EFFECTS OF CAUSAL EXPLANATIONS AND CONFIRMATION ON PROBABILISTIC JUDGEMENTS.

Can Konuk[1],Nicolas Navarre[2] Salvador Mascarenhas[1]

[1]School of Informatics, School of Philosophy, Psychology and Language Sciences, University of Edinburgh, Scotland, United Kingdom
[2]Ecole Normale Supérieure, Department of Cognitive Studies, Institut Jean-Nicod, ENS, EHESS, PSL University, CNRS, Paris France.

Theoretical background

Base-rate neglect (Kahneman & Tversky, 1973): when computing the probability of an hypothesis H given a body of evidence E, people tend to ignore the prior probability $P(H)$ rely exclusively on E. Two classes of explanations have been proposed:

Evidential Impact Perspective
People ignore Base-rate because they are more interested in the confirmation relation $c(H, E$ between E and H (How much learning E increases or decreases the probability of H) than in the posterior probability $P(H \mid E)$ (How probable is H, given E?).

Causal Reasoning View
People are interested in posteriors, but they ignore base-rate because in most ordinary cases, those do not bear a transparent causal connection to the hypothesis of interest.

Example: the Mammogram problem (Eddy, 1982; Krynski & Tenenbaum 2007)

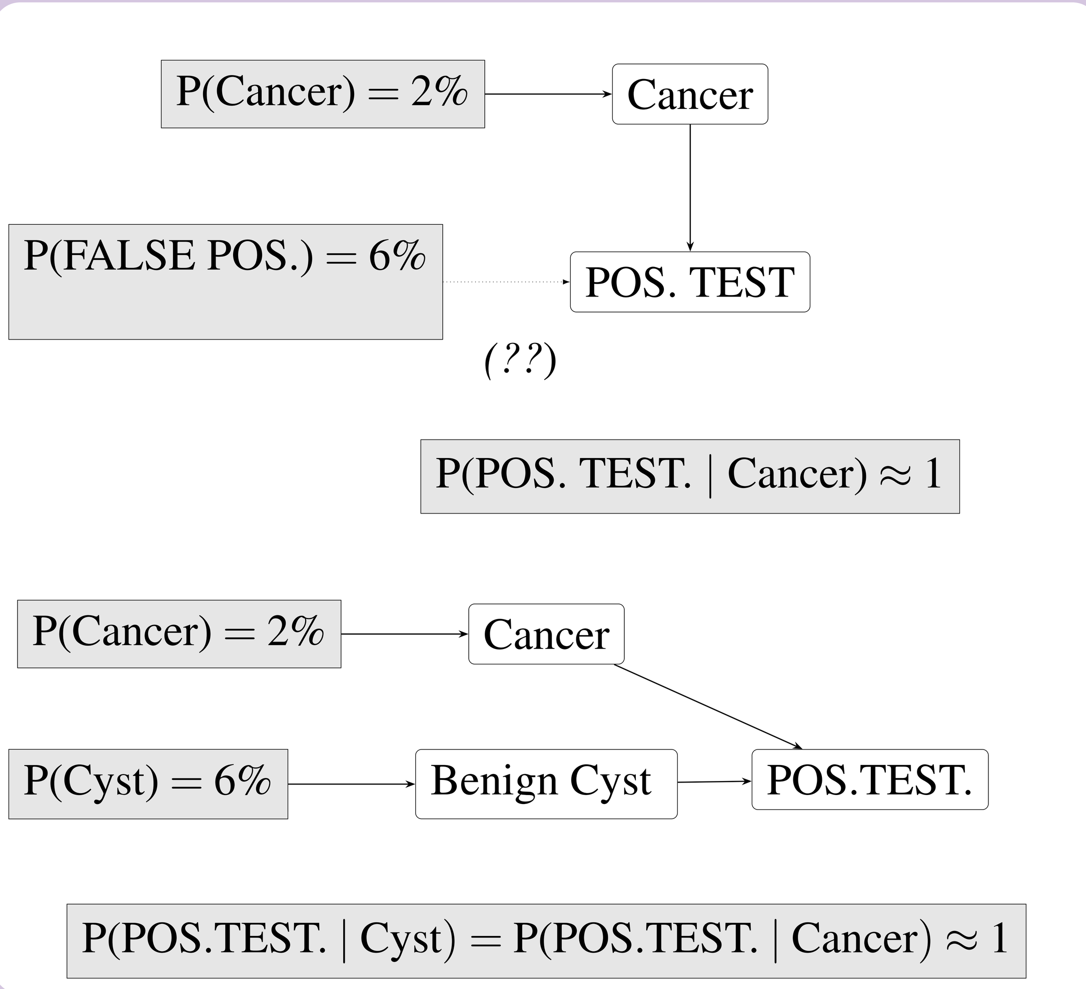

Teasing apart the two factors

Problem

In most cases that are likely to arise, confirmation and causal relations tend to be entangled. Wherever there is a causal relationship between two events A and B, there also tends to be a confirmation relation (positive or negative), and vice-versa.

- As a result, the two kinds of explanations above are hard to distinguish and often counfounding in existing experiments including Krynski & Tenenbaum (2007).
- In the Mamogram example presented above, the manipulation introduced by Krynski & Tenenbaum, the Cyst manipulation does not just introduce causality, it changes the evidential impact of the evidence, with $c(T+ \mid Cyst)$ being very high.

Proposal: To build a reasoning task, analogous to the original Lawyers and Engineers paradigm, where we specifically distinguish the two factors.

Selected references and acknowledgments

Kahneman, D., Tversky, A. (1973). On the psychology of prediction. Psychological Review, 80, 237–251.
Krynski, Tevye Tenenbaum, Joshua. (2007). The Role of Causality in Judgment Under Uncertainty. Journal of experimental psychology. General. 136. 430-50. 10.1037/0096-3445.136.3.430.
Eddy, D. M. (1982). Probabilistic reasoning in clinical medicine: Problems and opportunities. In D. Kahneman, P. Slovic, A. Tversky (Eds.), Judgment under uncertainty: Heuristics and biases (pp. 249—267). Cam- bridge, England: Cambridge University Press.

Try it yourself!

Research supported by Agence Nationale de Recherche grants ANR-17-EURE-0017 (FrontCog) and ANR-18-CE28-0008 (LANG-REASON). This work was also supported by a grant overseen by the French National Research Agency (ANR) as part of the "Investments d'Avenir" Programme (EURIPANR-17-EURE-0012) and by "Ecole Doctorale Frontières de l'Innovation en Recherche et Education - Programme Bettencourt" for financial support. Supported in part by the UKRI Centre for Doctoral Training in Natural Language Processing, funded by the UKRI (grant EP/S022481/1) and the University of Edinburgh, School of Informatics and School of Philosophy, Psychology & Language Sciences.

The experiment

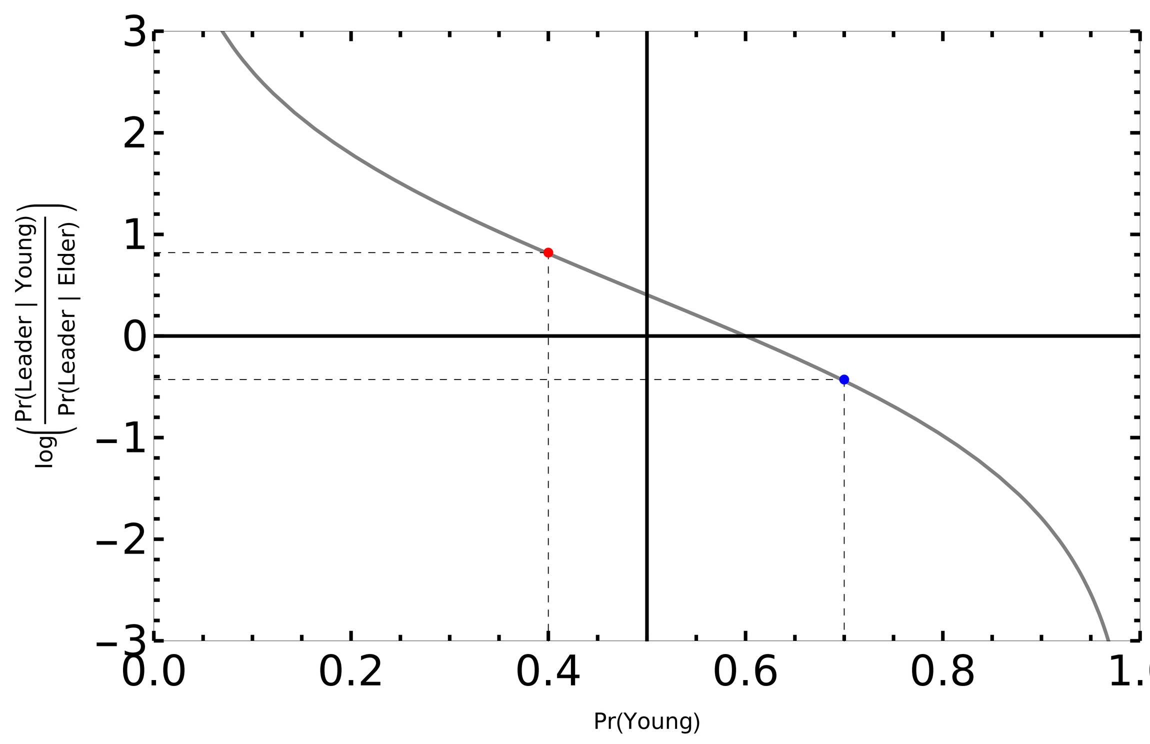

Figure 1: The $x-axis$ represents the proportion of Young people present in the set of villagers selected for the Parade, the y-axis, the proportion of Young villagers in the Parade that get selected to be Leader, in terms of the log-likelihood statistic. The curve represents the combinations of Base-rate and likelihood values that would give a posterior probability of Pr(Young|Leader) of 0.6. The dots represent the particular values of that conjunction that were chosen for our EV-POST (red) and PRIOR-POST (blue) conditions.

In the PRIORS-POST condition, the majority group in the first round of selection (corresponding to the priors or base-rate) matches the majority group in the posteriors, but **not** that of the second round.

- In this context, neglecting the base rate would lead to an underestimation of the proportion of Young People.

We also vary for each conditions:

- Whether or not the Base-rate information is causally connected to the target hypothesis.
- Whether or not the updating evidence is connected to the target hypothesis

Leading to a total of **2x2x2= 8 conditions.**

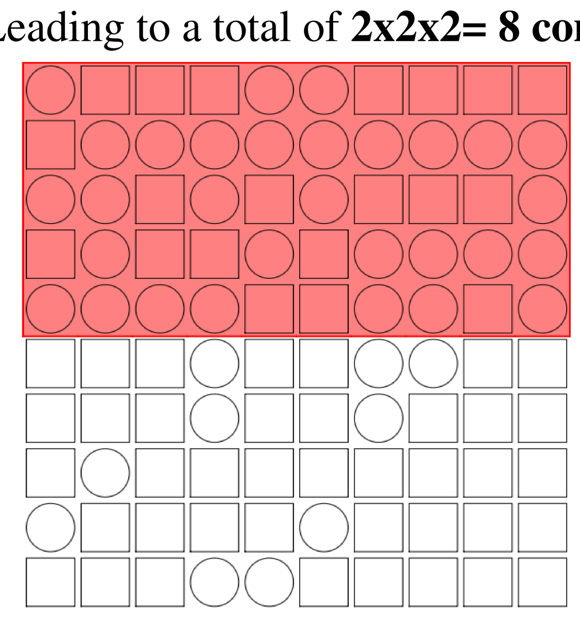

Figure 2: Images displayed along the story, in every condition. They helped people track the relevant proportions. The Youth are represented as circles, the Elder are represented as squares, and the red rectangle highlights the subset of villagers that got selected as Leaders.

Design of the experiment

1. Subjects were told to consider a fictional story about an ancient Mesopotamian village, in which a ritual Parade takes place every year. Only a limited number of men are allowed to participate each year, so a selection must occur, which involves two steps:
   (a) The first step selects a hundred men from the general population to be members of the Parade.
   (b) The second step selects a subset of 50 men from the members of the Parade, to be elected *Leaders* of the Parade. They wear a distinctive attire made out of a red mask.
2. At each round of selection, subjects are to told about the proportion of men from each one of two age brackets: "Youths" (aged 16 to 29) and "Elders" (aged 30 or older), that got selected.
3. Once subjects have heard the entire story, they are presented with a particular villager named 'Balthazar', and told that he wears the red mask distinctive of the Leaders of the Parade
4. From there we ask them and asked to estimate the probability that he'd be either an Elder or a Youth. This amounts to computing the posterior probability $P(Young \mid RedMask)$.

Across all conditions, the final distribution has 30 Youths, but only 20 Elders in the 50 Leaders of the Parade. So there is a *60% probability* that the villager is a Youth.

We also ask them about:

- The story that subjects have heard contains all of the information necessary to assess the probability:
  - The first round of selection provides them with **base-rate** information about the target hypotheses
  - The second round provides them with additional **updating evidence**.
- We then check how subjects' answer might differ from that correct answer in two kinds of conditions:
  - In the EV-POST condition, the majority group in the second round of selection (corresponding to the updating evidence) matches the majority group in the posteriors, but **not** that of the first round.
    - In this context, neglecting the base rate would lead to an overestimation of the proportion of Young People.

Results

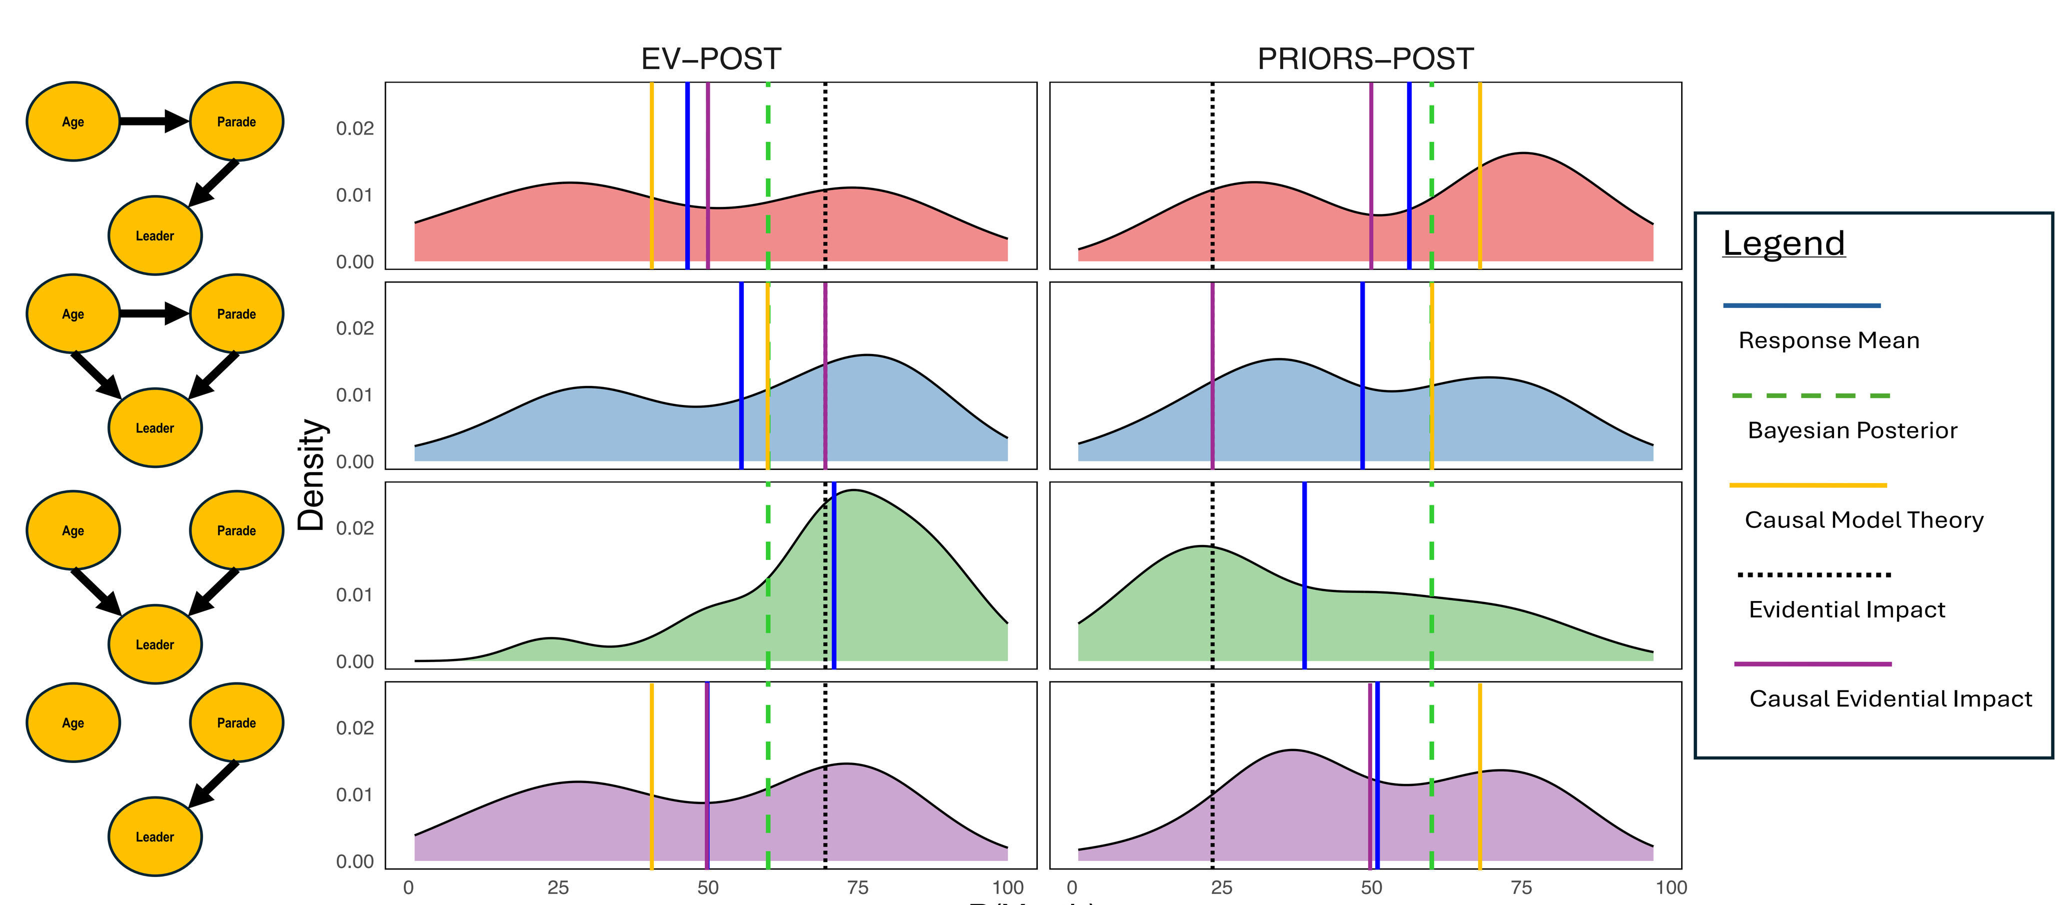

Figure 3: Distribution of Pr(Youth) estimates as a function of the causality of each type of information. The lines describe different model predictions. Note that the causal evidence condition all model predictions are in line with Evidential impact.

Summary

- People were more influenced by the impact of the evidence overall, regardless of whether the base-rate and or evidence were causal.
- Subjects' probability estimates were better predicted by their reported likelihood estimate $P(Red\ Mask \mid Youth)$ than by the posterior probability that came out of their Likelihood and Prior estimates (both of which were collected successively).
- Information that was framed causally was given more importance overall, but this mostly applied to updating evidence: where the evidence was causal, people's choices were more influenced by it than when it wasn't. On the other hand, the impact of causality on the base-rate was negligible.