Is It Rational to be Logical?

(Position Paper)

Robert C. Moore
Artificial Intelligence Center
SRI International
Menlo Park, CA 94025

Abstract

The topic of this panel is whether considerations of logicality or rationality should be paramount in modelling, or reasoning about, knowledge and belief. The first reaction one ought to have to such a question is to ask why there should be any conflict between these two types of consideration. Isn't one required to be logical in order to be rational? On closer examination, the issue is not as simple as that, but I believe nevertheless that there is no fundamental conflict between logicality and rationality. This issue is most sharply focused when one examines logics of knowledge or belief. At one level, a logic of knowledge (or belief) need not raise any special questions about logicality versus rationality. That is the level at which a logic of knowledge is simply a formal theory of knowledge. Just as one can have a formal theory of arithmetic, or physics, or anything else, one can have a formal theory of knowledge. In practice, however, logics of knowledge invariably make the claim that knowledge respects the deducibility relation of some "inner" logic, which is usually the same as the "outer" logic. That is, $K(P)$ implies $K(Q)$ in the outer logic if and only P implies Q in the outer logic.

This is where the issue is really joined, because it seems quite obvious that the knowledge of real agents is neither closed under, or limited to, the inferences that are validated by any of the logics of knowledge that have been proposed, or any logic that is likely to be developed using present methodology. How should we regard logics of knowledge, then? I believe they represent idealizations that are reasonable approximations to the truth for many purposes. While no actual agent's knowledge is closed under logical consequence, outside of mathematics there seem to be few cases where this significantly affects an agent's behavior. In everyday situations, the incompleteness and uncertainty of our information plays a much greater role in placing limits on our actions than does our inability to see all the consequences of our information. Similarly, while we do draw, and it is often rational to draw, inferences from our knowledge that are not logical consequences of our knowledge, it seems that these inferences can often be modelled as logical deductions from "appropriate" assumptions. Attempts to formalize an appropriate notion of appropriateness is what has led the development of nonmonotonic logics.

To improve on the idealizations made by existing logics of knowledge requires work in two directions. To deal with the lack of closure under logical consequence requires a substantive theory of the limits of human reasoning. To deal with inferences that go beyond logical consequence requires a substantive theory of plausible inference. Important first steps have been made in both these directions, but it seems to me that larger notions of rationality have limited relevance to either enterprise. While it is true that, in general, rationality requires that we act in ways that we believe will achieve our desires, the notion that we should therefore make ourselves believe what we desire to believe strikes me as at least slightly bizzare.