

Ethotic arguments and fallacies: The credibility function in multi-agent dialogue systems

DOUGLAS N. WALTON

In this paper, it is shown how formal dialectic can be extended to model multi-agent argumentation in which each participant is an agent. An agent is viewed as a participant in a dialogue who not only has goals, and the capability for actions, but who also has stable characteristics of types that can be relevant to an assessment of some of her arguments used in that dialogue. 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. In any particular case, an agent will begin a dialogue with a given degree of credibility, and what is here called the credibility function will affect the plausibility of the arguments put forward by that agent. In this paper, an agent is shown to have specific character traits that are vital to properly judging how this credibility function should affect the plausibility of her arguments, including veracity, prudence, sincerity and openness to opposed arguments. When one of these traits is a relevant basis for an adjustment in a credibility function, there is a shift to a subdialogue in which the argumentation in the case is re-evaluated. In such a case, it is shown how the outcome can legitimately be a reduction in the credibility rating of the arguer who was attacked. Then it is shown how the credibility function should be brought into an argument evaluation in the case, yielding the outcome that the argument is assigned a lower plausibility value.

In Aristotelian rhetoric, the audience is said to evaluate a speaker's argument at least partly by the *ethos*, or perceived character of the speaker (Brinton 1986). If the speaker is known to be a good person, her argument will be enhanced in its acceptability. If the speaker is seen as a bad or dishonest person, her argument will be judged to be less plausible. These kinds of assessments of an argument, which draw an inference from the credibility of the speaker to the acceptability of her argument, are quite common in everyday argumentation. But in logic, despite the useful literature on fallacies, an objective general framework for analyzing and evaluating them has been lacking.

In this paper it is shown how a new component has to be added to formal dialectic to model the kind of argumentation where the credibility of a participant in the dialogue as a collaborative arguer is a factor in the evaluation of the plausibility of her argument. In the previous systems of formal dialectic constructed by Hamblin (1970, 1971), Rescher (1977), Barth and Krabbe (1982), Mackenzie (1981, 1990), Walton and Krabbe (1995), and others, no resources are available within the formal system of dialectic that are fully adequate to deal with the evaluation of this kind of argumentation. To fill this gap, the notion of a credibility function is added to the system, of such a kind that an assessment of an arguer's personal credibility as a sincere collaborator in a dialogue can have some impact on how her argument is evaluated as more or less plausible.

One of the most obvious applications of this new structure is to certain forms of argumentation that have long been associated with informal fallacies in logic. In particular, the *ad hominem* and *ad verecundiam* types of arguments represent traditional fallacies where the credibility of the arguer plays an important role in the evaluation of her argument. It will be argued that not all kinds of arguments need to take the credibility function into account, and that it is only in certain kinds of cases (like arguments relating to these two traditional fallacies) where the credibility function needs to be used.

The way to carry out the addition of the credibility function to the current systems of formal dialectic, it will be shown, is to extend the notion of a participant in a dialogue to think of a participant as what is now called an 'agent' in multi-agent systems in computer science. Multi-agent systems are software engineering models in which two or more autonomous software systems called agents can interact verbally with each other, in a dialogue, in order to efficiently carry out a task requiring their intelligent collaboration

(Wooldridge and Jennings 1995). Such dialogues can be negotiations (Sandholm and Lesser 1995), but can take other forms as well.

1. Formal dialogue systems in logic

In formal logic, the idea of studying patterns of argumentation in a dialogue format, where two parties are reasoning with each other, is not widely accepted. Such a formal structure for dialectical reasoning was developed by the Erlangen School in Germany, but it never really caught on, and that branch of logic has never been widely influential, despite its further development as a set of formal dialogue systems in Barth and Krabbe (1982). However, dialogue logic did find an area of application when Hamblin (1970, 1971, 1987) advocated the use of formal dialectical structures as models for evaluating the kind of argumentation used in connection with the traditional subject of fallacies. Hamblin (1970:130) starts with a set P of participants (persons) and a set L of locutions. What he calls a *locution-act* (p, l) is a member of the set of participant-locution pairs. A *dialogue of length n* is then defined (p. 130) as a member of the set $(P \times L)^n$ of sequences of n locution-acts. For Hamblin (1971: 130), a dialogue is a member of the set

$$D = \bigcup_n (P \times L)^n$$

of dialogues of any length. A locution in a dialogue could be a question, for example, or a statement, which would be attributed to a participant, and then a number is attached to the participant-locution pair, representing its place in the longer sequence of participant-locution pairs that make up the dialogue. For example (p. 131), the following is a dialogue of length 3:

$$\langle 0, P_0, L_4 \rangle, \langle 1, P_1, L_3 \rangle, \langle 2, P_0, L_2 \rangle$$

We can see from this example how each move in a dialogue contains a single locution. The first move above, move zero, contains the locution L_4 put forward by the participant P_0 .

Hamblin doesn't tell us much about the central purpose or goal of each of the various systems of dialogue constructed in (Hamblin 1971). All he tells us (p. 133) is that they are "information-oriented", meaning that the purpose is to exchange information among the participants. The systems of

dialogue constructed in Hamblin (1971) seem to have a different purpose. But once again, it is not too clear exactly how the goal is to be defined. Hamblin describes them (1971:148) as systems which permit a participant "to develop an argument by securing assent to individual steps". Such systems, he writes (p. 148), are not information-oriented, but do involve a kind of "rationality" in the sense that once a participant has committed herself to a statement, she should not retract commitment to another statement that is an immediate logical consequence of the first statement. In the terminology of Krabbe and Walton (1995), such dialogue systems would be classified under the heading of persuasion dialogue.

Hamblin's proposal of using formal dialogue systems to analyze fallacies was a novel one at the time. But since then, the idea of using formal dialectical structures to study fallacies has revived an interest in formal dialogue logic among argumentation theorists and practitioners of applied logic. Recently dialogue logic has also come to be accepted as a new and important subject for study in computer science, especially in AI. The new field of computational dialectics investigates how a user and a software computer program would "reason together" in an interactive dialogue format, in order to answer a query, or solve some problem of the user. Dialogue logic has a natural application to fields like expert systems, where the user and the system can naturally be seen as engaging in a goal-directed dialogue where inferences are drawn and logical reasoning is used.

Hamblin (1971), and following him Mackenzie (1981, 1990) and Walton and Krabbe (1995), defines a dialogue not only as a set of moves, made by two participants, typically called the proponent and the respondent, but as a set of moves that are to be made according to certain rules. These rules, according to the account given in Walton and Krabbe (1995) are of four kinds. Locution rules define the types of moves allowed, like the asking of questions or the making of assertions (p. 149). Structural rules define what kind of move a participant may make at any particular point in the sequence of a dialogue, depending in many cases on what the prior move of the other party was (p. 150). Commitment rules determine which propositions are inserted into or deleted from a participant's commitment store at each type of move (p. 149). Win and loss rules define the goal of the dialogue, so that it is made clear what constitutes an achieving of the goal (winning), or a failure to achieve the goal (losing). Calling the latter rules win and loss rules implies a competitive aspect, but different types of dialogue can be more or

less competitive to a varying degree. The kind of dialogue most extensively treated in Walton and Krabbe (1995) is the persuasion dialogue. Each participant has the aim of winning by proving her designated thesis from the commitments of the other participant. In this respect, the dialogue is competitive, and partisan (advocacy) argumentation is used by both sides. However, the goal of the dialogue as a whole is to resolve a conflict of opinions, or at least to throw light on the issue discussed by considering the strongest arguments on both sides.

In this account of the formal structure of a dialogue, the concept of the commitment store of a participant is central. Each participant has a repository, a set of propositions indexed to that participant. As the dialogue proceeds, propositions (statements) are inserted into that set, or can be deleted from that set, depending on what happens at a particular move. For example, if the proponent asserts a particular proposition at a given move, then just after that move, that proposition she asserted will be inserted into her commitment store. The commitments of a participant operate approximately like a *persona* of her beliefs. But it is important to realize, as Hamblin (1970, p. 257) emphasized, that commitments are not exactly the same as beliefs. A commitment is something you have gone on record as accepting, whether you actually believe it or not. As van Eemeren and Grootendorst (1984, 1992) have also made clear, there is an important distinction between belief and acceptance. Belief is psychological, whereas acceptance is more like a voluntary action of communicating to another party that you stand behind a proposition, and are communicating that you are willing to defend it, if challenged to give your supporting reasons for undertaking to accept it. Just as commitment is taken on by voluntary acceptance, it can also be retracted or given up by an act of indicating that you no longer accept something.

2. Fallacies

Now we come to the point where it needs to be asked whether the dialectical structures outlined above are adequate to usefully model certain types of argument that are very important in the study of fallacies. In particular, the *ad hominem*, or personal attack argument, and the *ad verecundiam*, or appeal expert opinion argument (or more broadly, appeal to authority), are types of argument associated with traditional fallacies (Hamblin 1970; Walton 1995).

Can these types of argumentation be modeled adequately by the existing structures of formal dialectic, or do these structures need to be extended in a particular way, in order to do the job? Some recent research (Nuchelmans 1993, Walton 1998) has shown that although, since Locke (Hamblin 1970:159-160), the *argumentum ad hominem* has often traditionally been thought of as species of argument from the commitment of the other party in a dialogue, really it is more than that. It is a species of argumentation in which one party in a dialogue attacks the person of the other party - for example, attacks the character or reputation for veracity of that person - and then uses that attack to argue that the person's argument should not be accepted. In this account, the notion of the party being attacked, as a person, is fundamental. If this latter account is the right concept of the *ad hominem* argument needed for logic, then it would appear that we require the concept of an arguer (participant, person) as a distinct entity in the dialectical structure needed to analyze and evaluate *ad hominem* arguments. Recent research on the *ad verecundiam* type of argument (Walton 1997) also shows that, to analyze the appeal to expert opinion type of argument, we need the idea that an expert's credibility as an arguer is often attacked or questioned (for example, in legal argumentation), and then this attack is used to argue for non-acceptance of the expert's argument. What appears to be needed then, and what is lacking in the current systems of formal dialectic, is the modeling of the participant in a dialogue as a person or agent, in a way that can represent the idea of the person's character, or reputation for integrity for example, as being the target of an attack in an argument by the other party in the dialogue. One can see why formal logic, in the past, has not gone this far, however, for it seems to introduce a human or personal dimension into logic that is alien to the subject as an objective discipline. This element appears to be the psychologistic aspect that was firmly ruled out of logic in the late nineteenth century when it grew to be a purely formal and mathematical discipline. So one might now ask whether the concept of commitment, as introduced by Hamblin into formal dialectic, should be enough to model the kinds of arguments involved in fallacies. A commitment set is, after all, a set of propositions.

The concept of the commitment store, because it does act as a kind of *persona* of an arguer's personal stand or position on an issue in argumentation, does come some distance towards modeling a participant in a dialogue as a person or agent, with qualities of character that could persist over longer

sequences of argumentation in the dialogue. Indeed, throughout the history of informal fallacies, the *ad hominem* type of argument has frequently been identified with argumentation from the commitments of the other party in a dialogue, as noted above. But the question for us is whether this notion of participant in a dialogue goes far enough so that it adequately models (1) the concept of the person as arguer involved in the personal attack argument that is the basic type of *ad hominem* argument, and (2) the concept of credibility involved in the kind of argument where an expert's credibility is questioned, and then that questioning is used to criticize the appeal to expert opinion as an argument. In recent work on these two fallacies (Walton 1997,1998), it is concluded that both these types of arguments require a framework in which an arguer is seen as a participant in a dialogue who has a certain degree of credibility. It is the attack on this arguer's credibility as a participant in a collaborative dialogue that is the key to evaluating the argument.

In (Walton 1998: Chapter 6), five subtypes of the direct (or so-called abusive) type of *ad hominem* argument are recognized, according to the aspect of the arguer's character that is attacked : (1) veracity, (2) prudence, (3) perception, (4) cognition, and (5) morals. Each of these types of direct *ad hominem* is called ethotic, in the sense that the character (*ethos*) of the arguer is involved (Brinton 1986). In the appeal to expert opinion, an argument is judged to have greater plausibility on the grounds that the arguer has positive *ethos*, because she is supposedly an expert. In the *ad hominem* argument, the arguer is attacked on the grounds that she has negative *ethos* (bad character). The five subtypes of negative ethotic *ad hominem* arguments, representing the most common types of direct *ad hominem* arguments found in the case studies surveyed in (Walton 1998) are the following. Here *a* stands for the arguer, or agent, if you like.

Negative Ethotic Argument from Veracity

a has bad character for veracity

Therefore *a*'s argument should not be accepted

Negative Ethotic Argument from Prudence

a has bad character for prudent judgment

Therefore *a*'s argument should not be accepted

Negative Ethotic Argument from Perception

a has bad character for realistic perception of the situation

Therefore *a*'s argument should not be accepted

Negative Ethotic Argument from Cognitive Skills

a has a bad character for logical reasoning

Therefore *a*'s argument should not be accepted

Negative Ethotic Argument from Morals

a has a bad character for personal moral standards

Therefore *a*'s argument should not be accepted

The fifth scheme is more general than the first two, and may be taken to refer to personal moral characteristics other than veracity or prudence. All five of the argumentation schemes cited above presume a context of dialogue in which one party has already put forward some particular argument on an issue two parties are discussing, and the second party is criticizing the first party's argument by using one of the types of arguments cited above. In other words, the *ad hominem* argument is not just any personal attack. It must be the use of personal attack to criticize or refute an argument that has been put forward by the person who is the subject of the attack.

We see in both the *ad hominem* and *ad verecundiam* types of argument a central characteristic of using an evaluation of the *ethos* of a speaker (an *ethos* that can be positive or negative), in order to influence judgment of how the argument should supposedly be evaluated as strong or weak. It is this notion of the credibility of the speaker affecting the plausibility of the speaker's argument that is the framework in which both the *ad hominem* and the *ad verecundiam* fallacies need to be analyzed.

Notice also that the evaluation of these types of argument will depend on the context of dialogue in which the argument was used. In the *Federal Rules of Evidence*, for example, representing a legal framework, ethotic arguments from veracity may be judged relevant in cross-examining a witness, whereas negative ethotic arguments from morals might not be considered relevant. Or to cite another instance, negative ethotic argument from perception might be quite a relevant and strong argument in a deliberation dialogue, whereas the same argument could be quite a bit weaker, or even irrelevant, in a persuasion type of dialogue.

But how could this kind of framework be set up in a formal dialectical system? What is already present in the current systems of formal dialectic is a framework in which a proponent can put forward an argument, made up of a set of statements (propositions) in which one is designated as a conclusion, and the others are meant to be premises that support, or give reasons to the

respondent to accept that conclusion. The argument is evaluated on the basis of (a) how plausible the premises are, and (b) how strong is the inference from the premises to the conclusion. The two measures of evaluation of the acceptability, or strength of the argument are (a) the measure of premise plausibility, or worth as evidence, and (b) the measure of the strength of the inference link between the premise-set and the conclusion. Different standards for evaluating (b) are possible, including deductive validity, inductive strength, and abductive (or presumptive) strength. This sort of structure is already present in the various formal systems of dialectic that are already known. What, then, needs to be added to it, to deal with arguments of the special kind that depend on an arguer's credibility?

What needs to be added is some sort of way of evaluating an arguer's so-called credibility as a proponent of an argument, reflecting the common argumentation practice of some cases where the proponent has a certain *ethos*, that is, a character or reputation, of a kind that can influence how her argument is received in a dialogue (Brinton 1986). For example, if an arguer has a positive *ethos* - e.g., if she is an expert in the domain of knowledge into which the argument falls and has a personal reputation for honesty and integrity - that positive credibility will, quite appropriately, tend to make her argument more plausible - that is, more highly acceptable, to the respondent to whom it was directed. More credibility leads to more plausibility. On the other hand, if an arguer has a bad reputation for veracity, or shows by her performance in a dialogue, that she exaggerates, and uses all kinds of fallacies and tricky deceptions, a respondent to whom her arguments are directed in a dialogue will use that information to downgrade the plausibility of her argument, at least to some degree. Such practices are commonplace. Sometimes they are justifiable, and sometimes not. Sometimes *ad hominem* arguments are reasonable, and sometimes they are not.

To accommodate the possibility of such evaluations, what is required is some way of evaluating a person's character, insofar as such an assessment of character is relevant to an assessment of the person's argumentation in a dialogue. The relevance of such assessments will clearly vary with the type of dialogue that the participants are taking part in. For example, in a scientific discussion in a physics seminar, the character of the scientist putting forward an argument is not relevant to an evaluation of his argument on quantum theory, or some other topic in physics. But when the same scientist testifies as an expert witness in court, his reputation for veracity

would be relevant (at least, from the point of view of the rules of evidence). So while character is not always relevant, when it is relevant, what is presumed is that there can be some rational basis for evaluating an arguer's character, either positively or negatively.

First, it must be assumed that there is some way of evaluating an arguer's character in some relevant respects. Second, it needs to be assumed that this evaluation, can, at least in some cases, transfer over to an evaluation of that arguer's argument. What is needed is a credibility function that takes values of the arguer's *ethos*, or personal character for argument, and then transfers this evaluation to an assessment of her argument that is already in place. What is needed then is a function that takes the given plausibility of an argument as a prior value, and then updates or modifies this value, according to the character assessment.

But how can such a credibility function be introduced into the structure of formal dialectic? What needs to be done is to re-think the notion of a participant in a dialogue, so that the participant has some sort of characteristics as an arguer that are of a lasting sort, that persist over the sequences of argumentation in the dialogue. But what could these characteristics be? And how could we get an account of them that does not make the structure of the dialogue anthropomorphic or psychologistic, in a bad way, that would defeat the purpose of the formal dialectical structure as an objective normative framework that can be employed in logic?

The type of dialogue in which practical reasoning is most typically used in everyday conversations is that of deliberation. Deliberation arises from a need to take prudent action in a given situation. Deliberation can be a solitary procedure in some cases, where a single agent looks at the arguments on both sides of a dilemma, and plays devil's advocate for the side that does not represent her own commitment on the issue. But in many cases deliberation is a collaborative type of dialogue involving two or more agents.

Now, from the literature in philosophy on practical reasoning, it has been made clear what an agent is, at least to some extent, and how agents can not only reason, but reason together in a dialogue format. But this account is too narrow, in certain respects. For one thing, agents can engage in dialogue other than the deliberation type. They can also engage in negotiations, and in the persuasion type of dialogue that is so central to the study of informal fallacies. But even further, agents have qualities of character that are an important part of how we should evaluate their argu-

ments when they reason together in dialogue exchanges. We need to clarify the notion of an agent even further, with respect to how agents interact with each other in dialogues.

It might be noted here that with respect to the negative ethotic argument from prudence, and also the circumstantial *ad hominem* argument, where an agent is accused of acting in a way contrary to her stated goals, practical reasoning is vital to an analysis and evaluation of the argumentation in a given case. The preferred method of analysis in such cases is the action-state semantics of Hamblin (1987: Chapter 4) that has been refined for use in systems of dialogue in (Walton and Krabbe 1995: Appendix A). The key notion is the *partial strategy* of an agent (Hamblin 1987: 155-158), the set of paths of action from which the addressee of an imperative should select one individual action (deed) at a time in order to follow a strategy of action. This notion of a partial strategy is central to Hamblin's theory of imperatives, and it is vital to understanding how attributions of actions, especially those associated with circumstantial *ad hominem* arguments should be evaluated in dialogues. An agent needs to be seen for such purposes, as an autonomous entity which has a set of partial strategies in a case where the agent is seen as carrying out an action, or as having the capability of carrying out an action, or several alternative possible courses of action.

3. Multi-agent systems

It is here that the notion of an agent from the technology of multi-agent systems in computer science comes in. An agent, according to Franklin and Graesser (1996:22) is an entity that can perform autonomous execution of actions, can perceive its environment, including the effects of its own actions, through sensors, and can modify its actions in view of what it sees. Also, an agent is an entity that has goals, and that bases its actions on its goals. So far, this notion of an agent fits perfectly with the notion of an agent in practical reasoning as sketched out above. But how are two agents supposed to interact with each other?

Another characteristic of multi-agent systems is that an agent can "engage in dialogues and negotiate and coordinate transfers of information" (p. 23) with another agent. In engaging in such verbal actions in a dialogue with another agent, the agent must be able to have some capability of

estimating what the goals of the other agent are, and some capability of using that information in guiding its moves in the dialogue. This account gives us some idea of what the capabilities of an agent are, and how the agent is a goal-directed kind of entity. But for purposes of formal dialectic, how exactly should the concept of an agent be defined? What is important to recognize is that there can be different ways of defining the concept of an agent, for different purposes, in different types of dialogue contexts. But within AI, there is some basis for agreement on the range of ways in which the term 'agent' is used.

Wooldridge and Jennings (1995:116-117) distinguish between two usages in AI of the term 'agent', which they call the stronger and the weaker sense. According to the weaker sense, an agent is a computer system that has the following four properties (p. 116).

1. *Autonomy*: the agent has control over its actions and internal states.
2. *Social Ability*: an agent can interact linguistically with other agents.
3. *Reactivity*: an agent perceives its environment and reacts to changes in it.
4. *Pro-activeness*: 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 stronger sense, an agent has the following four additional properties (p. 117).

5. *Mobility*: an agent can move around an electronic network.
6. *Veracity*: an agent will not knowingly communicate false information.
7. *Benevolence*: an agent will do what is asked, and not have conflicting goals.
8. *Rationality*: 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).

Wooldridge and Jennings report that the weaker sense of the term 'agent' is relatively well accepted in computer science, while the stronger sense of the term is more contentious, and not so widely accepted.

To fully accommodate the notion of the credibility function, of the kind that is outlined above, it is the stronger notion of an agent that is needed. To deal with the evaluation of *ad hominem* arguments, for example, characteristics like veracity and rationality are the key. The characteristics that are

especially important are those that pertain to the collaborative qualities of sincerity and trustworthiness required to be presumed in the kind of cooperative conversational exchange described by Grice (1975). What is required is that both agents in a dialogue must have a set of these character traits, and that both agents can either be aware of the other party's traits, or at least come to know something about them, as a dialogue proceeds. Also, certain events, or kinds of moves in the dialogue made by a participant, need to be distinguished as identifiable patterns of argumentation that can be perceived by the other participant as evidence of the first participant's character traits. For example, if a participant becomes committed to both a proposition and its negation, and refuses to budge from this position even after the other party has explained the contradiction and challenged it, then that would be evidence of a kind of insincerity in the character of the party who was challenged. It is a relevant character fault in a critical discussion, for example, because it means that the offender is not collaboratively taking part in the dialogue in a way that helps it along towards its goal of resolving the initial conflict of opinions by rational argumentation.

4. Adding agents to formal dialectical structures

So the questions posed are : (1) How would the concept of an agent, as used in AI, be useful to add to formal dialectical structures of the kind used in applied logic to evaluate argumentation?, and (2) how could the concept of an agent be modeled in formal dialectical structures so that it could fit in with the existing formal systems of dialogue logic already developed by Hamblin (1970, 1971, 1987), Mackenzie (1981, 1990), and Walton and Krabbe (1995)?

One reason why it could be useful to model a participant in a dialogue as an agent (in at least the weaker sense) is that it would enable us to grasp how argumentation in a dialogue is a realistic kind of communication exchange between the participants in which the one party is really reacting to the moves of the other party, and not just following a pre-determined set of rules. This modeling would help us to understand cases where a participant in a dialogue violates the rules, in response to some prior move of the other party, in a way that is simply not comprehensible from the rules alone. Generally, to understand and evaluate the moves made in many realistic

sequences of argumentation, a participant needs to be seen as a goal-directed entity that will react to the moves of another party in a way that is in line with trying to carry out her own goals, whatever they may be. A participant in argumentation will have not only rules or conventions for a conversational kind of exchange that enable us to understand her moves in a dialogue. For us to understand the rationale of such moves well enough to be able to evaluate them as good or deficient arguments, we need to see how she personally reacted to a prior move of the other party. We need to appreciate how she is able to recognize a move made by the other party in the dialogue as significant - to grasp it as a move that affects her own position in the dialogue in a certain way, so that she can then react to it appropriately at the next move. To accomplish such a feat of recognition and reaction in a way we recognize as intelligible, a participant needs to be viewed more than just a commitment store. She needs to be seen as having the capability to grasp something in her dialogue environment as input, to be aware of it as a certain type of move, with potentially significant consequences for the outcome of the dialogue, and then to react to it appropriately at the next move (or at least be capable of doing that). It is exactly our awareness of the significance of sequences of such moves that enables us to recognize them as argumentation tactics - that is, as sequences of dialogue moves that are goal-directed, and are based on an arguer's own internal goals, and on her appreciation of how the sequences of argumentation moves of the other party represent the attempts of the other party to engage in goal-directed activity that is going in the opposite direction.

Judging from the case studies of interactive dialogues characteristic of examples where fallacies and other problems of argumentation evaluation are at issue (Walton 1995), how do the participants contestively interact with each other in sequences of verbal argumentation in dialogue exchanges? What is required is to see the process as one in which the one party is not only aware of the move the other party has made, she must also be able to recognize that move as being a certain type of move, like a question or an argument, and to have some grasp of its significance in the dialogue. To evaluate some cases, account may also have to be taken of the extent to which the participant has recognized, or failed to recognize, that the move made by the other party is inappropriate, at that particular point in the dialogue. If a move, in a given case, is inappropriate, it may be important to try to judge how the recipient of the move reacted to it, and how she saw a

deficiency in it. Such matters of argumentation tactics are vital in evaluating particular cases in applied logic, to judge whether or not a fallacy was committed in a given case.

5. Evaluating fallacies and blunders

To judge whether or not a fallacy was committed in a particular case, what is required is not necessarily to know the actual intention of the supposed fallacy committer. What is required is to have some grasp of how the sequence of argumentation in question was used by the proponent as a tactic to get the best of the other party. Typically there is an element of deception in such cases. We need to see how the arguer was trying to deceive the other party by the use of some sequence of moves that looks like it should be persuasive, or at least that looks like an appropriate sequence of moves for a dialogue. We need to use the given textual evidence to try to judge whether a systematic argumentation tactic was being used, or whether the deficiency in the sequence only represents a mistake or "blunder" (Walton and Krabbe 1995). In evaluating such cases, where mistakes and conversational miscommunications need to be understood and evaluated, the device of viewing both participants in the dialogue as agents is useful.

Grice (1975) showed us how various kinds of infelicities and floutings of maxims of polite communication in conversational exchanges can be insightfully explained as calculated violations of the collaborative rules of dialogue. Making a remark that is obviously irrelevant to the conversation, for example, could be a way of communicating the message that what the other party has just said is inappropriate or impolite. But to understand exactly how such a remark is irrelevant, you have to see it as a kind of tactic used for some conversational purpose, in reaction to the previous move in the conversation by the other party. To grasp the real import of such a sequence of moves, an observer must grasp the significance of the prior move, and must understand what would constitute an appropriate reply in that type of collaborative conversational exchange at that point in the conversation. The observer must also grasp that the each party in the dialogue has normal proactive expectations about how the other would appropriately move in this type of dialogue.

So it is to grasp this kind of evaluation of sequences of argumentation

that the concept of the arguer as agent would be useful. A lot of everyday argumentation is only comprehensible as argumentation of a kind that can be evaluated as logical or illogical because each party in a dialogue has expectations of how the other party should normally react to a certain move, at a particular point in an argument exchange. We expect the other party not to make replies that are clearly irrelevant, for example, and not to say things that both parties, quite clearly, already know. We expect another party not to say things that we know are false, and that she know we know are false. And if she makes such a move, an *ad hominem* or personal attack argument, based on questioning the arguer's character for veracity, can be quite appropriate and reasonable.

Generally, the type of circumstances in which *ad hominem* and *ad verecundiam* arguments need to be evaluated are cases in which there are two sides to an argument used in a dialogue. There may be some evidence for a proposition A, and also some evidence for the proposition $\neg A$, for example, and the problem is to decide which of these two propositions should be accepted. Let's consider a kind of case like that in (Gabbay 1996: 68), where the evidence for A is slightly stronger, but where it is possible to attack the source of A. In what is called by Gabbay "priority logic" such an attack could be a correct move. In a balance of considerations case, where the known objective evidence is insufficient to prove either A or $\neg A$, we may still need to decide, for practical purposes, which proposition to accept. We can't accept them both. But if the arguer who has vouched for A has some known defect of character that is a good reason for doubting his credibility, evidence of that would be reason for reducing the weight of acceptance attached to A. Therefore, on a balance of considerations, the rational move in the dialogue would be to reduce the plausibility value (the weight of acceptance) for A. Let's say that this reduction is sufficient to make it the case that, on balance, $\neg A$ now has the greater weight of acceptance. In this situation, the rational conclusion would be to retract commitment to A, and to commit to $\neg A$. At any rate, in this kind of situation it could be reasonable to accept an *ad hominem* argument, based on the credibility of the person who advocated that argument, and to use that argument to alter one's commitments in a dialogue on a balance of considerations case.

In such a case then, some known facts or evidence concerning the character of a participant in a dialogue could be relevant to the rational

evaluation of an argument in the dialogue, because part of the evidence that is relevant in the dialogue is the credibility of that person. The person is seen as an autonomous agent who has certain characteristics that can be relevant to weighing his credibility.

6. How should 'agent' be defined in formal dialectic?

The question now is how exactly should 'agent' be defined for purposes of formal dialectic. This question is tricky, because it may be that 'agent' needs to be defined differently in the different types of dialogue structures that have been identified. Hence no single, univocal definition give all the information needed to fully understand how the concept of an agent should work in formal dialectic. However, it is fairly clear from the discussion above that the concept of an agent needs to meet the following requirements.

1. It must be an entity that is capable of having a commitment store attached to it, so that statements can be inserted and removed from this store, and a record made of these changes, and of what is in the store, at any given point in a dialogue.
2. It must be an entity that has goals, and that can carry out action, based on these goals.
3. In particular, it must be able to carry out verbal actions in dialogue communications with another agent.
4. It must have the capability to be aware of the external circumstances in its environment, and an awareness of how its actions produce effects on those circumstances.
5. In particular, in a dialogue, it must have the capability to grasp the import and meaning of the communicative actions of another agent in the dialogue, so that it can react in an appropriate way.
6. It must have qualities of character, especially insofar as such qualities are necessary for two agents to engage in collaborative, goal-directed dialogues with each other. For example, it must be capable of having qualities like veracity and sincerity, insofar as such qualities need to be generally presumed in collaborative goal-directed dialogues of various kinds, as indicated by Grice (1975)

With respect to item 5, there needs to be some kind of evidence that is accessible to participants in a dialogue, that allows a participant make some estimate of another participant's character, insofar as this participant's character is relevant to an evaluation of his argumentation put forward in the dialogue. This evidence bears on the credibility of the agent, as a participant who can be relied on to take part in a collaborative dialogue in a constructive way. So, for example, if a participant puts forward statements that are known to be false, and that presumably, he knows are false, or if he persists in using arguments that are irrelevant to the issue being discussed, that should signal something about his character. In turn, this knowledge about the agent's character can yield relevant evidence relating to how his arguments in a dialogue should be evaluated. Some aspects of an arguer's character may be irrelevant to such an evaluation. For example, if we find out an arguer lacks generosity or courage, that finding may not tell us too much about the arguments on the abortion issue he has just put forward in a discussion on the issue. But if we find good evidence in the dialogue that the agent is a habitual liar, or persists with maintaining a position even though it has been pointed out that it contains a significant contradiction, this finding may certainly alter our assessment of some of the arguments he has put forward.

7. How to evaluate cases

The problem is one of evaluating when matters of an arguer's character are relevant when cited during the course of an argument. What is involved is a transition from a given sequence of argumentation in a dialogue on some subject being disputed to a subdiscussion of matters of the character of one of the participants. For example, in a trial, the problem could be one of a witness testifying in a case, where the cross-examining attorney starts to question the credibility of the witness by bringing his honesty into doubt. Or in ordinary conversational argumentation, the problem could be one of deciding how to judge an *ad hominem* argument that one party has used against another in a critical discussion of some disputed issue.

A clue to how to handle such cases has already been given in (Walton and Krabbe 1995), where this kind of case is treated as a dialectical shift from a permissive type of dialogue to a tightened up kind of dialogue where

the rules for the exchange are made more precise. This kind of case is treated in (Walton and Krabbe 1995:125) as involving a kind of shift to a type of dialogue in which one party probes towards a deeper clarification of the other party's position on the issue. For example, suppose two parties are engaging in a persuasion dialogue on the abortion issue, and one of them appears to have contradicted himself, or otherwise given some indication that he may not be sincerely taking part in the persuasion dialogue. To put it in Gricean terms, the suspicion is that he is not following the cooperative principles of the dialogue by collaborating in a constructive way. In a persuasion dialogue, there is a general presumption that participants will be collaborative, and follow the maxims of polite communication. But the conversational rules are not enforced precisely by any absolute rules. You could say that the cooperativeness of the participants is based on trust, or on a presumption of collaboration. So a dialogue of this sort is a permissive type of persuasion dialogue. The participants have some freedom, and the rules are not enforced in a rigid manner. However, if there is evidence that one participant may be violating the rules, then the dialogue may shift to a tightening up phase. In this phase, the other participant may question the sincerity and collaboration of the other party by engaging in a probing explanatory sequence of questioning. It is this phase that corresponds to the *ad hominem* attack. The presumed offender's prior commitments and moves in the dialogue may have to be probed, raising questions about his sincerity and veracity.

The exact technology for modeling this kind of shift has already been presented in (Walton and Krabbe 1995), by giving structures for both the initial permissive type of persuasion dialogue and the rigorous type of dialogue. In a permissive type of persuasion dialogue (**PPD**) (Walton and Krabbe 1995:135), a move is a six-tuple, in which a participant can choose to put forward any or all of the following six types of moves: (1) retractions, (2) concessions, (3) requests for retraction, (4) requests for concessions, (5) arguments, and (6) challenges. In this respect the **PPD** type of dialogue is more complex and more permissive than the Hamblin type of dialogue structure outlined above. But in the rigorous type of persuasion dialogue (**RPD**), the kind of response each player can make is tightly regulated (Walton and Krabbe 1995:154-163).

The precise rules are not so important for our purposes here. And in a legal case, the rules of evidence, which define relevance, and give precise

rules for judging when character is deemed relevant in a trial, are defined by codes like the *Federal Rules of Evidence*. What is important here is to see that there is a kind of shift involved, and that the problem is one of fitting the subconversation after the shift into the flow of the preceding dialogue. What is central is that the persuasion dialogue depends on Gricean presumptions about the collaborative character of the participants. When the collaborative aspect of a participant is called into question, bringing matters of the arguer's ethical character into question, such an *ad hominem* attack can be judged to be relevant in some cases. The reason it is relevant is that the persuasion dialogue is permissive in nature generally, and therefore depends for its success on the presumed honesty and collaborative willingness to participate in a constructive fashion. If this honesty is not manifested in the moves he makes, it is a serious impediment to the success of the dialogue. In principle then, *ad hominem* arguments are relevant, provided the dialogue after the shift fits into the main persuasion dialogue in a constructive way that contributes to the goals of this prior dialogue exchange. *Ad hominem* arguments should be evaluated retrospectively, from the viewpoint of the type of dialogue the participants were originally supposed to be engaged in.

In any kind of argumentation, like the giving of testimony in a case in court, where trust and an assumption of honesty is properly a part of the evaluation of the argument, *ad hominem* arguments will generally be relevant. The participant in the dialogue will start out with a given degree of credibility, even if there is no prior evidence of his honesty or dishonesty. But if there is some evidence of his dishonesty, and his honesty is brought into question, then that questioning will have an effect on the presumption that he is collaboratively taking part in the discussion. Any successful argument against his character as an agent will therefore lower his credibility rating (according to the credibility function). Such an interlude of questioning the arguer's character could be relevant, and could be reasonable generally as an argument, provided several conditions are met. But as stressed above, the evaluation of any particular case will depend, first of all, on the type of dialogue the participants were originally supposed to be engaged in. But it will also depend on the type of *ad hominem* argument used, on the evidence brought forward to support the premises for that type of *ad hominem* argument, and on how strongly that evidence supports the conclusion to be proved. All these matters need to be judged by examining the form of the *ad hominem* argument, in line with the various types of

arguments identified in (Walton 1998). You have to ask whether it is a direct type, a circumstantial type, a bias type, a poisoning the well type, and so forth, and then studying the details of the case to see whether the requirements for that type are met in the case.

8. The solution to the problem

The technical solution to the problem now advocated is to use the basic type of dialogue framework already presented in Hamblin, Mackenzie and Walton and Krabbe, where a participant is simply thought of as an individual, and nothing further is said about the make-up or content of the participant. What needs to be added is an additional, optional type of enriched system in which a participant is defined as an agent. This type of system can be called an agent dialogue. In an agent dialogue, both participants, the proponent and the respondent (or however many participants there may be), need to have certain characteristics relevant for different types of dialogue attached to them. For example, in a deliberation type of dialogue, an agent needs to be open to new, incoming information, and to alter her commitments based on this new information, once it become known to her. This characteristic could be called "reactivity to circumstances", or something of the sort.

The next important feature of an agent dialogue is that each agent has a credibility function. Going into the dialogue, each participating agent will have a particular credibility rating. If nothing relevant is known, the agent would have a rating that is "normal" for that type of dialogue. But then, as new evidence about that agent's relevant character is revealed by her performance in the dialogue, that evidence will bring the credibility function into operation, raising or lowering that agent's credibility by a degree appropriate for the evidence. For example, suppose a speaker strongly advocates the thesis that sexual harassment is a crime that should be strongly penalized. Then suppose it comes out in the dialogue that he admits that he continues to make unwanted sexual advances to the female employees in his own office, but thinks he should not be punished for these actions. In such a case, the situation would trigger the characteristic circumstantial *ad hominem* argument on the basis of the speaker's "hypocrisy". The positive character quality involved would be something like "being consistent in your personal actions with a policy you are advocating for everyone". Once the

character of the agent with respect to the negative feature of hypocrisy is revealed in relation to the agent's partial strategies in the dialogue, the credibility function comes into action. His personal credibility as an agent who is advocating a particular policy is downgraded, and therefore his argument for that policy is rated as correspondingly less plausible in the dialogue.

Or, to take another kind of case, suppose that a participant in a dialogue is rated as an expert in a domain of knowledge into which the claim she is making falls. Suppose she is highly rated as an expert in this domain, and that therefore her initial credibility is high. Suppose, for example, she is giving you financial advice on what investments to make. But then you find that she is making demonstrably false claims about the performance of certain stocks that she gets an unusually high fee from, if you buy them. If the evidence for these statements is quite convincing, it will show something about her character that is relevant for the kind of planning or information-seeking dialogue the two of you are engaged in. The resulting lowering of her personal credibility in the dialogue as an agent will bring the credibility function into play. You will, quite justifiably, begin to distrust what she says, and assign less plausibility to her arguments. The possibility of shifting to a subdialogue on such matters, however, will depend on how she reacts to your probing into her motives.

The technical solution works on the assumption that an agent's argument has a given plausibility value, representing the degree of strength by which the premises support the conclusion. Second, the agent herself has an initial credibility value. Once the negative ethotic argument against the agent comes into the picture, the agent's credibility value is lowered, to a greater or lesser degree. The credibility function is then applied. The outcome is that the plausibility value of the agent's argument is correspondingly reduced. The more the agent's credibility is reduced, the less plausible is her argument. The domain of the credibility function C is S , the set of credibility values that an agent can have in a given case. The codomain is the set T of plausibility values that the agent's argument can have in a case. The credibility function is a rule that assigns to each element in the domain S a single element in the codomain T . So, C is a function on S to T .

An agent needs to be seen as beginning a dialogue with a certain degree of credibility that would be normal. If an agent is an expert in the subject matter of the discussion, then her credibility will be higher than normal, by some appropriate degree. These initial credibility values are not determined

by any moves made in the dialogue itself. They are set as opening, given values. But as the dialogue proceeds, certain types of moves, or sequences of moves, will alter the initial value upwards or downwards, through the operation of the credibility function.

We let the credibility values of an agent in a dialogue be represented by a fraction between 0 and 1. The normal credibility value of an agent will be .5. If an agent is an expert in the subject being discussed, her credibility value could be very high - say .9. But then if she is shown to be biased, or to have bad character for veracity, her credibility value could be lowered to, say .3. We let the plausibility value of some given argument advocated by an agent also be represented as a fraction between 0 and 1. A high value, for example, would be assigned to an argument that not only has plausible premises, but has a structure or form of argument that is recognized as being structurally correct. The plausibility value of an argument is determined by not only these two factors, but also contextually, by its place in a dialogue. Each argument may only have a small degree of plausibility by itself, but when a mass of such arguments are weighed together in a dialogue, they can collectively swing the weight of a proof on an issue in the dialogue one way or the other. In general, worrying too much about assigning exact numbers to credibility and plausibility values is a mistake. Only some rough approximation is needed to judge whether a value is high or low.

The main thing is that a proponent is critically examining the argument of a respondent in a given case, and the two parties then engage in a dialogue sequence. Let's say, for example, the respondent is an expert witness in a trial, and has a high credibility, and that therefore her argument (expressed opinion on an issue on which she was questioned) has a high plausibility. Then, let's suppose, the proponent attacks her character by showing that she has lied in the past. Without analyzing the case in detail, let's say that we now have a subdialogue in which the proponent and respondent play their parts. She puts forward her expert testimony, he attacks her credibility, and she replies to the attacks. Let's say that as the outcome of the subdialogue, her credibility as an expert witness is reduced. By applying the credibility function, the plausibility value of her argument will also be reduced, correspondingly, depending on how successful the *ad x hominem* attack was. How successful the *ad hominem* attack was needs to be judged not only by its form, but also by the fit of the subdialogue into the larger context of dialogue in the given case. The assessment should be

contextual and dialectical, but it is mediated by the credibility function.

In relation to *ad hominem* argumentation, the agent needs to be seen as having a number of traits or qualities of character. The first set of five traits - truthfulness, prudence, perception, cognition, and morals - corresponds to the most common types of direct *ad hominem* arguments. But the circumstantial and the bias types of *ad hominem* arguments also need to be taken into account. Certain qualities of character need to be associated with these types of *ad hominem* arguments, and we need to know what kinds of evidence of performance in a dialogue bring the credibility function into operation to raise or lower an agent's credibility, based on some performance linked to one of these qualities.

The key character quality relating to the circumstantial *ad hominem* argument is that of hypocrisy. The evidence relating to this quality in an agent is an instance of the agent's argumentation in a dialogue where the agent advocates a certain policy, or course of action, that she says should be followed by all agents, though she herself turns out to be committed to the opposite course of action. The typical kind of case is one where the agent is criticized in the dialogue on the basis that she "does not practice what she preaches". For example, she gives a lecture saying how bad it is to take drugs, and concludes that nobody should take drugs, but then she admits that she has not only taken drugs in the past, but is still taking them. Such a speaker is open to the circumstantial *ad hominem* attack, on the grounds that since she says one thing but does the opposite, she is a hypocrite - an agent who is not herself personally committed to the very course of action she advocates for everyone.

It is fairly clear from the case studies of the circumstantial *ad hominem* argument I discussed elsewhere (Walton 1998) what sort of evidence corresponds to the finding that an agent is a hypocrite in this type of case. There is, first of all, the circumstantial inconsistency, which amounts to an inconsistency in the agent's commitments. Evidence of this kind of inconsistency is taken to indicate that the agent is not a sincere participant in the dialogue. Second, there is the evidence from the dialogue on how the agent reacts to the charge of inconsistency when challenged. In some cases, the apparent inconsistency can be admitted, but can also be explained in a way that rebuts the charge of hypocrisy. But in other cases, the arguer tries to squirm out of the inconsistency without really admitting it, even though the evidence is there. Particularly indicative in such a case is the evidence of a

closed attitude or "closed mind", where the agent won't budge in the dialogue, and refuses to try to explain or even admit the inconsistency. The agent may even try to use irrelevant *ad hominem* arguments to try to attack the critic. In such a case, the evidence comes from the dialogue sequence after which the charge is laid, in which the reaction to the charge is part of the dialogue exchange.

Another type of *ad hominem* argument that needs to be taken into account is the bias subtype, where one party accuses the other of having a particular bias. One type of evidence of this claim is the having of something to gain by the party who is accused of being biased. For example, in a debate on the environmental issue of acid rain, one party may accuse the other of being in the pay of a coal company. Such a bias type of *ad hominem* attack would have quite an impact, especially if the attacked party had failed to announce her affiliation at the outset of the debate. Once the attack was made, if it could be supported by convincing evidence, it would mean that the agent's credibility as an open-minded participant in the dialogue would be lowered. The suspicion would be that she is secretly always pushing for the one side, instead of collaboratively looking at the evidence on both sides of the issue. The quality of character involved in this type of *ad hominem* argument is one of openness to the evidence on both sides of a dispute, as opposed to having made up one's mind beforehand on one's conclusion. There is a presumption in a persuasion dialogue that an agent will honestly look at the evidence on both sides revealed in a dialogue, and not just take a one-sided approach.

We can see then that, in addition to the five qualities of character relating to the direct *ad hominem* argument, the qualities of sincerity (and its opposite, hypocrisy) and openness (in the sense of being capable of going with the evidence on either side, and not always sticking to one side) also need to be taken into account. What should happen in a dialogue, a conversational exchange of the sort described by Grice (1975), is that it is presumed at the outset that both participants have all these qualities of character or attitudes. Both participants begin the dialogue with a certain normal degree of credibility, on the grounds that both are presumed to follow Grice's Cooperative Principle, which requires that both will collaboratively contribute to the progress of the dialogue by making the appropriate kinds of moves at its various stages. Participants are expected to use relevant argument arguments, not to obviously contradict what they just said before, and

so forth. Any evidence of the breach of one of these presumptions will affect an arguer's credibility, downwards.

As shown in analyses of *ad hominem* argumentation, however, there are many shortcomings of this type of argument (Walton 1998). Sometimes the personal attack is not relevant. Sometimes there is insufficient evidence to support it, but because of the smear effect, the attack gains a great deal of plausibility by innuendo anyway. Sometimes, as in the case of the poisoning the well type of *ad hominem* argument, the attack is used as a sophistical tactic to close off the dialogue and to prevent the attacked agent from taking any further meaningful part in the dialogue. Each individual case needs to be evaluated on its merits. But now at least we have the framework needed for such evaluations.

In conclusion, the dialogue structures currently in use in the modeling of fallacies can be enriched in a useful way by introducing an additional set of families of dialogue types in which the participants are defined as agents in the sense in which this term is used in multi-agent systems. The deliberation type of dialogue in which the argumentation contained is typically that of practical reasoning yields the basic notion of what an agent is. An agent is a goal-directed entity that carries out actions based on its goals, and on incoming information on its situation, and in particular, feedback directing new actions on the basis of what is perceived about the consequences of the old ones. An agent, in particular, can carry out verbal actions in dialogues with another agent, and the two of them can discuss their goals, and plan on how to carry them out. They can not only engage in such joint deliberations, but can also engage in other types of dialogue, like negotiation, persuasion dialogue, information-seeking dialogue, investigative dialogue (inquiry), and even in some instances of eristic (quarrelsome) dialogue.

The University of Winnipeg

Acknowledgments

For research grants that supported the work in this paper, I would like to thank the Social Sciences and Humanities Research Council of Canada and the Canada-U.S. Fulbright Program.

References

- Barth, E. and Krabbe, E. C. W 1982. *From Axiom to Dialogue*. Berlin: de Gruyter.
- Brinton, A. 1986. "Ethotic argument". *History of Philosophy Quarterly* 3: 245-257.
- van Eemeren, F. H. and Grootendorst, R. 1984. *Speech Acts in Communicative Discussions*. Dordrecht: Foris.
- van Eemeren, F. H. and Grootendorst, R. 1992. *Argumentation, Communication and Fallacies*. Hillsdale, NJ: Erlbaum.
- Federal Rules of Evidence* 1986-1987. Binghamton, New York.
- Franklin, S. and Graesser, A. 1996. "Is it an agent, or just a program?: A taxonomy for autonomous agents". In J. P Muller, M. J. Wooldridge, and N. R. Jennings (eds), *Intelligent Agents III: Agent Theories, Architectures and Languages*. Berlin: Springer, 21-35.
- Gabbay, D. M. 1996. *Labelled Deductive Systems*, vol. 1. Oxford : Clarendon Press.
- Grice, H. P 1975. "Logic and conversation". In D. Davidson and G. Harman (eds), *The Logic of Grammar*. Encino, CA: Dickenson, 64-75.
- Hamblin, C. L. 1970. *Fallacies*. London: Methuen.
- Hamblin, C. L. 1971. "Mathematical models of dialogue". *Theoria* 37: 130-155.
- Hamblin, C. L. 1987. *Imperatives*. New York: Blackwell.
- Jennings, N. R. and Wooldridge, M. 1995. "Applying agent technology". *Applied Artificial Intelligence* 9: 357-369.
- Mackenzie, J. 1981. "The dialectics of logic". *Logique et Analyse* 94: 159-177.
- Mackenzie, J. 1990. "Four dialogue systems". *Studia Logica* 49- 567-583.
- Rescher, N. 1977. *Dialectics*. Albany, NY. State University of New York Press.
- Nuchelmans, G. 1993. "On the fourfold root of the *argumentum ad hominem*". In E. C. W. Krabbe, R. J. Dalitz, and P A. Smit (eds), *Empirical Logic and Public Debate*. Amsterdam: Rodopi, 37-68.
- Sandholm, T. and Lesser, V 1995. "Issues in automated negotiation and electronic commerce". In V Lesser (ed.), *Proceedings of the First International Conference on Multi-Agent Systems*. Menlo Park, CA: AAAI Press, 328-335.
- Walton, D. N. 1990. *Practical Reasoning*. Savage, MD: Rowman and Littlefield.
- Walton, D. N. 1995. *A Pragmatic Theory of Fallacy*. Tuscaloosa, AL: University of Alabama Press.
- Walton, D. N. 1997. *Appeal to Expert Opinion*. University Park, PA: Penn State University Press.
- Walton, D. N. 1998. *AdHominem Arguments*. Tuscaloosa, AL: University of Alabama Press.
- Walton, D. N. and E. C. W Krabbe. 1995. *Commitment in Dialogue: Basic Concepts of Interpersonal Reasoning*. Albany, NY : State University of New York Press.
- Wooldridge, M. and Jennings, N. R. 1995. "Intelligent agents : Theory and practice". *The Knowledge Engineering Review* 10: 115-152.