
Applying Labelled Deductive Systems and Multi-agent Systems to Source-based Argumentation

DOUGLAS WALTON, *Department of Philosophy, University of Winnipeg, Winnipeg, Manitoba, R3B 2E9 Canada.*
E-mail: *walton@io.uwinnipeg.ca*

Abstract

In this paper, it is shown how labelled deductive systems and multi-agent systems can be used to evaluate argumentation that is source-based and depends on a credibility function. When agents engage in argumentation in dialogues, each agent has a credibility function that can be adjusted upwards or downwards by certain types of arguments brought forward by the other agent in the dialogue. One type is the argument against the person or *argumentum ad hominem*, in which personal attack on one party's character is used to attack his argument. Another is the appeal to expert opinion, traditionally associated with the informal fallacy called the *argumentum ad verecundiam*. Both types of argument are frequently used in legal argumentation, sometimes fallaciously and sometimes non-fallaciously. The problem is to have some kind of structure that can be used to identify the form of the argument, and to appraise how the form has been used in specific cases, in order to determine which cases represent fallacious uses of argumentation, and which do not. Part of the system is dialectical, meaning that it views the argumentation in light of its use in a collaborative communicative exchange, or so-called dialogue, between two agents.

Keywords: Legal argumentation, credibility, dialogues, witness testimony, fallacies, evidence, *ad hominem* arguments, practical reasoning, *ad verecundiam* arguments, agent communication, expert system.

1 Introduction

A source-based argument is one in which evaluation of the argument depends not only on the structure of the inference used in the argument, but also on some assessment of the sources of the premisses [25, pp. 271–281]. By the 'source' of a premiss is meant the evidence on which it was supposedly based. For example, a source could be a set of medical files, or some person or organization that put the proposition forward, or it could be a chain of prior arguments. Two of the most common forms of source-based arguments are the appeal to expert opinion and the personal attack or *ad hominem* argument. These two types of argumentation are singled out for special attention in this investigation, because they have been acknowledged as especially important in the history of logic, as traditional informal fallacies, and because they are both extremely common and important in argumentation, and most notably in legal argumentation. But the problem of how source-based arguments should be evaluated in applied logic should be of general interest to researchers in the area of logic and computing.

The fact is that in many cases where conclusions are drawn from a database in everyday argumentation, the decision on what conclusion to accept is based to a considerable extent on the sources of the information used as the premisses of the argument, and on some weighting of the relative credibility of these sources. The evaluation of source-based reasoning has never been taken very seriously as a problem for logic in the past however, probably because

there just seemed to be no formal structure that could be applicable to this kind of problem. Now with the advent of new methods used in computing, that picture is changing. There are formal methods available that are at least applicable to some parts of the problem. In this investigation, the parts are put together in a sequence. The four methods brought to bear in this investigation are labelled deductive systems, systems of practical (goal-directed) reasoning, multi-agent systems, and formal dialectical systems. The result shown by the investigation is that these structures, currently in use in various areas of computing, are shown to be applicable at different stages of the evaluation process needed to deal with realistic cases of source-based argumentation. The method proposed is not one seamless formal structure, but a series of structures, some of which have been formalized. Other structures have been partly formalized, or can be formalized in different ways. The whole method shows promise of being precise enough to lead to formal models through the refinement of concepts currently in use in AI.

2 **Ad hominem and ad verecundiam arguments**

The appeal to expert opinion and the personal attack (*ad hominem*) types of argument are particularly important and common in legal argumentation, especially in witness testimony in a trial. What is common to both these types of argument is that evaluation of the strength of the argument depends crucially on an evaluation of the credibility of a source who backs up the premisses of the argument. One is the personal attack or *ad hominem* type of argument, where the one party in a dialogue mounts a personal attack against the other party, and then uses this personal attack as a basis on which to criticize or refute that other party's argumentation. This form of argumentation is very common in legal cases where the testimony of a witness is attacked on grounds of 'impeachment' of the witness, i.e. attacking the witness as an unreliable source [19, pp. 345–349]. Traditionally the *ad hominem* argument was classified in logic as a fallacy [9], but recent research has shown more and more that *ad hominem* arguments can often be reasonable, and in many instances are not fallacious at all [23]. The other type of argument is the appeal to expert opinion (or more broadly, appeal to authority), in which one party in a dialogue supports her conclusion by citing the opinion of an expert in a domain of knowledge, where the expert opinion is put forward as supporting the arguer's conclusion. This form of argument is extremely common in the use of expert testimony in legal trials. Its use in law has recently been a subject of controversy, as 'junk science' in expert testimony has been criticized [12]. This argument too was traditionally classified as a fallacy, as noted above. It was, and still is called the *ad verecundiam* fallacy, which can be (roughly, although intelligibly, to someone not familiar with this expression) translated as appeal to respect (for authority)—see [9, 23]. Both the appeal to expert opinion type of argument [24] and the *ad hominem* argument [25], however, can be quite reasonable in many cases. Both are in fact so widely used in common legal reasoning of many kinds that they are well worth studying from that viewpoint alone.

What is very important in the job of attempting to provide criteria for the evaluation of both these types of arguments is the recognition that both depend essentially on the credibility of one of the participants in the argument. In the case of the *ad hominem*, the argument is essentially a personal attack on the character of the one participant, often concentrating on her character for veracity, sincerity, or trustworthiness as a collaborative participant in the dialogue. Then the character attack is used to impugn the credibility of the arguer. And then this reduction in the credibility of the arguer is used as a basis for inferring that the plausibility

value of the participant's argument should be reduced. In other words, the *ad hominem* argument is a way of attacking somebody's argument by attacking that person's character, and then transferring the character evaluation to the evaluation of that person's argument. Clearly, what is central to this type of argument are the three assumptions that a participant in an argument has what is called credibility, that this credibility depends on the participant's personal character (or what that is taken to be by others), and that the participant's credibility is a factor in evaluating his argument.

What needs to be done to model these types of source-based arguments in formal dialectic is to introduce a credibility function from a participant in a dialogue to that participant's argument used in the dialogue, so that an evaluation of a participant's credibility (as high or low) can affect the evaluation of that participant's argument that she proposed in a given instance. As an example, suppose that a proponent has put forward an argument, and that argument is evaluated as strong, so that the premisses provide strong support to make the conclusion plausible. But then let's suppose that the respondent in the dialogue attacks the character for veracity of the proponent with some plausible allegations, mounting a plausible and well-supported *ad hominem* argument against the proponent. The next event in the dialogue should be a devaluation of the credibility of the proponent, which would, in turn, lead to a reduction in the plausibility value of the argument that the proponent has advanced. The credibility function takes a downward evaluation of the arguer's credibility, in this case, to a downward evaluation of the plausibility value of that participant's argument.

In other cases, the credibility function could take the evaluation of an argument the other way. Suppose that a proponent has put forward an argument that is weak, or not very plausible, because little in the way of evidence has been advanced to support the argument. But then suppose that the proponent's credibility is enhanced. This upward evaluation of credibility could occur for various reasons. It could be that the proponent is shown to be an expert in a domain of knowledge into which his argument falls. Or it could be that the proponent is shown to be a person with an established record of achievements of a kind that makes her character appear to have high ethotic value—in other words, she is seen to be a person of great integrity that we would respect as trustworthy and having *gravitas*, or seriousness of character. On the grounds of this high credibility rating, because of the credibility function, the plausibility value of her argument is then raised.

Formal deductive logic appears to offer no way of modelling the operation of the credibility function in a way that would be useful to help in evaluating appeals to expert opinion and *ad hominem* arguments. Of course, we have expert systems that model the reasoning of an expert in a domain of knowledge. But by themselves, such systems do not provide a method for using the evaluation of an expert as a source as part of the means for drawing conclusions on what to accept in a typical legal kind of case, for example, when the experts disagree. What would be useful would be some general method for reasoning from a database, where the propositions in the database come from different sources, and lead to different conclusions, and where there is a need to choose between these conclusions on some rational basis, even if it is inconclusive. We now turn to a general method of logical reasoning that is useful for this purpose.

3 Labelled deductive systems

A labelled deductive system [7, p. 67] begins with a database Δ , containing 'declarative units' of the form $t : A$, where t is a label and A is a formula. The label t is said to 'annotate'

the formula A . The use of annotations attached to formulas (propositions) can have various applications. One [7, p. 67] is to indicate what conclusion should be drawn from a database that contain several arguments leading to different, or even opposite conclusions. To cite a simple example given by [7, p. 67], suppose we have a database containing two propositions, A and B , and the database also contains the annotation t for A , and the annotation s for B . Suppose it is also known in the database that s has a higher evidential priority than t . Finally, suppose that the database also contains the two conditional propositions, $A \rightarrow \neg C$ and $B \rightarrow C$, and that both conditionals have equal priority. In this situation, there are two lines of reasoning available, represented by the following pair of valid arguments.

$$\frac{A \rightarrow \neg C}{A} \quad \frac{B \rightarrow C}{B}$$

The conclusion that should be derived is C . The reason is that s , which annotates B , has a higher priority than t , which annotates A . C is better supported than $\neg C$ because the greater weight of evidence that B has (greater than that attached to A) is transferred to C by the inference. In a kind of case where we have to choose between C or $\neg C$, the comparative weight of evidence available in the database, on balance, tilts towards acceptance of C .

What sort of information could be contained in a label? In the example above, it was a priority, indicating that one proposition was based on more weighty evidence than another. In another example [6, p. 311], it could be an annotation giving the source of evidence on the basis of which a proposition was accepted into a database, like a medical file, or the pronouncement of panel of social workers. In another case [6, p. 311], it could be the name of a person (source) who put the proposition in question forward, along with some indicator of the reliability of that person as a source of data. For example, a label could be (John, 0.7), where the figure is some comparative indication of how reliable John is as a source, compared to other sources annotated in the database. As indicated in the case above, these labels could be an important factor in deciding which conclusion to draw from the database,

LDS is a big step forward in the evaluation of *ad hominem* and *ad verecundiam* arguments, because it enables us to base our evaluation of such arguments on a label indicating a comparative assessment of the source of the propositions that were put forward. If an appeal to expert opinion, for example, was based on a source that is weak, this data can be used to assist with our evaluation of reasoning in which the appeal is a part. When an *ad hominem* attack is used to discredit a source as untrustworthy, at least we now have some basis for evaluating the argument in relation to the label that annotates the source. But is that the end of these matters? It would seem not. For the big problem with *ad hominem* arguments as fallacies is one of relevance [25]. For example, in legal argumentation in court, *ad hominem* arguments are sometimes relevant. But the problem is that they can be extremely powerful in persuading a jury (that someone is guilty of a crime, for example), even when they are not relevant. The big problem with *ad verecundiam* arguments is that juries are frequently too awed by experts, even intimidated by them, and a cross-examining attorney needs to ask the right critical questions so that assumptions are revealed, and the testimony of the expert is put in a right perspective, where important qualifications may need to be pointed out [24]. To get a grip on how these arguments can be used fallaciously, we need to better understand how they sometimes deceptively look persuasive, even when they should not be. We need to understand more about the source, and more about how that source should be questioned by a critic who has to deal with the *ad hominem* or *ad verecundiam* argument critically.

Experts who testify in court are often practitioners of some skill, and to examine the expert testimony, an attorney needs to ask many practical questions about how something was done, or how it should have been done. When an expert's credibility is attacked by an *ad hominem* argument, questions of competent execution of actions feature prominently in the examination. It would seem reasonable that, to supplement deductive logic, the kind of goal-directed reasoning called practical reasoning would need to be taken into account. Although practical reasoning has not yet been formalized in the way that deductive systems of reasoning have, it does have a kind of structure of a kind that shows its potential usefulness in modelling *ad hominem* and *ad verecundiam* argumentation.

4 Practical reasoning

Practical reasoning, according to the accounts given in [4, 2, 1] and [22] is a goal-directed, knowledge-based, action-guiding species of reasoning that combines stated goals with possible alternative courses of action, in relation to an agent's knowledge of its given circumstances. Practical reasoning is said to be carried out by an *agent*, an entity that is capable of intelligent action on the basis of observing its circumstances, and using this information to guide its actions [22, p. 191]. An agent can have some information incoming on the consequences of its actions, as these are observed to occur, and can adjust its subsequent actions accordingly, and this capability was described as feedback, in [22, pp. 144–158]. This conception of practical reasoning has long been known in philosophy, stemming from Aristotle's account of practical and theoretical reasoning. According to Aristotle, practical wisdom (*phronesis*) can be distinguished from the abstract kind of reasoning (*episteme*) used in theoretical science. Theoretical reasoning, in the Aristotelian view—an example would be Euclidean geometry—is based on necessary demonstrations (deductive, or in Aristotle's theory, syllogistic argumentation). Practical reasoning is goal-directed reasoning by an agent, aware of its circumstances in a particular given case, and using its knowledge of these circumstances to guide its actions intelligently. Practical reasoning works mainly by applying generalizations that are open to exceptions in special cases, and that require judgment skills to be used in fitting the particular to the general. But the importance of practical reasoning was neglected by subsequent generations of philosophers, as Aristotle's account of theoretical reasoning—his theory of the syllogism—became the dominant preoccupation of the teaching of logic, and philosophy became a more and more abstract and purely theoretical subject. With the supremacy of experimental science in the Enlightenment programme came a valuing of what was taken to be exact and objective scientific and theoretical reasoning as the only kind of correct reasoning worth paying serious attention to in matters of importance. The notion of practical reasoning was ignored, and was not taught as a kind of reasoning that was significant for logic. But now with the advent of AI and robotics, this imbalance has been corrected, and practical reasoning has finally been accorded a place as a kind of reasoning that is important in logical thinking. But what exactly is practical reasoning?

In the account presented in [22], practical reasoning is a chaining together by an agent of what are called practical inferences. A practical inference is defined as having two characteristic types of premisses. One states that the agent has a particular goal. The other states a means that the agent could use, in its given circumstances, to carry out that goal. According to the account given in [22, p. 85], a practical inference has the following general form. The letters *A, B, C, . . .*, stand for things brought about by agents, which may be thought of as contingent propositions that can be made true by an agent, and are called 'states of affairs'. For

purposes of exposition, the agent is referred to below by the use of the first- person pronouns ‘I’, ‘my’, and so forth.

(*PInf.*) *A* is my goal.
 To bring about *A*, I must first bring about *B*.
 Therefore I must bring about *B*.

The ‘must’ is intended to express a ‘practical ought’, conveying the idea that the agent, if it is to be practically rational, should become committed to bringing about *B*, once it knows that it is committed to bringing about *A*, and that bringing about *B* is necessary to bring about *A*. The accounts of practical inference given by Clarke [4] and Audi [1] are roughly similar to (*PInf.*), except that they use ‘wants’ or ‘intentions’ of the agent in the major premiss, and the beliefs of the agent in the minor premiss. Their accounts are examples of the ‘belief-desire’ model of practical reasoning that tends to be the dominant theory of practical reasoning, which could be classified as psychologistic. The commitment-based account of [22] is inherently different, because it is meant to be used to furnish a formal structure for the analysis and evaluation of arguments in applied logic. It could be classified as normative, as opposed to being psychologistic in nature. The normative type of theory is commitment-based (or acceptance based), rather than belief-based. It has arisen from Hamblin’s notion of commitment as a dialectical tool to be used for the evaluation of arguments associated with the traditional informal fallacies [9–11]. The normative theory does not use as premisses the intentions or wants of any real agent. Instead, the premiss of the practical inference represents the agent’s goals, as far as these can be determined from what the agent has gone on record as saying in a dialogue.

The form of inference (*PInf.*) represents the so-called necessary conditions schema for practical inference. The corresponding sufficient condition schema is comparable to (*PInf.*), except that the second premiss states a sufficient condition relation. In any actual case, the necessary and sufficient schemata are chained together in sequences that include instances of both of them at various points in the sequence. Once a sequence of practical reasoning has been put forward by a proponent in a dialogue, a weight of presumption is put in place making the conclusion of the practical reasoning conditionally plausible, based on the argument provided by the premisses. But by questioning the sequence of practical reasoning, using appropriate critical questions to match the reasoning used in the case, the other party in the dialogue can rebut a practical inference.

Matching the form of inference (*PInf.*) is the following set of five critical questions.

- CQ1. Are there alternative possible courses of action to *B*?
- CQ2. Is *B* the best (or most acceptable) of the alternatives?
- CQ3. Do I have goals other than *A* that ought to be considered?
- CQ4. Is it possible to bring about *B* in the given circumstance?
- CQ5. Does *B* have known bad consequences that ought to be considered?

As practical reasoning is used in given case, an argument of the form (*PInf.*) shifts a weight of plausibility from the premisses to the conclusion, indicating that the agent should go ahead with making it true that *B*. As noted above, practical reasoning is generally defeasible, and does not prove a proposition absolutely, ending a dialogue. Instead, it makes a conclusion tentatively plausible, subject to further knowledge coming in that might indicate that the original conclusion should be reconsidered, or even retracted. Plausible reasoning [17] is

seen as inherently presumptive in nature. Asking any one of the above five critical questions, in a given case, shifts the weight of presumption back onto the other party. Until the question is answered adequately, the weight of presumption in favour of the conclusion of the original (*Pinf.*) type of inference is suspended. But if the question is answered adequately, the weight of presumption in favour of the original argument is restored (or perhaps even increased, if the reply presents strong evidence supporting the argument).

In practical reasoning, an agent needs to be aware of incoming messages from another agent, and needs to modify the conclusions of her reasoning in accord with the received information. An agent also needs to respond to such messages in light of the new information. But what exactly is an agent, anyway. It is this fundamental question that is so vital to gaining further insight into how the *ad hominem* and *ad verecundiam* arguments work, and should be evaluated. In such source-based arguments, an attack on a source, or the assessment of a source, is typically based on the characteristics of the source as an agent. If the source has been shown to be incompetent, or a liar, that is a basis for downgrading the source's credibility. If a source advocates a goal, but then apparently acts contrary to that goal, that is evidence that the source is of questionable reliability. In all such cases, the basis of how we should annotate the source as credible or not appears to lie in a prior assessment of certain kinds of characteristic qualities of the source as an autonomous agent.

5 What is an agent?

The development of software agents in computer science has provoked the asking of the definitional and philosophical question, 'What is an agent?' Franklin and Graesser [5] have surveyed a number of proposed definitions offered by workers doing agent research. Among the characteristics stressed in various of these definitions are the abilities of agents to perform 'autonomous execution' of actions (p. 22), to 'perform domain oriented reasoning' (p. 22), to 'perceive its environment through sensors' (p. 22), to 'act on its environment' (p. 22), to 'realize a set of goals and tasks' (p. 22), to 'act autonomously' (p. 22), to perceive, affect and interpret 'dynamic conditions in the environment' (p. 22), to 'employ knowledge of the user's goals or desires' in carrying out 'some set of operations' (p. 23), to 'engage in dialogs and negotiate and coordinate transfers of information' (p. 23), to have 'some degree of independence and autonomy' (p. 23), to carry out 'autonomous, purposeful action in the real world' (p. 24), to be 'autonomous, goal-oriented, collaborative, flexible, self-starting, and to have character, adaptiveness, mobility and communicative skill' (p. 24). This list appears to be quite extensive and varied, but the central ideas in the field have been expressed in a systematic analysis of eight characteristics of an agent by Wooldridge and Jennings.

Wooldridge and Jennings [27, pp. 116–117] distinguish between two usages of the term 'agent', a stronger and a weaker use. According to the weak notion, an agent is a computer system that has the following four properties (p. 116).

1. *Autonomy*, meaning the agent has control over its actions and internal states.
2. *Social Ability*, meaning that an agent can interact linguistically with other agents.
3. *Reactivity*, meaning that an agent perceives its environment and reacts to changes in it.
4. *Pro-activeness*, meaning that an agent can take the initiative in its goal-directed actions, so that it is not just responding to these changes in its environment.

According to the strong notion, an agent has the following four additional properties (p. 117).

5. *Mobility*, meaning that an agent can move around an electronic network.
6. *Veracity*, meaning that an agent will not knowingly communicate false information.
7. *Benevolence*, meaning that an agent will do what is asked, and not have conflicting goals.
8. *Rationality*, meaning that an agent will act in order to achieve its goals, and not prevent its goals from being achieved (in line with its beliefs about these matters).

According to Wooldridge and Jennings [27], the weak usage of the term ‘agent’ is ‘relatively uncontentious’ in computer science, whereas the strong usage is ‘potentially more contentious’.

For the purpose of this investigation, the characteristics of an agent can be grouped under two general headings. One is the reasoning used by the agent as it perceives its external circumstances, and its ability to take into account its knowledge of these circumstances as it carries out goal-directed actions (and perceives the effects of these actions on the changing external circumstances). The second group of characteristics have to do with communication with other agents. The same abilities are used, but instead of acting on ‘circumstances’ the agent is acting (mainly linguistically) on other agents, who respond in dialogues to these linguistic actions (speech acts). The first group of characteristics has been studied in philosophy under the headings of action theory and practical reasoning, as outlined above. The second group has been studied in the area called argumentation theory, although the central kind of speech act studied has been that of argument. Clearly agents can interact in other ways than by arguing with each other. For example, explanation is an important form of communicative interaction for agent technology. But argumentation is important, and in many ways central.

Multi-agent reasoning poses a number of philosophical questions, and also suggests a number of directions in which the field of applied (practical) logic (argumentation theory) needs to be extended. At the same time, research on multi-agent reasoning can profit from investigations of closely allied subjects currently being studied in work on argumentation and informal fallacies. According to Jennings and Wooldridge [13, p. 364], a major problem with multi-agent systems is that ‘the overall system is unpredictable and nondeterministic: which agents will interact with others in which ways to achieve what cannot be determined in advance’. What is needed is ‘a sophisticated means of dealing with incomplete and conflicting viewpoints’ so that agents can ‘help with decision support tasks’ [13, p. 365]. In fact what is needed in multi-agent technology is a systematic taxonomy of the different types of dialogue in which argumentation is most often used, so that argumentation and practical reasoning between agents can be understood and evaluated as a contribution to a goal-directed collaborative dialogue exchange. As shown below, the new pragmatic field of applied logic can offer such a systematic framework, based on case studies of argumentation.

6 Types of dialogue

To more deeply understand how source-based reasoning needs to be evaluated, the idea has to be introduced that an agent can engage in reasoned dialogue with another agent. A typical type of dialogue that agents engage in with each other is deliberation, where a decision needs to be made as a basis for some kind of collaborative action of a group to solve a common problem. But practical reasoning is used in other types of conversational exchanges between two parties called ‘dialogues’ or types of dialogue as well—see Table 1—and it needs to be evaluated within a context of use specified by the normative requirements for the kind of dialogue it was supposedly part of. Deliberation, as noted above, arises out of the need to

take action in a given situation, and the purpose of the deliberation is to indicate the most prudent course of action in that situation. Deliberation can be a solitary procedure in some cases, where the same agent looks at the arguments on both sides of a dilemma, or choice between two courses of action, but in many cases it can be seen as a collaborative type of dialogue where two agents (or several) look at two (or several) sides of a problem to decide what to do.

As noted in Section 4, a practical inference does have logical form, and defects in the form of the reasoning can be a kind of logical failure. But what has been stressed here is the evaluation of practical reasoning in a context of use. A given instance of the use of practical reasoning is pragmatically evaluated as correct or not, in a given case, insofar as it was used in such a way that it contributed to the type of dialogue that the participants were supposed to be engaging in, in that case. An example of a deliberation type of dialogue exchange would be a town hall meeting, let's say in a case where the citizens in a neighbourhood have called a meeting to discuss whether they should go ahead and pay for a new sewer system in their district, or wait for another year.

Recent developments in applied logic of argument evaluation have taken a pragmatic turn towards evaluating arguments as used in a given case with respect to how well the argument contributes to a goal-directed type of dialogue or conversational exchange in which two parties are attempting to reason together. This approach is called 'dialectical', in the ancient Greek sense (not the Marxist–Hegelian sense), meaning that the two parties are supposed to be collaboratively (although partly contestively) taking part in a rule-governed exchange of viewpoints in a dialogue where there is a difference of opinions or viewpoints, so that one party argues for the one side of an issue, and the other party argues on the opposed side. The dialectical approach to the evaluation of arguments has distinguished a number of different types of dialogue in which two parties may argue. The six basic types are indicated in Table 1. These six types of dialogue represent conventional kinds of goal-directed conversational exchanges representing different contexts in which an argument can be used, in a given case, to prove a point.

These types of dialogue can contain adversarial, partisan argumentation, but they also require a certain amount of cooperativeness in following conventions of polite conversation if the participants are to have a successful argument that contributes to the goal of the dialogue. Participation in a dialogue is therefore to be in accord with the cooperativeness principle (*CP*) of Grice [8, p. 67]: 'Make your conversational contribution such as is required, at the stage at which it occurs, by the accepted purpose or direction of the talk exchange in which you are engaged'. Each type of dialogue [20, 21] has four stages—an opening stage, a confrontation stage, an argumentation stage, and a closing stage. An argument is evaluated as reasonable or not in a given context of dialogue, to the extent that it contributes or not to the goal of the dialogue at the particular stage where it was used. An argument is fallacious if used in such a way that it blocks or goes against the goals of the dialogue of which it was supposed to be a part [23]. This normative framework of the types of dialogue is used to evaluate an argument as used in a given case. First, you have to look at the context of the case, and make an assumption about what type of dialogue it is supposed to represent. Then you have to identify the argument used in the case (its premisses and conclusion), and then judge how well that argument was used to contribute to the goals of that type of dialogue.

The purpose of the *critical discussion*, is to resolve the initial conflict of opinions, one way or the other. In this type of persuasion dialogue, in other words, the dialogue is only successful if the conflict of opinions is resolved, so that one party is the winner and the other

TABLE 1. Types of dialogue

TYPE OF DIALOGUE	INITIAL SITUATION	PARTICIPANT'S GOAL	GOAL OF DIALOGUE
<i>Persuasion</i>	Conflict of opinions	Persuade other party	Resolve or clarify issue
<i>Inquiry</i>	Need to have proof	Find and verify evidence	Prove (disprove) hypothesis
<i>Negotiation</i>	Conflict of interests	Get what you most want	Reasonable settlement
<i>Information-seeking</i>	One party lacks information	Acquire or give information	Exchange information
<i>Deliberation</i>	Dilemma or practical choice	Co-ordinate goals and actions	Decide best course of action
<i>Eristic</i>	Personal conflict	Verbally hit out at opponent	Reveal deeper conflict

is the loser. But of course, in many persuasion dialogues, this goal is not so strict, for the dialogue can be successful if real light is thrown on the issue, even if the conflict is not resolved and neither party ‘changes her mind’ as a result of the discussion. So in this type of persuasion dialogue, the goal is that of throwing light on the positions of both sides. In the *maieutic function* of dialogue [26], probing arguments in a dialogue clarify a participant’s commitments and strengthen her arguments so that she gains a deeper insight into, not only the other party’s position, but her own as well. In either event, the participants in a persuasion dialogue need to use relevant arguments that are based on good evidence that supports their contentions, and have to allow and even encourage the other side to present its strongest arguments to support its thesis. Only by this interaction of strong partisan arguments coming from both sides will the issue really be tested, so that the goal of the dialogue is achieved.

The goal of the inquiry is to prove that a particular proposition is true (or false) by using only premisses that are very well established and verified, so that (ideally) there should be no need to go back and have to retract them. Whereas the persuasion type of dialogue is fairly free in allowing retractions [26], the inquiry goes to great lengths to eliminate the need for retraction altogether. The principal characteristic of the argumentation in an inquiry is cumulateness, meaning that the set of commitments used as premisses is always getting larger and larger as the line of argument proceeds. The semantics for intuitionistic logic given by Kripke [14] is a very good model of the kind of cumulative argumentation characteristic of the inquiry. The nodes in the tree structure represent ‘evidential situations’ or points of time where propositions are ‘verified’ or ‘not verified’. As a line of argumentation goes along the tree, away from the root, more and more propositions are verified, and none are ever retracted

or ‘unverified’.

An example of an inquiry would be an investigation into an air crash disaster. The goal is to assemble all the available evidence, being as complete as possible, and then in a later stage of the inquiry, draw conclusions only from propositions that have been carefully established, using a high standard of proof.

The purpose of negotiation dialogue is, to use the colloquial expression, ‘to make a deal’. The initial stage giving rise to the negotiation dialogue is that both parties have certain interests at stake (it could be financial interests, but other interests, like prestige, could also be at stake), and they can only maximize their share of these interests by making trade-offs with the other party. The goal of each participant should be to get what is most important to them, while sacrificing some interests that they see as less important in their order of priorities. The goal of negotiation overall as a type of dialogue is to reach a settlement that both parties are satisfied with, as what they want most. But it is important that neither side feel cheated, or that they got ‘a bad deal’. The settlement should ideally be one that both sides ‘can live with’. So negotiation should not be wholly adversarial in nature, although it is, of course, as so often stressed, a highly adversarial kind of argumentation. Negotiation has been extensively studied in recent research in multi-agent systems, and recognized as important for fields of application like electronic commerce [18].

The aims of negotiation are quite different from those of persuasion dialogue. In persuasion dialogue, the correct or successful type of argument is one that is used to prove that some proposition is true (false). Negotiation has little to do with what is true or false, and if you think of it that way, you will do badly in it. The purpose is not to present evidence to persuade. It is to bargain, and strike a deal. It is about interests. It is sometimes stressed that a negotiation should be ‘principled’, but what may be indicated here is that an inquiry or persuasion dialogue can often be joined to a negotiation dialogue (in a manner indicated below, under ‘dialectical shifts’). At any rate, it is important to see that the goals of persuasion dialogue are quite different from those of negotiation dialogue.

The goal of information-seeking dialogue is to effect a transfer of information from the one party to the other. This type of dialogue arises from an initial situation in which one party is privy to certain information that the other party lacks. A simple example, would be a person who asks another person on the street information on where a certain building is located. Another different kind of example would be a celebrity interview on television, where the interviewer tries to get the celebrity to reveal personal information of a kind that would be of interest to fans or to viewers generally.

The information-seeking type of dialogue is characteristic of user-machine interaction in searching through a database. A framework for evaluating moves made in this kind of information-seeking dialogue in electronic environments has been presented in [18]. Important problems here are how to formulate questions, and how to use them, and follow up with sequences of questions, in various kinds of information-seeking strategies.

In deliberation, two parties, representing two points of view, or two proposals on how to proceed, confront a practical problem. Typically, the problem can be cast in the form of a dilemma, a kind of choice between two opposed possible courses of action that represent the best ways to proceed, or perhaps the only two ways to proceed, in the face of some difficulty. The thread of reasoning that runs through deliberation dialogue and forms the fabric of argumentation is called practical reasoning (see Section 4). The problem in deliberation is always one of what to do, what course of action to take in a situation that demands that some kind of choice be made. What is sought in deliberation is a prudent or practical line of action

that represents the best way to proceed, choosing from the available alternatives in a given situation at a particular point of time.

Eristic dialogue is an almost purely adversarial type of dialogue where each party tries to win out over the other by producing arguments that devastate the contrary arguments of the other side, or even preferably if these arguments make the other side look foolish, confused, or generally incompetent in argumentation. The best known and most familiar type of eristic dialogue is the quarrel (or personal quarrel), where both parties verbally ‘hit out’ at each other, trying to humiliate that person, and make him look guilty for having committed some culpable actions in the past. However, the quarrel, and eristic dialogue generally, is not completely adversarial or anarchical in nature. Some degree of cooperation is required. For example, it is not really possible to conduct a quarrel unless both participants take turns making accusations to each other, and replying to these accusations.

The purpose of the quarrel as a whole is to reveal hidden grievances that both parties have had, usually for a long time, but have been reluctant to voice. The quarrel is inappropriate in polite conversation, and it generally ‘bursts out’, often accompanied by impolite language and excessive dramatics and emotional outpouring. The quarrel is usually portrayed in a wholly negatively light, as a deterioration of argumentation. But, in some cases, the quarrel can be a beneficial or constructive type of dialogue. In a successful quarrel, the hidden grudges that have been bothering the participants for a long time are brought to the surface and articulated, in a cathartic effect. If the quarrel has been successful, the participants ‘make up’ at the closing stage, and resolve to be more sensitive in the future to the things that have emerged as being bothersome to the other party.

The above classification of types of dialogue is by no means complete. Some other types of dialogue are mixed, or hybrids. For example the forensic debate is a mixture of critical discussion and eristic dialogue, where a judge, referee or audience decides who had the strongest argument. Some other types of dialogue are tied to specific institutional rules of procedure. For example, the kind of legal argumentation used in a criminal trial is regulated by legal rules of procedure and evidence, and although it partly has the structure of a persuasion dialogue, you can’t evaluate their arguments used unless you relate them to the codified body of laws and procedures that makes up the legal system. The six types of dialogue listed in Table 1 have proven to be fundamental for the purpose of evaluating arguments of the kind commonly used in everyday discourse (especially in relation to the analysis of informal fallacies of the kind studied traditionally in logic).

7 Commitment in dialogue

The central idea introduced by Hamblin [9, 10] was that each participant in a dialogue has a repository, a so-called commitment store, into which statements can be inserted as the dialogue proceeds. Statements can also be erased or taken out of a commitment store. The basic idea is that when a participant makes a particular type of move in a dialogue, then the insertion or deletion of commitments is dictated by certain rules, called commitment rules. For example, if a participant asserts a particular proposition at any given point in a dialogue, then that proposition is inserted into her commitment store (sometimes also called a commitment set), and it stays there until such time that it may be retracted. The idea of having a commitment store represents an ideal of rational argumentation that is frequently not met in real argumentative exchanges. Arguers often forget what they said before, or dispute what they supposedly said, or disagree on whether what they supposedly said commits them to such-

and-such a proposition. But ideally, according to Hamblin's model, the commitment store would always be there, and fully visible to both parties at all times. So in the Hamblin model of dialectic, there is a rationality assumption that may not be met in much commonplace argumentation.

A key feature of all the types of dialogue is that the argumentation in them is commitment-based. This feature is particularly evident in the persuasion dialogue. What a participant in a persuasion dialogue must do to have a persuasive argument is to base that argument exclusively on premisses that are commitments of the other party. Or if they are not already commitments of the other party, then they must be propositions that the other party can be persuaded to accept, at some future point in the dialogue. Then she must use these premisses in a chain of argumentation that leads to her own thesis as the ultimate conclusion of the argument. Sometimes other kinds of premisses can be used—for example, statements that are said to be true by an expert, who has been consulted on some factual question that is relevant to the argument. But these statements are only useful in the persuasion dialogue because both parties have agreed that they will accept as commitments whatever can be extracted by reasonable inference from what this expert says. So persuasion dialogue is essentially commitment-based. What matters most is that to have a persuasive (good, successful) argument, you must base it on the commitments of the other party.

The concept of commitment used in formal dialectic seems to be thought of as pertaining to an individual participant in argumentation. However, it has been indicated in studying fallacies and other kinds of criticisms commonly made in arguments that an individual is criticized (for example by using an *ad hominem* argument) because the individual belongs to a certain group, or has commitments that are attributed to her in virtue of her being known to belong to a certain identifiable group. Thus one desirable way of extending the concept of commitment would be to take this group commitment aspect into account. Research of this nature has already been carried out in AI [3], and could be usefully applied to formal dialectic and applied logic generally. Castelfranchi [3, p. 42] has introduced a definition of social commitment that is agent-based, and has also studied certain properties of agents that depend on relationships between individual and social commitment, including the attribute of honesty, which is shown in Section 8 below to have fundamental importance in applied logic in the study of informal fallacies.

The historical problem with *ad hominem* arguments is that following Locke's account [9, p. 160], they have been seen as being identical to argumentation from commitment. Locke, as quoted by Hamblin (p. 160), described *argumentum ad hominem* as 'to press a man with consequences drawn from his own principles or concessions'. This account of the *ad hominem* is too broad, making it identical to the type of argument used when one participant in a dialogue uses the commitments of the other participant as premisses in her (the first party's) argument. The *ad hominem* argument is partly based on argumentation from commitment, but the two types of argumentation are not identical. The *ad hominem* argument, of the kind central to fallacies, and typically featured in the logic textbooks, is essentially personal attack. More exactly, it is the attack on the character of the other party in a dialogue, used to argue that this party is not a credible source, and thereby used to cast doubt on the conclusion put forward by that party. This mixup in terminology is a long historical story, sorted out in [25]. Suffice it to say here that although *ad hominem* has often been taken in logic to be equivalent to argumentation from commitment, this supposed equivalence is a fundamental confusion. The kind of *ad hominem* argument that is so important from the viewpoint of fallacies and legal argumentation is the use of personal attack (along with its variants, the circumstantial

and bias subtypes).

Accordingly, what is needed in formal dialectic to model *ad hominem* arguments must go beyond commitment sets. Commitment is an important concept in formal dialectic, but by itself, it is not sufficient to model the evaluation of source-based arguments like the *ad hominem* and *ad verecundiam*.

8 What needs to be added to formal dialectic

Following the outline of how the structure of formal dialectic was set up by Hamblin [9, 10, Chapter 8], the usual assumption is that in a framework of dialogue, there will be two participants, called the proponent and the respondent. As each of these two participants makes moves in the dialogue, propositions will be inserted into or deleted from their respective commitment stores, according to the commitment rules. But nothing is really said about the internal makeup of these two participants. Hamblin only requires of them, ‘they speak in turn in accordance with a set of rules or conventions’ (p. 255). The various kinds of moves made by the two participants, and the rules they need to follow in making these moves are described in some detail by Hamblin [9, 10]. But very little is said about the participants themselves, by way of describing any properties they might have that might be important for evaluating the kinds of move they make when arguing with each other in a dialogue exchange. The same lack of citing any specific properties of the participants in dialogues was still true of work in this area quite recently—for example, in [26], where sets of rules for different types of dialogue are described in detail, but where comparatively little appears to be said about the properties of the participants in the dialogues.

This lack was understandable, because (1) these works represent the beginning stages of serious work in the field of dialectical logic, and the basic types of dialogue and the rules were important to establish, and these factors seemed like the most basic elements to concentrate on, (2) any attempt to define the qualities or characteristics of the participants would have made the dialogue concept appear anthropomorphic, especially to critics who, from the standpoint of the traditional formal logic, were even dubious about the whole dialogue idea altogether, because it already seemed too personalistic and subjective to bring into the field of logic, and (3) it appeared that there was no real need to bring in constraints on the characteristics of the proponent and respondent, because there was no evidence that the analysis of such characteristics was useful or necessary for the purpose of evaluating argumentation.

The biggest factor in showing the need for a shift in the field of dialectical logic towards consideration of the characteristics of the participants in a dialogue is the evidence from work designed to evaluate source-based arguments that require a credibility function. What then needs to be added to formal dialectic to make it possible to build in a component that would accommodate the credibility function attached to a participant in a dialogue? What needs specifically to be added to the present framework of formal dialectic is the modelling of a participant as not just a set of commitments, but as an agent, and moreover, as an agent that can be said to have a certain degree of credibility. Even more, what needs to be added is the kind of LDS structure in which the credibility assessment can play a role in how the reasoning in a case is evaluated, so that it can function as a rational basis for selecting one conclusion over another. Also, certain characteristics of agents, like their veracity, or the practical consistency of their actions, need to be seen as part of the basis for the assessment of credibility. All this begins to look like a complicated set of structures needed to be brought to bear on the evaluation of cases of source-based arguments. What is centrally needed is

to expand the concept of a participant beyond being a mere repository of commitments in a dialogue, and define a participant as being an agent in the sense outlined by Wooldridge and Jennings. In particular, it is necessary to see an agent as an entity that has the property of veracity, the sixth characteristic of an agent in the list given by Wooldridge and Jennings above.

Some dialectical requirements also need to be added. There needs to be a presumption that a participant in a dialogue will not knowingly communicate false information. This presumption needs to be considered, not only with the other maxims of collaborative dialogue coming under the Gricean (*CP*), but as a characteristic built into the participant herself, by defining the participant as an agent. Thus a participant in a dialogue should not only be seen as having the four characteristics of the weaker usage of the term ‘agent’, but also as an entity that has the four characteristics of the stronger sense of ‘agent’. Other social attributes of commitment could come in here as well, like the property of honesty defined by Castelfranchi [3]. What has been revealed to be particularly important, from the discussion above, is the three attributes of veracity, benevolence, and rationality. But clearly with respect to the *ad hominem* and *ad verecundiam* arguments, the attribute of veracity is especially important. An agent participating in a dialogue should be seen as an entity that not only has commitments, but can generally be presumed to be trustworthy, or to have veracity, when it puts forward an argument or opinion based on those commitments, or when it makes those commitments apparent in a dialogue. Accordingly, the respondent in the dialogue will evaluate the proponent’s arguments and other moves in accordance with his (the respondent’s) estimate of the veracity of the proponent. And this estimate, in fact, should be carried out in accordance with the way the credibility function works. If a proponent shows lack of veracity, the respondent should react by lowering the credibility value he attaches to the proponent’s argument. If the proponent gives evidence to support a presumption of high veracity, the respondent should make an upwards credibility adjustment.

Finally, to get the deepest understanding of how the *ad hominem* and *ad verecundiam* arguments work as fallacies, we need to come to grips with the notion of dialectical relevance. An *ad hominem* argument or appeal to expert opinion argument that is quite reasonable in a persuasion dialogue, might be irrelevant in an inquiry, where character is not part of the issue, and where hard scientific evidence that can be directly verified is all that should count as relevant. The deepest problem with the *ad verecundiam* and *ad hominem* as fallacies is that they are so powerful when used in the right context of dialogue that we may overlook the shift to their use in a different type of dialogue, where they really are not relevant.

9 Dialectical shifts and fallacies

In everyday arguments, there can be a dialectical shift, or movement from one type of dialogue to another, during the same sequence of argumentation. For example, during a negotiation dialogue, where a homeowner is negotiating the cost of basement renovations with a contractor, the two may temporarily switch to an information-seeking type of dialogue, where the homeowner asks the contractor about how concrete work is done, or what the city regulations on the depth of a concrete basement floor are, and so forth. In this kind of case the shift to the information-seeking dialogue could be beneficial to the satisfactory progress of the negotiation dialogue. So, in such a case, we would say that the dialectical shift is licit. Illicit dialectical shifts can occur where the advent of the second dialogue is obstructive to the progress of the first one. For example suppose that two parties are supposed to be having

a critical discussion on the issue of abortion. But then as the discussion gets more heated, the pro-choice advocate one switches to an eristic kind of dialogue where she uses a personal attack against the other party, saying things like, ‘Well you’re a man, a one-sided arguer who is just not objective about the issue, and therefore your argument is not worth much!’ In this case, the dialectical shift would be illicit if the discussion was originally supposed to be a critical discussion, but then the one party unilaterally shifted to a quarrelsome kind of dialogue. The *ad hominem* attack would be said to be irrelevant, for this reason, and could be judged to be a fallacious argument. This phenomenon of the dialectical shift in *ad hominem* argumentation was clearly observed by Aristotle (*Topica* 8.II—see [15, p. 167]), when he remarked that sometimes the questioner in a dialogue is forced to argue against the respondent, instead of arguing against his thesis, if the respondent ‘takes every means of thwarting him with unscrupulous effrontery’. Aristotle comments that such an arguer is a ‘bad associate’ as a partner in the dialogue, and that his ‘perversity’ makes the argumentation ‘eristic’ [15, p. 167].

Many of the traditional informal fallacies are failures of relevance of a kind associated with an illicit dialectical shift during the course of an argument. One of the most interesting fallacies of this type is the *argumentum ad consequentiam*, or argument from consequences. The following textbook example is from [16, p. 82].

The United States had justice on its side in waging the Mexican war of 1848. To question this is unpatriotic, and would give comfort to our enemies by promoting the cause of defeatism.

Presumably, in this case, the context is that of a critical discussion of the issue of which side ‘had justice’ on it in the war of 1848. When one of the participants uses argument from consequences, citing supposed bad consequences of the other side’s maintaining its view, there is a shift to a practical type of discourse (deliberation). Because of the shift the argument from consequences is not really relevant in the critical discussion. Rescher (p. 82) diagnoses the fallacy as a failure of relevance. If the original dialogue had been a deliberation on how to avoid defeat in a current war by avoiding defeatism, the argument from consequences would have been relevant (even though, in fact, it does not seem very plausible, in any case). What makes the argument irrelevant is that the dialogue was evidently supposed to be a critical discussion. So from that perspective, the argument from consequences is beside the point.

Many of the other traditional fallacies can be explained, in many instances, as being arguments in which a dialectical shift occurred. For example, the *ad baculum* argument, or appeal to threat, could be a relevant argument, in some instances, where the context of use is that of a negotiation dialogue. In union–management negotiations for example, threats of strikes and slowdowns, or threats of wage cuts, are often a normal part of the bargaining process. But suppose that during a philosophy seminar where there is a critical discussion of some topic in ethics underway, one party threatens the other party. Such a move would be transparently seen as irrelevant and inappropriate (in that context). So the concept of a dialectical shift is very useful in helping to explain why and how many of the traditional fallacies represent failures of argumentation.

The concept of a dialectical shift also points up some anomalies in the traditional treatment of fallacies however. Appeal to expert opinion has often been treated as a fallacious type of argument in the past (see below, and also [9, pp. 42–44]). But now, with the advent of expert systems, it tends to be judged more as a reasonable type of argument. In fact, a shift from another type of dialogue to an expert consultation dialogue (a subtype of information-seeking

dialogue) can be a licit shift that is highly beneficial to the original dialogue. For example, during the course of an inquiry, experts may be called in to give testimony and advice, and such an interval of expert consultation dialogue could vastly improve the progress of the inquiry, and even be essential to it. The concept of a dialectical shift is clearly necessary to understand how the appeal to expert opinion works as a kind of argument that can be quite reasonable in many instances, but can be abused or used ineptly in other cases.

10 Conclusions

In this investigation, the goal was to find some useful way of assisting with the evaluation of reasoning in source-based arguments. Such arguments are common in legal argumentation, and rulings are made on when they are relevant and when not in the Federal Rules of Evidence; but what kind of formal structure could be brought to bear, of a kind that might be useful, for example, in applying artificial intelligence to legal reasoning? In the investigation, we began with one formal method that is applicable, and then moved along to other formal structures that also appear to be needed, if a deeper modelling of source-based argumentation is to become feasible.

We started out with LDS, because it is basic to the formal modelling of source-based arguments, and because it has a formal structure well investigated by Gabbay [7]. Here, the annotation of the source of declarative units used in an argument is the key. This method is applicable to the kind of argumentation used in a trial setting in law, where a conflict of opinions needs to be resolved, and where most of the argumentation used for this purpose is source-based (witness testimony). But then, the next problem is how to evaluate the characteristics of the sources, leading to a credibility assessment that is the basis of the annotation. Taking this step requires seeing the proponent of an argument as having certain characteristics that relate to her credibility as a source. The standard resources of deductive logic offered no way of solving this problem. So then we turned to practical reasoning. Even there, the resources had to be expanded. Another tool was needed. For this purpose, we used agent-based systems as the tool of choice.

But then, to evaluate aspects like relevance and questioning of a source, we had to work in a framework of multiple agents interacting with each other through argumentation. But source-based arguments are used for different purposes in different types of dialogues. To handle this aspect of the evaluation, we had to move to yet another structure, that of formal dialectic. And even at that, we had to introduce an expansion into formal dialectic by modelling the participants as agents, in a way that goes beyond the usual approach to formal dialectic. The formal structure resulting is now quite complex. But the applicability of this expanded framework to problem of the evaluation of *ad hominem* and *ad verecundiam* types of arguments is highly promising, in relation to modelling the kinds of parameters that are needed to deal in a useful way with the evaluation of cases studies of these types of arguments. The new framework can even go so far as to help explain, through the notion of the dialectical shift, how these powerful kinds of argumentation can be used as fallacies of irrelevance.

Note

I have attempted to meet the requirements for the use on non-sexist language, without having to use confusing and redundant circumlocutions, by following the rule of generally making the proponent in a dialogue 'she' and the respondent 'he'. In some cases, however, it is easier to designate an agent as 'it', thus circumventing the need for equitable gender locutions.

Acknowledgements

I would like to thank the Social Sciences and Humanities Research Council of Canada for a Research Grant (Pragmatics of Argumentation) that funded this work, and the University of Oregon, where I was a Distinguished Visiting Research Fellow in the Oregon Humanities Center in 1997.

References

- [1] R. Audi. *Practical Reasoning*, Routledge, London, 1989.
- [2] M. E. Bratman. *Intentions, Plans, and Practical Reason*, Harvard University Press, Cambridge, MA, 1987.
- [3] C. Castelfranchi. Commitments: from individual intentions to groups and organizations. In *Proceedings of the First International Conference on Multi-Agent Systems*, V. Lesser, ed. AAAI Press, Menlo Park, 1995.
- [4] D. S. Clarke, Jr. *Practical Inferences*, Routledge and Kegan Paul, London, 1985.
- [5] S. Franklin and A. Graesser. Is it an agent, or just a program?: a taxonomy for autonomous agents. In *Intelligent Agents III : Agent Theories, Architectures and Languages*, J. P. Muller, M. J. Wooldridge and N. R. Jennings, eds. pp. 21–35. Springer, Berlin, 1996.
- [6] D. M. Gabbay. Labelled deductive systems and the informal fallacies: a preliminary analysis. In *Analysis and Evaluation : Vol. II of the Proceedings of the Third ISSA Conference on Argumentation*, pp. 308–318. Sicsat, Amsterdam, 1995.
- [7] D. M. Gabbay. *Labelled Deductive Systems, Vol. I*, Clarendon Press, Oxford, 1996.
- [8] H. P. Grice. Logic and conversation. In *The Logic of Grammar*, D. Davidson and G. Harman, eds. pp. 64–75. Dickenson, Encino, CA, 1975.
- [9] C. L. Hamblin. *Fallacies*, Methuen, London, 1970.
- [10] C. L. Hamblin. Mathematical models of dialogue. *Theoria*, **37**, 130–155, 1971.
- [11] C. L. Hamblin. *Imperatives*, Blackwell, New York, 1987.
- [12] P. Huber. *Galileo's Revenge : Junk Science in the Courtroom*, Basic Books, New York, 1991.
- [13] N. R. Jennings and M. Wooldridge. Applying agent technology. *Applied Artificial Intelligence*, **9**, 357–369, 1995.
- [14] S. Kripke. Semantical analysis of intuitionistic logic I. In *Formal Systems and Recursive Functions*, J. N. Crossley and M. Dummett, eds. pp. 92–113. North-Holland, Amsterdam, 1965.
- [15] E. Poste. *Aristotle on Fallacies or the Sophistici Elenchi*, Garland, New York, 1987.
- [16] N. Rescher. *Introduction to Logic*, St. Martin's Press, New York, 1964.
- [17] N. Rescher. *Dialectics*, State University of New York Press, Albany, 1977.
- [18] T. Sandholm and V. Lesser. Issues in automated negotiation and electronic commerce: extending the contract net framework. In *Proceedings of the First International Conference on Multi-Agent Systems*, V. Lesser, ed. pp. 328–335. AAAI Press, Menlo Park, 1995.
- [19] K. W. Saunders. Informal fallacies in legal argumentation. *South Carolina Law Review*, **44**, 343–382, 1993.
- [20] F. H. van Eemeren and R. Grootendorst. *Speech Acts in Communicative Discussions*, Foris, Dordrecht, 1984.
- [21] F. H. van Eemeren and R. Grootendorst. *Argumentation, Communication and Fallacies*, Erlbaum, Hillsdale, NJ, 1992.
- [22] D. Walton. *Practical Reasoning*, Rowman and Littlefield, Savage, Maryland, 1990.
- [23] D. Walton. *A Pragmatic Theory of Fallacy*, University of Alabama Press, Tuscaloosa, 1995.
- [24] D. Walton. *Appeal to Expert Opinion*, Penn State Press, University Park, 1997.
- [25] D. Walton. *Ad hominem Arguments*, University of Alabama Press, Tuscaloosa, 1998.
- [26] D. N. Walton and E. C. W. Krabbe. *Commitment in Dialogue : Basic Concepts of Interpersonal Reasoning*, State University of New York Press, Albany, 1995.
- [27] M. Wooldridge and N. R. Jennings. Intelligent agents: theory and practice, *The Knowledge Engineering Review*, **10**, 115–152, 1995.

Received 14 July 1997