Causal Reasoning in a Prediction Task with Hidden Causes

Pedro A. Ortega, Daniel D. Lee and Alan A. Stocker
University of Pennsylvania

Motivation

Causal knowledge

Humans guide their decisions using **causal knowledge**.
Causal knowledge predicts **what the world will do** when we **interact** with it.
The processing of causal information is deeply embedded in cognition.

But how causal knowledge is:
- **represented**,
- **learned**,
- and **used**,

is currently not **well understood**.

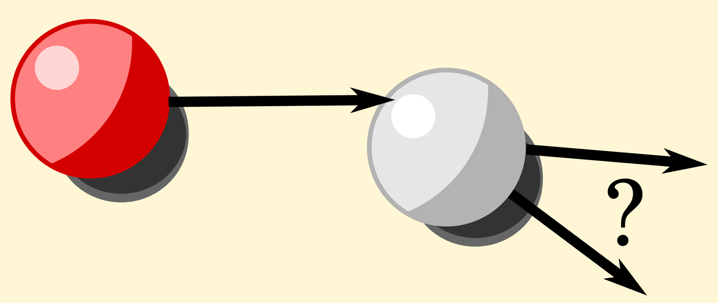

Predictions in observations *vs.* interventions

Causal knowledge affects the way we:
- **interpret evidence** and
- **make predictions**:

Complex Causal Model

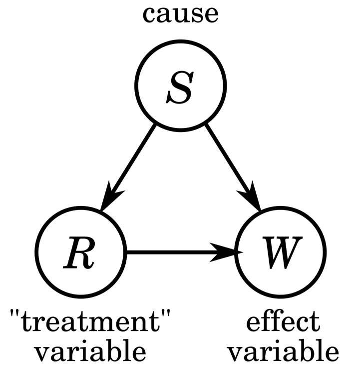

This is the **most general** causal dependency. It captures the idea of **partial control.**

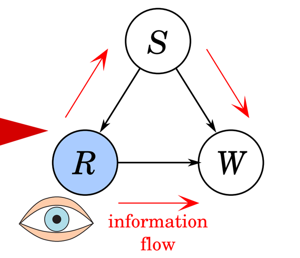

Observational Regime

$$P(W|R) = \sum_s P(W|S=s,R)P(S=s|R)$$

The observation provides information about the underlying hidden cause.

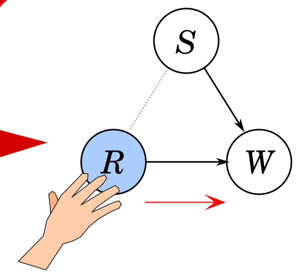

Interventional Regime

$$P(W|\mathrm{do}(R)) = \sum_s P(W|S=s,R)P(S=s)$$

The information doesn't provide information about the hidden cause.

How do we learn complex causal dependencies?

Hypothesis

Is it sufficient to experience **both regimes** to learn a complex causal dependency?

Method

We let people play a repeated **betting game** that they can **intervene** half ot the time.
We infer their beliefs from their bets and compare them to the causal model.

Subjects must rely on their **sense of agency** to interpret their experience.

Experiment

The betting game

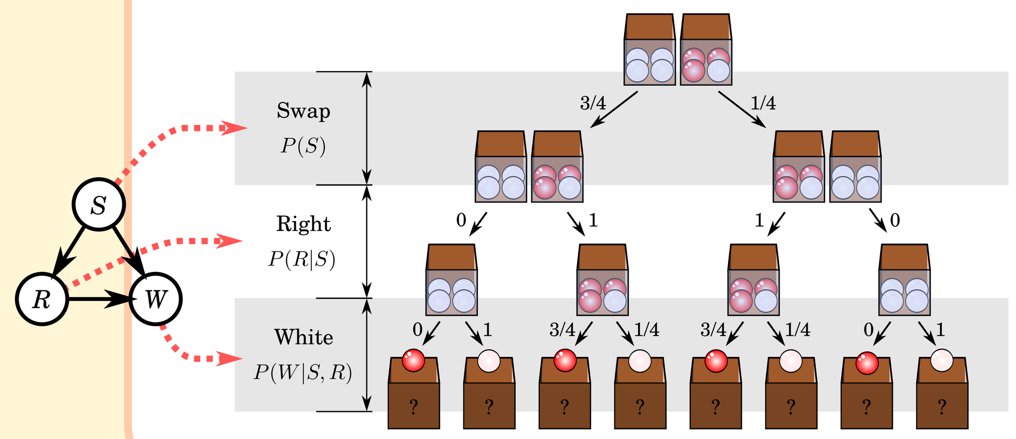

There are **two boxes** containing red and white balls.

Contents are **hidden**.

Subjects must bet on the colour of a **randomly drawn** ball.

Game and trials

Subjects **must complete** 40 blocks (levels) of 10 trials each.
They are allocated an **initial budget** at the beginning of each block.
Each **bet reduces** the budget.
Their goal is to **keep** as much as possible of the initial budget.
If they reach zero, they **must repeat** the block.

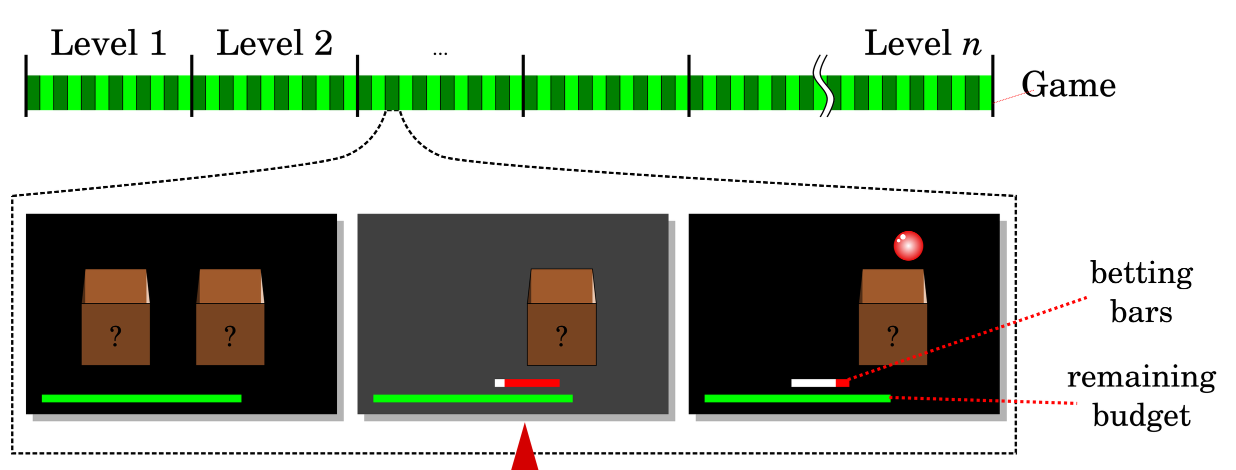

Bets are placed here by adjusting the length of the **coupled bars**.

Betting system

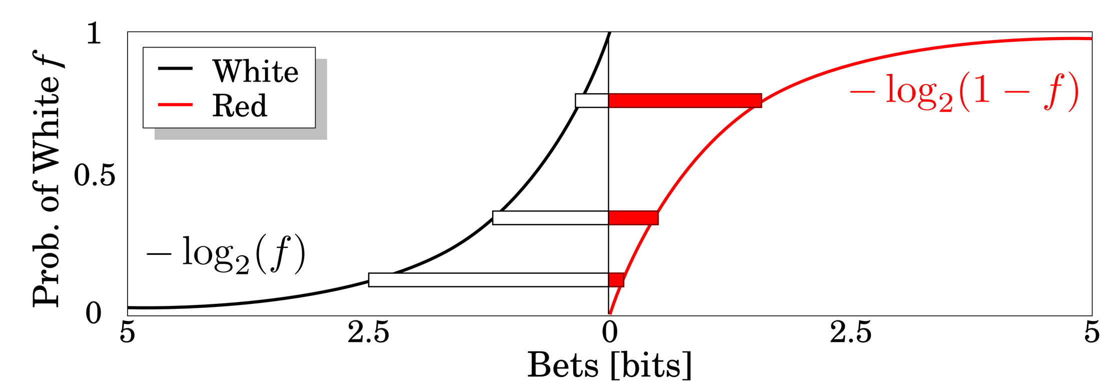

Log-loss scoring rule encourages reporting the true beliefs.
The system allows measuring beliefs on a **trial-by-trial** basis.
Confident bets are **too risky**.
The initial budget is set so that **conservative guesses** cannot be successful.

Results

Final prediction probabilities

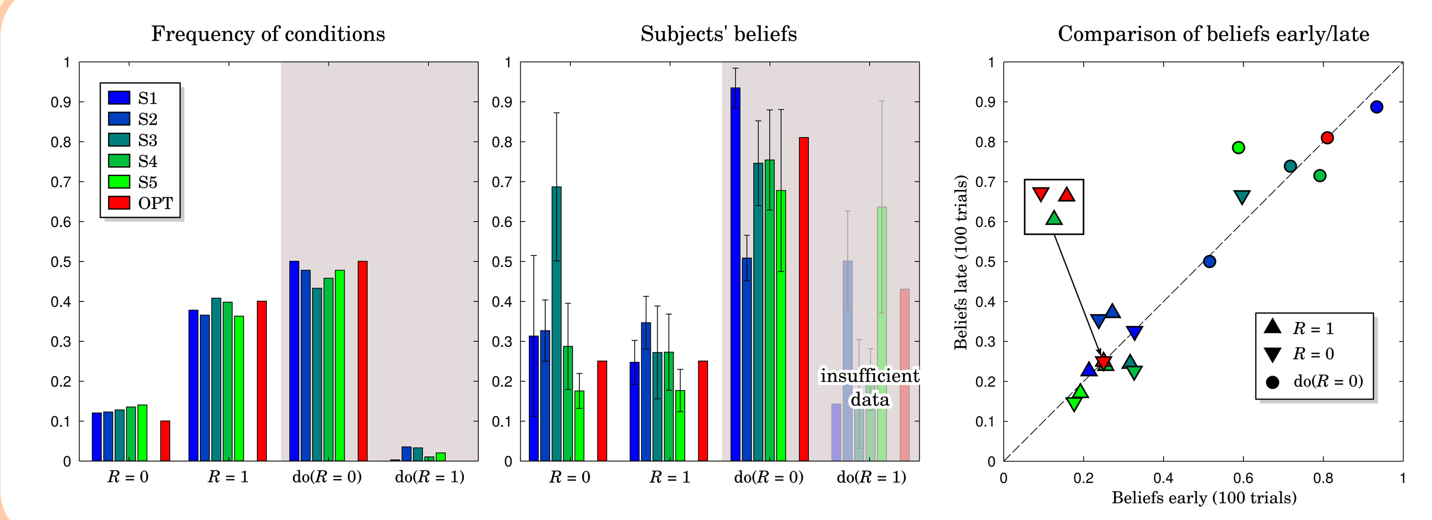

4 out of 5 subjects learned to predict correctly **right from the start**.
The results show expected utility, Bayes, and causality in action.
S3 treated every condition as interventional.

Game	Levels	Transparent	Interventions
Training 1	10	yes	no
Training 2	10	yes	yes (50%)
Test	40	no	yes (50%)

Learning curves

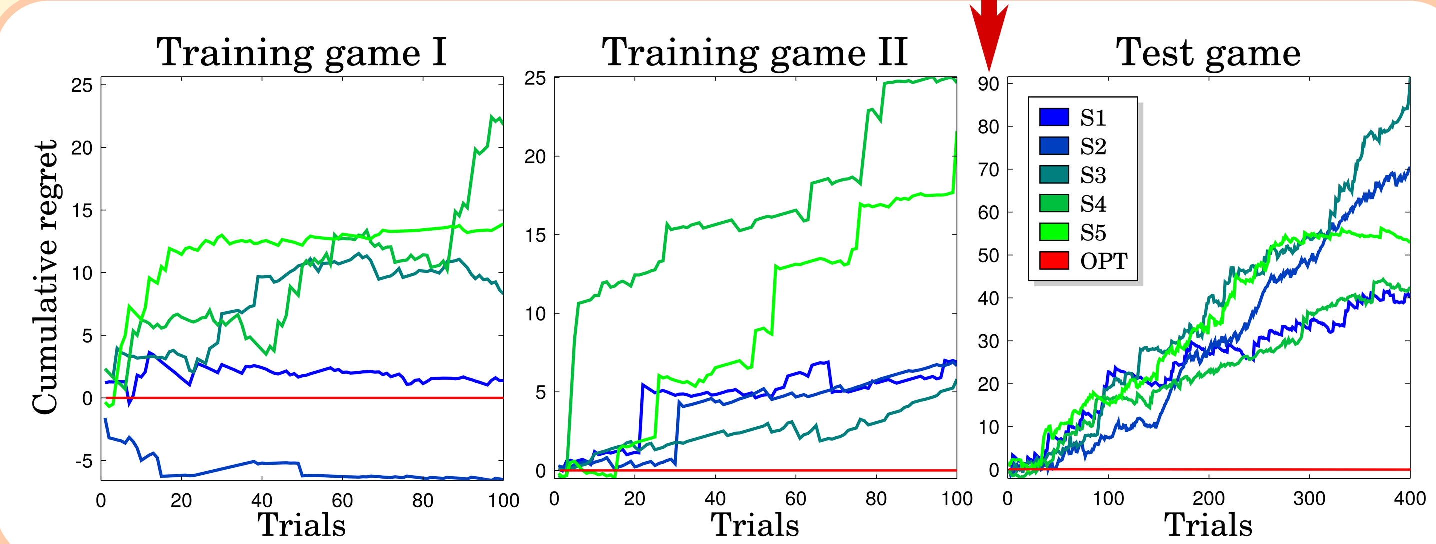

Training games: learning is very quick (< 40 trials)
Test game: little to **no learning**, yet positive slope: noisy beliefs?
S3 performs pretty well during the training games: smaller hypothesis space?

Summary

Excepting S3, the subjects made bets that were **consistent** with the **causal model**'s predictions.

Hence, they **learned** the causal model, **marginalised** over hidden causes, and **distinguished** between actions and observations.

Subjects appear to rely on a **sense of agency** to interpret their experience as either interventional or observational, even though they do not need to do so to perform well.

References

Hagmayer, Y. and Sloman, S. (2009). People conceive of their choices as intervention. *Journal of Experimental Psychology: General*, 138:22-38.
Hagmayer, Y. and Meder, B. (2013). Repeated Causal Decision Making. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 39(1).
Meder, B., Hagmayer, Y., and Waldmann, M.R. (2009). The role of learning data in causal reasoning about observations and interventions. *Memory & Cognition*, 37(3):249-264.
Dawid, A. P. (2006). Probability Forecasting. *Encyclopedia of Statistical Sciences*.