

Research seminar
 CUNY 2011


Meaning shifts and Conditioning

★

Jan-Willem Romeijn
 Faculty of Philosophy
 University of Groningen

Conditioning and meaning shifts


In the Bayesian model, beliefs over sentences like s and r are represented with probabilities over propositions S and R .



Within the set of possible worlds consistent with the sentence that is learnt, the probability is kept unchanged.

A shift in meaning

A shift in the meaning of a sentence r is represented by a change in the associated collection of possible worlds R .



Such meaning shifts lead to a conflict between conditional and updated probability.

Examples of meaning shifts

Meaning shifts can be read into several well-known philosophical problem domains.

- Violations of the principle of reflection, as discussed in van Fraassen (1989) and Maher (1993): after drinking a bottle of whiskey, the meaning of “being fit to drive home” changes.
- Vagueness in meaning, as considered by Williamson (1994) and, in terms of conceptual covers, by Aloni (2000): learning that Bill is a monk impacts on the meaning of “Bill is rich”.
- Reasoning about knowledge, as formalised in dynamic epistemic logic: Bob’s being in doubt changes what Alice means with “I am not sure”.


I focus on the last domain, in the hope of stimulating a rapprochement between probabilistic epistemology and dynamic epistemic logic.

Contents

1 A problem case from DEL	6
2 A Bayesian model?	10
3 Knowledge structures	14
4 Dempster-shafer belief functions	18
5 Discussion	20

1 A problem case from DEL


Van Benthem (2003) shows that some belief changes cannot be modelled by Bayesian conditioning. His example involves Alice, Bob, and three worlds.



If Alice announces that she does not know r , or $\neg K_A r$ for short, we find that Bob knows s . The Bayesian model gives this a probability half.

A re-representation


We can represent the same semantics slightly differently, showing these two epistemic propositions.



This representation is instrumental to making the fallacy of the naive Bayesian explicit.

Updating by Bayes' rule


The Bayesian model of belief change simply eliminates worlds inconsistent with the information provided.



If we announce $\neg K_A r$, we can eliminate W_3 . This leads to $P(\neg K_B s | \neg K_A r) = \frac{1}{2}$.

Updating on epistemic repercussions



A complete update requires that we also operate on epistemic relations between worlds, leaving $\neg K_B S$ with zero probability.



The second step in the update is effectively a shift in meaning: the extension of the sentence $K_B S$ changes.

2 A Bayesian model?


To construct a Bayesian model of meaning shifts, we provide worlds with an internal structure that captures the epistemic relations.



All relevant aspects of the new information are made explicit in the possible worlds semantics, so that operations on it can be kept simple.

Meaning shifts as conditioning


Specifically, we can model the the entire update by means of a conditioning operation.



This first stage corresponds to standard Bayesian conditioning.

Operating within worlds


In the second stage, instead of operating on relations, we operate within worlds.



In the remaining possible worlds semantics there are no worlds left in which Bob does not know s .

Probabilities over epistemic states?

We might think that the proper Bayesian model may simply take epistemic states as the units of analysis.



But in that case we arrive back at square one. To capture the meaning shift, probabilities need to pertain to worlds, as sets of states.

3 Knowledge structures

The foregoing uses so-called knowledge structures, first discussed in Fagin et al (1984). They can be defined inductively, assuming relations among worlds like R_A and R_B .

$$W_k = k \times \{i : \langle k, i \rangle \in R_A\} \times \{j : \langle k, j \rangle \in R_B\} \times \dots$$

Possibility structures within worlds, like $W_k = \langle k, i, j, \dots \rangle$, can be made as rich as necessary, involving any depth, any number of n -ary relations, and any interpretation of the relations.

Information versus propositional content

Sets of possibilities may cut across worlds. This creates room for distinguishing propositional and informational content.


- The propositional content $\ulcorner u \urcorner \in \mathcal{P}(\mathcal{W})$ of sentence u consists of worlds W_k for which u is true.
- The informational content of u , written as $[u]$, is a set of possibilities, typically included in $\ulcorner u \urcorner$, that may cut across worlds.

Assuming an epistemic frame and using only the first-step possibility structure, the propositional and informational content are related according to

$$[u] = \{\langle k, i, j \rangle : W_k, W_i, W_j \in \ulcorner u \urcorner\}.$$

Possibility-dependent propositions


Whether a world belongs to the propositional content of a sentence may depend on the presence of particular possibilities within that world.



The propositional content of the sentence $\neg K_B s$ is $\{W_2, W_3\}$, because in both these two worlds we find a possibility in which Bob thinks s and one in which he thinks $\neg s$.

Application to the problem case

In moving from $P(\cdot)$ to $P_{\neg K_A r}(\cdot)$, we condition not on the propositional content $\ulcorner \neg K_A r \urcorner$ but on the information $\lceil \neg K_A r \rceil$.



Moreover, the proposition $\ulcorner \neg K_B S \urcorner$ is possibility-dependent. After conditioning on $\lceil \neg K_A r \rceil$, we find that $\ulcorner \neg K_B S \urcorner = \emptyset$.

4 Dempster-shafer belief functions

Knowledge structures require a weaker notion than probability: Dempster-Shafer belief functions. They are defined by a mass function m on worlds.

- We have $m(U) \in [0, 1]$ for all members of $U \in \mathcal{P}(W)$, and $\sum_k m(W_k) = 1$
- Any other set of possibilities V receives a minimal probability $\underline{P}(V)$ determined by the maximal mass among those $U \subset V$.
- And it receives a maximal probability $\bar{P}(V)$ determined by the minimal mass among those U for which $V \subset U$.

Belief functions are basically interval-probabilities $[\underline{P}(V), \bar{P}(V)]$. If we only consider worlds, they coincide with probabilities.

Conditioning on informational content

A special case of Dempster's rule of combination, itself a generalization of Bayes' rule, covers conditioning on information V :

$$m_V(W_k \cap V) = \frac{I_k(V) \times m(W_k)}{\sum_k I_k(V) \times m(W_k)}.$$

Here $I_k(V) = 0$ if $W_k \cap V = \emptyset$ and else $I_k(V) = 1$. For any world W_k not intersecting with V , we obtain $P_V(W_k) = 0$. For worlds W_k intersecting with V , we have

$$P_V(W_k) = P \left(W_k \left| \bigcap_{k:I_k(V)=1} W_k \right. \right).$$

So P_V retains the proportions among the probabilities $P(W_k)$ for all worlds W_k intersecting with V .

5 Discussion

Dempster's rule provides a Bayes-style model for belief changes framed by knowledge structures. But there are many loose threads.

- The model hinges on a thick notion of possible worlds. What determines the trade-off between thick worlds, or thin ones with operations over them?
- Knowledge structures are very similar to the Harsanyi type spaces used in game theory. How exactly do these two match up?
- Dempster's rule also allows for Jeffrey-style updating over knowledge structures. How do such belief changes relate to what is modelled by DEL?
- It is tempting to apply the above model to other philosophical problems that seem to involve meaning shifts. But once you have a hammer. . .

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.