Title
How Well Do Humans Learn Conditional Probabilities?

Journal
Proceedings of the Annual Meeting of the Cognitive Science Society, 45(45)

Authors
Strößner, Corina
Hahn, Ulrike

Publication Date
2023

Peer reviewed

How Well Do Humans Learn Conditional Probabilities?

Corina Strößner (corinastroessner@posteo.de)
Department of Psychological Sciences, Birkbeck College, University of London,
Malet Street, London WC1E 7HX, UK

Ulrike Hahn (u.hahn@bbk.ac.uk)
Department of Psychological Sciences, Birkbeck College, University of London,
Malet Street, London WC1E 7HX, UK

Abstract

Although there is a great deal of interest in conditional probabilities in Bayesian cognitive science, there is still little understanding of how well human agents can learn them. This paper addresses the issue by theoretical and experimental means. In the theoretical part, we distinguish between cases of vacuous learning, where learned probabilistic information is not new, and belief revision is unwarranted, from cases with truly new information. In the experimental part, we investigate how well participants can distinguish these cases and how well they respect the probabilistic norms, thus adding new insights to the long-standing question of the extent to which the human mind is adapted to probabilistic norms.

Keywords: Bayesianism; Conditional probabilities; Belief revision; Coherence

How probabilistic is the human mind?

Are humans good at probabilistic thinking? Before Tversky and Kahneman (1983) famously demonstrated flaws of human reasoning in statistical inferences, many psychologists were quite optimistic about humans' ability to reason with probabilities. In a review of the debate at the time, tellingly named 'Man as an intuitive statistician', Peterson and Beach (1967, p. 42–43) conclude: 'Experiments that have compared human inferences with those of statistical man [the probabilistic norm] show that the normative model provides a good first approximation for a psychological theory of inference'. The extensive work of the time, focused mainly on belief revision in light of new evidence, noted some biases, but no general misalignment. The fallacies noted by Tversky and Kahneman (1983), however, severely conflict with probabilistic norms. Their work shaped the widely influential image of human reasoning as not even close to living up to the rules of probability theory.[1]

With the rise of Bayesian cognitive science, probabilities again became a frequently applied framework for understanding the human mind, for instance in the study of reasoning: Oaksford and Chater (2007) argue that probability theory rather than logic provides the standard of human reasoning, ushering in 'the new paradigm' in the psychology of reasoning.

A major role in the new paradigm is played by *conditional probability*.[2] It has been studied as a potential candidate for providing semantics (or, at least, a quantity implied by the semantics) for conditionals (Over & Cruz, 2017; Krzyżanowska, Collins, & Hahn, 2021; Skovgaard-Olsen, Singmann, & Klauer, 2016). In the context of this research, it has been found that human intuitions on conditionals are better explained in terms of probabilistic coherence than classical logic (Pfeifer, 2021; Cruz, Baratgin, Oaksford, & Over, 2015). While our work is closely related to this research, our main interest in this paper is not on the study of conditionals, their semantics or their role in reasoning, but on the learning of conditional probabilities, often expressed by a conditional statement such as 'If A, it is likely that C'.

Conditional probabilities arguably provide the backbone for Bayesianism, and with that Bayesian cognitive science. They figure centrally in Bayes's theorem and determine the impact of learning new evidence. Despite this, research examining learning from a probabilistic perspective has focused on factual evidence.[3] For a more complete picture, it is important to understand how conditional probabilities are learned.

From a mathematical standpoint, the learning of new conditional information is quite problematic. There is no unique probability distribution Q, that results from a prior distribution P upon learning a conditional probability $Q(C|A) = q$. This problem has prompted intensive debate within formal epistemology (Fraassen, 1981; Douven & Romeijn, 2011; Eva, Hartmann, & Rad, 2019). We will not go into the details of this discussion here. For the present paper, it suffices to note that a new probability distribution can be determined by stipulating further parameters that should remain fixed. The resulting normative model and its assumptions are explained in the next section. This model provides the basis of our experimental work that is described in the third section of this paper.

A simple normative model

The probability distribution for two atomic propositions A and C can be defined by determining $P(A)$, $P(C|A)$ and $P(C|\neg A)$, which allow calculating the probability of all potential states as follows:

[1] For a more comprehensive illustration of the bias debate, see Hahn and Harris (2014).

[2] The conditional probability of C given A is usually defined as $P(C|A) =_{def} \frac{P(A\&C)}{P(A)}$. Some scholars, e.g. Pfeifer (2021), take it as a foundational undefined notion.

[3] For a notable exceptions see Zhao, Shah, and Osherson (2009)

In M. Goldwater, F. K. Anggoro, B. K. Hayes, & D. C. Ong (Eds.), *Proceedings of the 45th Annual Conference of the Cognitive Science Society*.

$$
\begin{aligned}
P(A\&C) &= P(A)*P(C|A)\\
P(A\&\neg C) &= P(A)-P(A)*P(C|A)\\
P(\neg A\&C) &= (1-P(A))*P(C|\neg A)\\
P(\neg A\&\neg C) &= 1-P(A)-(1-P(A))*P(C|\neg A)
\end{aligned}
$$

If the agent learns a conditional probability $Q(C|A)$, a new probability distribution can be determined if we assume $Q(A) = P(A)$ and $Q(C|\neg A) = P(C|\neg A)$. Such an update procedure corresponds to a minimisation of the inverse Kullback-Leibler divergence (IKL), which is advocated to be normatively privileged for such a context within the philosophical literature (Douven & Romeijn, 2011; Eva et al., 2019). We will not restate the arguments in favour of it, but simply characterise the two resulting constraints $Q(A) = P(A)$ and $Q(C|\neg A) = P(C|\neg A)$ intuitively:

No influence beyond the condition: $Q(C|\neg A) = P(C|\neg A)$ The agent should not change her belief about events or cases that are not covered by the conditional probability. For instance, assume you learn that Ahmed will likely go by bike if he visits a library. This information should not influence your beliefs about whether Ahmed takes a bike if he goes somewhere else. This constraint maintains continuity with factual probabilistic learning (i.e., by Jeffrey conditionalisation), in which conditional probabilities are not changed (Eva et al., 2019). It also keeps continuity with the way one generally applies conditional probabilities. If understood as conditional probability, conditionals are generally void if the condition is not fulfilled. This also corresponds to betting behaviour: If I bet on Ahmed taking the bike given he goes to the library but he never goes there, then the bet is cancelled and I neither win nor lose money (de Finetti, Galavotti, & Hosni, 2008; Politzer, Over, & Baratgin, 2010). Accordingly it is highly plausible that the learning of a conditional probability should not influence states in which it is void.

Nevertheless people might *infer* additional information from this learning experience, for instance that Ahmed generally likes biking. Moreover, some natural language conditionals have a bi-conditional 'if and only if' interpretation due to pragmatic factors. For instance, '*If* Sam studies really hard, it is likely that she passes the exam' may imply that the consequence is likely *only* if the antecedent is fulfilled but unlikely otherwise. In these cases the agent does not just learn a single conditional probability, but infers additional information beyond it.

Fixity of the antecedent likelihood: $Q(A) = P(A)$ The second constraint is the fixity of the antecedent's probability. For many conditionals, the intuitive appeal of this constraint is beyond doubt. For example, the sentence 'Jasmine likely takes a bus if it rains' is not informative about whether it will rain or not. Pragmatic factors make such a statement useful and informative only in as far as there is a real chance, but no certainty, about whether it will rain (Krzyżanowska et al., 2021). The utterance would be pointless to make if the antecedent is false and needlessly complicated if it is (almost) certainly true.

	Case 1A	**Case 1B**	**Case 2A**		**Case 2B**	
	$P=Q$	$P=Q$	P	Q	P	Q
A	0.5	0.5	0.5	0.5	0.5	0.5
$C\|A$	0.9	0.1	0.1	0.9	0.9	0.1
$C\|\neg A$	0.9	0.1	0.1	0.1	0.9	0.9
$A\&C$	0.45	0.05	0.45	0.05	0.05	0.45
$A\&\neg C$	0.05	0.45	0.05	0.45	0.45	0.05
$\neg A\&C$	0.45	0.05	0.45	0.45	0.05	0.05
$\neg A\&\neg C$	0.05	0.45	0.05	0.05	0.45	0.45
C	0.9	0.1	0.9	0.5	0.1	0.5

Table 1: Exemplary prior (P) and posterior (Q) probability distributions. In case 1A/1B, the learned probabilistic information is already implicitly believed, and no revision is needed. In case 2A/2B, the received information is inconsistent with the prior and a change is necessary. A revision that keeps the probability of A and of A given $\neg C$ fixed leads to a revision of the probability of C.

Beyond such pragmatic constraints, it seems reasonable to make revisions to the antecedent where we hesitate to change our belief about C. Assume you are learning that if Alan committed a fraud, he likely did it with Cathy. According to the constraint, you will not change your belief about Alan's guilt, but raise the likelihood that Cathy was involved. However, if you strongly disbelieve that Cathy was involved and you are entrenched to this belief, you may alternatively decrease the likelihood that Alan committed a fraud (Eva et al., 2019; Günther & Trpin, 2022).[4]

Let us now look at some exemplary probabilities that will play a central role in our experimental work (see also Table 1):

Case 1A: The agent is neutral about A, but believes C to be the case, independently of whether A is true or false. She then learns that C is likely given A. A corresponding prior probability is $P(A) = 0.5$, $P(C|A) = P(C|\neg A) = 0.9$ and the learned constraint is $Q(C|A) = 0.9$.
Case 1B: The agent is again neutral about A, but disbelieves C. She learns that C is unlikely given A. A corresponding prior probability is $P(A) = 0.5$, $P(C|A) = P(C|\neg A) = 0.1$ and the learned constraint is $Q(C|A) = 0.1$.
Case 2A: The agent is neutral about A, but believes C to be the case, independently of whether A. She then learns that C is unlikely given A. A corresponding prior probability is $P(A) = 0.5$, $P(C|A) = P(C|\neg A) = 0.9$ and $Q(C|A) = 0.1$ is the learned constraint.
Case 2B: The agent is again neutral about A and disbelieves C. She learns that C is likely given A. Here, the

[4] This updating strategy is reminiscent of modus tollens: 'If A, then C. Not C. Thus not A'. Also note that the updating by a material conditional 'C or not A' involves a decrease of $P(A)$, but this has been viewed as counter-intuitive (Douven, 2012).

prior probability is $P(A) = 0.5$, $P(C|A) = P(C|\neg A) = 0.1$ and the learned constraint is $Q(C|A) = 0.9$.

Cases 1A and 1B are instances of *vacuous* learning. The conditional probability in question was implicitly already part of the agent's prior belief. The normative model thus recommends keeping all probabilities as they are. In cases 2A and 2B, however, the probability distribution has to be changed because $Q(C|A) \neq P(C|A)$. The posterior Q, given that the above two constraints are met, can be seen in table 1. One should no longer believe or disbelieve C but come to a more neural attitude, that is, $Q(C) = 0.5$. Unlike for 1A/1B, this result depends on the above constraints. An alternative posterior Q' for case 2A, which violates the second constraint, would be $Q'(A) = 0.1$, $Q'(C|A) = 0.1$, $Q'(C|\neg A) = 0.9$, with $Q'(C) = 0.82$. A violation of the first constraint, namely that $P(C|\neg A)$ is not changed, can only exert its influence numerically if the agent interprets the conditional probability as an unconditional probability of C.

Experimental Study

In consideration of the normative model, we address the following empirical questions:

1. Do people's intuitions respect the difference between vacuous information in cases 1A/1B and new information in case 2A/2B?
2. Do people's intuitions respect the symmetry between the cases 1A/1B and 2A/2B? Do they deal with the concept of unlikely as they do with the concept of likely?
3. To which extent do their judgements depend on explicit conservativity constraints such as the fixity of antecedent probability?

While our study is obviously descriptive, it involves a comparison to a normative framework (probability theory). The material asks subjects how the individuals in a presented story *should* change their belief considering new information, that is, their intuitions on normative matters. The methodology resembles third-person evaluation, which has been extensively applied in research on argument strength (Hahn & Oaksford, 2007; Oaksford, 2014). Regarding concerns against working on the interface of descriptive and normative models (Elqayam & Evans, 2011), it should be noted that the study does not involve fallacious ought-is inferences. We do not debate the status of probability theory as a normatively valid account of learning, which has been done elsewhere (Hájek, 2008; Pettigrew, 2019). We also do not *infer* descriptive facts about human intuitions from the expectation of the normative model, but empirically test them.

The two experiments reported here are exploratory and aimed at a first understanding regarding the above questions.[5]

Our materials embed the probabilistic information in stories where individuals hold some prior probabilistic beliefs and are then confronted with new conditional probabilities. For instance, in one story, Pierre wonders whether his friend Maria, who went either on a boat or bike trip (both are equally likely) will be back by dinner. From reading tourist information, he infers new information about whether she will likely be back if she is on a bike trip. Across conditions, the stories manipulated the probabilistic information, e.g., the prior belief and the learned conditional probability. Pierre's story thus varied as follows:

1A He [initially] thinks she will probably be back for dinner, regardless of what she does. [...] he concludes that Maria will probably be back if she is on a bike tour.

1B He [initially] thinks she will probably not be back for dinner, regardless of what she does. [...] he concludes that Maria will probably not be back if she is on a bike tour.

2A He [initially] thinks she will probably not be back for dinner, regardless of what she does. [...] he concludes that Maria will probably be back if she is on a bike tour.

2B He [initially] thinks she will probably be back for dinner, regardless of what she does. [...] he concludes that Maria will probably not be back if she is on a bike tour.

Condition was varied between subjects, that is, each participant was assigned to one experimental group and completed all target items in this condition. After each story, subjects were asked several questions. This included questions of understanding: choosing the learned conditional from four alternatives, and determining the prior belief in C (*Maria is back on time*). Participants rated whether the individuals should revise their belief in C and in which direction, using a five-point scale, ranging from 'much less likely' to 'much more likely'. Finally, they were asked what the agent's new belief about C should be on a five point scale from 'very unlikely' to 'very likely'. Finally, we asked them how confident they were in their answers for each of these questions.

The experiment was run in two versions. In Exp. 1, it was explicitly stipulated that the protagonists respect the conservativity constraints above. In Pierre's story, it was said: '[...] he still considers it as just as likely that she has taken a boat tour instead. He has not found any information about paddle boat tours. He accordingly still suspects her to be *on time / too late* for dinner if she is on a boat tour.' In Exp. 2, no such explicit constraints were given.

Experiment 1

Procedure After informed consent and a warm-up task, participants completed four target items and four filler items that were used to break the repetitious pattern of the target items. Finally, participants provided information about prior

[5] The study was preregistered (`https://osf.io/uy7zv`). Further data of the project as well as documents with the complete analysis can be found at the project page (`https://osf.io/aqhxw/?view_only=925e76760e5644ed8c13800012e314e8`) with the following coding of conditions: 1A=A, 1B=D, 2A=B, 2B=C.

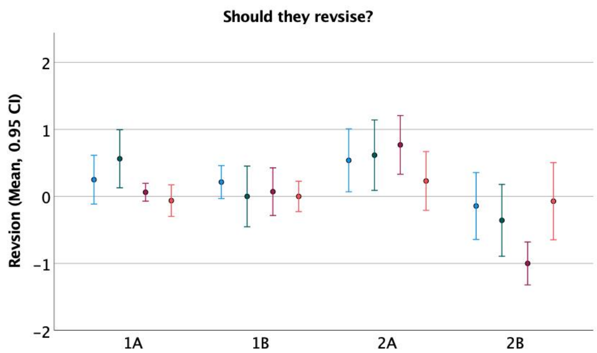

	Mean	SE	0.95 CI
1A	0.20	0.11	[-0.03, 0.44]
1B	0.07	0.12	[-0.17, 0.31]
2A	0.54	0.12	[0.29, 0.79]
2B	-0.39	0.12	[-0.64, -0.15]

Figure 1: Revision rating in Exp. 1 with -2 meaning 'much less likely', 2 'much more likely' and the value 0 meaning 'no change': The graphic shows the mean and 0.95 confidence intervals for all items. The table below displays least square means over all items. From a normative perspective, it is expected that the rating is close to 0 in conditions 1A/1B, because they are cases of vacuous learning. The rating should be higher in 2A and lower in 2B: A revision in the direction of the learned conditional is required. The expectation is largely met. The least square means do not diverge significantly from 0 in 1A/1B but they do in 2A/2B, even though the graph shows notable difference across items.

knowledge of Bayesianism as well as demographic information (optional). The questionnaire was programmed and hosted on SoSciSurvey (`https://www.soscisurvey.de`).

Participants Participants were recruited and rewarded via Prolific. Overall, 117 persons participated in the experiment. Their age was between 19 and 64 (mean 32, SD: 12). 38.5% of them identified as female, all others as male.

Before running the experiment, it was planned that data will be excluded if participants failed more than two of the overall eight reading checks (two per item) or if they claimed to have prior knowledge about Bayesianism (six participants). However, many participants failed more than two reading checks. Overall, only 57 data sets were thus eligible according to the above-mentioned criteria. Since the high selectivity may distort the results, we performed the analysis also with the full data set and point out notable differences in the discussion.

Results The main variables of interest were the judgement about whether agents in the stories should revise their belief about *C*, and in what direction (revision rating), and the final belief in *C* (the posterior).[6]

Figure 1 displays answers to the revision question of whether the agents should revise the rating and, if so, in which direction.[7] The rating tendentiously follows probabilistic expectations, though with notable distortions. The mean rating in condition 1A was slightly higher than the normatively expected value 0 (no change).[8] This deviation can be interpreted as the result of a relevance bias: If agents understand the conditional 'If *A*, then it is likely that *C*' as expressing a positive relevance of *A* for *C* rather than just a restatement of the belief that *C* is likely anyway, it seems intuitive to further raise belief in *C*, even though not very much. Notably, such a relevance bias was not observed in condition 1B.

In condition 2A, it is expected that the belief is revised, namely *C* becomes more plausible. The pattern is indeed seen in the rating of the items, though to a lower extent than expected. In condition 2B, the rating should be clearly below 0, e.g., *C* should be revised and become less likely. The cleaned data set satisfies this norm, even though only item 3 is significantly lower. Moreover, the analysis of all data revealed a different pattern, where subjects failed to see that *C* becomes less likely (-0.02 [-0.25, 0.21], SE: 0.10).

To summarise, the revision ratings of our participants followed the normative probabilistic model to some extent. However, there was an effect of polarity: high likelihoods had a stronger impact relative to low likelihoods. In 1A subjects wrongly judged vacuous high likelihood information to be influential but no such effect was observed for vacuous low likelihood information in 1B. Furthermore, participants failed to consider new low likelihood information as relevant in 2B. In contrast, participants in 2A had no problem to view high likelihood information as relevant.

The probabilistic expectation for the posterior likelihood was that the rating is close to 0 for case 2A and 2B, where the prior and the learned conditional 'cancel each other out', but higher in 1A and lower in 1B. The overview of the results in figure 2 shows that these expectations were largely met. According to the mixed effect model, the rating is significantly above 0 in condition 1A and significantly below 0 in condition 1B. In contrast, there is no significant deviation from 0 in condition 2A and condition 2B. Note, however, that a model with all data revealed a small but significant positive deviation from 0 in condition 2A and 2B.[9]

In comparison to the revision rating, the answers for the posterior were even more consistent with the normative model. We found less strong effects of polarity than in the revision rating, but the subjects still tended to be more influenced by high likelihood. That is, in case 1A/1B, the subjects seemed to deviate more strongly from the neutral value

[6] We do not report the results of the confidence rating, which had no interesting variations over the conditions and items.

[7] All analyses are based on a mixed effect model run on R with the packages lme4, lmertest, and performance (R Core Team, 2018; Bates, Mächler, Bolker, & Walker, 2015; Kuznetsova, Brockhoff, & Christensen, 2017; Lüdecke, Makowski, Waggoner, & Patil, 2020) with item and subject entered with random intercepts.

[8] In the analysis of all data, the deviation from 0 became significant (mean with 0.95 Confidence interval: 0.28 [0.03, 0.53], SE:0.12).

[9] Estimates for 2A: 0.31 [0.11, 0.51], SE: 0.10, for 2B: 0.23 [0.03, 0.42], SE: 0.09.

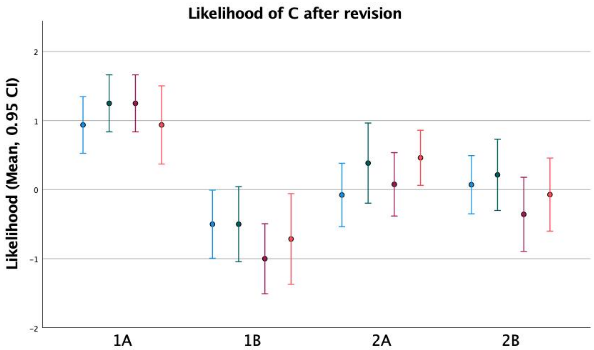

	Mean	SE	0.95 CI
1A	1.09	0.12	[0.84, 1.35]
1B	-0.68	0.13	[-0.95, -0.41]
2A	0.21	0.14	[-0.07, 0.49]
2B	-0.04	0.13	[-0.31, 0.23]

Figure 2: Posterior rating in Exp. 1 with -2 meaning 'very unlikely', 2 'very likely' and the value 0 meaning 'neither likely nor unlikely': The graphic shows the mean and 0.95 confidence intervals for all items. The table below displays least square means over all items. From a normative perspective, it is expected that the rating is close to 0 in conditions 2A/2B, because the prior and learned probability 'cancel each other out'. The rating should be high in 1A and low in 1B, where there should be no change in comparison to the prior belief. This expectation is largely met: the mean values in 2A/2B are much closer to 0 than in 1A/1B

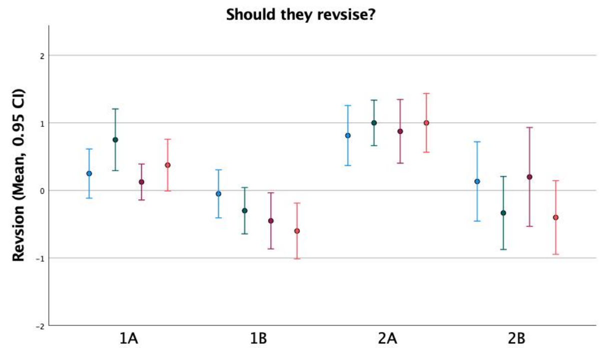

	Mean	SE	0.95 CI
1A	0.38	0.14	[0.09, 0.66]
1B	-0.35	0.13	[-0.61, -0.09]
2A	0.92	0.14	[0.63, 1.21]
2B	-0.10	0.15	[-0.40, 0.20]

Figure 3: Revision rating of Exp. 2, where -2 means 'much less likely', 2 'much more likely' and the value 0 'no change': The graphic shows the mean and 0.95 confidence intervals for all items. The table below displays least square means over all items. As in Exp. 1, the normative model predicts 0 for 1A/1B but much higher values for 2A and much lower ones for 2B. The rating diverged from this expectation in 1A/1B. In 2A the rating fits to the expectations while it was not met in 2B, where subjects failed to appreciate the impact of low likelihood information on revision.

0 if they had to process high likelihoods (1A) than when they were in the low likelihood group (1B). In true learning scenarios, subjects were also more strongly influenced by the new learned conditional probability in 2A ('likely') than 2B ('unlikely').

In a nutshell, we found a partial compliance with the normative model and quite notable polarity effects, especially in the revision rating. In the posterior rating, the compliance with the normative model was even better, and polarity effects were much weaker. This addresses the first two questions that we extracted from the normative model above. To evaluate the final question, namely the role of explicit conservative constraints in this experiment, we run a study that did not include these constraints for comparison.

Experiment 2

Procedure The procedure and material were the same as in Exp. 1, but the items had no explicit conservative constraint stating that antecedent likelihood as well as the likelihood given a false antecedent remain fixed.

Participants Participants were recruited and rewarded via Prolific. Overall, 101 persons participated in the experiment. Their age was between 18 and 68 (mean 31, SD: 10). One person identified as non-binary and 38.6% as female. Three participants were removed because they claimed to have prior knowledge of Bayesianism. We again planned to remove data of participants, who failed more than two of the eight reading checks. This was the case for 32 participants, which is still a large number, but remarkably less than in Exp. 1. The answers of 67 participants were in the cleaned data set, which we used in the analysis. As before, we also analysed the full data set and point out notable differences in the discussion of the results.

Results The main focus of the second experiment is the comparison with the results from Exp. 1, but we also compare the results to expectations from the normative model.

Figure 3 shows the results of the mixed effect model from the revision rating. For comparison with Exp. 1, we also ran an ANOVA of both experiments with item as (fixed) within-subjects factor and condition and experiment as between-subjects factors. This revealed a significant effect of condition ($F(3,116) = 23.0$, $p < 0.001$) and a significant interaction of condition and experiment ($F(3,116) = 4.1$, $p = 0.008$) as well as a significant interaction of item with condition and experiment ($F(9,348) = 2.73$, $p = 0.004$). The experiment type (with or without explicit conservativity) thus influenced the rating in the groups differently. The normatively unexpected positive rating for 1A is even slightly higher. The rating in 1B is considerably lower than in the previous experiment, and now deviates significantly from the normative model. These results are hard to explain in terms of a violation of the two above discussed conservativity constraints. This result fits more to a relevance effect (Skovgaard-Olsen

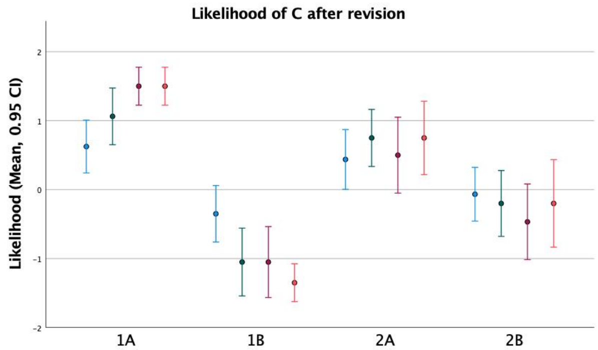

	Mean	SE	0.95 CI
1A	1.17	0.13	[0.92, 1.43]
1B	-0.95	0.11	[-1.18, -0.72]
2A	0.61	0.13	[0.35, 0.87]
2B	-0.23	0.13	[-0.50, 0.03]

Figure 4: Posterior likelihood rating of Exp. 2, where -2 means ‘very unlikely’, 2 ‘very likely’ and the value 0 ‘neither likely nor unlikely’: The graphic shows the mean and 0.95 confidence intervals for all items. The table below displays least square means over all items. As in Exp. 1, the normative model predicts values above 0 for condition 1A and below 0 for 1B. In comparison, values in 2A and 2B are expected to be closer to 0. This expectation is largely met, with the exception of 2A, where subjects weighted the new high likelihood information too strong.

et al., 2016), that is, the reading of conditionals as implying a relevant connection between the likelihood of the consequent and the antecedent.[10] The explicit conservativity constraints in Exp. 1 may have moderated the effect because the additional information led the focus away from the learned conditional probability.

Regarding 2A/2B, where a revision is expected, the positive effect of learning an unexpected high conditional probability in condition 2A is slightly stronger than in Exp. 1. In condition 2B, there is no clear tendency: the effect of learning a low conditional probability seems slightly less strong, but the inclusion of the full data does not confirm this trend (Exp. 1: -0.02 [-0.25, 0.20], SE: 0.10; Exp. 2: -0.13 [-0,36 0,10], SE: 0.12).

The posterior likelihood rating in this second experiment was quite similar to the rating in Exp. 1. Figure 4 shows an overview of the results. The rating was almost identical for 1A, slightly lower in 1B and 2B, and slightly higher in condition 2A. The ANOVA with item as repeated measure, and condition and experiment as between subject factors revealed a significant interaction between condition and item ($F(9,348) = 3.15$, $p = 0.001$), a highly significant effect of condition ($F(3,116) = 90.59$, $p < 0.001$), but only a mild interaction between condition and experiment ($F(3,116) = 2.85$, $p = 0.04$). A series of t-tests on the rating of all items in all conditions showed a significant deviation only for item 4 in condition 1B ($-0.64[-1.24, -0.03]$, $t(32) = -2.13$, $p = 0.041$). This value is not corrected for multiple comparisons, and thus within the range of what could be expected coincidentally. However, it is noteworthy that the ratings in Exp. 2, even though barely different, tend to deviate from Exp. 1 in the direction of the learned conditional probability, which is high in condition 1A and 2A, but low in condition 1B and 2B. This effect can again be explained pragmatically. Without the conservativity clause, which followed the conditional learning in Exp. 1, the learned conditional probability became more salient. This deviation cannot be explained in terms of a violation of the conservativity constraints (for instance, lowering $Q(A)$ rather than lowering $Q(C)$ in condition 2B), which would moderate the influence of the new learned conditional on the posterior probability of C rather than increasing it.

Taken together, the results of the second experiment revealed no violation of conservative constraints. There were notable variations in the revision rating that indicate a decisive pragmatic role of the explicit conservativity constraints in Exp. 1. Thus, the moderate agreement with the normative model in the first experiment must be evaluated critically. However, the assessment of the posterior probability corresponded to the expectations from the normative model in both experiments. This makes one suspect that people are much better in updating their belief than we would expect from their own revision rating.

Conclusion

We described two exploratory experiments examining an under-researched question in the otherwise extensive research literature on human probabilistic reasoning: how people revise related probabilities upon learning a (new) conditional probability. This sheds much needed light on this cornerstone of reasoning with probabilities.

In our studies, participants did not find it straightforward to distinguish cases of vacuous ’learning’ and cases where updates were genuinely required in order to avoid inconsistency. However, the posterior probabilities proved surprisingly close to the normative standard imposed by probability theory supplemented by conservativity constraints. To the extent that there were deviations, these were not interpretable as violations of those conservativity constraints. This suggests that minimisation of IKL—the best candidate for a normative model of learning conditional probabilities (Douven, 2012; Eva et al., 2019)— provides an intuitive approach for capturing belief revision upon receipt of conditional probabilities, and may have some cognitive plausibility as well.

[10]This effect holds for ‘If-then’ conditionals but not for ‘Even if-then’ conditionals.

Acknowledgement

This work was carried out in the UK-German Collaborative Research Project 'Descriptive Accounts in the Philosophy and Psychology of Reasoning and Argumentation: Tension or Productive Interplay?' (Ulrike Hahn and Stephan Hartmann), kindly funded by the Arts and Humanities Research Council (AHRC) of the UK Research and Innovation (UKRI) and the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation). We thank our team members Marko Tešić, Borut Trpin, Kirsty Phillips, Anita Keshmirian, Ammar Jalalimanesh, and Rafael Fuchs for discussing the material and the results with us. Furthermore, we thank Annika Schuster, Joanna Kuchacz and Maria Sekatskaya for pretesting the experiment. Finally, we are grateful to three anonymous reviewers who provided helpful comments for improving the paper.

References

Bates, D., Mächler, M., Bolker, B., & Walker, S. (2015). Fitting linear mixed-effects models using lme4. *Journal of Statistical Software*, *67*(1), 1–48. doi: 10.18637/jss.v067.i01

Cruz, N., Baratgin, J., Oaksford, M., & Over, D. E. (2015). Bayesian reasoning with ifs and ands and ors. *Frontiers in Psychology*, *6*. doi: 10.3389/fpsyg.2015.00192

de Finetti, B., Galavotti, M. C., & Hosni, H. (2008). *Philosophical lectures on probability* (A. Mura, Ed.). Dodrecht, NL: Springer.

Douven, I. (2012). Learning conditional information. *Mind & Language*, *27*(3), 239-263. doi: 10.1111/j.1468-0017.2012.01443.x

Douven, I., & Romeijn, J.-W. (2011). A new resolution of the Judy Benjamin problem. *Mind*, *120*(479), 637–670. doi: 10.1093/mind/fzr051

Elqayam, S., & Evans, J. S. B. T. (2011). Subtracting 'ought' from 'is': Descriptivism versus normativism in the study of human thinking. *Behavioral and Brain Sciences*, *34*(5), 251–252. doi: 10.1017/s0140525x1100001x

Eva, B., Hartmann, S., & Rad, S. R. (2019). Learning from conditionals. *Mind*, *129*(514), 461-508. doi: 10.1093/mind/fzz025

Fraassen, B. C. V. (1981). A problem for relative information minimizers in probability kinematics. *The British Journal for the Philosophy of Science*, *32*(4), 375–379. doi: 10.1093/bjps/32.4.375

Günther, M., & Trpin, B. (2022). Bayesians still don't learn from conditionals. *Acta Analytica*. doi: 10.1007/s12136-022-00527-y

Hahn, U., & Harris, A. J. (2014). What does it mean to be biased. In *Psychology of learning and motivation* (pp. 41–102). doi: 10.1016/b978-0-12-800283-4.00002-2

Hahn, U., & Oaksford, M. (2007). The rationality of informal argumentation: A Bayesian approach to reasoning fallacies. *Psychological Review*, *114*(3), 704–732. doi: 10.1037/0033-295X.114.3.704

Hájek, A. (2008). Arguments for, or against, probabilism? *The British Journal for the Philosophy of Science*, *59*(4), 793–819. doi: 10.1007/978-1-4020-9198-8_9

Krzyżanowska, K., Collins, P. J., & Hahn, U. (2021). True clauses and false connections. *Journal of Memory and Language*, *121*, 104252. doi: 10.1016/j.jml.2021.104252

Kuznetsova, A., Brockhoff, P. B., & Christensen, R. H. B. (2017). lmerTest package: Tests in linear mixed effects models. *Journal of Statistical Software*, *82*(13), 1–26. doi: 10.18637/jss.v082.i13

Lüdecke, D., Makowski, D., Waggoner, P., & Patil, I. (2020). Performance: Assessment of regression models performance. *CRAN*. Retrieved from `https://easystats.github.io/performance` (R package) doi: 10.5281/zenodo.3952174

Oaksford, M. (2014). Normativity, interpretation, and Bayesian models. *Frontiers in Psychology*, *5*. doi: 10.3389/fpsyg.2014.00332

Oaksford, M., & Chater, N. (2007). *Bayesian rationality: The probabilistic approach to human reasoning*. Oxford, UK: Oxford University Press.

Over, D. E., & Cruz, N. (2017). Probabilistic accounts of conditional reasoning. In *The Routledge international handbook of thinking and reasoning* (p. 434-450). London, UK: Routledge.

Peterson, C. R., & Beach, L. R. (1967). Man as an intuitive statistician. *Psychological Bulletin*, *68*(1), 29–46. doi: 10.1037/h0024722

Pettigrew, R. (2019). Epistemic utility arguments for probabilism. In E. N. Zalta (Ed.), *The Stanford encyclopedia of philosophy.* Metaphysics Research Lab, Stanford University.

Pfeifer, N. (2021). Probability logic. In M. Knauff & W. Spohn (Eds.), *The handbook of rationality.* MIT Press.

Politzer, G., Over, D. E., & Baratgin, J. (2010). Betting on conditionals. *Thinking & Reasoning*, *16*(3), 172-197. doi: 10.1080/13546783.2010.504581

R Core Team. (2018). R: A language and environment for statistical computing [Computer software manual]. Vienna, Austria. Retrieved from `https://www.R-project.org/`

Skovgaard-Olsen, N., Singmann, H., & Klauer, K. C. (2016). The relevance effect and conditionals. *Cognition*, *150*, 26–36. doi: 10.1016/j.cognition.2015.12.017

Tversky, A., & Kahneman, D. (1983). Extensional versus intuitive reasoning: The conjunction fallacy in probability judgment. *Psychological Review*, *90*(4), 293–315. doi: 10.1037/0033-295X.90.4.293

Zhao, J., Shah, A., & Osherson, D. (2009). On the provenance of judgments of conditional probability. *Cognition*, *113*(1), 26–36.