A Dialogical Argument for the Usefulness of Logic in MAS

Frank Dignum

Institute for Information and CS
Utrecht University
dignum@cs.uu.nl

Liz Sonenberg

Dept. of Information Systems
The University of Melbourne
l.sonenberg@unimelb.edu.au

Abstract. Using an ancient form of argumentation in dialogue style, we respond to the article of Bruce Edmonds. We hope to achieve two things. First some refutation of claims that seem to be exaggerated. Further, to give a taste of motivations for using logic based formal methods in MAS and to note some successes.

Introduction

This joint response grew out of discussions between us on the merits of Bruce's paper, and the possible advantages and disadvantages of producing a public response. Before offering our reaction to the substantive content of Bruce's article, Frank explains why at first he did not want to reply.

"As will be evident later, it was not because I had nothing to state in defense of the use of formal logic. However, I was wondering where this discussion would ultimately lead us. There are two possible outcomes. First Bruce's arguments could be so powerful that any reply would not stand up to it. It might then lead to a situation where people would actually start following his advice and papers using formal logic to model problems in MAS would all be rejected on the basis that they are to premature. Would this stop these people from writing papers using logic? I don't really think so. However, it would probably, after some time, stop them from submitting their papers to workshops such as the RASTA 2002 workshop. Therefore the community of persons doing research in the area of MAS and social science would effectively be split up into two (even) smaller communities. These communities would miss any opportunity of cross-fertilization between those who use simulations as their main tool and those who use formal methods as their main tool. Personally I would see that as a defeat of what doing science and especially organizing workshops is all about. The exchange of ideas between all kinds of people, all with their own backgrounds, intuitions, favorite tools and capabilities, and all looking at roughly the same problems from their own perspective.

What would happen if people would write a response against Bruce's statement that would be accepted by most people in the community and thus papers based on

formal logics would still be presented? Possibly it would lead to frustration of Bruce and lead to alienation of the rest of the community. We would lose a very valuable contributor to this field."

It seemed there were some problems with responding, and others with not responding. On further reflection, however, it seemed that in the process of responding to Bruce we could shed some light on the reasons why people that use formal systems feel encouraged, or maybe even compelled, to do so. As in many situations, understanding the motivations of other people brings people closer and creates more tolerance for differences. Indeed, just as Bruce proposes that discussion between like-minded academics is not as likely to provoke progress as testing intuitions by "rubbing them against a real problem" we hope to demonstrate that a dialogue exposing some different perspectives of contributors to the same community can also be useful.

Having commented on our motivations, let us move on to the substance of Bruce's claims. We leave the floor to Liz.

An analysis of Bruce's paper and a personal story

"I read Bruce's article as having (at least) three intertwined threads: constructive criticism of paper selection criteria in MAS conferences and workshops, cogitations on the development processes within a scientific subdiscipline, and inadequately motivated or defended propositions. I'm commenting on the first two threads, leaving the reader to reflect on other elements as they choose.

Taken literally, the title of the paper "How Formal Logic Can Fail to be Useful…" suggests that the article will point to situations where the operation of a formal logic in an MAS context has not been useful – useful in the particular sense of assisting the ongoing progress of modelling or designing systems. At a surface level, this is not a surprising, or controversial, claim. Author motivations vary widely, and exploration of ideas, albeit preliminary ones, can be as viable a motivation as the perhaps more pragmatic ones of assisting system modelling or design. The observation that in a relatively new subdiscipline, research is carried out that seems to make small and not obviously useful progress, is certainly not confined to MAS! So I have no real difficulty in agreeing with the claim that no all formal logic is useful for MAS."

What do you think Frank?

"Of course it all depends on what one might define as useful in this context. As said above, I think the most important purpose of this response is to illustrate the motivations of people using formal logics and thus give an idea to those people that are not familiar with their use, what they are all about. So, let me illustrate with an example what I consider to be some useful progress by the use of formal logic.

The first time I came in contact with logic (besides the one course in logic I barely passed during my studies in computer science) was during my PhD research about 15 years ago. I was wrestling with the problem of representing a deductive database with

a number of constraints in a way such that the whole system would not collapse once one of the constraints was violated. This was a very practical problem. The systems used in that time would check whether the database complied to the constraints by checking whether the conjunction of the two parts would be consistent. This is not the case if a constraint is violated!

Once the system becomes inconsistent it is possible to derive any conceivable fact (due to the fact that one can derive anything starting from “false”). Therefore the system would become unusable and thus it would halt. In many cases, it is possible to control the environment in such a way that constraints are always fulfilled in the database. E.g. if there is a constraint on the database of bank customers that states that their accounts have to have a positive balance, then one could check whether the balance becomes negative when a customer withdraws money from his account. Whenever the balance would become negative, the withdrawal would be prevented. However, there are also constraints that cannot be enforced by the system. E.g. suppose a customer is allowed to have a negative balance of at most $100 and has to pay a monthly interest over this negative amount. How can the system enforce both constraints at the same time if the customer has a negative balance of $100 and has to pay the interest? If the system automatically takes the interest out of the account of the customer his balance is getting to low. However, the system has no way to force the customer to pay otherwise!

It is at that time I first learned about deontic logic. Using deontic logic it is possible to state that the constraints *should* be fulfilled, but the system does not become inconsistent once the constraints are not fulfilled. Instead the system enters a state that is not ideal. It is now possible to specify explicitly which actions should be undertaken to get back to a situation where all constraints are satisfied again. (e.g. go to the customer to get the money). The way this could be implemented was in principle quite straightforward and could be done with standard means.

However, this story is not meant to educate people how to solve database problems as well as to show how the introduction of deontic logic helped to model the situation in a way that led to an elegant and implementable model. Does this example give you some idea of how broadly useful a formal focus can be?"

"Frank, although you haven't expressed it in this way, this anecdote seems to provide a nice counterpoint to a claim that Bruce makes in several places – that focusing on *expressiveness* of a new language (or feature) is a red herring from an applied modelling perspective. The point is that knowing that first order predicate logic is generally expressive does not help you find an elegant, or if you prefer an appropriately compact, representation of a given set of concepts and hence does not assist the modeller. The importance of naturalness and ease of use of modelling notations is certainly widely accepted now in software engineering, and it's surely just as relevant here? Indeed Bruce himself seems to agree (at least partially) with this point in an earlier discussion of "implicit content" of a formal system (Edmonds, 2000).

Of course the next question in this context is, of course, whether you really *needed* deontic logic to model the situation. How do you respond to that?"

"Well Liz, formally this question should be answered negatively. It is possible to express most modal logics in (sorted) first order predicate logic. So, we could have

modeled the situation using first order predicate logic. However, this would not have given me the concepts of *ideal state*, *violation*, etc. that ultimately helped me solve the problem. So, in this respect I agree completely with Bruce that the choice of a certain formal tool is an entirely pragmatic choice. It should support you to easily model those aspects of the problem that are most important for your focus of research. Moreover, after modeling the problem using this formal tool it should alert you to problems that you were unaware of yet, to incompleteness of the specification of the problem and hence to suppositions that might have to be made explicit or are even wrong."

"Usually modelling a realistic problem using some form of logic does not only enhance your insights in the problem domain, but also makes one aware of some shortcomings of the logical tool that is used. In my case the solution of the problem raised several new issues, which were not so easily solved using the deontic logic available at that time. (In terms of the saying that Bruce cited: we finally got our screwdriver, which worked OK for some screws, but we discovered that we actually wanted a Philips screwdriver). This related to the fact that I wanted to not just register that the system is in a non-ideal state, but also what caused this state. Moreover, I should be able to distinguish non-ideal states where only one constraint was violated (or was violated one time) and non-ideal states where several constraints were violated (possibly more than one time). In order to solve these issues some extension would be needed to the deontic logic that I used (by introducing separate non-ideal states for all cases). This is not an isolated incident, but happens very often. The logics that are available can model a large part of the problem, but are inadequate for the last bit. Therefore, one needs to spend some time to work on the logic and find out whether the tool can be extended in a way to be helpful for the last bit as well. "

"That is a nice story Frank, and it forms some counter to Bruce's claim that researchers tend to "change the problem to suit the tool" because you have illustrated that explorations with alternative formulations play a role in helping us to see the boundaries of the usefulness of various tools. Indeed we've done this a bit ourselves looking at integrating deontic concepts into a (moderately) pragmatic BDI setting (Dignum et al, 2000). But can you add some more to respond to Bruce's arguments?"

Some things learned

"Above I discussed my first encounter with the use of logic. Since that time I have moved into research in MAS, agent communication and e-commerce. Now my exercises in the field of deontic logic helped me as well. In deontic logic one tries to capture obligations, permissions and prohibitions. However, the proposed logics are still riddled with paradoxes. Although the logic itself did thus not directly provide me with solutions for the fields I was studying, the very fact of the difficulties in the logic made me sensitive to the type of problems that one could expect. All of them have to do with a few issues, such as the balance between social norms/rules and personal autonomy, the effect of social relations on personal behaviour and how to describe

protocols/procedures without knowing in advance all possible situations in which they are executed.

I am not claiming that one should study logic in order to get these insights. It just illustrates that using logic can also have indirect effects. It can sometimes give you some insight in what are the really hard problems to solve or which are the essential concepts to model.

To illustrate the last point, let me refer to some work done in e-commerce. One of the issues in e-commerce is how trade procedures have to be modeled when the trade is performed by electronic means. Should they mimic the traditional paper-based procedures or can we leave out or change parts? In order to check the trade procedures some audit principles have been used that check whether they comply to some essential norms (such as exchanging product and money at the same time). By analyzing these audit rules again using a multi-modal logic we could pinpoint a few essential principles behind the rules. E.g. the fact that both parties should have equal knowledge about the state of the procedure and cannot repudiate such knowledge. By pinpointing this essential purpose of the steps and/or documents used in the trade procedures it became possible to determine whether this purpose could also be achieved in alternative ways and whether performing steps electronically would make an essential difference. So Liz, this is just one example of some very practical use of logic including epistemic operators!"

"Frank, this story might show that logic can be useful in some ways and so provides a counter to the claim that logic fails to be useful for modeling or designing MAS. I'd like to add to this analysis some observations from joint work we have undertaken ourselves (Dignum, Kinny & Sonenberg, 2002).

Remember this work started from the observation that in MAS, interactions between individuals need to be carefully captured (i.e. represented so they can be explicitly manipulated) and that certain relationships have received inadequate attention in the past. Specifically we argued that in open systems where independently developed agents will be cooperating and competing, traditional MAS architectures are handicapped by not distinguishing between desires and goals, and also that such architectures are not rich enough to capture the various influences on goals and intentions that arise in applications where obligations and norms are significant. In the paper we revisited the relationships between goals, desires and intentions in the BDI model with an eye to the kinds of situations in which both goals and desires should be represented, gave a brief formal analysis of desires and goals, and offered reasons why it is useful to represent them both explicitly in MAS settings.

We noted that while desires and goals are rather similar, the way in which an agent's desires influence, determine or perhaps are elevated to become goals is an important issue even within an individual agent, and one which becomes even more significant in MAS where other sources of motivation and influence on intentions arise. For example, obligations arise from interactions between pairs of agents as a result of visible, explicit commissives such as promises, or agreements to requests or contracts, but may also arise due to the "rules of the game" which apply to an interaction, i.e., they may be explicit or implicit elements of an interaction protocol. Norms, as manifest in human societies and organizations, assist in standardising the behaviour of individuals, making it easier to cooperate and/or interact within that society. Similarly,

if agents are able to be designed so they tend to follow norms, several writers have observed that knowledge of these norms can allow for easier coordination, as certain behaviours of others can be anticipated with some degree of reliability. These observations, and some logical modelling that followed in the paper formed an initial attempt to be precise about these complex and subtle relationships, and to provide some tentative first steps towards guidelines for system implementers, were on the basis of a conceptual analysis alone. We were able to document the intuitions informally first, and used the logic (albeit tentatively) to add precision, and to provide a focus for further analysis.

If one were simply to ignore the need for formalism and focus on developing and refining implementations, there would be a risk of losing key intuitions and insights that could guide further elaboration of the agent paradigm. Subsequent development of agents with concepts such as emotions, and obligations would more than likely be increasingly ad hoc, leaving no reproducible trail for subsequent analysis of run-time behaviour, other than system runs themselves - an approach long discarded in other communities. So while conceptual analysis supported by implementation and evaluation has its place, I think this is a useful illustration that so can conceptual analysis alone.

However the cogitations in Bruce's paper put a much stronger position. Bruce's argument seems to be captured in his attempt to apply to MAS the claim "In the past, premature `armchair theorising' has not helped the eventual emergence of useful theory, but rather impeded it." The suggestion then is that certain published works in logics not only "fail to be useful" but actually are *harmful*.

Bruce goes on to say "choosing the wrong kind of formal system will bias our attempts and make our task more difficult." Confined to a reflection on a given paper, or group of related papers, as might be supposed from the title, the assertion that progress is difficult if you choose a poor tool is an easy one to agree with. But Bruce wants this to be a more potent argument – indeed that working with the "wrong kind" of formal logic will bias and hence impede progress of the MAS field, a claim that he has also made elsewhere (Edmonds, 2000). In this paper we find: "there has grown up a tradition [in MAS] which discusses and compares different axiomatisations of logic and logical systems based purely on plausibility and the ability to encode particular examples (i.e. its expressive power). It is this approach that I am arguing against on the grounds that it will not be useful in either understanding or building MAS." Further, in the context of discussing different approaches to successful developments in science, Bruce talks about the importance of using data drawn directly from the phenomena under question to constrain and guide research – as contrasted with the view from the "philosopher's armchair" (Edmonds, 2000).

So Frank, as someone who has published papers on different axiomatisations of logical systems in an MAS context, how do you feel about this strong claim?"

"I think you can counter the argument that using the "wrong" logic might impede the development of MAS research by giving some examples and history of MAS. The development of KARO (Linder, 1996), a logical framework to describe agents, can very well be seen as a kind of reaction on the inadequacies of the well-known BDI

logic. In this framework agents are described using their knowledge, abilities and opportunities. The goals of the agents are described in terms of these concepts. One of the developments was that KARO is based on dynamic logic, thus introducing a temporal and causal element in the formal system. At this moment we are developing an agent programming language (3APL) directly based on the logic developed in the KARO framework (Hindriks, 2001). The first version is up and running and available through the WWW (www.cs.uu.nl/3apl). I should add that 3APL is, of course, not the only agent programming language that is directly based on formal logic. The first attempt to implement the BDI logic was made in dMARS (see (d'Inverno, 1998) for a formal description of dMARS) in the beginning of the '90s. Another agent programming language based on formal methods is ConGolog developed in Toronto (Giacomo, 2000).

This example shows two points. First it shows that, although BDI logics had many shortcomings and thus it was impossible to implement agents directly based on BDI logic alone, the development of this logic was a starting point for many others. Therefore they served as a reference point and inspiration for other formal systems that are being developed at the moment and of which many are (directly or indirectly) being implemented. Thus, whatever the shortcomings of the BDI logic may be, it has had a great positive impact on the development of the agent research.

The second point this example shows is that there are now several groups working on agent programming languages that are in the one hand directly based on some logical framework, but in the other hand are practical and implementable.

So, based on a first attempt to specify agents using a formal modal logic, several years later a number of systems are build based on improved models. All of these systems are basically single agent systems. My feeling is that we still need to go through the same development for MAS. We are now experimenting with all types of logical systems that might explain small parts of MAS. Based on the findings and all the inadequacies in a few years we will maybe build systems based on improved models.

I think that in general we can state that the BDI logic might have been prominent, or even dominant in agent theory for many years, but has never been an impediment for further developments in agent research. As shown above, there have been people who have developed theory based on BDI, and we have contributed to this ourselves of course. But there have also been many groups that developed agent platforms that were not based on BDI theory at all. So, it seems that the sheer number of different agent platforms that are only vaguely based on BDI type of theories is by itself a refutation of the position that formal logic theories might be an impediment of further developments, or indeed that a surge of interest in different formal systems might somehow signal a stagnation of the discipline."

Controversy or misunderstanding?

The above can be seen as the reflections of researchers in MAS who have been using logic for several years. Although an individual might feel that the research has been useful for the discipline, and also at a more personal level because such research

is used to progress one's career, it does not prove that the use of formal logic in MAS (and publishing about it) actually helps the field advance or helps other people to get a better understanding of it. The relevance, illusionary or otherwise, of the above examples of the use of logic in MAS depends on the importance that one might accord to the particular problems that are addressed in the examples.

Surely this point is central to this whole discussion. What is the purpose of the research people are doing in MAS? A not so daring claim is that people in this field of MAS and people in the field of Multi-Agent Based Social Simulation have widely differing interests and goals. For the people coming from social simulation the primary goal is to simulate a large group of persons in a social setting and check the influence of different factors on the society as a whole. Issues like emergent behavior are important in this setting.

People coming from computer science are often more interested in the properties of the MAS by itself. For example, can we prove that it will always terminate in a correct way or that certain features are preserved during the run of the system. Another important point is how the agents should be programmed. Can we develop a new paradigm that makes it natural and efficient to program MAS and which still allows us to prove some properties as is done with traditional programs?

Given these widely differing interests and backgrounds it is unsurprising that the preferred tools also differ widely. Coming from computer science it is clear that one would like to use formal methods to describe the models developed. Implementations connected to systems described formally are most likely to be able to exhibit some controlled behaviors. This does not necessarily mean that one would have to use logic. Nor does it mean that a run of the system should be captured within a derivation of the logic. One might use the logic to check a static description of the system against some constraints and have a separate formal (maybe transition based) model to describe the development of the system over time.

As a social scientist one probably does not want to specify the MAS up to the level where it performs a completely controlled set of actions. In that case the simulations would not add much to the theory. One needs agents that learn from experience and adapt their behavior. Concepts such as learning and evolutionary algorithms are more appropriate for this purpose. Only after extensive experimenting and recording all results it will be possible to formulate some general theses about the societies behavior in terms of the types of agents. So, we agree with Bruce that the use of logical formalisms from this viewpoint seems premature.

However, while focusing on comparisons between different research paradigms, while accepting the importance of using empirical data to help shape intuition, one can argue that the field of adaptive systems is also still in its infancy. Of course one can use it, but the features of such systems are often not well understood. For example, often it is unclear whether the outcomes of the system represent an interesting phenomenon or are merely a direct result from the choices made at the modeling or input stage. So, from this point of view one might argue that papers using these techniques are also premature, because it is still unclear how the results from the experiments can be explained.

In other words, we propose "*No experimentation without explanation*" as an adaptation of Shoham's exhortation "No notation without exploitation" (Cousins & Sho-

ham, 1994), itself a descendant of a relevant plea[1]. To take this further, we note that Bruce's final point could perhaps be captured as "logic without exploitation is like simulation without execution" whereas we suggest framing the accusation as "logic without exploitation is like simulation without explanation."

The way forward?

In the section of Bruce's paper with this same title, he states that he expects there "will be hundreds of essentially different 'species' of MAS". We completely agree with this statement. We also completely agree with his subsequent statement that there will probably not be any set of easily accessible universal principles covering all 'species'. The set of systems covered by the term MAS is simply to broad to describe it by a simple set of general principles. However, we do not agree that therefore it is no use to formulate intuitions (even vague ones) using formal logic. One should however specify explicitly for which types of MAS these intuitions would be interesting! E.g. open or closed MAS (where an open MAS is a system where agents may come from different parties and agents may enter and leave the system at will). Probably formal logics are most useful for MAS that are being designed for a specific task, such as information management, workflow management or managing business transactions.

Is there only one way forward? No, we don't think so. In the natural sciences such as physics and chemistry, practical and theoretical streams develop in parallel. Sometimes the practical work produces the most successful results, while at other times the theoretical work inspires people to try new avenues. One simply cannot say that because practical experimental work is needed, one should stop doing all work on the logical formalism until the moment it can be grounded completely in practice. Neither should one claim that we should wait with the experimental work until we have a complete formal theory with which to test theses and perform experiments.

We should give both sides some room to develop their theories and standards. Although some of the formal logics might not be applicable right away they should also not be dismissed on forehand. By keeping an open mind towards people that use a different viewpoint to research the same problem area we might get new ideas and might be able to point to those parts where verification of the logic is possible through simulation. It may also indicate new interesting theses that seem to be implied by the logic and should be verified by logic. In the other hand the experimental work might lead to new intuitions about connections between concepts in the logical framework, or important new concepts that were still missing.

Whether any of the above types of results count as positive contributions to the development of MAS research depends of course on the problem discussed and the perceived importance for MAS. Bruce seems to want us to accept that the only positive contributions would be by implemented systems based on the formal system (though

[1] McDermott, D. (1978), "Tarskian semantics, or no notation without denotation!", *Cognitive Science* 2(3), 277--282.

elsewhere notes that both "foundational" and "empirical" approaches have their part to play (Edmonds, 2000). While we are enthusiastically in agreement with this last observation, Frank's examples show that some progress can also be made on a conceptual level only. Sometimes waiting for the formal systems to be actually exploited and hence to check the results would actually impede progress. If an implementation is not very easily done it may take considerable time before experimental results supporting or attacking the formal model would become available. In the mean time it might be possible to discuss and point to some possible fallacies of the model on a conceptual level without having an implementation. This is similar to the observation that one does not have to wait for a system to be completely implemented before trying to validate (some of) the requirements based on the conceptual specifications.

Concluding remarks

In essence we read much of Bruce's commentary and ten suggestions for "ways forward" as addressing the choice of criteria by which papers should be selected for publication in various settings. The key perspective is that of "the audience." Bruce makes the point that the author having had their own good reasons for writing the work, does not ensure that good reasons abound for the reader to invest *their* time or effort to learn about it. Although the above seems very simple, the difficulty lays in the fact that people have different conceptions of the audience of their paper. Due to their different background they expect a different context to suffice for their audience. We should all be aware of this problem and try to be explicit as possible. One does not only need an account of what has been done, but also an account of why this is relevant for the development of the field. Sometimes the contribution is just that the problem can be formally specified and thus analyzed using other formal tools. Sometimes the contribution will be the signaling of new problems based on a formal analysis of an area. Sometimes, a way for resolving some issues is proposed based on the formal model. Such explicit discussions will enable the reader to evaluate how relevant this paper is for his own research.

There is no doubt that many papers are written that cause many reviewers to regret their existence and the time required to proffer a constructive response. Indeed, in reviewing such papers I have from time to time found myself tempted to adopt the "sublimely evocative phrase '*...this paper fills a much needed gap in the literature*' " (Jackson, 1997)[2]. But admitting that more papers are written than are worth reading, and claiming that an entire strand of work is not only useless, but harmful, are widely differing positions.

So, can we conclude with definite recommendations for the field? We don’t believe so. That is, we don’t believe we are influential enough to determine the direction of

[2] On gaps, see also: The Society for Preservation of Gaps in the Literature, http://www.econ.uiuc.edu/~roger/gaps.html (accessed September 2002)

this complete community. People have many motivations to write and submit their papers. Those motivations will only be marginally effected by our efforts. However, we do give one recommendation for judging any submitted paper. As a reviewer, ask yourself the question: “Is this a much needed paper filling a gap in the literature or is it a paper filling a much needed gap in the literature?” If the first applies accept it, otherwise reject it.

Acknowledgement

This response was developed while Frank Dignum was visiting The University of Melbourne partially supported by a grant from the University of Melbourne International Collaborative Grant Scheme. The paper would not have taken its present shape without us having the type of dialogues in person as are represented in a stylized form in this paper.

References

1. R. Bons, F. Dignum, R. Lee, and Y.-H. Tan. *A Formal Analysis of Auditing Principles for Electronic Trade Procedures*. In International Journal of Electronic Commerce, Vol.5, Nr.1, 2000, pages 57-82.
2. Frank Dignum, David Morley, Liz Sonenberg and Lawrence Cavedon, "Towards socially sophisticated BDI agents," in Proceedings of the Fourth International Conference on MultiAgent Systems (ICMAS 2000) Boston, USA, pp 111-118
3. Frank Dignum, David Kinny, Liz Sonenberg, "From Desires, Obligations and Norms, to Goals" *Cognitive Science Quarterly*, 2(3/4) 2002 *in press*
4. B. Edmonds, "The Purpose and Place of Formal Systems in the Development of Science," CPM Report 00-75, September 2000, accessed at http://cfpm.org/cpmrep75.html (September 2002)
5. G. de Giacomo, Y. Lesperance and H. Levesque. ConGolog, a concurrent programming language based on the situation calculus. *Artificial Intelligence*, 121(1-2), pp.109-169, 2000.
6. K.V. Hindriks, *Agent Programming Languages. Programming with Mental Models*. PhD Thesis, Utrecht University, 2001.
7. M. d'Inverno, D. Kinny, M. Luck, and M. Wooldridge. A Formal Specification of dMARS. In M. P. Singh, A. S. Rao, and M. Wooldridge, (eds.): *Intelligent Agents IV S*pringer-Verlag LNAI 1365, 1998.
8. Jackson, "Chinese Acrobatics, and Old-Time Brewery, and the `Much Needed Gap': The Life of Mathematical Reviews", *Notices of the American Mathematical Society*, vol 44, no. 3, March 1997, pages 330-337; accessed at http://www.ams.org/notices/199703/comm-mr.pdf (September 2002)
9. B. van Linder, *Modal Logics for Rational Agents*, Ph.D. thesis, Utrecht University, 1996.
10. Y Shoham & S B Cousins, "Logics of mental attitudes in AI," In G. Lakemeyer and B. Nebel, editors, Foundations of Knowledge Representation and

Reasoning, volume 810 of Lecture Notes in Computer Science (subseries LNAI), pages 296-309. Springer-Verlag, 1994.