

Effects of causal structure and evidential impact on probabilistic reasoning

Can Konuk (can.konuk@ens.fr)

Department of Cognitive Studies, Institut Jean-Nicod
Ecole Normale Supérieure
29 rue d'Ulm, 75005 Paris, France

Nicolas Navarre (n.s.navarre@sms.ed.ac.uk)

School of Informatics, University of Edinburgh
10 Crichton St, Newington, Edinburgh, EH8 9AB, UK

Salvador Mascarenhas (salvador.mascarenhas@ens.fr)

Department of Cognitive Studies, Institut Jean-Nicod
Ecole Normale Supérieure
29 rue d'Ulm, 75005 Paris, France

Abstract

We compare two perspectives on *base-rate neglect* (Kahneman & Tversky, 1973) in probabilistic judgment. The *evidential impact* perspective derives it from humans' focus on the impact of evidence on belief, rather than conditional probabilities. The *Causal Models* perspective derives it from humans' inability to integrate information that is causally opaque, as base-rates often are in such experiments. Because causal and evidential-impact relations are often concomitant and confounded, we designed an experiment that specifically teases apart their respective influence on probabilistic judgment. Our results support a combination of the two perspectives, with causal transparency influencing the degree to which one engages in evidential impact reasoning strategies.

Keywords: causal reasoning; rationality; probability; evidential impact; confirmation;

Introduction

Suppose there is a particular used car you're considering buying, and you're wondering if it is in as good a condition as the dealer advertises it to be. Several kinds of data might be relevant for you to consider. For one thing, you might want to consider the proportion of cars on the used-car market which are in good shape. For another, you might consider the trustworthiness of the dealer. Do they have a good reputation? And how does that affect the chances that the car they're selling you is in good shape?

Classical (Bayesian) probability theory provides a normative standard for how to evaluate and combine these different kinds of information. The first example above would define a prior probability or base rate, while the second would provide additional evidence, to be combined into a posterior probability for the hypothesis h 'This car is in good shape', given the evidence e 'The dealer is trustworthy'. Bayes' theorem provides a formula for combining these two kinds of information in a rational way.

Yet, psychological research has repeatedly found that human judgment appears to deviate from these standards significantly and often. In particular, when evaluating hypotheses, people tend to place more importance on the immediate impact of the evidence e , and tend to neglect the base rate or prior probability of the hypothesis h (Kahneman & Tversky, 1973).

Over the past twenty years, a theoretical perspective on these questions with great descriptive adequacy and promising explanatory potential has been taking shape. The idea is that humans are especially interested in the *impact* a piece of evidence has on belief in a hypothesis, rather than in the probability of the hypothesis conditional on the evidence. This view can be cashed out in terms of the Bayes factor or any of a number of measures of evidential impact in the formal-epistemology literature (Fitelson, 2007). The perspective has been applied to great effect to various reasoning problems from the representativeness literature, in particular the conjunction fallacy (Tentori, Crupi, & Russo, 2013; Crupi, Fitelson, & Tentori, 2008; Tenenbaum & Griffiths, 2001).

Other authors propose to shift the focus away from the probabilistic information itself (Krynski & Tenenbaum, 2007; Bar-Hillel, 1983). Krynski and Tenenbaum (2007) propose that humans tend to disregard information that does not lend itself to a *causal* interpretation, that does not seem to fit a causal model of the situation at hand. The reason base rates are often neglected in psychological experiments is because they rarely bear a causal connection to the hypothesis that is being considered. The proportion of cars in good condition on the market in general, however relevant and normatively useful, does not bear a transparent *causal* connection to the condition of the particular car I am considering. It does not *explain* it, in the way that the trustworthiness of the dealer might do.

But suppose you know that of late many people have been selling their cars out of a desire to buy an electric car instead of their traditional one, so that most vehicles on the used-car market right now are actually in pretty good shape. This might offer a causal connection to link the base rate information with the main hypothesis concerning the state of the particular car you're considering. Were a salient causal connection of this form available, Krynski and Tenenbaum argue, subjects would be far better at handling the available information in a Bayesian way. The authors back this proposal with a series of experiments on probabilistic reasoning that control for the extent to which a causal narrative is available for the base-rate information (Krynski & Tenenbaum, 2007).

But there is a huge risk with such experiments, not countenanced by Krynski and Tenenbaum: in many if not most situations, *causality* and *evidential impact* as discussed above are closely related and often confounded. If two events A and B are causally connected, then — in most cases likely to arise — learning about one will change my degree of belief in the other. Conversely, whenever there is an evidential relationship between two events, there is often a causal one as well. For this reason, the existing experimental literature is insufficient to distinguish between these two conflicting theories of base-rate neglect.

Thankfully, the link between the two notions is not inextricable. In this manuscript we introduce a new experiment designed specifically to tease apart the respective role of evidential impact and causal heuristics in subjects' judgments. Our findings reveal a more complex relationship between causality and probabilistic judgment than previously suspected. There is indeed an imbalance in the treatment of probabilistic information, in line with the predictions of evidential-impact approaches. Subjects overall tend to underestimate the importance of base-rates, even when they have a causal interpretation. But on top of that, the causal status of the updating evidence *does* bear an influence on subjects' proclivity to engage in evidential-impact reasoning. They appear to do so more readily when that evidence is causal. We speculate that this is so because evidential-impact reasoning strategies stem from a search for *causal explanations* for phenomena.

Theoretical background

We are interested here in two classes of interpretations of apparent violations of Bayesian principles.

On an *evidential-impact perspective*, the error typically ascribed to base-rate neglect actually stems from a tendency to prioritize the *evidential impact* of information over the posterior probability of hypotheses.

A common way to cash out the evidential impact of a certain piece of evidence e upon a given hypothesis h , is via confirmation measures $c(h, e)$, which provide various ways of computing the extent to which learning the evidence e increases one's degree of belief in h (Fitelson, 1999). A commonly used metric is the likelihood ratio, defined for a given hypothesis h and piece of evidence e as $L(h, e) = \log \frac{P(e|h)}{P(e|\neg h)}$.

By contrast, on a *causal-reasoning view*, base-rate neglect is less about a nonstandard use of information and more about how the information is integrated into causal models (Ali, Chater, & Oaksford, 2011; Krynski & Tenenbaum, 2007).

According to Krynski and Tenenbaum (2007), judgment under uncertainty starts by building a causal model of the situation at hand, where all relevant pieces of information are connected causally to the hypotheses under consideration, and then proceeds to read probabilities off this graph and combine them, using Bayes' rule.

A consequence of this two-step process is that only information that can indeed be integrated into a causal model will

effectively factor into the posteriors. In many cases, however, base-rates do not bear a transparent causal interpretation. Only once such an interpretation is made available will humans integrate base-rate information in the appropriate Bayesian manner. Crucially, this interpretation entails that humans' default way of handling probabilistic information at its core follows the standards of Bayesian rationality. Mistakes arise from subjects' difficulty to represent non-causal information, but the information that subjects do represent is processed in a normative manner. This is in clear contrast with the evidential perspective, in which one's strategy for processing information is itself non-normative.

Krynski and Tenenbaum (2007) present a number of behavioral experiments meant to support this view: One such experiment involves a reformulation of the mammogram problem, a classic problem in probabilistic reasoning (Eddy, 1982; Gigerenzer & Hoffrage, 1995). In the traditional version, participants are provided with three key facts: (1) the prior probability of breast cancer for any woman getting a routine screening is 1%; (2) 80% of women with cancer get a positive mammogram test (true positive rate); (3) 9.6% of women without cancer also get a positive mammogram test (false positive rate).

Participants are then asked to judge the probability that a woman with a positive test has cancer. They often give answers in the 70%–90% range, close to the true positive rate, whereas Bayes' theorem prescribes a much lower probability of 7.8% integrating all available information.

Krynski and Tenenbaum's reformulation consists in reframing the false positive rate in point (3) above in a causal fashion: they tell subjects that about 9.6% of the time, women tested can present with a benign cyst, which would also yield a positive result on the test, while being completely unrelated to cancer. As they predict and observe, this reframing brings subjects' estimates much closer to Bayesian standards.

This finding is intriguing. Krynski and Tenenbaum interpret it as evidence that making a base rate causally understandable reduces base-rate neglect, by allowing subjects to connect the base-rate information with the rest in their causal model of the situation. But these results can also be viewed through the lens of evidential impact. While in the original version the confirmation relation $c(\text{cancer}, \text{positive test})$ is weak, the relation $c(\text{cyst}, \text{positive test})$ is strong. To put it more plainly, the moment the possibility of a benign cyst and its correlation with positive test results become salient, *irrespective of a causal story connecting the two*, learning of a positive test result raises the probability of a benign cyst to a significantly greater extent than the same evidence raised belief that cancer was absent in the traditional version. Krynski and Tenenbaum's reformulation in the mammogram scenario does more than simply introduce causality. It highlights a special case (benign cyst) of the original hypothesis (not cancer) that significantly increases the confirmation value of the evidence.

If we want to make sure that the causal status of informa-

tion plays a role in how subjects use it in probabilistic judgments — above and beyond evidential impact — we need to reliably tease apart the two factors. This is what the experiment we present next accomplishes.

Experiment

Design and materials

Our experiment involves comparing different versions of a task analogous to one proposed by Krynski and Tenenbaum (their “CIA” experiment), itself based on Kahneman and Tversky (1973). It varies across conditions the direction and evidential impact of the evidence, the causal status of the base-rate information, and the causal status of the evidence. This allows us to compare the merits of the Causal Bayesian Reasoner theory, the Evidential Impact Reasoner theory, and a third possibility, prompted by the suggestion that confirmation-theoretic reasoning strategies are triggered by a search for causal explanations of phenomena (Mangiarulo, Pighin, Polonio, & Tentori, 2021), as we detail shortly.

Subjects were told to consider a fictional story about an ancient Mesopotamian village.¹ In this village, a great ritual parade takes place every year. Only one hundred men get to participate each year, so a selection must occur. The selection involves two steps: one meant to decide which 100 men will be invited to participate, the other to decide which subset of 50 out of these 100 men will be proclaimed the “Leaders” of the parade and wear a distinctive attire comprising a red mask and robe. At each step, subjects are 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 made it through the selection.

After having read the story, participants are told that one villager randomly taken from the 100 participants in the parade (the base-rate) wears the distinctive red mask of the Leaders of the parade (the evidence), and asked to estimate the probability that he is an Elder/Youth.

Using a fictional story in an exotic setting allowed us to have maximal control over our stimuli by minimizing the interference of world knowledge. This way we could manipulate the probabilities presented to participants and the different causal and non-causal scenarios used at each step of the story’s selection procedure.

The story was presented in eight different versions, between participants. These eight conditions varied around three dimensions, which we detail presently.

The majority group at each step. In all conditions, the proportions after the two rounds of selection were such that there were 30 Youths and 20 Elders selected as Leaders of the Parade. A rational Bayesian reasoner should thus give an individual taken at random among the Leaders a 60% chance of being a Youth (40% of being an Elder). But the conditions

¹Readers can take the experiment themselves at the following link: <https://web-risc.ens.fr/~smascarenhas/experiments/cc-2/dispatcher.php>.

Condition	PRIORS-POST		EV-POST	
	Youths	Elders	Youths	Elders
Base-Rate	70	30	40	60
Likelihood	43%	66%	75%	33%
Posteriors	30/50 Youths (60%), 20/50 Elders (40%)			

Table 1: Base-rates and likelihoods for each condition.

varied with respect to whether the Youths had this higher posterior probability because of a favorable first round of selection (i.e. base-rate) or because of their success at the second round of selection (likelihood).

In the first set of conditions, the first round of selection (base rate) would have 60 Elders and only 40 Youths partaking in the parade, while the second round would favor the Youths, with 75% (30 out of 40) of them being selected as Leaders, for only 33% (20 out of 60) of the Elders. In what follow, we refer to this condition as EV-POST, because the likelihood of the evidence given the target hypothesis $P(e | h)$ and the posteriors $P(h | e)$ align: the group (Youths) that gives the highest likelihood to the evidence is also the group whose posterior probability is highest. In the second set of conditions, the first round of selection would have 70 Youths, for only 30 Elders partaking in the parade. The second round of selection would favor the Elders, with 66% (20 out of 30) of them being selected as Leaders, for only 43% (30 out of 70) of the Youths. We refer to this condition as PRIORS-POST, because the group (Youths) that has the highest prior probability is also the group whose posterior probability is highest. Across selections, the numbers we gave participants were accompanied by a schema representing the numbers of people from each group that got selected. This visual aid was added so that the difficulties in multiplying the numbers wouldn’t be an obstacle for computing posteriors probabilities, as subjects could simply count on the schema the number of Elders and Youths after the final round of selection. The numbers used at each step, as well as the corresponding proportions, are summarized in Table 1.

The causal status of the base-rate information. The first round of selection is overseen by the Priestess of the village, who chooses which 100 men from the village will participate in the parade.

In one set of conditions, the age class you belong to would have a causal impact on your chance of being selected. The Priestess picks a majority of Elders since, being elderly herself, she has more friends and longer acquaintances among the older members of the village. Alternatively, she selects more Youths because she wants to favor their integration in the adult society of the village.

In another set of conditions, the selection explicitly follows a random process. The Priestess puts together amulets in a chalice, each bearing the name of a villager, and a blindfolded child is asked to pick 100 of them at random. The proportion

of Youths and Elders selected at random merely reflects the demographic base-rate of Youths and Elders in the general population.

The causal status of the updating evidence. The second round of selection brings the 100 men selected for the Parade to an arena outside the village, where they engage in a sacred ritual involving a bull.

In one set of conditions (causal), the 100 participants gather in the arena, each equipped with a single dart with his name written on it. A bull is then released into the arena, and each participant must try to plant his dart in the bull's back. Every participant whose dart stayed stuck on the bull's back after 20 minutes gets chosen as a Leader of the parade. The Elders tend to fare better at this game, because of their superior wisdom and experience with the ritual. Alternatively, the Youths fare better, because of their superior athleticism.

In another set of conditions, (noncausal), the 100 participants are spread across the arena and each is provided with an identical staff with his name written on it. All of them are invited to plant their staffs into the ground and leave the arena, whereupon the bull is released from an undisclosed location. The bull is left free to trample every staff it finds on its way. After 20 minutes, every man whose staff is left standing after the bull's rampage is declared a Leader of the parade.

Procedure

We recruited 400 participants on Prolific from the United States, United Kingdom and Canada. Each participant was randomly assigned to one of the 8 conditions. They were explicitly invited to pay particular attention to the details of the story relative to age.

After participants had gone through the story, they were asked a question regarding the age of a man ("Balthazar") participating in this year's parade, who had the red mask distinctive of the Leaders — meaning he had made it through both selection rounds. Participants responded on a sliding scale which we recorded as a number between 1 and 100, corresponding to their estimate for the probability that the man is a Youth.

They answered a series of questions meant to check their understanding of the story. The first pair of questions asked whether they thought age was a causal factor in each of the selection events. A second set of questions asked them about the proportion of Youths and Elders in the original selection for the parade (querying their base-rate recall), and about the probability that a member selected for the parade was also selected to be a Leader: (a) given that he is an Elder; (b) given that he is a Youth (querying their likelihood recall).

Finally, participants answered two more comprehension questions about non-essential aspects of the story, completed a brief demographic questionnaire, and were redirected to Prolific for payment. The experiment was coded in the jsPsych library (De Leeuw, 2015).

Predictions

Causal Bayesian Reasoner theory. Krynski and Tenenbaum's Causal Bayesian Reasoner theory proposes that subjects only take into account the information that they can connect causally to the target hypothesis (here the age class of an individual). When a piece of information is used, it is used following the process of Bayesian updating. In this theory, subjects are maximally rational when both base-rates and evidence are given a causal interpretation and give similar answers in PRIORS-POST and EV-POST conditions, as the posterior probabilities for each condition are the same. When only one of base-rate or evidence is causal however, they are expected to underuse the non-causal information and give an estimate that is biased towards the hypothesis most favored by the causal one.

Evidential impact theory. On this view, the main distinction relevant here is that between base-rate and updating evidence. Rather than computing posterior probabilities using Bayes' rule, we expect subjects to track the confirmation values of the evidence. This predicts that they should give answers biased in favor of the group that is most represented in the evidence, in a way that is independent from the causality of either the base-rate or the evidence information.

Causal evidential impact theory. A third option amounts to combining the two perspectives. In this theory, subjects do indeed represent information preferentially in a causal format. Yet the process they use to integrate information is not Bayesian, but instead relies on evidential-impact heuristics. Causality has an impact on subjects' reasoning by virtue of the fact that subjects read off a causal model the probabilities with which they apply those heuristics. As such, it does not necessarily make them more rational; instead subjects might engage in even more evidential-impact reasoning when causal structures are transparent.

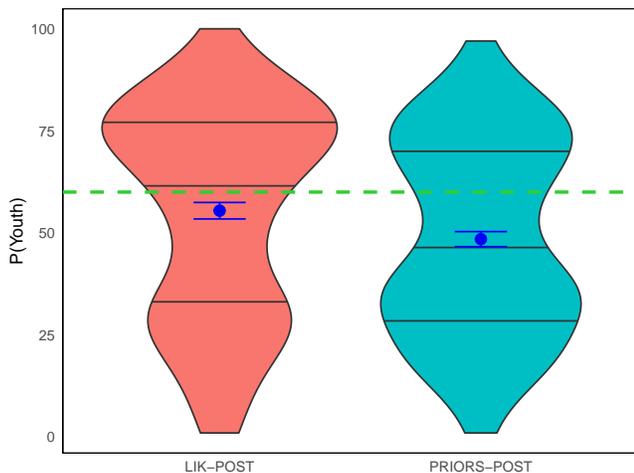
Results

We excluded a total of 48 participants who failed at least one of the comprehension questions, leaving us with a total of 352 participants included in all analyses below.

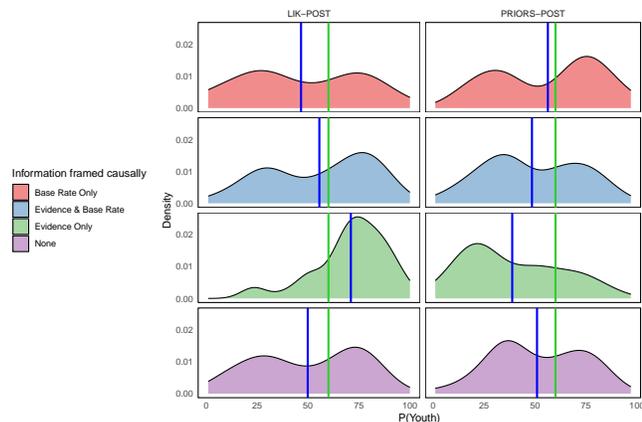
Subjects give more importance to the evidence than to the base-rate. The posterior probability distribution over the relevant hypotheses was the same in all conditions: Balthazar had a 60% probability of being a Youth and a 40% probability of being an Elder. This means that if subjects make a balanced use of the information available, giving equal importance to the base-rate information and to the updating evidence, they should arrive at similar answers in both conditions. Conversely, an imbalance between the two conditions is a sign that one kind of information is given more importance than the other. If subjects estimate that the probability that Balthazar is a Youth is greater in the PRIORS-POST than in the EV-POST condition, this would mean that they are giving more importance to the base-rate information than to the later evidence, as base-rates feature more Youths in

PRIORS-POST, but more Elders in EV-POST. If subjects estimate that the probability that Balthazar is a Youth is lower in the PRIORS-POST than in the EV-POST condition, this would mean that they are giving more importance to the evidence than to the base-rates, as the updating evidence features more Youths in EV-POST, but more Elders in PRIORS-POST.

Results are summarized in Figure 1a. Subjects gave significantly higher $P(\text{Youth})$ rating in the EV-POST conditions. The significance was confirmed by comparing a linear model that used the condition as predictor for subjects' responses to a null model with just a random intercept, which revealed that introducing the condition as a predictor significantly increased model fit ($\chi^2 = 6.4722, p < 0.011$).



(a) Distribution of Youth estimates across conditions. Black lines represent the median and quartiles, blue dots represent the means per condition, along with their standard error. The dotted green line represents the correct posterior value.



(b) Distribution of Youth estimates as a function of the causality of each type of information. Blue lines represent the means, green lines mark the correct posterior value.

Figure 1: Plots of responses by subconditions

The causality of the evidence matters more than that of the base-rate. We looked at whether the causality of both the base-rate and evidence had an effect on the responses

given by participants. Results are summarized in Figure 1b. For almost all sub-conditions, responses tend to follow a bi-modal distribution, one mode corresponding to subjects' thinking that Balthazar is more likely to be a Youth, the other to subjects thinking he is more likely an Elder.

The height of the modes varied depending on the causal status of information. In *unbalanced conditions*, where one piece of information was causal but the other not, subjects tended to favor the hypothesis whose causal role was transparent. This effect is particularly marked in the *evidence only* condition (green curve in Figure 1b), where the distribution takes a uni-modal shape, whose peak favors the group faring better in the second round of selection for each condition (Youths for the EV-POST condition, Elders for the PRIORS-POST condition). It is less marked in the *base-rate only* condition, which mostly just shows an effect for the PRIORS-POST condition, with an increase of the Youth mode. In *balanced conditions*, where either both pieces of information are causal or neither is, we still observe a higher mode for the group that is favored by the second round of selection, in line with the hypothesis that evidential-impact strategies are at play across the board, regardless of the causal status of information.

These results were confirmed by comparing different linear models, using subjects' responses as predictors and progressively removing parameters to test their effect on model fit, reported in Table 2. We took as reference a model M3 using all three factors: condition, causal status of the evidence and base-rates and their interactions to predict subjects' responses. We then looked at the effect of removing selected variables: when removing evidence causality and its interaction with the condition (model M2), the log-likelihood of the model saw a highly significant decrease, and the BIC of the model increased, indicating a poorer fit to the data. When removing base-rate causality however, the log-likelihood decrease was smaller, and the BIC score decreased, suggesting that the extra predictive power provided by the base-rate causality was not proportionate to the penalty incurred for extra degrees of freedom. We found similar results looking at the impact of subtracting either base-rate causality or evidence causality from models that had already subtracted the other variable: while that subtraction was highly significant in the case of evidence causality, it was barely significant in the case of base-rate causality ($p < 0.024$), and it also led to a decrease in the BIC score compared to the model before subtraction.

These results point to the causal status of the base-rate's being a rather poor predictor of subjects' responses in this dataset, especially compared to the causal status of the evidence.

Additionally, we also performed a model comparison of a similar shape, but this time using the absolute difference Δ between subjects' reported posteriors and the correct value as dependent variable, to check what effect the causality of either the evidence or the base-rate had on subjects accuracy. The causality of the evidence proved to have a negative effect

Model	Predictors	LogLik	Df	χ^2	p -value	BIC
M3	$CN * (EVC + BRC)$	-1568.1	7			3177.104
M2	$M3 - EVC - EVC : CN$	↓ -1579.7	5	23.01	1.008e-05 ***	↑ 3188.45
M2'	$M3 - BRC - BRC : CN$	↓ -1573.4	5	10.525	0.005182 **	↓ 3175.965
M1	$M2' - EVC - EVC : CN$	↓ -1583.4	3	19.992	4.558e-05 ***	↑ 3184.293
M1'	$M2 - BRC - BRC : CN$	↓ -1583.4	3	7.5068	0.02344 *	↓ 3184.293

Table 2: Comparison of different models for subjects’ responses, with M3 the full model, where *CN* stands for condition (PRIORS-POST or EV-POST), *EVC* and *BRC* stand respectively for evidence causality and base-rate causality. We use R syntax for the * and : operators. The χ^2 comparisons, associated p -values, and the arrows indicating the evolution of the log-likelihood and BIC metrics, correspond to model comparisons with the model that each was built from, via the subtraction of predictors (i.e. M2 and M2’ are compared to M3, M1 to M2’, M1’ to M2).

on the accuracy (increasing the Δ values), while neither the causality of the base-rate nor its interaction with the causality of the evidence proved to have a significant effect.

Subjects responses estimates aligned with their likelihood estimates, rather than with expected posteriors.

Since we collected subjects’ own estimates for the base-rates $P(\text{Youth})$ and $P(\text{Elder})$, as well as for the likelihood factors $P(\text{Leader} | \text{Youth})$ and $P(\text{Leader} | \text{Elder})$, it was possible to get a measure of the *internal coherence* of subjects’ estimates: how much can a subject’s reported posterior probability $P(\text{Youth} | \text{Leader})$ be predicted from their reported $\frac{P(\text{Youth})}{P(\text{Elder})}$ and $\frac{P(\text{Leader} | \text{Youth})}{P(\text{Leader} | \text{Elder})}$ using Bayes theorem? Across conditions, we compared four predictors for subjects’ individual reported posteriors: (a) their expected posteriors, reconstructed from their individually reported base-rate and likelihood estimates; (b) their reported base-rates; (c) their reported likelihood estimates for $P(\text{Leader} | \text{Youth})$; (d) their reported likelihood ratios $\frac{P(\text{Leader} | \text{Youth})}{P(\text{Leader} | \text{Elder})}$ (recoded into a value between 0 and 1).

The results are in Table 3. The poorest predictor of subjects’ reported posteriors was the base-rate estimate. Both subjects’ likelihood and likelihood-ratio estimates proved to be better predictors of their posterior estimates than subjects’ expected posteriors.

We also investigated whether the causal status of the evidence had an impact on the match between subjects’ likelihood and posterior estimates. To do so, we looked at the impact of the causality of the evidence on the absolute value of the difference Δ_R between reported posteriors and reported likelihoods/likelihood ratios, by comparing to an intercept-only null model. The comparison showed a trend in the expected direction, with causal evidence decreasing Δ_R , (estimate: -3.876), but the difference with the intercept-only model was only barely significant ($\chi^2 = 3.8631$, $p = 0.049$).

Discussion

Our results provide substantial evidence against the notion that base-rate neglect comes solely from subjects’ inability

Models	LogLik	χ^2	p -value	Estimate	p -value
intercept	-1638.3				
base rate	-1637.9	0.8388	0.3597	0.05507	0.361
posteriors	-1614.8	47.062	6.877e-12	0.38798	8.15e-12
likelihood	-1603.2	70.193	< 2.2e-16	0.48372	< 2e-16
likelihood ratio	-1603.2	70.215	< 2.2e-16	0.53289	< 2e-16

Table 3: Comparison of different predictors for subjects’ posteriors. The first row is an intercept-only linear model serving as a null model. Each subsequent row corresponds to a different addition to this intercept model. *Base rate* adds the reported prior probabilities; *posteriors* adds subjects’ expected posterior estimate, as reconstructed from their reported priors and likelihood estimates; *likelihood* the reported probability of the evidence conditional on the hypothesis; and *likelihood ratio* the ratio of the likelihood for the hypothesis and the likelihood for its negation. All values are considered for each subject individually. All χ^2 reports concern likelihood-ratio tests comparing with the intercept-only model. We report the estimate for the added predictor in each model.

to process information when it isn’t causally framed, contra Krynski and Tenenbaum (2007). The influence of updating evidence on subjects’ estimates was more aligned with the perspective of evidential-impact theory than with the Causal Bayesian reasoner approach. However, these findings do not completely conform to the traditional evidential-impact theory, which posits that the evidential impact is paramount, regardless of the causal connections between variables. They indicate that causality does have an effect on subjects’ estimates, albeit not quite in the manner predicted by Krynski and Tenenbaum. Rather than making subjects’ reasoning overall more rational by integrating otherwise neglected base-rate information, it affected subjects’ use of the evidence, making them rely even more heavily on evidential-impact in cases where evidence was causal.

The critical question that emerges is why causality amplifies such tendencies. One plausible explanation is that subjects simply pay more attention to causally framed information. On this view, subjects universally apply evidential-impact heuristics but only effectively use information they can easily track. One important factor for this kind of tracking is a causal framing, perhaps by more sharply capturing participants’ attention than equivalent information without the causal framing. This interpretation is reminiscent of Krynski and Tenenbaum’s two-step information-processing proposal, but it suggests that the second step involves applying nonstandard heuristics rather than engaging in normative Bayesian computation. An alternative interpretation posits that subjects’ reliance on evidential-impact heuristics is driven by an inherent interest in causal relations. Subjects prioritize evidential relations as those are good proxies for causal ones, which is what they are ultimately after, hence they engage more visibly in such heuristics when the presence of a causal connection is provided by the setup.

Acknowledgments

We thank Katya Tentori, Thomas Icard, and the LANG-REASON team at Ecole Normale Supérieure for invaluable feedback. This work was supported by *Agence Nationale de la Recherche* grants ANR-18-CE28-0008 (LANG-REASON; PI: Mascarenhas) and EURIPANR-17-EURE-0012 (PhD funding: Konuk), and by a grant from *Ecole Doctorale Frontières de l'Innovation en Recherche et Education—Programme Bettencourt* (PhD funding: Konuk).

References

- Ali, N., Chater, N., & Oaksford, M. (2011). The mental representation of causal conditional reasoning: Mental models or causal models. *Cognition, 119*, 403–418. doi: 10.1016/j.cognition.2011.02.005
- Bar-Hillel, M. (1983). The base rate fallacy controversy. *Advances in Psychology, 16*, 39–61. doi: 10.1016/S0166-4115(08)62193-7
- Crupi, V., Fitelson, B., & Tentori, K. (2008). Probability, confirmation, and the conjunction fallacy. *Thinking & Reasoning, 14*(2), 182–199. doi: 10.1080/13546780701643406
- De Leeuw, J. R. (2015). jspsych: A javascript library for creating behavioral experiments in a web browser. *Behavior research methods, 47*(1), 1–12. doi: 10.3758/s13428-014-0458-y
- 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). Cambridge University Press.
- Fitelson, B. (1999). The plurality of Bayesian measures of confirmation and the problem of measure sensitivity. *Philosophy of Science, 66*(Supplement: Proceedings of the 1998 Biennial Meetings of the Philosophy of Science Association), 362–378. Retrieved from <https://www.jstor.org/stable/188784>
- Fitelson, B. (2007). Likelihoodism, Bayesianism, and relational confirmation. *Synthese, 156*(3), 473–489. doi: 10.1007/s11229-006-9134-9
- Gigerenzer, G., & Hoffrage, U. (1995). How to improve Bayesian reasoning without instruction: frequency formats. *Psychological Review, 102*(4), 684–704. doi: 10.1037/0033-295X.102.4.684
- Kahneman, D., & Tversky, A. (1973). On the psychology of prediction. *Psychological Review, 80*(4), 237–251. doi: 10.1037/h0034747
- Krynski, T. R., & Tenenbaum, J. B. (2007). The role of causality in judgment under uncertainty. *Journal of Experimental Psychology: General, 136*(3), 430–450. doi: 10.1037/0096-3445.136.3.430
- Mangiarulo, M., Pighin, S., Polonio, L., & Tentori, K. (2021). The effect of evidential impact on perceptual probabilistic judgments. *Cognitive Science*. doi: 10.1111/cogs.12919
- Tenenbaum, J. B., & Griffiths, T. L. (2001). The rational basis of representativeness. In J. D. Moore & K. Stenning

(Eds.), *Proceedings of the 23rd annual conference of the cognitive science society* (pp. 1036–1042).

- Tentori, K., Crupi, V., & Russo, S. (2013). On the determinants of the conjunction fallacy: probability versus inductive confirmation. *Journal of Experimental Psychology: General, 142*(1), 235–255. doi: 10.1037/a0028770