A New Approach to Testimonial Conditionals

Stephan Hartmann (S.Hartmann@lmu.de)
Munich Center for Mathematical Philosophy, LMU Munich, 80539 Munich (Germany)

Ulrike Hahn (U.Hahn@bbk.ac.uk)
Department of Psychological Sciences, Birkbeck, University of London, WC1E 7HX (UK)

Abstract

Conditionals pervade every aspect of our thinking, from the mundane and everyday such as 'if you eat too much cheese, you will have nightmares' to the most fundamental concerns as in 'if global warming isn't halted, sea levels will rise dramatically'. Many decades of research have focussed on the semantics of conditionals and how people reason from conditionals in everyday life. Here it has been rather overlooked how we come to such conditionals in the first place. In many cases, they are learned through testimony: someone warns us about the ill-effects of cheese. Any full account of the conditional must consequently incorporate such learning. Here, we provide a new formal account of belief change in response to a testimonial conditional.

Keywords: Indicative conditional reasoning; testimony; belief change; probability

Introduction

Conditionals figure centrally in our everyday discourse. On hearing them, we must interpret them and adjust our beliefs.

Decades of research have sought to understand exactly what conditionals are, that is, provide an account of their *semantics*. However, as pointed out recently by (Collins, Krzyżanowska, Hartmann, Wheeler, & Hahn, 2020), it remains somewhat mysterious how we change our beliefs when we encounter a conditional. There is a vast body of work on reasoning with conditionals that probes the inferences people will draw from a given conditional. But other aspects of interaction with conditionals have been overlooked: virtually no work has examined how beliefs change simply on hearing a conditional asserted by a testimonial source. This is not just an intriguing gap in its own right, but as Collins et al. argue, how beliefs change upon encountering an assertion of a conditional constrains theories of what conditionals are.

With respect to the meaning of the conditional, a wealth of research has demonstrated that people do not view the conditional as the material conditional of classical logic, and that reasoning performance is better characterised by a probabilistic model which takes conditional reasoning to be Bayesian belief revision with the conditional premise expressed as the conditional probability $P(\mathrm{C}|\mathrm{A})$ (Oaksford & Chater, 2007). Furthermore, judgments of the conditional probability correlate well with judgments of the probability of the conditional (Evans, Handley, Neilens, & Over, 2007; Evans, Handley, & Over, 2003; Over, Hadjichristidis, Evans, Handley, & Sloman, 2007; Oberauer & Wilhelm, 2003; Politzer, Over, & Baratgin, 2010; Over et al., 2007). Studies have shown the association with the conditional probability for indicative conditionals (Evans et al., 2003); causal conditionals (Over et al., 2007); conditional promises (Ohm & Thompson, 2006); conditional tips, threats, and warnings (Evans, Neilens, Handley, & Over, 2008); and counterfactual conditionals (Over et al., 2007). This evidence suggests that whatever is learnt from a conditional is likely to at least include something about the conditional probability. However, recent data also suggest that for the conditional probability and the probability of the conditional to correspond well, the antecedent may need to be positively relevant for the consequent (Skovgaard-Olsen, Singmann, & Klauer, 2016; Krzyżanowska, Wenmackers, & Douven, 2013; Krzyżanowska, Collins, & Hahn, 2017; Krzyżanowska, 2013). In other words, there may need to be some link between antecedent and consequent for a conditional to be assertable.

That the conditional probability is implicated in the indicative conditional suggests that integrating the conditional with testimony should be possible in a probabilistic framework. After all, the uncertainty raised by less than full reliability of a testimonial source seems naturally captured by a probabilistic perspective. However, such a merger turns out to be technically less straightforward than one might assume.

Some Stylised Facts

How, then, might beliefs change on hearing a conditional and how *should* they change? Collins et al. (2020) used the following simple context. Imagine that you are looking at cars at a large car dealership, and someone tells you "If a car on this lot is a Mercedes, then it's black".

If there is an association between the conditional and $P(\mathrm{C}|\mathrm{A})$, a recipient receiving a conditional should increase this conditional probability. However, intuitively it should also matter who uttered the conditional: rationally, one might change one's belief more if the conditional is uttered by a source with relevant expertise, such as the manager of the car dealership, as opposed to an accompanying child.

This putative change to the conditional probability is the most obvious intuition, but likely not the only change to take place (but see (Douven, 2012)). Imagine believing today will be a fine, sunny day, but as you prepare to leave your home, someone says "If it rains today, then you'll get wet." This may raise your belief in the probability of rain. Imagine additionally that you are heading out to the bakery, and the same person also says "If the bakery is open, can you get me a croissant." Here, you might actually decrease your judgment of the probability of the bakery being open. Finally, it seems harder to generate clear intuitions about the probability of the consequent: but there would seem to be contexts on which it would likely rise, given that one has just learned about a new way of bringing it about. Collins et al. (2020) provide experimental evidence in line with these intuitions. We do not review that evidence here, but rather focus on the question of how best to model such changes.

Modelling the Belief Change

In addition to backing up basic intuitions about belief change in response to testimonial conditionals, Collins et al. (2020) also probe the range of extent approaches to conditionals and demonstrate how they fail with the acquisition of conditionals via testimony. Again, space limits prevent full discussion of their results, hence we focus on the challenge for probabilistic accounts such as the suppositional theory of the conditional. Here, the difficulty lies in the fact that conditionals express a relationship between variables, not (just) new information about the state of one or more of those variables. Thus, the standard tool for modelling belief revision within a probabilistic framework–Bayesian conditionalization (or "Bayes' rule")–is at best indirectly applicable. Conditionalization normatively prescribes belief change on observing evidence for the state of a variable, and Bayesian networks (as in Figure 1) provide computational tools for propagating that evidence to other connected variables (Bovens & Hartmann, 2003; Sprenger & Hartmann, 2019; Hartmann, 2020). On the suppositional theory, a conditional such as 'if A, then C' implies something about the conditional probability $P(\mathrm{C}|\mathrm{A})$. For a Bayesian network, however, this is not a claim about the variables A and C per se, but rather about links between them.

From the perspective of Bayesian networks, a model with a link between A and C is simply a different model than A and C without that link. This suggests that capturing such changes requires thinking in terms of model transition. This is exactly the novel approach to learning conditionals taken in (Eva & Hartmann, 2018; Eva, Hartmann, & Rafiee Rad, 2020). These authors seek to identify rationality constraints on the transition. Specifically, they identify formal means that make the transition *conservative*: learning about a previously unknown connection in a model should not lead one to abandon all other knowledge the model embodies. What is desired is the minimal adjustment needed to accommodate the new belief.

However, additional challenges are posed by the uncertainty concerning the reliability of the source where knowledge of that link is acquired via testimony (Coady, 1992). Even the most well-intentioned sources make mistakes and are less than fully reliable as a result. Hence, a central challenge for humans is determining the reliability of our sources in order to normatively factor in that partial reliability. Considerable research in the last 15 years has sought to identify normative constraints on source reliability (Olsson, 2011; Bovens & Hartmann, 2003; Hahn, Merdes, & von Sydow, 2018) and examined laypeople's descriptive responses (A. J. L. Harris, Hahn, Madsen, & Hsu, 2016; Jarvstad & Hahn, 2011; Collins & Hahn, 2019; Collins, Hahn, von Gerber, & Olsson, 2018; P. L. Harris & Koenig, 2007). Here, we adopt the spirit of one such normative approach (Bovens & Hartmann, 2003) in order to capture the uncertainty associated with sources in terms of model uncertainty.

Finally, adequately capturing the way one's beliefs change on hearing the assertion of a conditional by a testimonial source, will likely need to factor in general considerations of natural language pragmatics (Mey, 2001): speech acts carry a presumption of relevance (Grice, 1989; Wilson & Sperber, 2004), hence the fact that someone sees fit to raise a topic at all may alter our beliefs about the issues involved.

In this paper, we combine these ingredients into a novel approach for formalizing belief change in response to conditional assertions. We next detail the formal implementation of these basic ideas.

The Baseline Model

An agent considers the propositions A (= the antecedent) and C (= the consequent) and learns the conditional 'If A, then C' from a partially reliable information source. She assigns a reliability $r \in (0,1)$ to the source. How shall she update her (partial) beliefs about A and C? And how should the conditional probability $P(\mathrm{C}|\mathrm{A})$ change?

To address these question, we proceed in two steps. In the first step, we assume that the agent has a prior probability distribution P over the propositional variables A (with values A and $\neg$A) and C (with values C and $\neg$C). To represent P, we consider the Bayesian network in Figure 1 and assign the prior probability of the antecedent,

$$P(\mathrm{A}) = a, \tag{1}$$

and the conditional probabilities of the consequent C, given the values of its parent,

$$P(\mathrm{C}|\mathrm{A}) = p \quad , \quad P(\mathrm{C}|\neg\mathrm{A}) = q. \tag{2}$$

With this, the joint prior probability distribution P over the variables A and C is given by

$$\begin{aligned} P(\mathrm{A},\mathrm{C}) &= a\,p \qquad , \qquad P(\mathrm{A},\neg\mathrm{C}) = a\,\overline{p} \\ P(\neg\mathrm{A},\mathrm{C}) &= \overline{a}\,q \qquad , \qquad P(\neg\mathrm{A},\neg\mathrm{C}) = \overline{a}\,\overline{q}, \end{aligned} \tag{3}$$

where we have used the shorthand notation $P(\mathrm{A},\mathrm{C})$ for $P(\mathrm{A} \wedge \mathrm{C})$ which we will use throughout this paper. We also use the shorthand $\overline{x}$ for $1-x$ and assume that $a, p, q \in (0,1)$.

In the second step, we ask how the prior probability distribution changes once the agent learns the conditional 'If A, then C'. Let Q denotes the posterior probability distribution

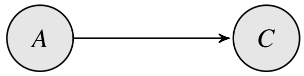

Figure 1: The Bayesian network representation of the relation between A and C.

after learning the conditional from a partially reliable information source. To determine Q, we consider two extreme cases: (1) If the source is fully reliable, i.e. if $r = 1$, then the agent updates by conditioning on the corresponding material conditional. P'_{rel} can be parameterized in the same way as P with

$$a' = \frac{ap}{ap + \overline{a}} \quad , \quad p' = 1 \quad , \quad q' = q. \tag{4}$$

It is interesting to note that one obtains the same result if one uses the more general distance-based approach to Bayesianism and determines P'_{rel} by minimizing some f-divergence between P'_{rel} and P, taking the constraint $P'_{rel}(\mathrm{C}|\mathrm{A}) = p' = 1$ into account. See (Eva & Hartmann, 2018; Eva et al., 2020; Stern & Hartmann, 2018) for details. From eqs. (4) it is easy to see that

$$a' < a. \tag{5}$$

Defining $c := P(\mathrm{C}) = a\,p + \overline{a}\,q$ and $c' := P'_{rel}(\mathrm{C}) = a' + \overline{a'}\,q$, one finds that

$$c' > c. \tag{6}$$

At this point, one might wonder whether the utterance of the conditional by a partially reliable information source could be modeled by simply minimizing some f-divergence between Q and P, taking the constraint $Q(\mathrm{C}|\mathrm{A}) = p' < 1$ into account, where the value of p' depends on the reliability of the source (e.g. $p'(r) \to 1$ as $r \to 1$).[1] This proposal has some plausibility, but it faces serious difficulties when confronted with empirical data. Here is why: If one follows the described procedure, then the new probability of the antecedent and the new probability of the consequent only depend implicitly on r, viz. via the functional relationship between p' and r. That is, once p' is fixed (and an f-divergence is chosen), then the posteriors of A and C do not show any explicit dependence on r. However, the data suggest something else: If the prior probability distribution over the variables A and C is fixed and a value of p' is chosen, then the posterior distribution will still depend on r. Hence, the simple proposal does not work.

(2) Let us now consider the case where the source is fully unreliable, i.e. $r = 0$. Then a rational agent should disregard the new information completely. Hence, the posterior probability distribution P'_{unrel} equals the prior probability distribution: $P'_{unrel} = P$.

To obtain the posterior distribution Q for all values of the reliability r, we take the convex combination of P'_{rel} and P'_{unrel}, weighted with r or $\overline{r}$ respectively, i.e. $Q = r\,P'_{rel} + \overline{r}\,P$ (as $P'_{unrel} = P$). Proposition 1 summarizes the features of Q:

Proposition 1 *An agent considers the propositional variables A and C with a prior probability distribution P defined over them as in eqs. (3). The agent then learns the conditional 'If* A, C*' from a source to which she assigns a reliability r and updates her (partial) beliefs as described above. Then the posterior probability distribution Q has the following features: (i)* $Q(\mathrm{A}) < P(\mathrm{A})$*, (ii)* $Q(\mathrm{C}) > P(\mathrm{C})$*, and (iii)* $Q(\mathrm{C}|\mathrm{A}) > P(\mathrm{C}|\mathrm{A})$.

Q accounts for some of the stylized facts described above. For example, it accounts for the observation that the conditional probability of C given A always increases. We also see Collins et al.'s (2020) data that the probability of the conclusion (C) increases for a large range of prior distributions. However, the baseline model cannot explain that the probability of the conclusion sometimes decreases. It can also not explain that the probability of the antecedent (A) typically increases in experiments. Modeling this latter observation is a major challenge for all accounts of learning conditionals which aim to be normatively plausible and empirically adequate since Collins et al. show that extant accounts entail that the probability of the antecedent always decreases after learning the corresponding conditional. We address this issue next.

The Extended Model

Consider the following situation: A friend tells you, totally unexpected to you, that "If there is an earthquake, then there will be a considerable amount of air pollution in the area". You are surprised by the remark. On the one hand you find it plausible as an earthquake will probably cause air pollution. At the same time you do not expect an earthquake at all and so your prior of it is rather low. But why does your friend mention an earthquake? Does she have special information which you do not have? There must be a *reason* why your friend mentions the earthquake. Pondering the issue, you increase the probability of the antecedent. This seems rational.

Let us now formalize this insight. To begin with, we disregard the propositional variable C and focus on A. We assume that the agent also believes that things are somehow normal and unsurprising. (Earthquakes and the like are really unlikely in the area we live in.) Let us denote the corresponding propositional variable by N (with values N: "Things are normal." and ¬N: "Things are not normal.") and include it in the Bayesian network which represents the relevant beliefs of the agent (see Figure 2).

We furthermore assume that $n := P(\mathrm{N})$ is large, that $p_N := P(\mathrm{A}|\mathrm{N})$ is small and that $q_N := P(\mathrm{A}|\neg\mathrm{N})$ is large. In particular, we assume that $q_N > p_N$. Hence, $P(\mathrm{A}) = n\,p_N + \overline{n}\,q_N$ will be fairly small. Now, the fact that

[1] The f-divergence is defined as follows: Let $S_1, \ldots, S_n$ be the possible values of a random variable S over which probability distributions P (= the prior distribution) and Q (= the posterior distribution) are defined. The f-divergence between Q and P is then given by

$$D_f(Q||P) := \sum_{i=1}^{n} P(S_i)\, f\left(\frac{Q(S_i)}{P(S_i)}\right), \tag{7}$$

where f is a convex function such that $f(1) = 0$. For details and applications in the psychology of reasoning, see (Eva & Hartmann, 2018; Eva et al., 2020).

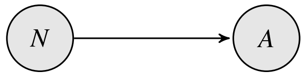

Figure 2: The Bayesian network representation of the relation between A and N.

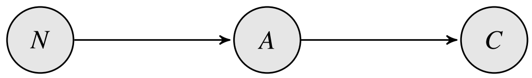

Figure 3: The Bayesian Network representation of the relation between A, C and N.

your friend mentions the possibility of an earthquake (i.e. A) makes you doubt that things are as normal as you previously thought. Hence, you decrease the probability of N and set $n' := P'(\mathrm{N}) < n$. To find the full posterior probability distribution P' you use Jeffrey conditionalization and obtain $P'(\mathrm{A}) = n'\, p_N + \overline{n'}\, q_N$. It is easy to see that $P'(\mathrm{A})$ >P(A) as $P'(\mathrm{A}) - P(\mathrm{A}) = (n - n')(q_N - p_N) > 0$.

This mechanism increases the probability of A. If the agent then updates on the conditional, it will decrease, but it may well be that the resulting posterior probability of A is larger than the prior probability of A. It may, however, also reach a value lower than the prior probability of A. The details depend on the prior probability distribution over the propositional variables A, C and N and on the value of n'. Let us now work this idea out in detail.

The full Bayesian network over the three propositional variables A, C and N is given in Figure 3. The corresponding prior probability distribution P is given by

$$P(\mathrm{N}) = n \tag{8}$$

and

$$\begin{aligned} P(\mathrm{A}|\mathrm{N}) = p_N \quad &, \quad P(\mathrm{A}|\neg\mathrm{N}) = q_N \\ P(\mathrm{C}|\mathrm{A}) = p_A \quad &, \quad P(\mathrm{C}|\neg\mathrm{A}) = q_A. \end{aligned} \tag{9}$$

As before, we first consider how the agent updates if she considers the information source to be fully reliable. In this case the agent learns two pieces of information: (i) The probability of N shifts from $P(\mathrm{N}) = n$ to $P'_{rel}(\mathrm{N}) = n'$. (ii) The indicative conditional $\mathrm{A} \to \mathrm{C}$. To update, we use the distance-based approach to Bayesianism and specify the following two constraints on P'_{rel}:

$$P'_{rel}(\mathrm{N}) = n' \quad , \quad P'_{rel}(\mathrm{C}|\mathrm{A}) = p'_A = 1 \tag{10}$$

With this, P'_{rel} can be represented in the same way as P.

Proposition 2 *An agent considers the propositional variables A, C and N with a prior probability distribution P defined over them as in eqs. (8) and (9). The agent then learns the conditional 'If* A, C*' from a source to which she assigns a reliability r and updates her (partial) beliefs by minimizing an f-divergence between P'_{rel} and P, taking the constraints (10) into account. Then $q'_A = q_A$ and*

$$p'_N = \frac{p_N\, p_A}{p_N\, p_A + \overline{p_N}} \quad , \quad q'_N = \frac{q_N\, p_A}{q_N\, p_A + \overline{q_N}}, \tag{11}$$

with $p'_N < p_N, q'_N < q_N$ and $q'_N > p'_N$ if and only if $q_N > p_N$.

Eq. (11) are interesting: If p_A is small, then p'_N and q'_N are also small and hence $P'_{rel}(\mathrm{A}) = n'\, p'_A + \overline{n'}\, q'_A$ is small and therefore likely to be smaller than $P(\mathrm{A})$. To proceed, let us define $\Delta_{rel}(\mathrm{A}) := P'_{rel}(\mathrm{A}) - P(\mathrm{A})$ and $\Delta_{rel}(\mathrm{C})$ accordingly. Figure 4 shows these functions for some typical values.

Panels 4(c) and 4(d) suggest that the probability of the conclusion (C) typically increases. However, in numerical studies we find that it decreases under certain conditions, e.g. if p_A is large (i.e. if learning the conditional does not provide much new information) and if $n' > n$, i.e. if the agent comes to belief, after learning the conditional, that things are actually more normal than expected. This phenomenon needs a more detailed experimental and theoretical analysis.

As in the case of the baseline model, we can now calculate the posterior probability distribution $Q = r\,P'_{rel} + \overline{r}\,P$. Defining $\Delta(\mathrm{A}) := Q(\mathrm{A}) - P(\mathrm{A})$ (and $\Delta(\mathrm{C})$ and $\Delta(C|A)$ accordingly), the following proposition summarizes our results.

Proposition 3 *An agent considers the propositional variables A, C and N with a prior probability distribution P defined over them as in eqs. (8) and (9). The agent then learns the conditional "If* A, C*" from a source to which she assigns a reliability r and updates her (partial) beliefs as described above. Then (i) $\Delta(\mathrm{A}) = r \cdot \Delta_{rel}(\mathrm{A})$, (ii) $\Delta(\mathrm{C}) = r \cdot \Delta_{rel}(\mathrm{C})$, and (iii) $\Delta(C|A) > 0$.*

Hence, our model explains why the (absolute values of the) differences $\Delta(\mathrm{A})$ and $\Delta(\mathrm{C})$ increase linearly with the reliability r. The plots in Figure 4 show the maximal (absolute) differences, i.e. for $r = 1$.

Conclusions

We have presented a new approach to an overlooked, but central, problem of human cognition. Human thought and language are unthinkable without the conditional, yet little attention has been given to how we acquire knowledge of conditionals through testimony from others. The approach taken here is simple yet powerful: when we acquire knowledge of a new conditional relationship, our model of the world changes. Where that knowledge is uncertain, because our source is less than fully reliable, model uncertainty remains. Arguably, both model transition and model blending are subject to rationality constraints. Future research working out these constraints in increasing detail is likely to not only provide a missing puzzle piece for theoretical understanding of conditional reasoning and communicating with conditionals, but also to provide a blue print for other contexts of model transition and model uncertainty in human cognition.

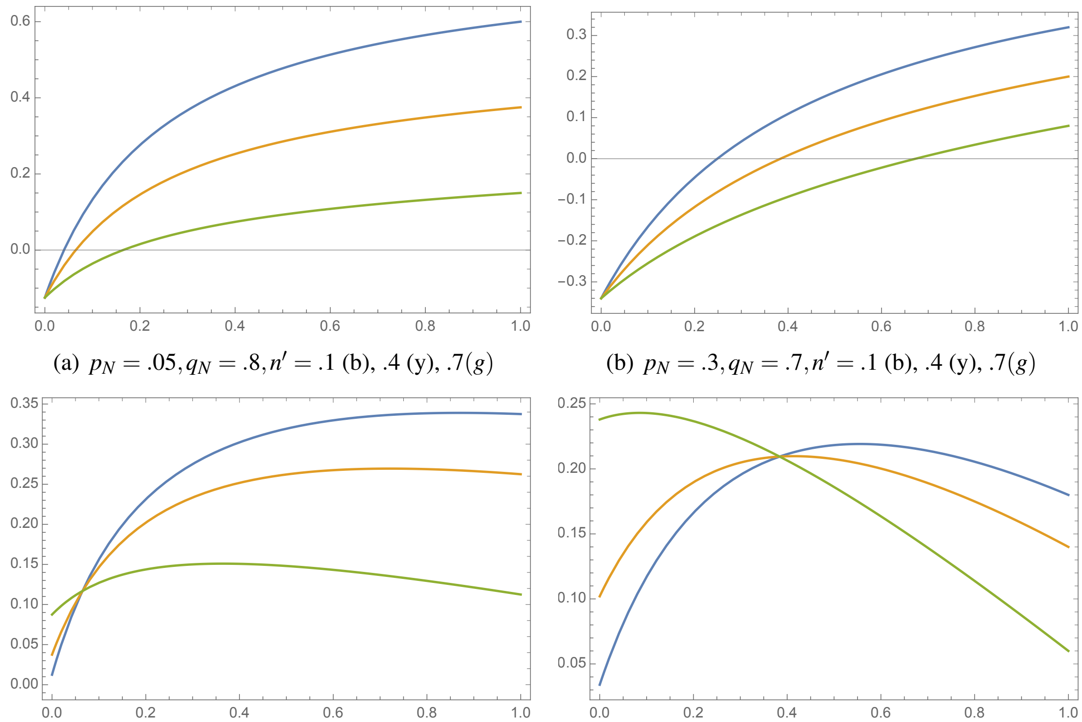

(a) $p_N = .05, q_N = .8, n' = .1$ (b), .4 (y), .7(g)
(b) $p_N = .3, q_N = .7, n' = .1$ (b), .4 (y), .7(g)
(c) $p_N = .05, q_N = .8, n' = .4, q_A = .1$ (b), .4 (y), .7(g)
(d) $p_N = .3, q_N = .7, n' = .4, q_A = .1$ (b), .4 (y), .7(g)

Figure 4: The differences $\Delta_{rel}(\mathrm{A})$ (top) and $\Delta_{rel}(\mathrm{C})$ (bottom) as a function of p_A for $n = .9$. b: blue, y: yellow, g: green

Proof of Proposition 1

The posterior probability distribution Q is given by

$$\begin{aligned} Q(\mathrm{A},\mathrm{C}) &= a'\,r + a\,p\,\overline{r} \quad , \quad Q(\mathrm{A},\neg\mathrm{C}) = a\,\overline{p}\,\overline{r} \qquad (12)\\ Q(\neg\mathrm{A},\mathrm{C}) &= \overline{a'}\,q\,r + \overline{a}\,q\,\overline{r} \quad , \quad Q(\neg\mathrm{A},\neg\mathrm{C}) = \overline{a'}\,\overline{q}\,r + \overline{a}\,\overline{q}\,\overline{r}, \end{aligned}$$

with a' given in eq. (4). Using eqs. (12), we calculate $Q(\mathrm{A}) = a'\,r + a\,\overline{r}$. Hence $Q(\mathrm{A}) - P(\mathrm{A}) = (a' - a)\,r < 0$ as $a' < a$ (see eq. (5)). Similarly, we calculate $Q(\mathrm{C}) = c'\,r + c\,\overline{r}$. Hence $Q(\mathrm{C}) - P(\mathrm{C}) = (c' - c)\,r > 0$ as $c' > c$ (see eq. (6)). Finally, $Q(\mathrm{C}|\mathrm{A}) = Q(\mathrm{A},\mathrm{C})/Q(\mathrm{A}) = (a'\,r + a\,p\,\overline{r})/(a'\,r + a\,\overline{r})$. Hence, $Q(\mathrm{C}|\mathrm{A}) - P(\mathrm{C}|\mathrm{A}) = a'\,\overline{p}\,r/(a'\,r + a\,\overline{r}) > 0$. It is interesting to note that $Q(\mathrm{C}|\mathrm{A})$ is not given by $p'' := r + p\,\overline{r}$, as one might have expected in analogy to the expressions for $Q(\mathrm{A})$ and $Q(\mathrm{C})$. In fact, $Q(\mathrm{C}|\mathrm{A}) \leq p''$.

Proof of Proposition 2

Using eq. (7), we calculate $F := D_f(P'_{rel}||P)$:

$$\begin{aligned} F = {} & n\,p_N\,p_A \cdot f\left(\frac{n'\,p'_N}{n\,p_N\,p_A}\right) + n\,\overline{p_N}\,q_A \cdot f\left(\frac{n'\,\overline{p'_N}\,q'_A}{n\,\overline{p_N}\,q_A}\right) \\ & + n\,\overline{p_N}\,\overline{q_A} \cdot f\left(\frac{n'\,\overline{p'_N}\,\overline{q'_A}}{n\,\overline{p_N}\,\overline{q_A}}\right) + \overline{n}\,q_N\,p_A \cdot f\left(\frac{\overline{n'}\,q'_N}{\overline{n}\,q_N\,p_A}\right) \\ & + \overline{n}\,\overline{q_N}\,q_A \cdot f\left(\frac{\overline{n'}\,\overline{q'_N}\,q'_A}{\overline{n}\,\overline{q_N}\,q_A}\right) + \overline{n}\,\overline{q_N}\,\overline{q_A} \cdot f\left(\frac{\overline{n'}\,\overline{q'_N}\,\overline{q'_A}}{\overline{n}\,\overline{q_N}\,\overline{q_A}}\right) \qquad (13) \end{aligned}$$

To find the values of p'_N, q'_N and q'_A which minimize F, we first differentiate F by q'_A and then set the resulting expression equal to zero. Hence,

$$\begin{aligned} & n'\,\overline{p'_N} \cdot \left[f\left(\frac{n'\,\overline{p'_N}\,q'_A}{n\,\overline{p_N}\,q_A}\right) - f\left(\frac{n'\,\overline{p'_N}\,\overline{q'_A}}{n\,\overline{p_N}\,\overline{q_A}}\right)\right] \\ & + \overline{n'}\,\overline{q'_N} \cdot \left[f\left(\frac{\overline{n'}\,\overline{q'_N}\,q'_A}{\overline{n}\,\overline{q_N}\,q_A}\right) - f\left(\frac{\overline{n'}\,\overline{q'_N}\,\overline{q'_A}}{\overline{n}\,\overline{q_N}\,\overline{q_A}}\right)\right] = 0. \end{aligned}$$

As this equation has to hold for all values of n', we conclude that the expressions in the square brackets have to vanish. Using the convexity of f, we obtain $q'_A = q_A$. Inserting this result into eq. (13), we obtain

$$\begin{aligned} F = {} & n\,p_N\,p_A \cdot f\left(\frac{n'\,p'_N}{n\,p_N\,p_A}\right) + n\,\overline{p_N} \cdot f\left(\frac{n'\,\overline{p'_N}}{n\,\overline{p_N}}\right) \\ & + \overline{n}\,q_N\,p_A \cdot f\left(\frac{\overline{n'}\,q'_N}{\overline{n}\,q_N\,p_A}\right) + \overline{n}\,\overline{q_N} \cdot f\left(\frac{\overline{n'}\,\overline{q'_N}}{\overline{n}\,\overline{q_N}}\right). \qquad (14) \end{aligned}$$

Next, we differentiate F by p'_N and set the resulting expression equal to zero:

$$n' \cdot \left[f\left(\frac{n'\,p'_N}{n\,p_N\,p_A}\right) - f\left(\frac{n'\,\overline{p'_N}}{n\,\overline{p_N}}\right)\right] = 0$$

Hence, again using the convexity of f, we obtain

$$p'_N = \frac{p_N\,p_A}{p_N\,p_A + \overline{p_N}}$$

and with this

$$p'_N - p_N = -\frac{p_N\,\overline{p_N}\,p_A}{p_N\,p_A + \overline{p_N}} < 0.$$

Similarly for q'_N. Finally, we calculate

$$q'_N - p'_N = \frac{q_N - p_N}{(p_N\,p_A + \overline{p_N})(q_N\,p_A + \overline{q_N})},$$

from which the last statement in the proposition follows.

Proof of Proposition 3

$\Delta(\mathrm{A}) := Q(\mathrm{A}) - P(\mathrm{A}) = r \cdot P'_{rel}(\mathrm{A}) + \overline{r} \cdot P(\mathrm{A}) - P(\mathrm{A})$ Hence, $\Delta(\mathrm{A}) = r \cdot (P'_{rel}(\mathrm{A}) - P(\mathrm{A}))$. Accordingly for $\Delta(\mathrm{C})$. Next, note that $P'_{rel}(\mathrm{A}, \neg\mathrm{C}) = 0$ and hence $P'_{rel}(\mathrm{A},\mathrm{C}) = P'_{rel}(\mathrm{A})$. Setting $a := P(\mathrm{A})$ and $a' := P'_{rel}(\mathrm{A})$, we finally find

$$\begin{aligned}\Delta(\mathrm{C}|\mathrm{A}) &= \frac{Q(\mathrm{A},\mathrm{C})}{Q(\mathrm{A})} - P(\mathrm{C}|\mathrm{A}) \\ &= \frac{a'\,r + a\,p_A\,\overline{r}}{a'\,r + a\,\overline{r}} - p_A \\ &= \frac{a'\,\overline{p_A}\,r}{a'\,r + a\,\overline{r}} > 0.\end{aligned}$$

References

Bovens, L., & Hartmann, S. (2003). *Bayesian Epistemology*. Oxford: Oxford University Press.

Coady, C. (1992). *Testimony: A Philosophical Study*. Oxford: Oxford University Press.

Collins, P. J., & Hahn, U. (2019). We might be wrong, but we think that hedging doesn't protect your reputation. *Journal of Experimental Psychology: Learning, Memory, and Cognition*.

Collins, P. J., Hahn, U., von Gerber, Y., & Olsson, E. J. (2018). The bi-directional relationship between source characteristics and message content. *Frontiers in Psychology*, *9*, 18.

Collins, P. J., Krzyżanowska, K., Hartmann, S., Wheeler, G., & Hahn, U. (2020). *Conditionals and testimony.* (submitted)

Douven, I. (2012). Learning conditional information. *Mind and Language*, *27*(3), 239–263.

Eva, B., & Hartmann, S. (2018). Bayesian argumentation and the value of logical validity. *Psychological Review*, *125*(5), 806–821.

Eva, B., Hartmann, S., & Rafiee Rad, S. (2020). Learning from conditionals. *Mind*, *129*(514), 461–508.

Evans, J. S. B. T., Handley, S. J., Neilens, H., & Over, D. E. (2007). Thinking about conditionals: A study of individual differences. *Memory & Cognition*, *35*(7), 1772–1784.

Evans, J. S. B. T., Handley, S. J., & Over, D. E. (2003). Conditionals and conditional probability. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *29*(2), 321–335.

Evans, J. S. B. T., Neilens, H., Handley, S. J., & Over, D. E. (2008). When can we say 'if'? *Cognition*, *108*(1), 100–116.

Grice, H. P. (1989). *Studies in the Way of Words*. Cambridge, MA: Harvard University Press.

Hahn, U., Merdes, C., & von Sydow, M. (2018). How good is your evidence and how would you know? *Topics in Cognitive Science*, *10*(4), 660–678.

Harris, A. J. L., Hahn, U., Madsen, J. K., & Hsu, A. S. (2016). The appeal to expert opinion: Quantitative support for a Bayesian network approach. *Cognitive Science*, *40*(6), 1496–1533.

Harris, P. L., & Koenig, M. A. (2007). The basis of epistemic trust: Reliable testimony or reliable sources? *Episteme*, *4*(3), 264–284.

Hartmann, S. (2020). Bayes nets and rationality. In M. Knauff & W. Spohn (Eds.), *Handbook of Rationality.* Cambridge, MA: MIT Press.

Jarvstad, A., & Hahn, U. (2011). Source reliability and the conjunction fallacy. *Cognitive Science*, *35*(4), 682–711.

Krzyżanowska, K. (2013). Belief ascription and the Ramsey test. *Synthese*, *190*, 21–36.

Krzyżanowska, K., Collins, P. J., & Hahn, U. (2017). Between a conditional's antecedent and its consequent: Discourse coherence vs. probabilistic relevance. *Cognition*, *164*, 199–205.

Krzyżanowska, K., Wenmackers, S., & Douven, I. (2013). Inferential conditionals and evidentiality. *Journal of Logic, Language and Information*, *22*(3), 315–334.

Mey, J. L. (2001). *Pragmatics: An introduction*. Hoboken, N.J.: Wiley.

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

Oberauer, K., & Wilhelm, O. (2003). The meaning(s) of conditionals: Conditional probabilities, mental models, and personal utilities. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *29*(4), 680–693.

Ohm, E., & Thompson, V. A. (2006). Conditional probability and pragmatic conditionals: Dissociating truth and effectiveness. *Thinking & Reasoning*, *12*(3), 257–280.

Olsson, E. J. (2011). A simulation approach to veritistic social epistemology. *Episteme*, *8*(02), 127–143.

Over, D. E., Hadjichristidis, C., Evans, J. S. B. T., Handley, S. J., & Sloman, S. A. (2007). The probability of causal conditionals. *Cognitive Psychology*, *54*, 62–97.

Politzer, G., Over, D. E., & Baratgin, J. (2010). Betting on conditionals. *Thinking & Reasoning*, *16*(3), 172–197.

Skovgaard-Olsen, N., Singmann, H., & Klauer, K. C. (2016). The relevance effect and conditionals. *Cognition*, *150*, 26–36.

Sprenger, J., & Hartmann, S. (2019). *Bayesian Philosophy of Science*. Oxford: Oxford University Press.

Stern, R., & Hartmann, S. (2018). Two sides of modus ponens. *The Journal of Philosophy*, *115*(11), 605–621.

Wilson, D., & Sperber, D. (2004). Relevance theory. In L. R. Horn & G. Ward (Eds.), *The Handbook of Pragmatics* (pp. 607–632). Oxford: Blackwell.