The Value of Uncertain Evidence

A classical result in Bayesian decision theory (Raiffa and Schlaifer 1961, ch. 4.5; Good 1967, Ramsey 1990), known as the value of evidence theorem (VET), says that, under certain conditions, a rational agent would postpone a decision in order to acquire cost-free evidence. The original VET, however, is limited to cases where the agent learns a proposition for certain, and hence may update her credences by Bayesian conditionalization. But is *uncertain* evidence worth waiting for in advance of making a decision?

The above question has not gone unnoticed in the literature. Graves (1989) showed that we can extend VET so that it holds for cases where we expect to obtain cost-free and uncertain evidence, and update our credences by using a rule called Jeffrey conditionalization. This rule requires uncertain evidence to be specified as a redistribution of the agent's credences over the propositions in some partition of the set of worlds she considers possible, without assigning absolute certainty to any particular proposition (hereafter, Jeffrey shift). To accommodate this type of uncertain evidence, Graves's argument is mobilized by two conceptual moves. The first one is that any Jeffrey shift can be specified as a sort of propositional certainty, i.e., as a proposition that receives posterior credence 1 in an enriched subjective probability space. This enrichment is achieved by adding to the original smaller space propositions about posterior probabilities attached to the members of a given partition. The second key move is to show that, under certain conditions, Bayesian conditionalization on propositions specifying the posterior credences over some partition in the enriched space is equivalent to Jeffrey conditionalization in the original smaller space.

After challenging the above key moves in Graves's argument, this paper offers an alternative extension of VET to the case of learning uncertain evidence. Instead of recasting uncertain evidence as certain in an enriched subjective probability space, the proposed view retains the uncertainty of evidence in the original smaller space, and provides a specification of this evidence by using the method of virtual evidence proposed by Pearl (1988) and developed in Chan and Darwiche (2005). According to this method, uncertain evidence can be specified as a set of likelihood ratios, where each likelihood ratio tells us how likely it is that some virtual evidence would be true given the truth of some proposition in a given partition as compared to what another proposition in that partition says about the likelihood of this virtual evidence. The virtual evidence here is meant to represent an outcome of ineffable learning experience that bears directly on the truth of propositions in that partition. For example, suppose that although after reading Tom's review Ann is not capable of expressing with certainty the proposition 'Tom's review says that her thesis's quality is good', she is nevertheless capable of assessing the likelihood of this proposition being true given that the thesis's quality is good as compared to the likelihood of this proposition given that the thesis's quality is not good.

Armed with the method of virtual evidence, I show how VET can be extended to the context of uncertain evidence. Two basic ideas underpin this extension. First, under a fairly plausible condition, updating on a piece of virtual evidence the agent considers possible, albeit cannot express as propositional certainty, can be reduced to an update method which accommodates uncertain evidence as a likelihood ratio and does not require calculating prior credences conditional on that piece of virtual evidence. Following Pearl, I call this update method 'virtual conditionalization'. Second, once we assume that the pieces of virtual evidence the agent considers possible form a partition, we can show that the expected worth of accommodating cost-free uncertain evidence by virtual conditionalization cannot be negative. And this expectation is calculated relative to the agent's prior credences over the possible pieces of virtual evidence. I then argue that using this method rather than the widely accepted Jeffrey conditionalization enables us to provide a much less cognitively demanding extension of VET to cases involving uncertain evidence.

References

Chan, H. and Darwiche, A. (2005). On the revision of probabilistic beliefs using uncertain evidence. *Artificial Intelligence*, 163(1):67 – 90.

Good, I. J. (1967). On the principle of total evidence. *British Journal for the Philosophy of Science*, 17(4):319–321.

Graves, P. R. (1989). The total evidence theorem for probability kinematics. *Philosophy of Science*, 56(2):317–324.

Pearl, J. (1988). *Probabilistic Reasoning in Intelligent Systems*. San Mateo, CA: Morgan Kaufmann.

Raiffa, H. and Schlaifer, R. (1961). *Applied Statistical Decision Theory*. Cambridge, MA: Harvard University Press.

Ramsey, F. P. (1990). Weight or the value of knowledge. *British Journal for the Philosophy of Science*, 41:1–4.