university of groningen

NWO
Netherlands Organisation for Scientific Research

CWI 2011

A new resolution of the Judy Benjamin problem

⋆

Igor Douven
Faculty of Philosophy
University of Leuven

Jan-Willem Romeijn
Faculty of Philosophy
University of Groningen

Judy Benjamin

In an example by van Fraassen [1981], Judy Benjamin is dropped in an area divided into Red (R) and Blue ($\neg R$) and into Second Company (S) and Headquarters ($\neg S$) sections. She assigns equal probability to all quadrants Q_i.

	S	$\neg S$
R	Q_1	Q_2
$\neg R$	Q_3	Q_4

Then she learns that if she is in Red territory, the odds are 3 : 1 that she is in Headquarters area. How probable is it now that she is in Blue territory?

Relative entropy distance minimization

The information imposes a specific constraint on the probability assignment over the segments Q_i. Using a relative entropy distance between probability assignments,

$$RE(P, P_{\text{old}}) = \sum_i P(Q_i) \log \frac{P(Q_i)}{P_{\text{old}}(Q_i)},$$

we can look for the closest new probability assignment that satisfies the constraint:

$$\Gamma = \left\{ P : \frac{P(Q_2)}{P(Q_1)} = 3 \right\}, \qquad P_{\text{new}} = \{ P \in \Gamma : RE(P, P_{\text{old}}) \text{ minimal} \}.$$

Conditional as material implication

Surprisingly, if we determine the new probability in this way, the probability of being in Blue increases!

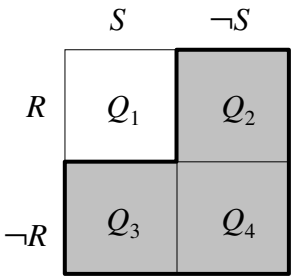

Van Fraassen [1989] explains this by reference to the limiting case. If Judy learns “If in Red, then in Headquarters, period”, the increase in the probability of Blue is a matter of course. Or is it?

Contents

1 Updating on conditionals

We can object to van Fraassen's line of argument for a number of reasons, having to do with the semantics of conditionals and with the event of learning a conditional. Consider the story of Sarah and Marian having sundowners at the Westcliff hotel.

Conditionalization on the material conditional entails a decrease in the probability of the antecedent, which is at variance with the example.

More than conditionalization

The example suggests that learning a conditional by conditionalizing on its truth conditions misses out on the context dependent implications of the conditional, including its assertability conditions.

- If conditionals have truth conditions: by asserting a conditional we convey more than truth-conditional content. They also have pragmatic implicatures, arguably of a probabilistic nature.
- If conditionals do not have truth conditions: why take the material implication as a limiting case to begin with?

The upshot is that we cannot simply defend the consequences of a minimum relative entropy update in the Judy Benjamin case by referring to conditionalization on a conditional.

What colour is the coat?

2 Alternative update mechanisms

Information does not always come in neat propositional packages. Richard Jeffrey devised a rule for updating a probability assignment on new information captured by a probability assignment over a partition of possible events.

$$P_{\text{new}}(C) = \sum_i P_{\text{new}}(U_i) P_{\text{old}}(C|U_i).$$

Jeffrey's rule does not tell us how we can obtain this probability assignment over the partition of U_i, other than that it stems from our observation and experience.

Employing Jeffrey conditionalization

Say that we learn "If R, then the odds for $\neg S : S$ are $q_1 : q_2$", and that we do not want to adapt our degree of belief $P(R) = r$.

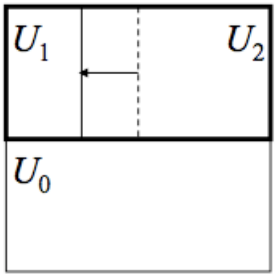

We can achieve this by applying Jeffrey conditionalization to the partition of events $\mathcal{U} = \{U_0, U_1, U_2\} = \{\neg R, R \wedge \neg S, R \wedge S\}$ using the odds, $(1-r)/r(q_1 + q_2) : q_1 : q_2$.

Adams conditioning

In the context of preference kinematics, Bradley [2005] proposes an update rule in which the invariance of the probability of the antecedent remains implicit: Adams conditioning.

> Given a partition $\{U_0, U_1, \ldots U_n\}$, and supposing we obtain new probabilities $P_{\text{new}}(U_i)$ for $i = 1, \ldots, n$, the new probability P_{new} must be as follows:

$$P(C) = P_{\text{old}}(C|U_0)P_{\text{old}}(U_0) + \sum_{i=1}^{n} P_{\text{new}}(U_i)P_{\text{old}}(C|U_i).$$

Clearly this is a special case of Jeffrey's rule of updating: the only difference is that in Adams conditioning, the probability of one of the elements is hardwired to be invariant.

Context as input, or implicit to rule

Applying Adams conditionalization to the cases of Judy and Sarah, we find the intuitively correct results: the probability of the antecedent is not affected by the update. We can choose to. . .

- take the invariance of the probability of the antecedent as an explicit part of the input to the update rule, as for Jeffrey's rule. We may then derive the required constraint from the context of the example cases.
- take the invariance of the probability of the antecedent as implicit to the update rule itself. Based on the context we may then decide that Adams conditioning is applicable.

The difference between these two ways of updating is of little consequence. The boundary between criteria for applicability and input seems vague.

3 A distance function for Adams conditioning

An attractive feature of Jeffrey's rule is that its results are replicated by a distance minimization procedure. This holds for a number of different distance functions.

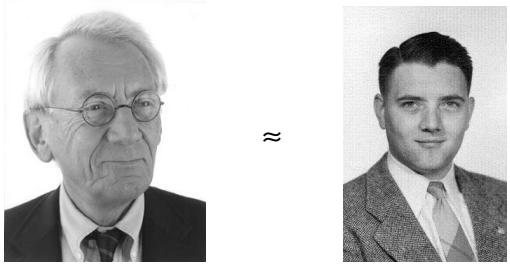

If the new probability assignment is not constrained by all elements in the partition, distance minimization leads to changes in the probability of those elements not involved.

Inverse relative entropy

Is there also a distance function that yields the results of Adams conditioning? It turns out that minimizing the inverse relative entropy distance exactly yields the required results.

$$IRE(P, P_{\text{old}}) = \sum_i P_{old}(Q_i) \log \frac{P_{old}(Q_i)}{P(Q_i)}$$

Imposing the constraint that the odds $P(Q_1) : P(Q_2)$ are $q_1 : q_2$, we find that $P_{\text{new}}(Q_0) = P_{\text{old}}(Q_0)$. Translated to the case of Judy: after learning the radio message and updating by inverse relative entropy minimization, her probability of being in Blue is not affected.

4 Distance minimization generalized

There are other cases than those of Sarah and Judy. In the example about Patricia, learning the conditional should not affect the consequent.

In fact Patricia's case can be accommodated by a variant of Adams conditioning, using a fixed probability for the consequence. But what if the update leads to a conflict between the probabilities of the antecedent and consequent?

A large class of distance functions

To accommodate a trade-off between antecedent and consequent, we may use a Hellinger distance and supplement it with weights $w_i > 0$ for the quadrants,

$$HEL_w(P, P_{\text{old}}) = \sum_{i=1}^{4} w_i \left(\sqrt{P(Q_i)} - \sqrt{P_{\text{old}}(Q_i)} \right)^2 .$$

The higher w_i, the more resistance to deviations in the probability $P(Q_i)$. This idea can be generalized to more complicated conditional statements and different kinds of dependence. Adams conditioning is a limiting case.

Epistemic entrenchment

We can model any trade-off between adapting the probability of the antecedent and the consequent by varying the w_i.

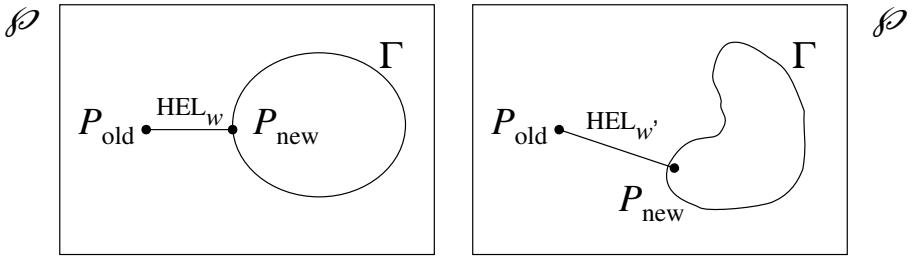

The values of the weights express epistemic entrenchment. Or in terms of a Lewisian imaging operation, they effectively determine the “closest possible world”.

Numerical example

Setting the odds $P_{\text{new}}(Q_1) : P_{\text{new}}(Q_2)$ to 3 : 1 and to 50 : 1 respectively, fixing the weight $w_2 = w_4 = 1$, and varying $w = w_1 = w_3$ from 1 to 100, we obtain the following updated probability assignments.

		Probability			
Odds	Weight	$Q_1 = R \wedge S$	$Q_2 = R \wedge \neg S$	$Q_3 = \neg R \wedge S$	$Q_4 = \neg R \wedge \neg S$
-	-	0.10	0.70	0.10	0.10
3	1	0.53	0.18	0.15	0.15
	5	0.21	0.07	0.13	0.60
	100	0.10	0.03	0.10	0.76
50	1	0.47	0.01	0.26	0.26
	5	0.15	0.00	0.13	0.72
	100	0.10	0.00	0.10	0.79

5 Discussion

- Conditionalization on the material implication is not necessarily the limiting case of updating by relative entropy minimization under the constraint of conditional odds.
- Hence, the fact that relative entropy minimization affects the probability of the antecedent cannot be defended by reference to this conditionalization.
- If we gather the constraints imposed by the Judy Benjamin story, they pin down a complete probability assignment over a partition, and we can apply Jeffrey's rule of updating.
- Alternatively, we can apply Adams conditioning, using an incomplete probability assignment over a partition as input. The further constraint then appears as a condition of applicability.

Discussion (continued)

- The distance function IRE provides an underpinning for Adams conditionalization: minimizing it under the constraint of an incomplete probability assignment gives the same results.
- We can define a whole class of distance functions, each of them associated with different epistemic entrenchments for the probabilities of the elements of the partition.
- In the face of this plethora of update rules, capturing the dynamics of belief in a single update rule seems unrealistic. We apply to update rules more generally what Richard Bradley says of conditionalization:

> "it should not be thought of as a universal and mechanical rule of updating, but as a technique to be applied in the right circumstances, as a tool in what Jeffrey terms the 'art of judgment'. "

Thank you

The slides for this talk will be available at http://www.philos.rug.nl/~romeyn. For comments and questions, email j.w.romeijn@rug.nl.