

Using Multi-Agent Systems to Represent Uncertainty

Joseph Y. Halpern
IBM Almaden Research Center
San Jose, CA 95120
email: halpern@almaden.ibm.com

Abstract

I consider a logical framework for modeling uncertainty, based on the use of possible worlds, that incorporates knowledge, probability, and time. This turns out to be a powerful approach for modeling many problems of interest. I show how it can be used to give insights into (among other things) several well-known puzzles.

Introduction

Uncertainty is a fundamental—and unavoidable—feature of daily life. In order to deal with uncertainty intelligently, we need to be able to represent it and reason about it. This invited talk describes one systematic approach for doing so.

Reasoning about uncertainty can be subtle. Consider the following well-known puzzles. (These puzzles are presented under the assumption that the uncertainty is quantified in terms of probability, but the issues that they bring out arise whatever method we use to represent uncertainty.)

The second-ace puzzle (Bar-Hillel & Falk 1982; Freund 1965; Shafer 1985): Suppose we have a deck with four cards: the ace and deuce of hearts, and the ace and deuce of spades. After a fair shuffle of the deck, two cards are dealt to Alice. It is easy to see that, at this point, there is a probability of $1/6$ that Alice has both aces, probability $5/6$ that Alice has at least one ace, probability $1/2$ that Alice has the ace of spades, and probability $1/2$ that Alice has the ace of hearts: Out of the six possible deals of two cards out of four, Alice has both aces in one of them, at least one ace in five of them, the ace of hearts in three of them, and the ace of spades in three of them. Alice then says “I have an ace”. Conditioning on this information, Bob computes the probability that Alice holds both aces to be $1/5$. This seems reasonable: The probability of Alice having two aces goes up if we find out she has an ace. Next, Alice says “I have the ace of spades”. Conditioning on this new information, Bob now computes the probability that Alice holds both aces to be $1/3$. Of the three deals in

which Alice holds the ace of spades, she holds both aces in one of them. As a result of learning not only that Alice holds at least one ace, but that the ace is actually the ace of spades, the conditional probability that Alice holds both aces goes up from $1/5$ to $1/3$. Similarly, if Alice had said “I have the ace of hearts”, the conditional probability that Alice holds both aces would be $1/3$.

But is this reasonable? When Bob learns that Alice has an ace, he knows that she must have either the ace of hearts or the ace of spades. Why should finding out which particular ace it is raise the conditional probability of Alice having two aces?

The Monty Hall Puzzle (Savant 1990/91; Morgan *et al.* 1991): Suppose you're on a game show and given a choice of three doors. Behind one is a car; behind the others are goats. You pick door 1. Before opening door 1, Monty Hall, the host (who knows what is behind each door), opens door 2, which has a goat. He then asks you if you still want to take what's behind door 1, or to take what's behind door 3 instead. Should you switch?

There is certainly far more to representing uncertainty than dealing with puzzles such as these. Nevertheless, the analysis of these puzzles will give us deeper insight into the process of reasoning under uncertainty and the problems involved with getting a good representation.

So how do we represent and reason about uncertainty? I shall use the *possible-worlds* framework. This is the standard approach for giving semantics to modal logic. The intuition is that besides the true state of affairs, there are a number of other possible states of affairs or “worlds”, that an agent considers possible. We can view the set of worlds that an agent considers possible as a qualitative way to measure her uncertainty. The more worlds she considers possible, the more uncertain she has as to the true state of affairs, and the less she knows. This is not quite enough for dealing with the puzzles above. We need to add two more features to the picture: time and probability. To add time, we need to have possible worlds describing not only the current state of affairs, but the state of

affairs at each time point of interest. As we shall see, it is also useful to assume that these states have some internal structure. This gives us the *multi-agent systems* framework of (Fagin *et al.* 1995). To add probability, we need to associate with each possible world a probability distribution over other possible worlds; this issue is discussed in detail in (Fagin & Halpern 1994; Halpern & Tuttle 1993).

The resulting multi-agent systems provide a powerful framework in which we can represent, in a natural way, time, knowledge, and probability. But where does the system come from? Typically, it is generated by a *protocol*. An important theme in the talk is the importance of specifying clearly the protocol generating the system. In particular (as already pointed out by Shafer 1985), this is the key to understanding puzzles such as the second-ace puzzle.

The material in this talk is largely covered in (Halpern 1995).

References

- Bar-Hillel, M., and Falk, R. 1982. Some teasers concerning conditional probabilities. *Cognition* 11:109–122.
- Fagin, R., and Halpern, J. Y. 1994. Reasoning about knowledge and probability. *Journal of the ACM* 41(2):340–367.
- Fagin, R.; Halpern, J. Y.; Moses, Y.; and Vardi, M. Y. 1995. *Reasoning about Knowledge*. Cambridge, Mass.: MIT Press.
- Freund, J. E. 1965. Puzzle or paradox? *American Statistician* 19(4):29–44.
- Halpern, J. Y., and Tuttle, M. R. 1993. Knowledge, probability, and adversaries. *Journal of the ACM* 40(4):917–962.
- Halpern, J. Y. 1995. A logical approach for reasoning about uncertainty. Research Report RJ 9972, IBM.
- Morgan, J. P.; Chaganty, N. R.; Dahiya, R. C.; and Doviak, M. J. 1991. Let's make a deal: the player's dilemma (with commentary). *The American Statistician* 45(4):284–289.
- vos Savant, M. 1990/91. Ask Marilyn. *Parade Magazine* (Sept. 9, 1990; Dec. 2, 1990; Feb. 17, 1991).
- Shafer, G. 1985. Conditional probability. *International Statistical Review* 53(3):261–277.