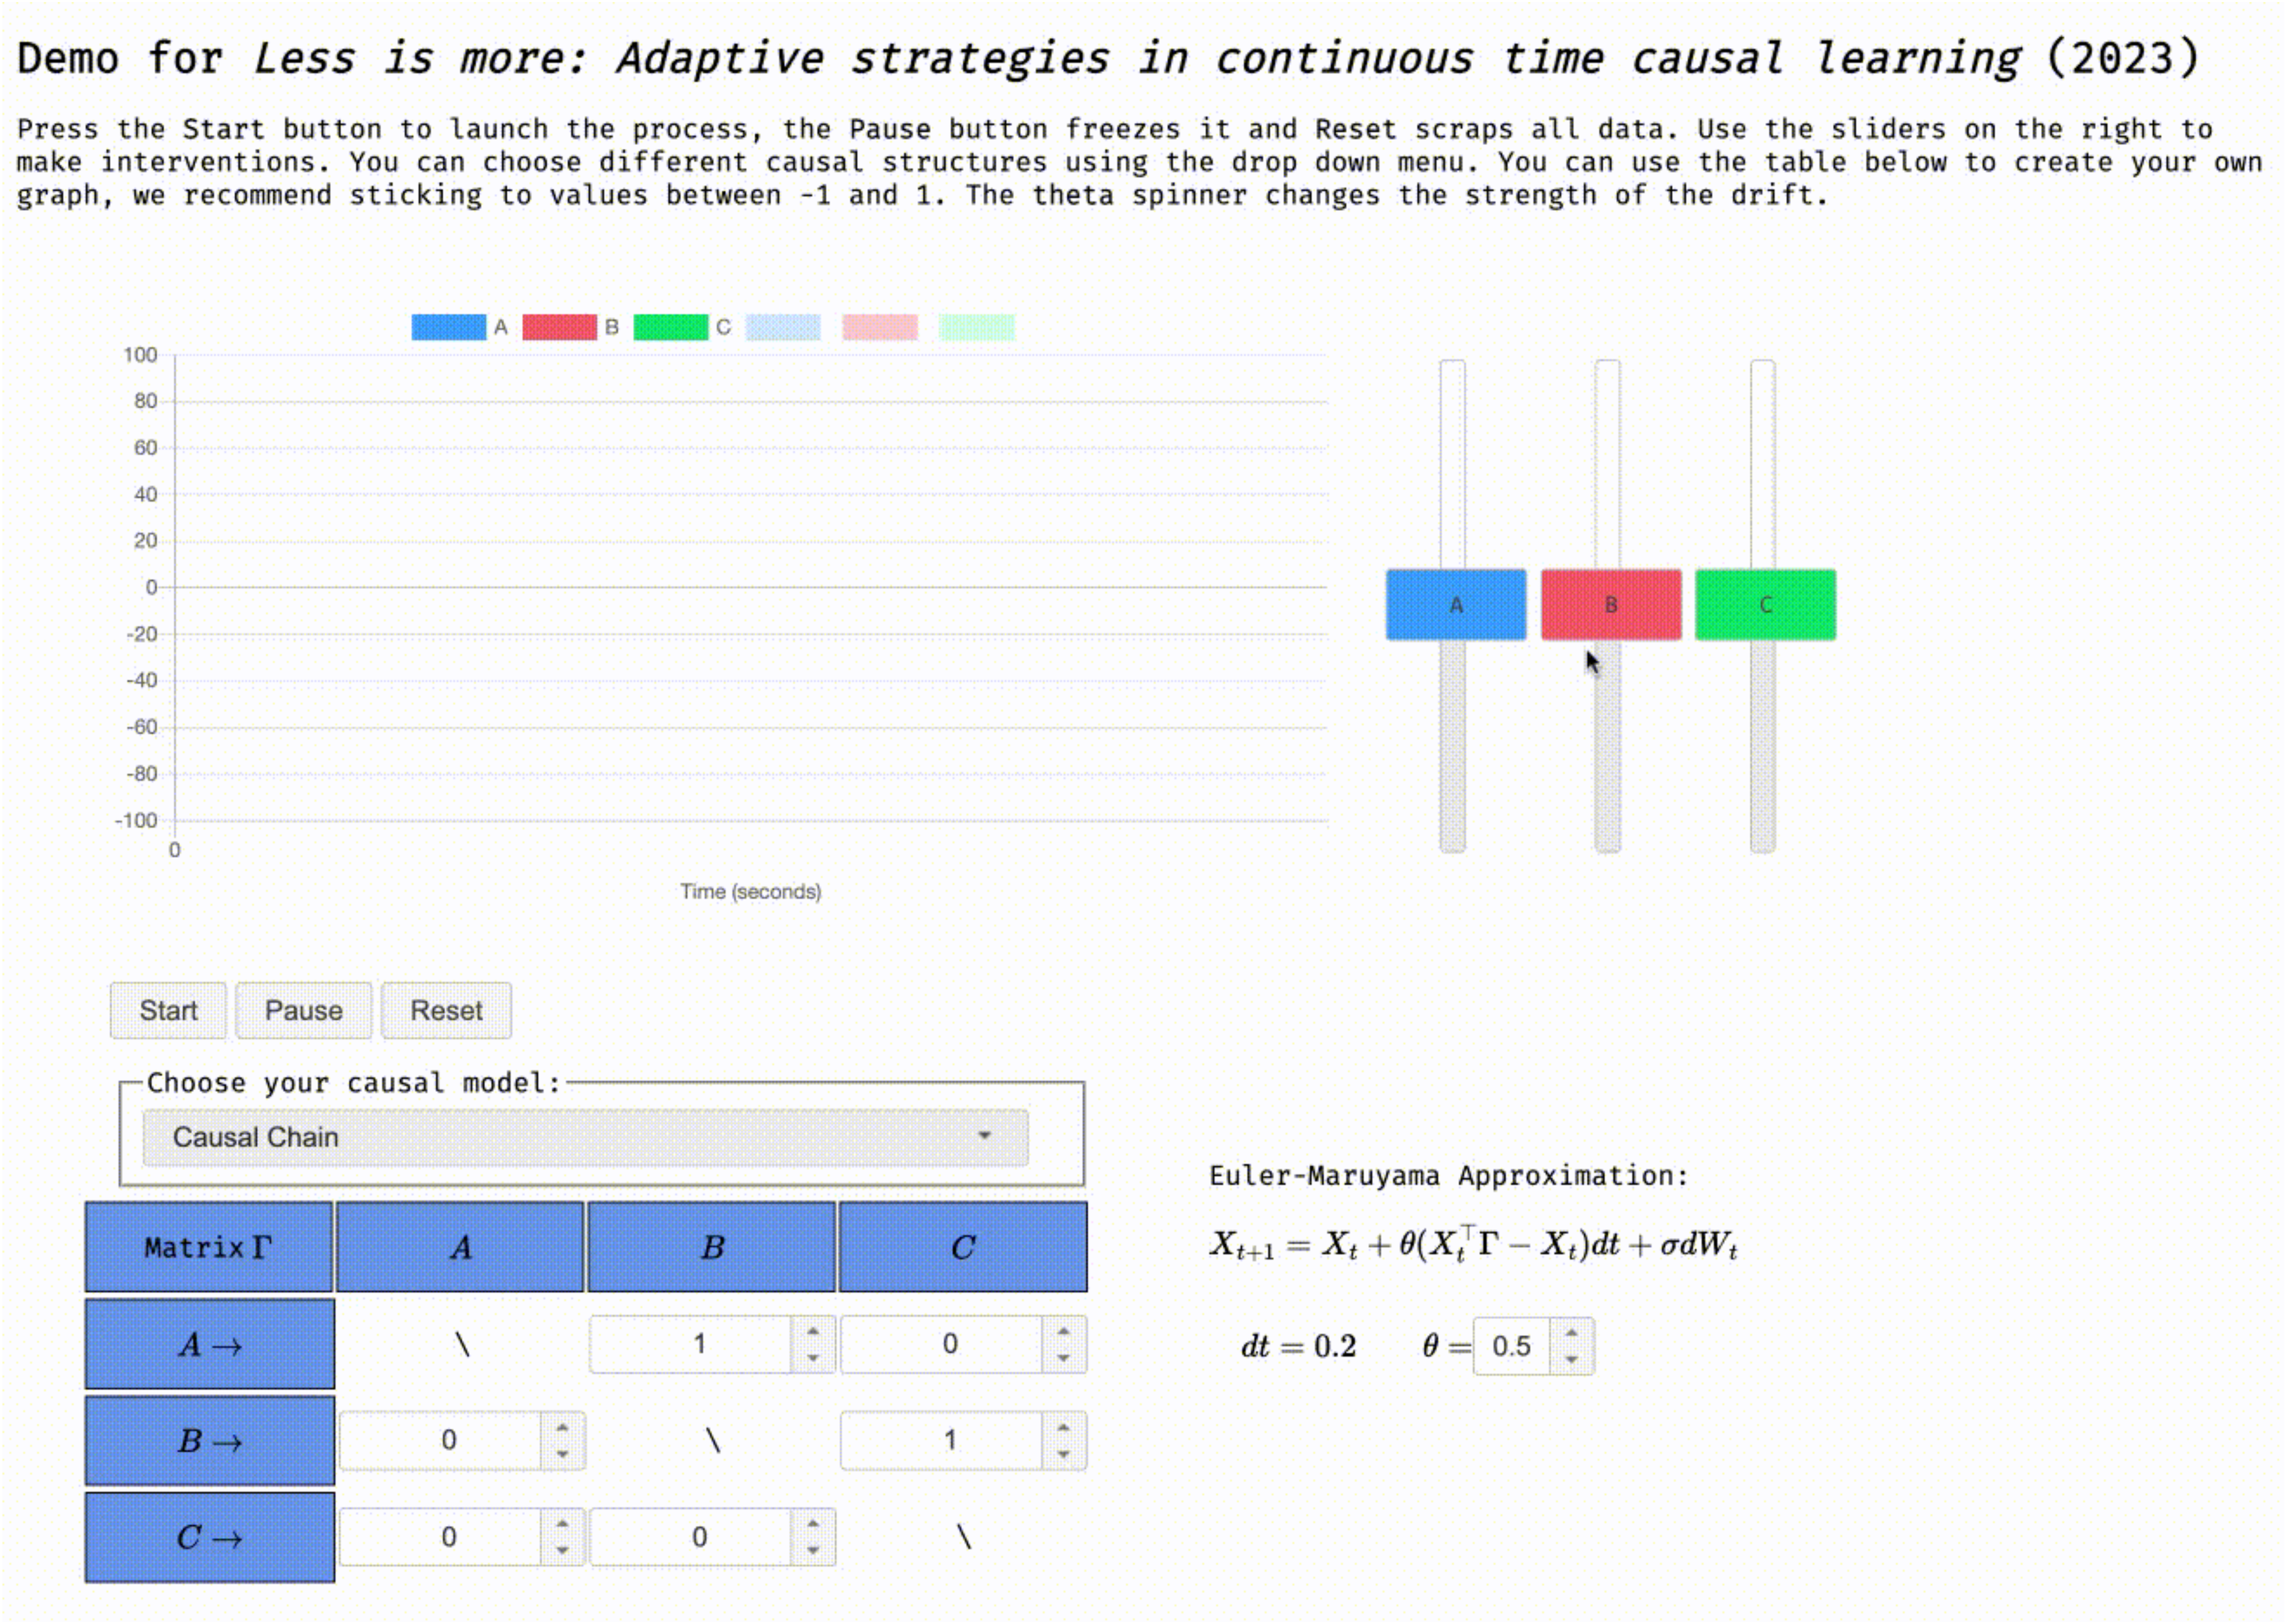

- 3 sliders represent 3 variables
- The values of the variables are computed according to Ornstein-Uhlenbeck processes

$$p\left(x_{t+dt} \mid x_t, y_t, \theta, \sigma, dt\right)$$

$$= \mathcal{N}\left(x_t + \theta\left(\mu_x\left(x_t, y_t\right) - x_t\right)dt, \sigma dt\right)$$

$$\mu_x\left(x_t, y_t\right) = \beta\left(x_t\right) + \gamma_{xx}x_t + \gamma_{yx}y_t$$

Paradigm

Reminder of the rules:

- ***0***: *no arrow from variable A to variable B means no effect of variable A on variable B*
- ***1***: *one positive arrow from A to B means that **MORE of A** results in **MORE of B***
- ***2***: *two positive arrows from A to B means that **MORE of A** results in **a lot MORE of B***
- ***-1***: *one negative arrow from A to B means that **MORE of A** results in **LESS of B***
- ***-2***: *two negative arrows from A to B means that **MORE of A** results in **a lot LESS of B***

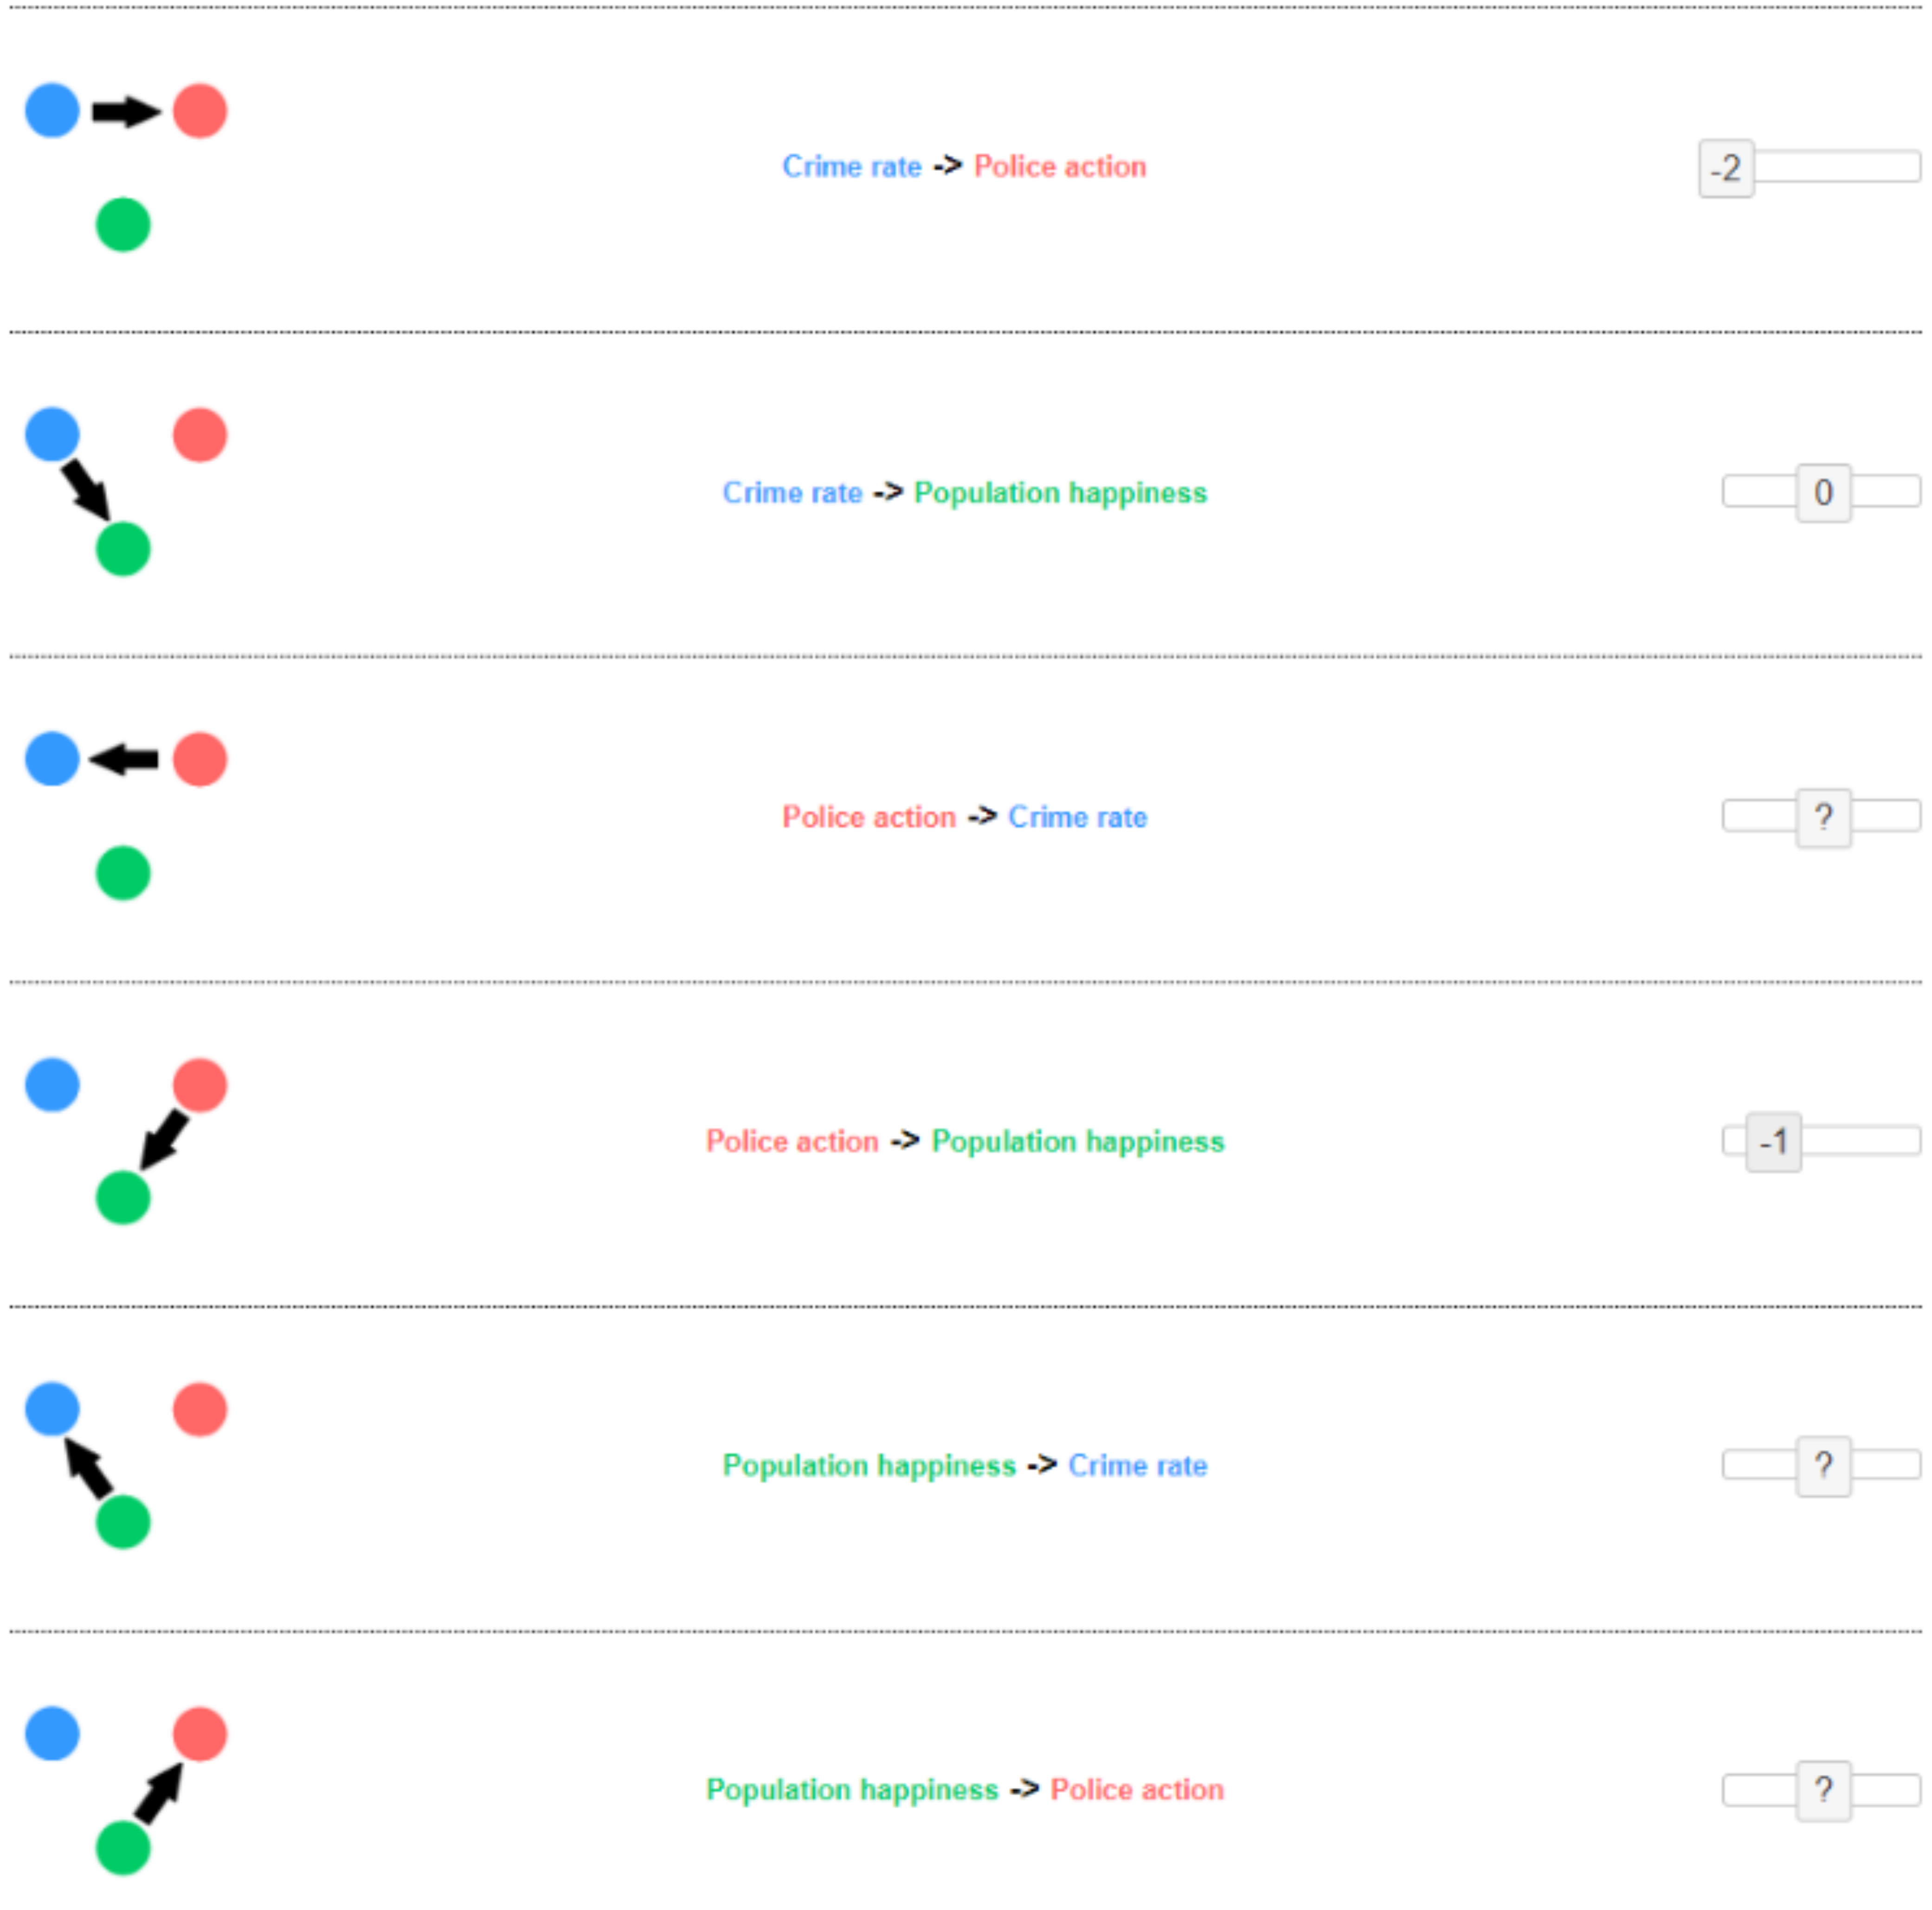

Experiment 1

Impact of semantic labels

- Test whether participants would perform differently if the variables were
  - labelled meaningfully, i.e. crime rate, police action and population happiness,
  - or generically, i.e. blue, red and green.
- See if the agreement between elicited prior and ground truth graphs would predict participants' ability to recover the correct structure.

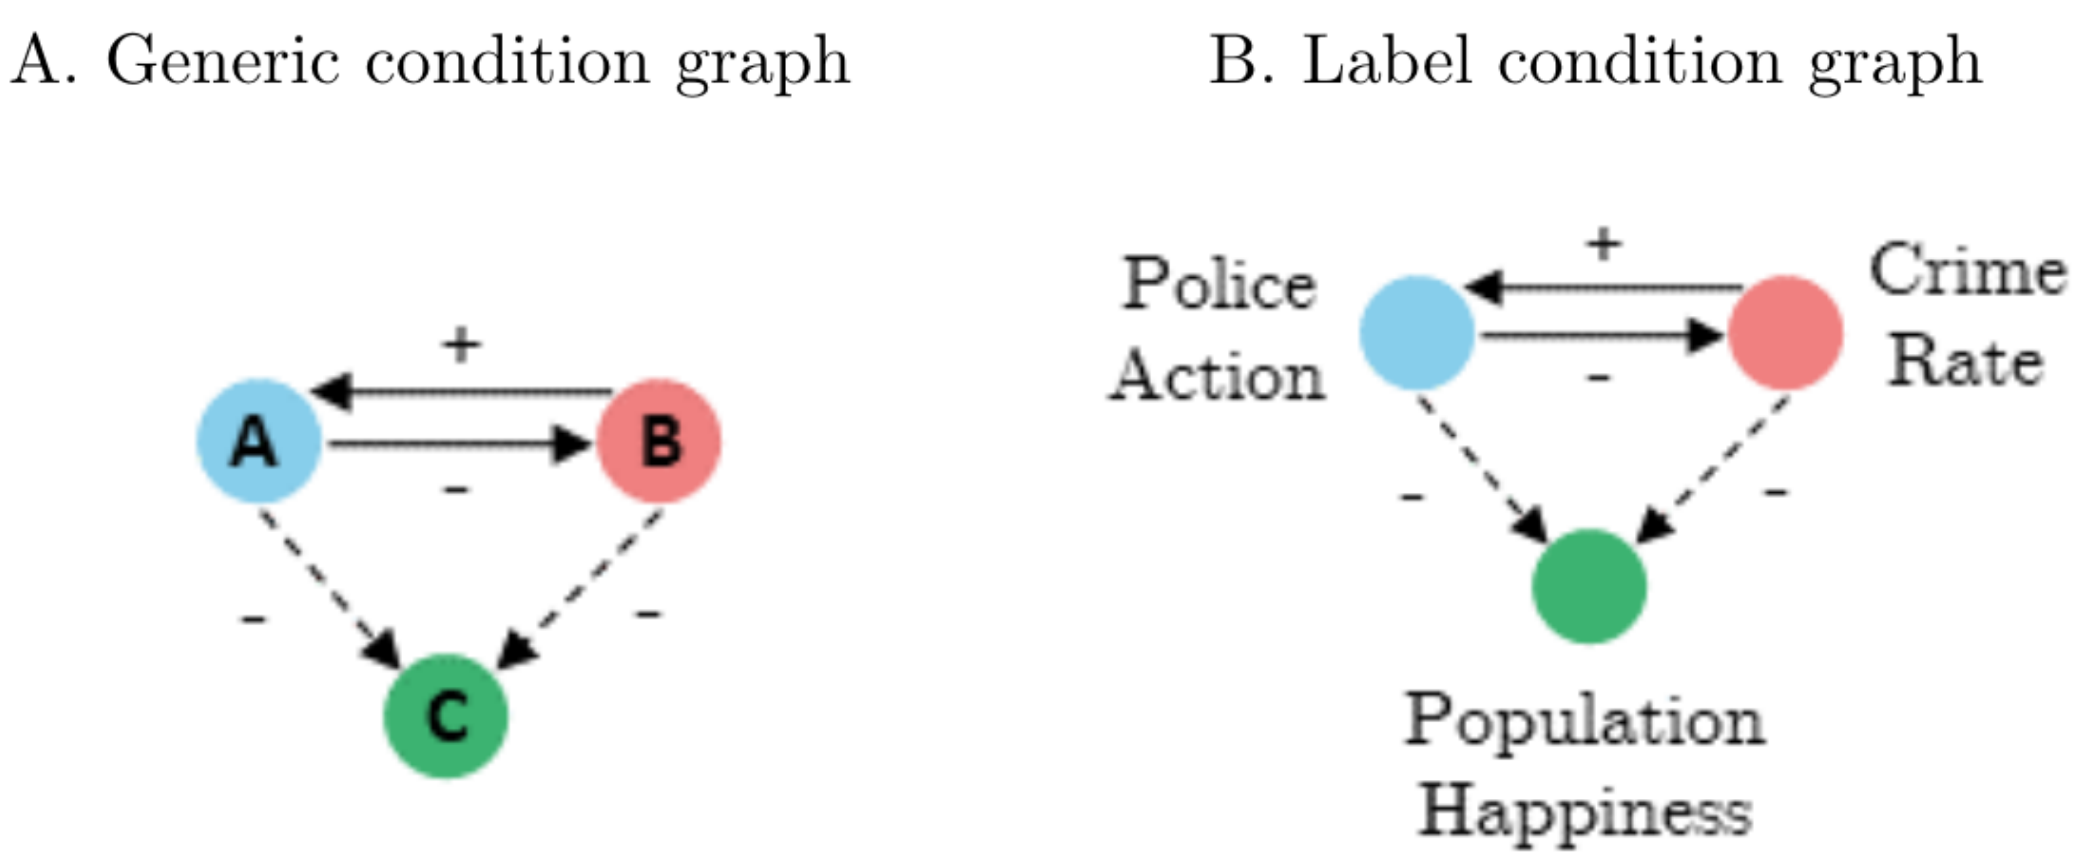

Experiment 1

Impact of semantic labels

- Participants were randomly allocated to a *generic* or a *label* condition.
    - Label: crime rate, police action and population happiness
- 3 generic causal learning trials + 1 controlled trial.

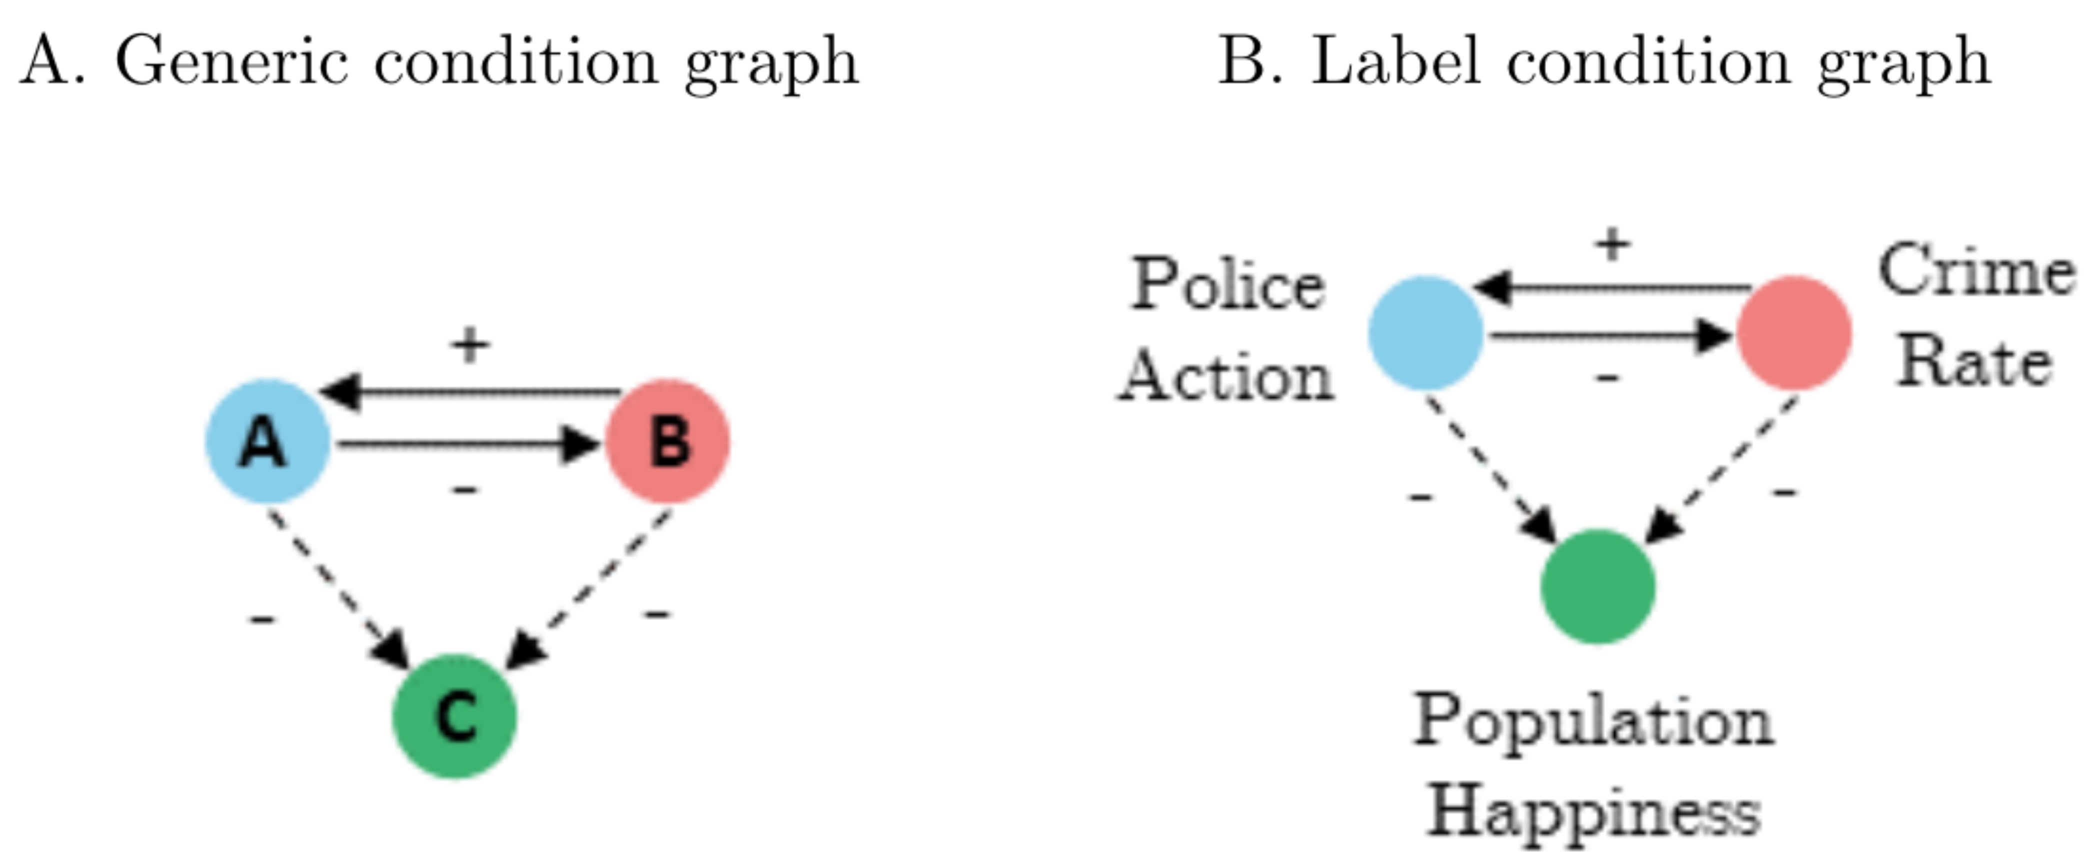

Experiment 1 results

Impact of semantic labels and prior

- **Correct Link Recovery**: accurately reported correct causal links 63% of the time, surpassing the chance level (33%)
    - **Difficulty with Link Strengths:** recovered weak (40%) vs. strong links (50%)
    - **Difficulty with Negative Links:** recovered positive links (87%) vs. negative links (73%)
- **Label Influence:** Labeling *did not* significantly impact overall accuracy, with meaningful (67%) vs. generic labels (66%)
- **Prior Influence:** Participants whose prior were closer to the ground truth were on average more accurate
- **Intervention Usage:** The average number of interventions was 7.18/trial with 3.38 (±3.96) seconds
- **Structure Complexity:**
    - Participants performed better on simpler graphs (e.g., collider and common cause structures) with accuracies of 78% and 76%, respectively, versus 68% for complex chain graphs.
    - In complex chain graphs, participants often failed to recognize certain direct links and reported non-existent ones, showing difficulties with complex causal structures.

Experiment 2

Impact of ground truth-prior congruence

- Three scenarios:
  - Criminality: crime rate, police action and population happiness
  - Real estate: house prices, desirability and population density
  - Finance: virus cases, stock prices and lockdown measure
- Difficulty level - (inverse) similarity between ground truth and prior:
  - Congruent: [0.8, 0.9]
  - Incongruent: [0.45, 0.55]
  - Implausible: [0.05, 0.15]

Similarity between two graphs

$$\mathrm{acc}(\boldsymbol{\gamma}) = 1 - \frac{\| \boldsymbol{\gamma}_{gt} - \boldsymbol{\gamma} \|}{\| \boldsymbol{\gamma}_{gt} - \boldsymbol{\gamma}_{*} \|}$$

Ground truth γ_{gt} and estimated γ are 6-d vectors of the strength of each link

Furthest possible graph γ_* from γ_{gt}

Experiment 2

Impact of ground truth-prior congruence

- 1 simpler generic graph (chain, common cause or collider) + 3 labeled trials
  - Each participant sees exactly one of each scenario and one of each prior condition

Experiment 2 results

Impact of ground truth-prior congruence

- **Scenario Impact:** the finance scenario was more challenging
- **Difficulty Impact:** difficulty level had a significant effect on accuracy

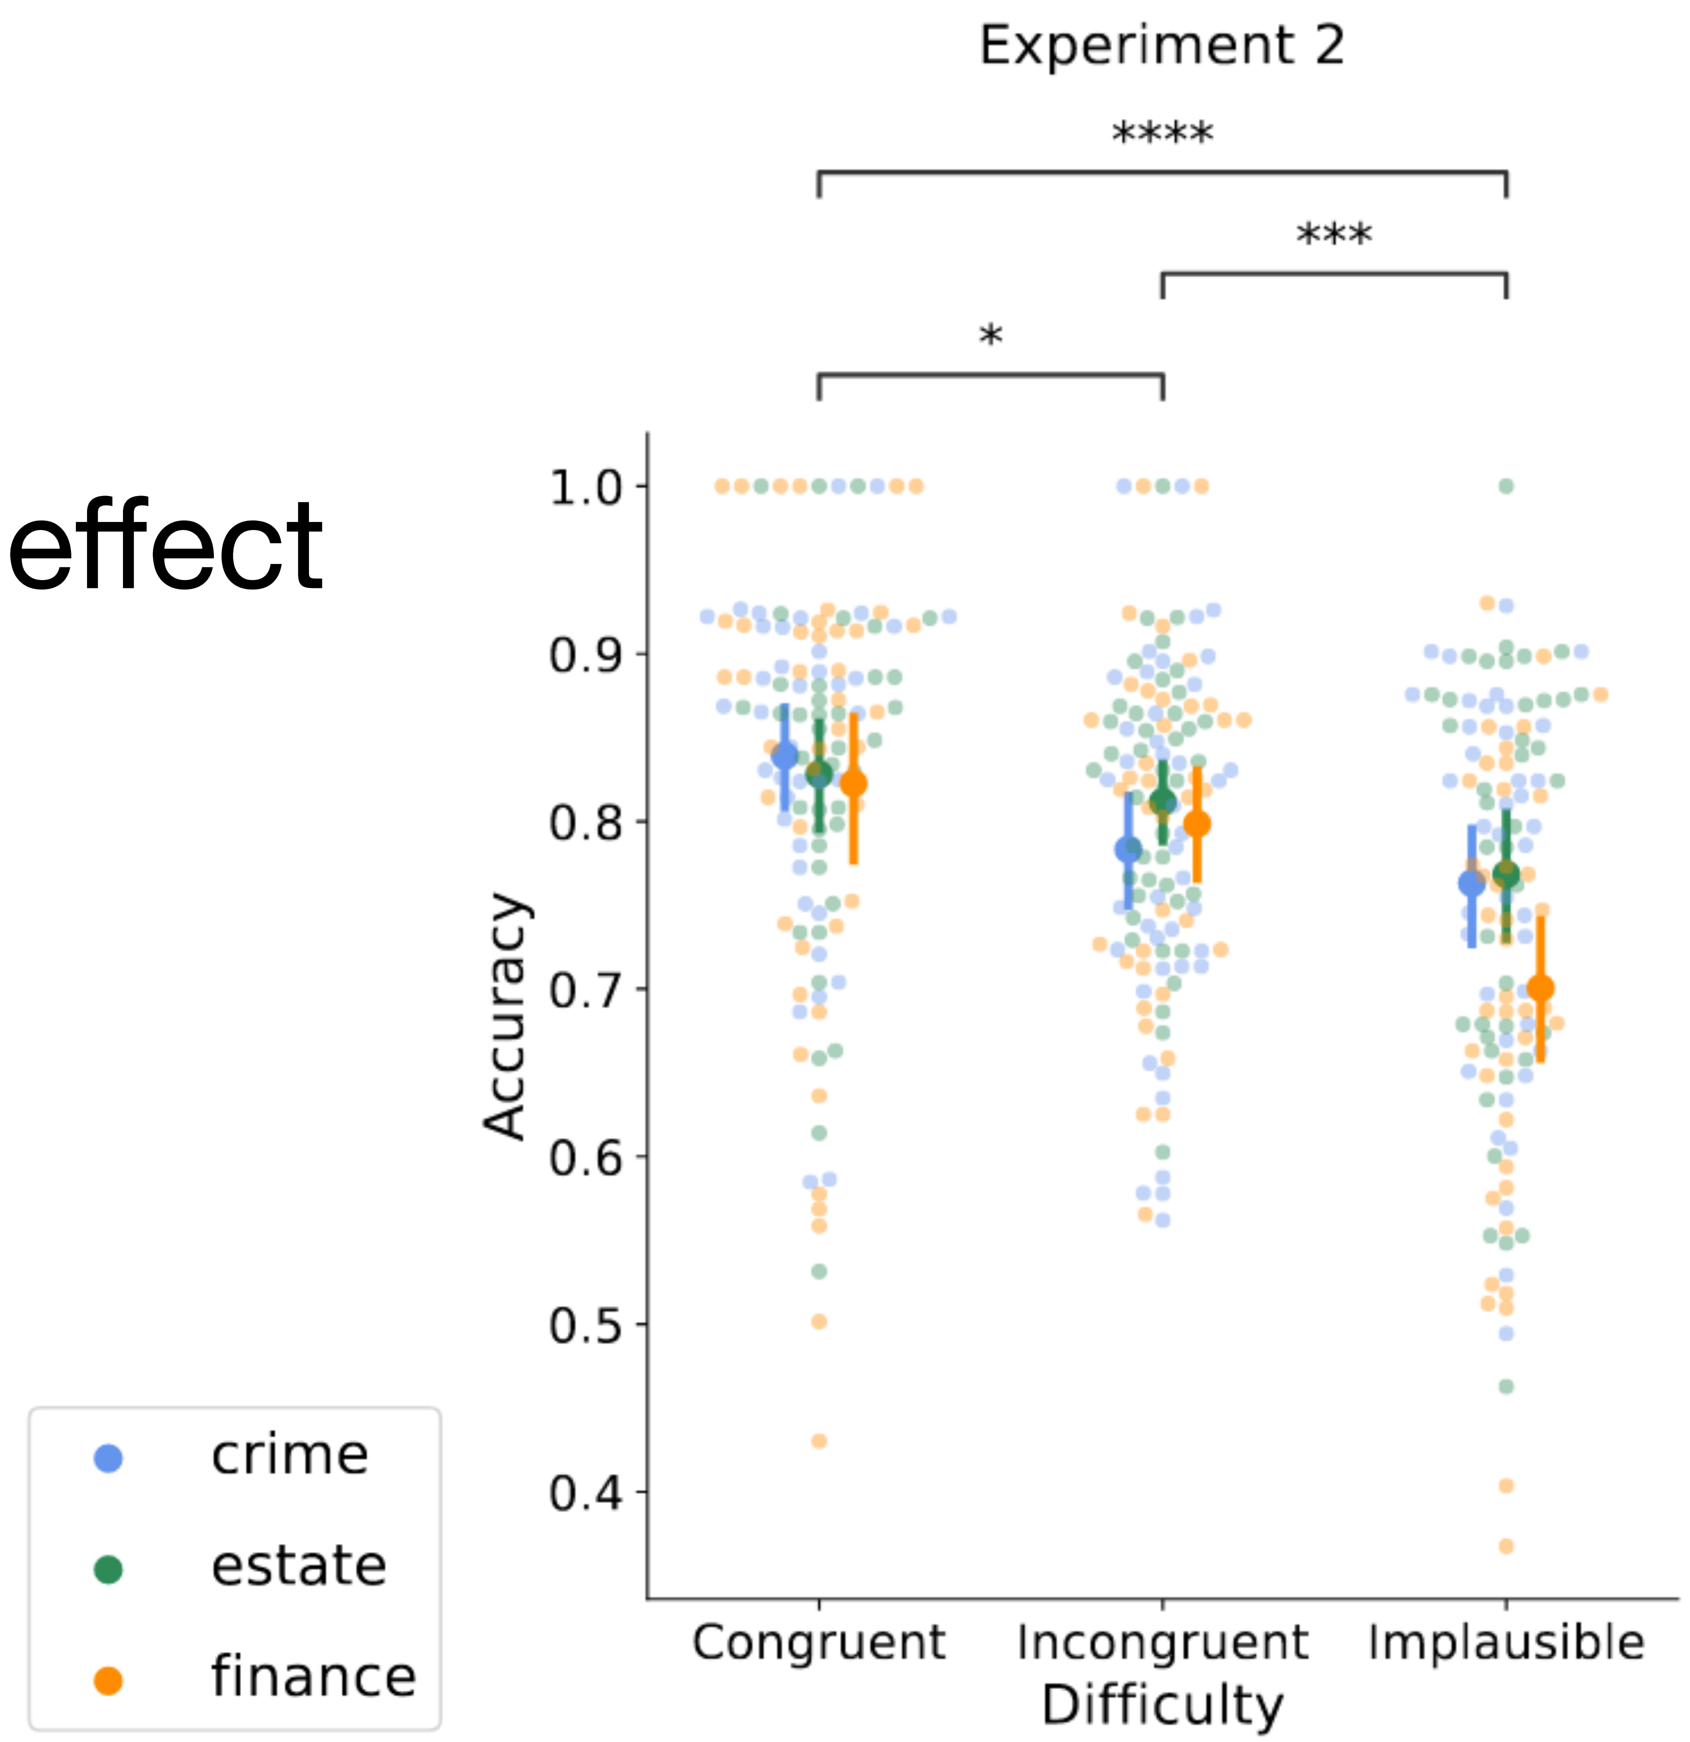

Experiment 3 & 4

Impact of *constrained* ground truth-prior congruence

- Make the manipulation more aligned with how people actually reason about causality (focusing on link direction rather than abstract distance)
- Avoid unintended structural complexities, like indirect effects, which made some graphs harder to learn and confounded the previous results.

Experiment 3 & 4

Impact of *constrained* ground truth-prior congruence

- Difficulty condition
  - Congruent: the ground truth is exactly the prior judgement.
  - Incongruent: the ground truth is the prior judgement graph but the direction of all links has been inverted, e.g. a chain A → B, B → C is now C → B, B → A.
- As only two difficulty conditions are left, we dropped the estate scenario.

Experiment 3 results

Impact of *constrained* ground truth-prior congruence

- **Effect of Scenario:** no
- **Effect of Difficulty:** yes

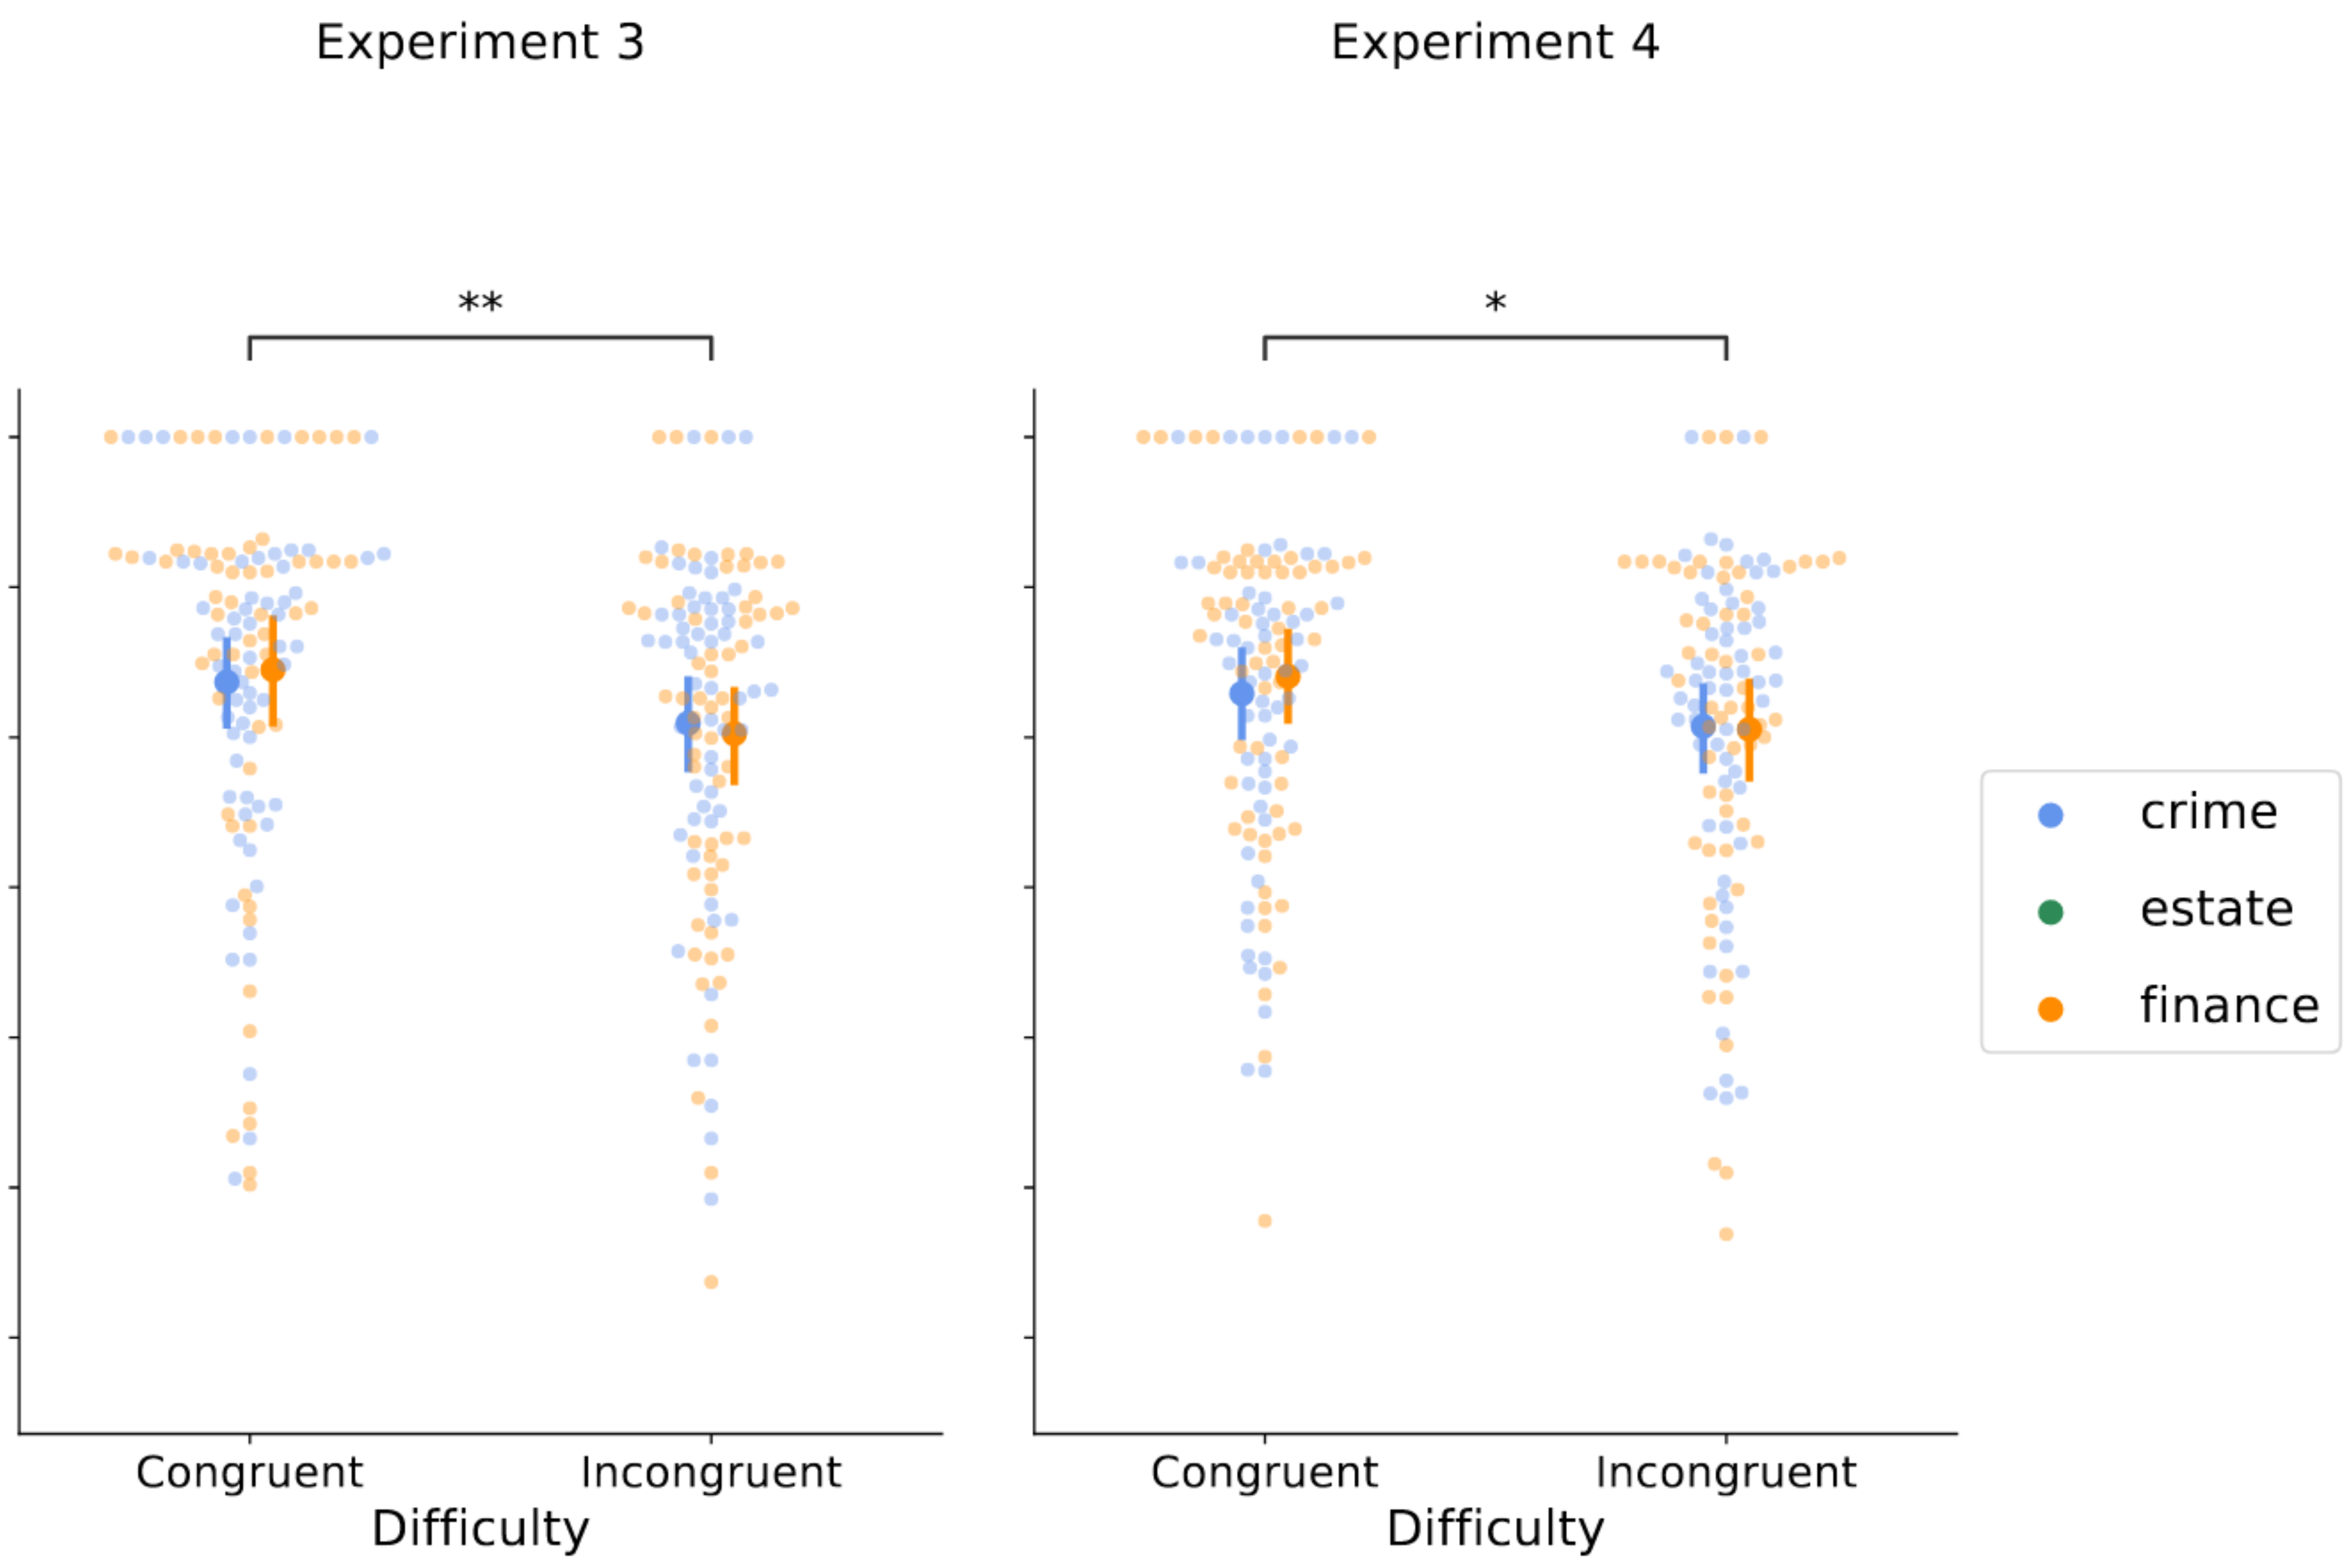

Modeling

All comparison models

- Inference over
  - graph structure Γ,
  - system parameter θ,
  - noise σ,
- Fit inverse temperature parameter τ

- Normative model: full inference $p(\Gamma, \theta, \sigma \mid X_{1:t})$
- Local computation: $\Gamma = \gamma_1 \cup \ldots \cup \gamma_6 \mid X_{1:t}$
- Causal event abstraction

$$\text{Causal Event } = \begin{cases} 1, & \text{if } X \geq T_X \text{ and } Y \geq T_Y \\ 0, & \text{otherwise} \end{cases}$$

- Adaptive selective LC: temporal locality & link focus

$$p\left(x_{t+dt} \mid x_t, y_t, \theta, \sigma, dt\right)$$
$$= \mathcal{N}\left(x_t + \theta\left(\mu_x\left(x_t, y_t\right) - x_t\right) dt, \sigma dt\right)$$

$$\mu_x\left(x_t, y_t\right) = \beta\left(x_t\right) + \gamma_{xx} x_t + \gamma_{yx} y_t$$

Model fitting results

- AS LC model provides a better fit than the Causal Event Abstraction (CEA) model both at individual and group levels.
- AS LC achieves similar or better performance than basic LC, even when only using data from interventions.
- Some participants showed a bias towards prior judgments, aligning with a model that has a non-flat prior distribution.

Model fitting results

- The AS LC model provided the best fit for participants' behavior at both group and individual levels.

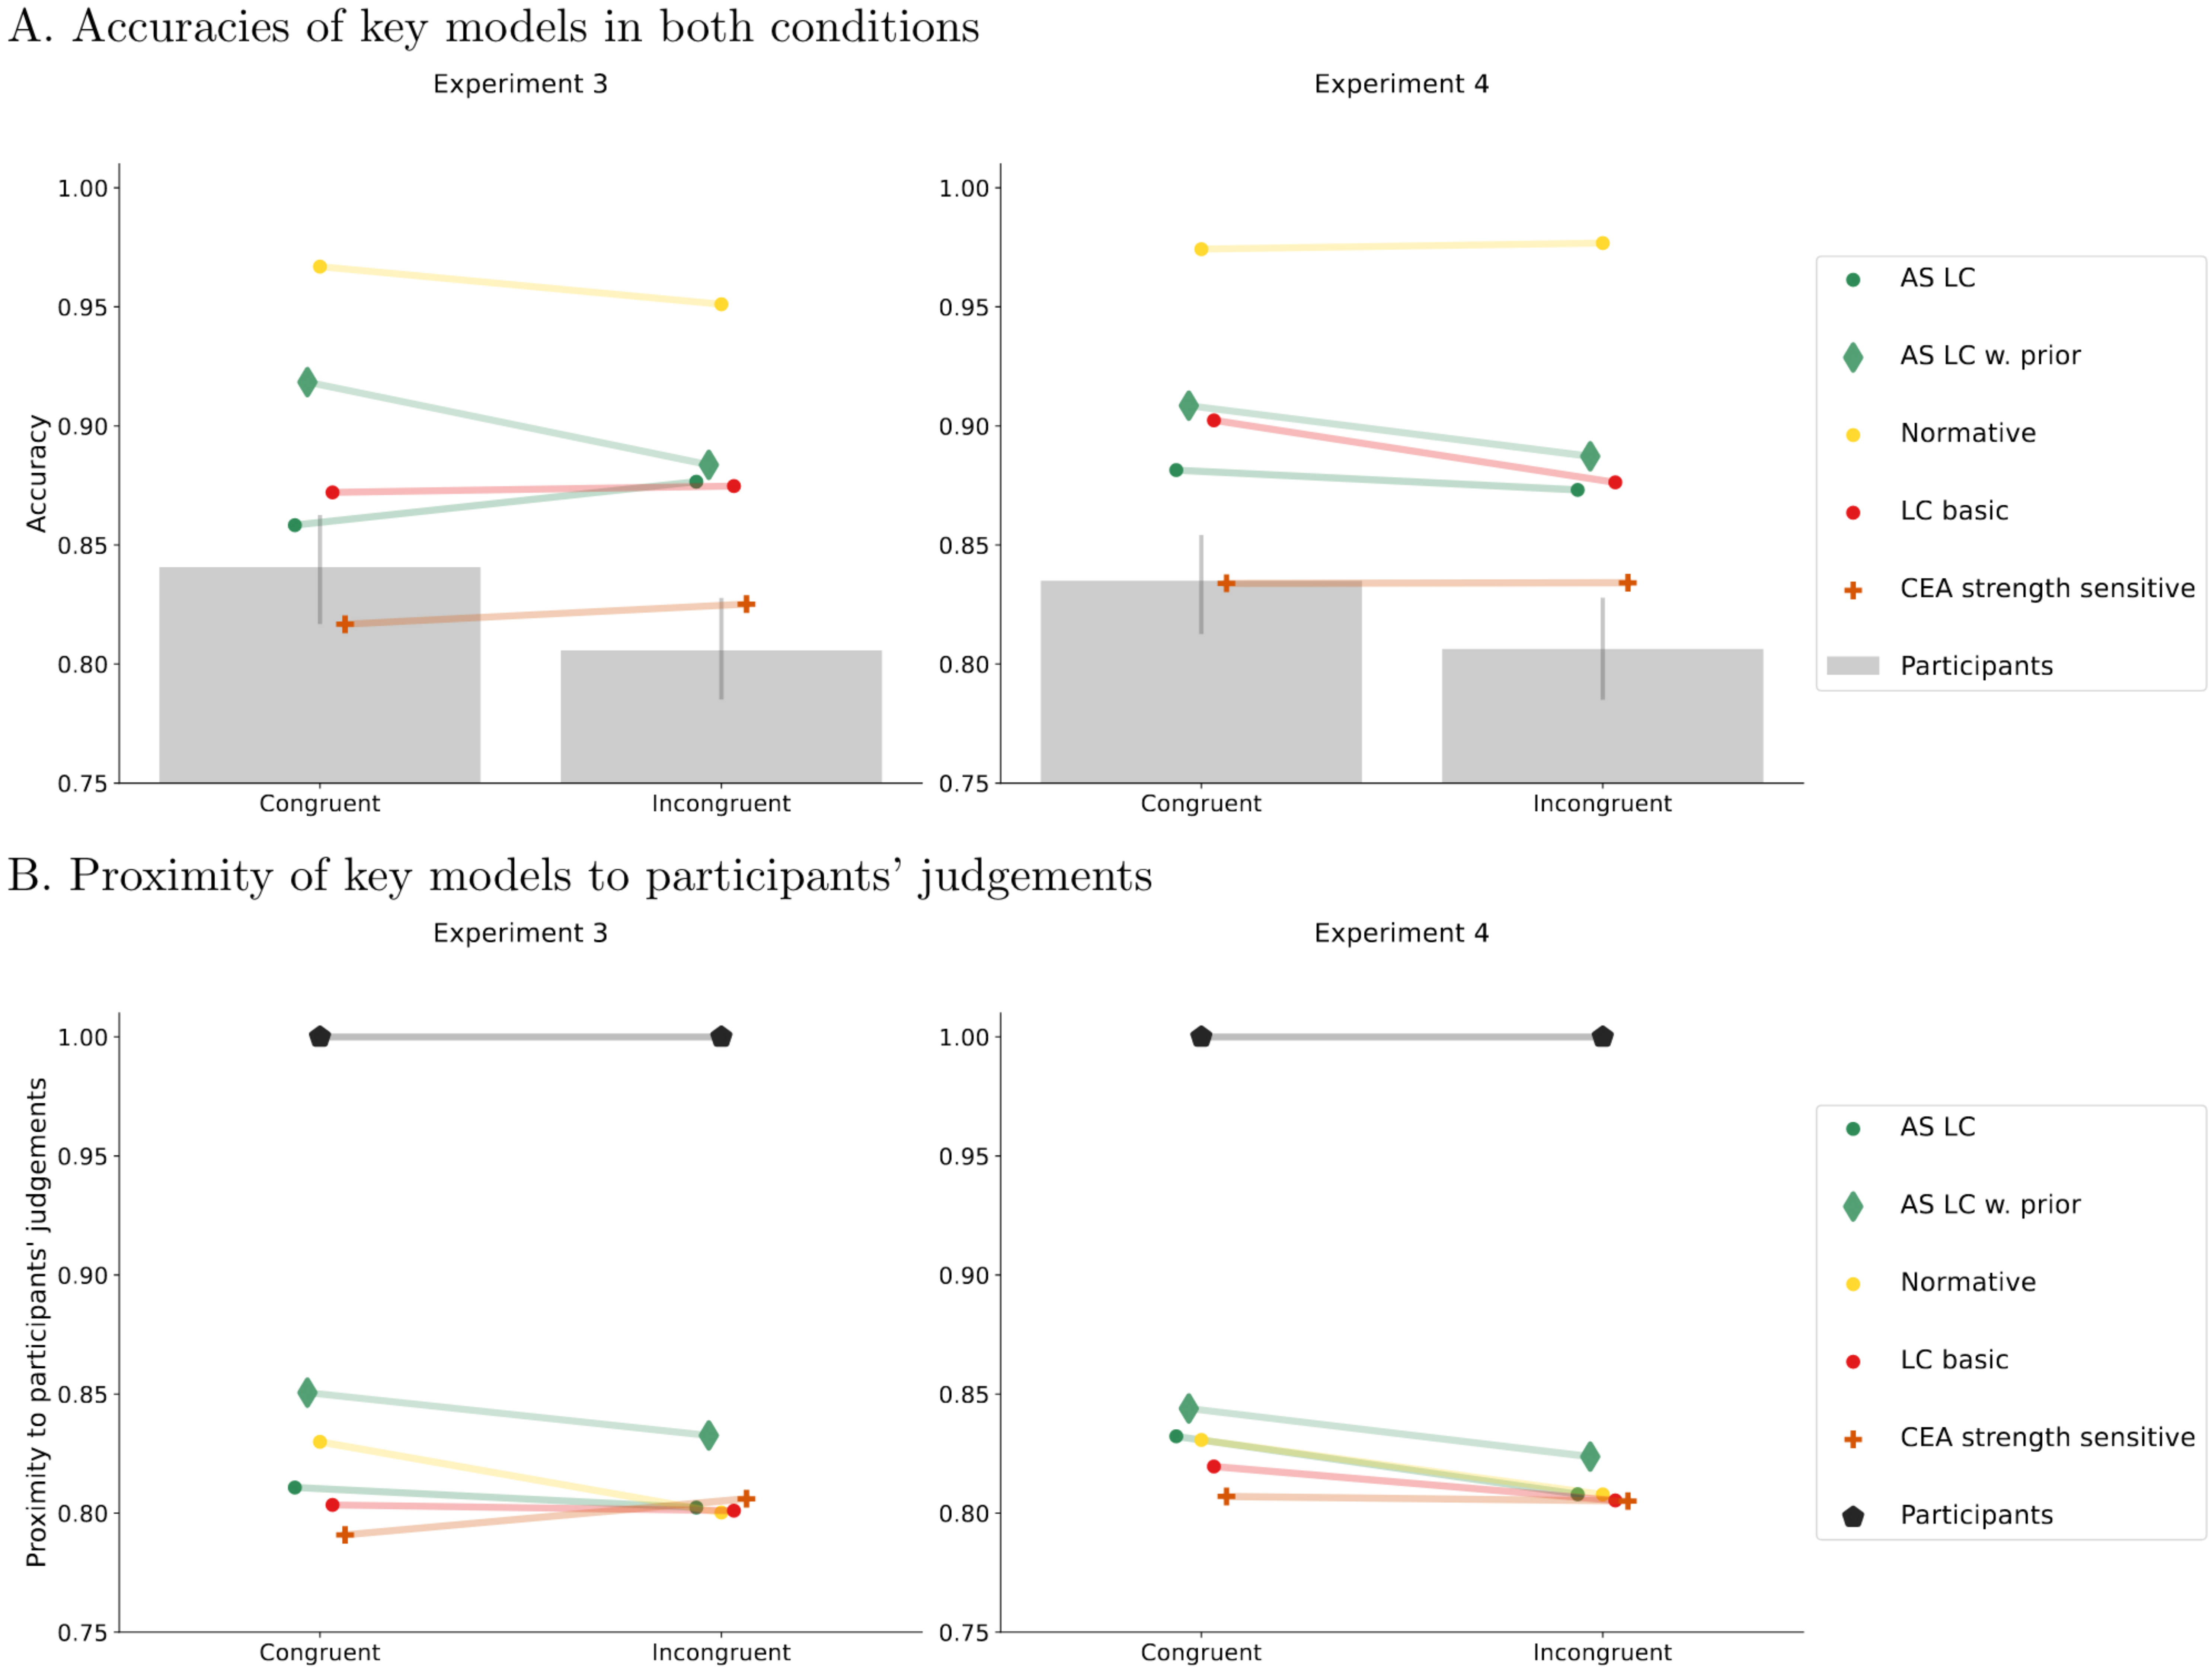

The prior shapes the inference process

- Participants were categorized based on whether their behavior was best fit by the Adaptive Selective LC (AS LC) model with a non-flat prior or with a flat prior (uniform).

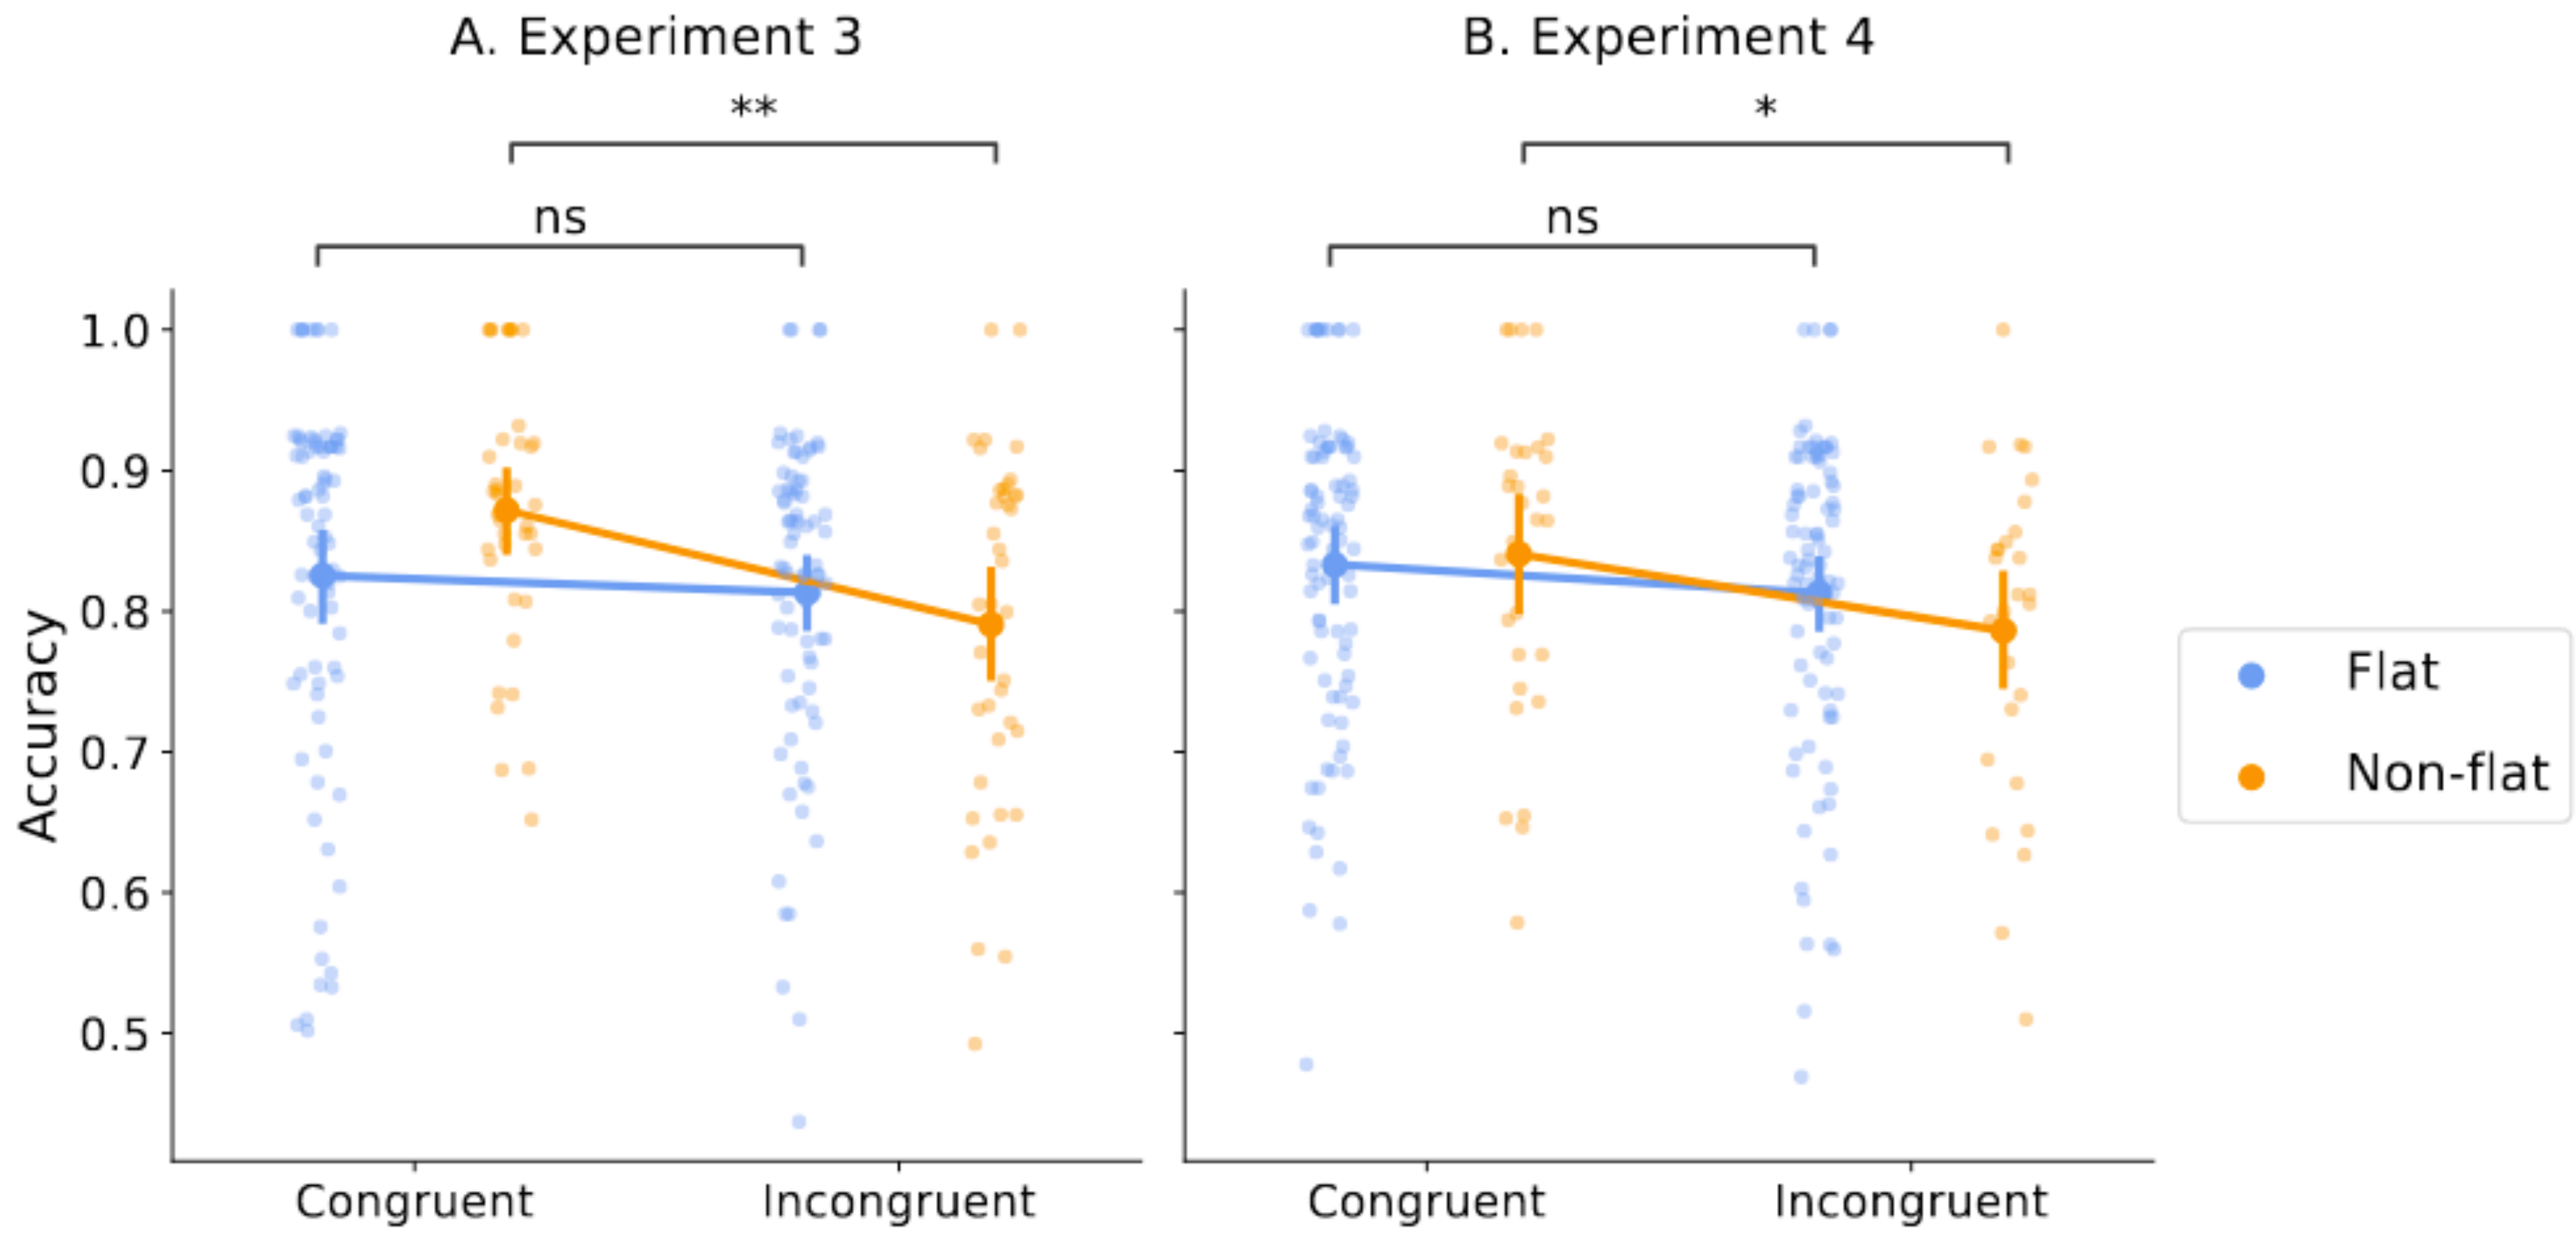

Figure 8: Mean differences in accuracy between condition for participants best fit by AD LC with a non-flat prior parameter and those best fit with assuming a flat prior and interactions with trial difficulty for experiment 3 and 4.

The prior shapes the intervention process

- Participants with a non-flat prior in incongruent trials
  - Have higher posterior entropy in actions
  - Favor interventions on the root node

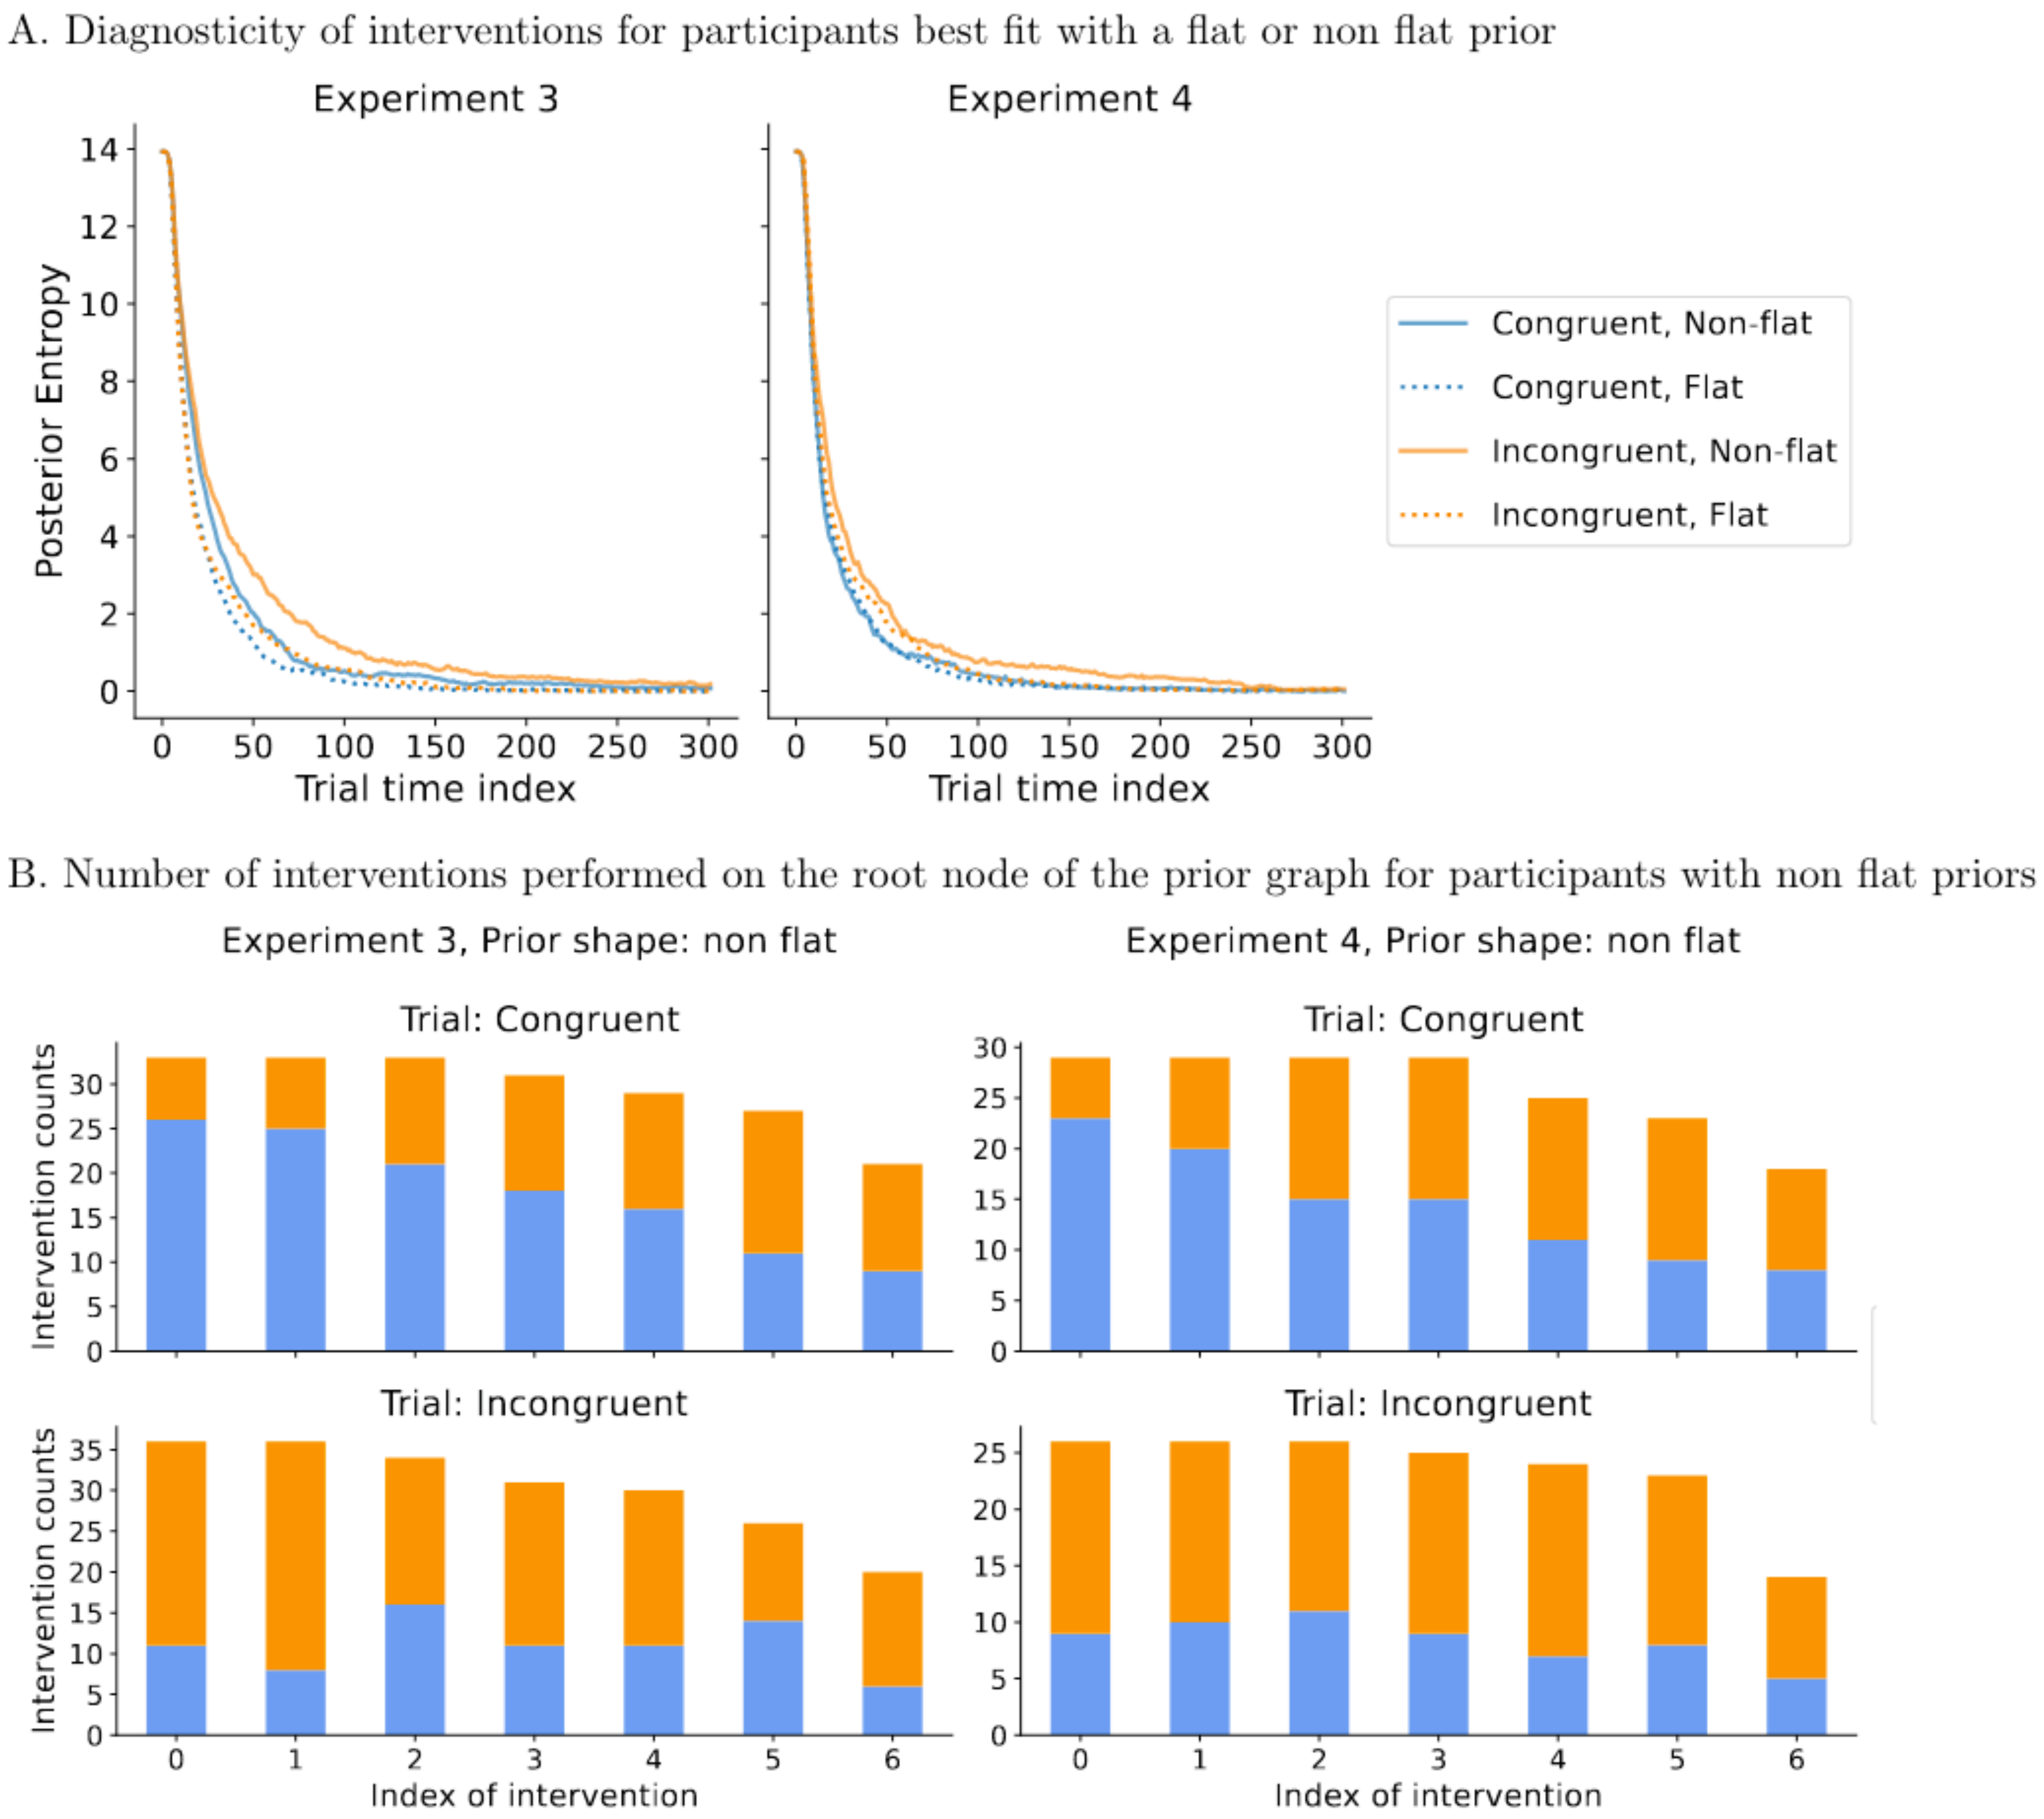

Summary

Summary of findings

- Behaviour
  - Participants who relied on prior beliefs performed better when these beliefs aligned with the ground truth, reducing errors associated with chain bias (mistaking indirect effects as direct).
  - When the ground truth contradicted prior beliefs, participants' intervention policies were less diagnostic, leading to worse performance.
- Modeling
  - An adapted Local Computations (LC) model, which considers data only during and shortly after interventions, closely matched participant behavior.