Reset reproduction of article published in *Computational Intelligence*, Vol. 10, No. 1 (February 1994), pp. 46-48. Written December 11, 1992. Reprinted July 1994.

Inference and Acceptance*

Jon Doyle DOYLE@LCS.MIT.EDU
Massachusetts Institute of Technology, Laboratory for Computer Science
Cambridge, Massachusetts 02139, U.S.A.

In [8], Henry Kyburg describes the approach he intends to pursue in studying believing and reasoning—an approach based on the agent accepting conclusions on the basis of their probability relative to a body of background knowledge and evidence—and sketches some arguments for and against this approach in the course of describing his intent. My comments here relate to some of these arguments and other points raised in his text.

Kyburg distinguishes between two views of belief and reasoning. One view hedges conclusions but posits sound grounding of these in knowledge and evidence, while the other view hedges the grounding relationship instead of (but perhaps in addition to) conclusions. He starts by observing that people tend to give answers in line with each of these views in different settings, but indulges the "human lust for uniformity and generality" by attempting to find a view that covers both rather than using each one where it seems most appropriate. His favored alternative (elaborated elsewhere) hedges inference steps rather than conclusions, but acknowledges the role of probability by making acceptance depend on probabilities of conclusions exceeding certain thresholds in certain probability measures.

Although I sympathize with Kyburg's view that a realistic theory of belief should give a large role to the idea of acceptance, I regard other elements of his approach with some suspicion. One qualm concerns his framework for using threshold probability levels to determine acceptance. Kyburg's view of acceptance permits the threshold of acceptance to vary with the context, so that, for example, one might require a higher threshold while designing nuclear power plants or excising a brain tumor than while deciding which door to use when entering a building. But Kyburg's examples all suggest that a single threshold suffices for each context. I find this implausible; different potential conclusions surely have different significances, which to my mind should influence the threshold needed for acceptance just as do differences in context. Perhaps Kyburg means for contexts to include the set of possible conclusions under discussion, but this did not seem to be his intent. Moreover, since the background knowledge and evidence also would appear to influence significance, they too might be necessary components of contexts if contexts determine thresholds. But in that case, the division of the attribution of belief into global threshold determination followed by individual acceptance or avoidance seems poorly motivated.

To put this qualm in a more general perspective, I find the notion of thresholds uniform within contexts unappealing because I think acceptance should depend on the significance or importance of beliefs as well as probability (and possibly other factors). Put in the broadest terms, I think that to justify acceptance we should interpret it as the result of a rational decision to believe based, as usual, on the utilities of the consequences of the act or condition

*Jon Doyle is supported by the USAF Rome Laboratory and DARPA under contract F30602-91-C-0018, and by National Institutes of Health Grant No. R01 LM04493 from the National Library of Medicine.

of believing and the probabilities of these different consequences. These utilities may reflect many different considerations, from computational or psychological costs "internal" to the agent to physical, monetary, societal, or other costs "external" to the agent. But whatever we take these utilities to reflect, different beliefs may have different utilities, which means that two beliefs may have the same probability with only one being rationally acceptable. Of course, the decision to believe may concern sets of beliefs, or sets of circumstances, and these may lead to acceptance policies of the sort Kyburg describes: in the circumstances at hand, accept any belief if its probability exceeds the specified threshold. But in the general case, one may decide to accept or refrain from each belief separately.

This view offers the virtue of uniformity motivating Kyburg's investigation, subsumes as special cases each of the views of reasoning he places in opposition, and explains without difficulty the pragmatic approach he mentions of using each one when it is best suited to the task at hand. Putting this last point in another way, using rationality as the guide for how one reasons explains (at least to my mind) Kyburg's opposing approaches without the need to hypothesize any particularly new criteria for acceptance.

Kyburg mentions something like this view at the very end of his paper when he refers to pragmatic efficiency as the strongest argument for acceptance. Pragmatic efficiency strikes me as too narrow a criterion unless we interpret efficiency in its most general sense as economic rationality taking into account the impact of beliefs on the agent's situation as well as computational costs.

My papers [1, 2, 3, 4, 5] provide more details on my view of rational belief and inference, generally taking a "coordinate-free" approach in which one may consider alternative ways of reading beliefs out of mental states. This permits, for example, interpreting specific beliefs and inference in any of the ways discussed by Kyburg, with no supposition that any of these enjoys some privileged status. Any privileged status would follow from the constitutional constraints reflected in the mental state space rather than the nature of the systems of interpretation. This approach suggests evaluating Kyburg's arguments for and against acceptance in terms of what difference acceptance really makes and what costs it really incurs (after the example of [6]) relative to the standard probabilist approaches.

In particular, I doubt that accepting beliefs and periodically revising them always improves on probabilistic revision. Kyburg seems to share this doubt in at least one passage of the paper, but the arguments he presents later concerning this question suffer greatly from a lack of appreciation of recent work on representing probabilistic information efficiently using general closed-world independence assumptions and qualitative representations, as exemplified by [9] and [10]. Such work casts doubts on Kyburg's claims that probabilistic representations require explicitly representing much more information than logical ones (a view true of earlier representations, but not of modern ones). The approach of [7] in addition casts doubts on Kyburg's conjecture that only three orders of magnitude suffice for describing evidential certainty. All of this work suggests that the effort involved in reasoning within a probabilistic framework need not exceed that involved in reasoning within a logical framework (and especially not that involved in reasoning within some nonmonotonic logics).

Kyburg correctly notes that while hedged reasoning may produce conclusions about probabilities, these probabilities need not have any bearing on the uncertainty of the inference. But the difficulty of probabilistic reasoning means that practical approaches may not calculate probabilities accurately in all cases. This need not affect the utility of these methods much, but it turns practical probabilistic reasoning into a system of hedged inference to hedged conclusions, with the inference hedge a general one about the accuracy of the algorithm in the case at hand. In such a system, however, the hedges in the conclusions may almost always actually express the uncertainty of the inference, making the lack of a necessary connection a philosophical rather than practical worry.

References

[1] J. Doyle. A model for deliberation, action, and introspection. AI-TR 581, Massachusetts Institute of Technology, Artificial Intelligence Laboratory, 545 Technology Square, Cambridge, MA, 02139, 1980.

[2] J. Doyle. Some theories of reasoned assumptions: An essay in rational psychology. Technical Report 83-125, Department of Computer Science, Carnegie Mellon University, Pittsburgh, PA, 1983.

[3] J. Doyle. Artificial intelligence and rational self-government. Technical Report CS-88-124, Carnegie-Mellon University Computer Science Department, 1988.

[4] J. Doyle. Constructive belief and rational representation. *Computational Intelligence*, 5(1):1–11, Feb. 1989.

[5] J. Doyle. Rational belief revision (preliminary report). In R. E. Fikes and E. Sandewall, editors, *Proceedings of the Second Conference on Principles of Knowledge Representation and Reasoning*, pages 163–174, San Mateo, CA, 1991. Morgan Kaufmann.

[6] J. Doyle. Reason maintenance and belief revision: Foundations vs. coherence theories. In P. Gärdenfors, editor, *Belief Revision*, pages 29–51. Cambridge University Press, Cambridge, 1992.

[7] M. Goldszmidt and J. Pearl. Rank-based systems: A simple approach to belief revision, belief update, and reasoning about evidence and actions. In *Proceedings of the Third International Conference on Principles of Knowledge Representation and Reasoning (KR'92)*, pages 661–672, 1992.

[8] H. E. Kyburg, Jr. Believing on the basis of the evidence. *Computational Intelligence*, 9, 1993. To appear.

[9] J. Pearl. *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference.* Morgan Kaufmann, San Mateo, CA, 1988.

[10] M. P. Wellman. Fundamental concepts of qualitative probabilistic networks. *Artificial Intelligence*, 44(3):257–304, 1990.