On the relevance of utterances in formal inter-agent dialogues

Simon Parsons[1] Peter McBurney[2] Elizabeth Sklar[1] Michael Wooldridge[2]

[1]Department of Computer & Information Science
Brooklyn College, City University of New York
Brooklyn NY 11210 USA
{parsons,sklar}@sci.brooklyn.cuny.edu

[2]Department of Computer Science
University of Liverpool
Liverpool L69 7ZF UK
{p.j.mcburney,m.j.wooldridge}@csc.liv.ac.uk

ABSTRACT

Work on argumentation-based dialogue has defined frameworks within which dialogues can be carried out, established protocols that govern dialogues, and studied different properties of dialogues. This work has established the space in which agents are permitted to interact through dialogues. Recently, there has been increasing interest in the mechanisms agents might use to choose how to act — the rhetorical manoeuvring that they use to navigate through the space defined by the rules of the dialogue. Key in such considerations is the idea of relevance, since a usual requirement is that agents stay focussed on the subject of the dialogue and only make relevant remarks. Here we study several notions of relevance, showing how they can be related to both the rules for carrying out dialogues and to rhetorical manoeuvring.

Categories and Subject Descriptors

I.2.11 [**Artificial Intelligence**]: Distributed Artificial Intelligence: *Coherence & co-ordination; languages & structures; multiagent systems.*

General Terms

Design, languages, theory.

Keywords

Argumentation, dialogue, relevance

1. INTRODUCTION

Finding ways for agents to reach agreements in multiagent systems is an area of active research. One mechanism for achieving agreement is through the use of *argumentation* — where one agent tries to convince another agent of something during the course of some *dialogue*. Early examples of argumentation-based approaches to multiagent agreement include the work of Dignum *et al.* [7], Kraus [14], Parsons and Jennings [16], Reed [23], Schroeder *et al.* [25] and Sycara [26].

The work of Walton and Krabbe [27], popularised in the multiagent systems community by Reed [23], has been particularly influential in the field of argumentation-based dialogue. This work influenced the field in a number of ways, perhaps most deeply in framing multi-agent interactions as *dialogue games* in the tradition of Hamblin [13]. Viewing dialogues in this way, as in [2, 21], provides a powerful framework for analysing the formal properties of dialogues, and for identifying suitable protocols under which dialogues can be conducted [18, 20]. The dialogue game view overlaps with work on conversation policies (see, for example, [6, 10]), but differs in considering the entire dialogue rather than dialogue segments.

In this paper, we extend the work of [18] by considering the role of *relevance* — the relationship between utterances in a dialogue. Relevance is a topic of increasing interest in argumentation-based dialogue because it relates to the scope that an agent has for applying strategic manoeuvering to obtain the outcomes that it requires [19, 22, 24]. Our work identifes the limits on such *rhetorical manoeuvering*, showing when it can and cannot have an effect.

2. BACKGROUND

We begin by introducing the formal system of argumentation that underpins our approach, as well as the corresponding terminology and notation, all taken from [2, 8, 17].

A dialogue is a sequence of messages passed between two or more members of a set of agents $\mathbf{A}$. An agent α maintains a knowledge base, Σ_α, containing formulas of a propositional language $\mathcal{L}$ and having no deductive closure. Agent α also maintains the set of its past utterances, called the "commitment store", CS_α. We refer to this as an agent's "public knowledge", since it contains information that is shared with other agents. In contrast, the contents of Σ_α are "private" to α.

Note that in the description that follows, we assume that $\vdash$ is the classical inference relation, that $\equiv$ stands for logical equivalence, and we use Δ to denote all the information available to an agent. Thus in a dialogue between two agents α and β, $\Delta_\alpha = \Sigma_\alpha \cup CS_\alpha \cup CS_\beta$, so the commitment store CS_α can be loosely thought of as a subset of Δ_α consisting of the assertions that have been made public. In some dialogue games, such as those in [18] anything in CS_α is either in Σ_α or can be derived from it. In other dialogue games, such as

AAMAS'07 May 14–18 2007, Honolulu, Hawai'i, USA.

those in [2], CS_α may contain things that cannot be derived from Σ_α.

DEFINITION 2.1. *An* argument A *is a pair* (S, p) *where* p *is a formula of* $\mathcal{L}$ *and* S *a subset of* Δ *such that (i)* S *is consistent; (ii)* $S \vdash p$*; and (iii)* S *is minimal, so no proper subset of* S *satisfying both (1) and (2) exists.*

S is called the *support* of A, written $S = \text{Support}(A)$ and p is the *conclusion* of A, written $p = \text{Conclusion}(A)$. Thus we talk of p being *supported* by the argument (S, p).

In general, since Δ may be inconsistent, arguments in $\mathcal{A}(\Delta)$, the set of all arguments which can be made from Δ, may conflict, and we make this idea precise with the notion of *undercutting*:

DEFINITION 2.2. *Let* A_1 *and* A_2 *be arguments in* $\mathcal{A}(\Delta)$*.* A_1 undercuts A_2 *iff* $\exists \neg p \in \textit{Support}(A_2)$ *such that* $p \equiv \textit{Conclusion}(A_1)$*.*

In other words, an argument is undercut if and only if there is another argument which has as its conclusion the negation of an element of the support for the first argument.

To capture the fact that some beliefs are more strongly held than others, we assume that any set of beliefs has a *preference order* over it. We consider all information available to an agent, Δ, to be stratified into non-overlapping subsets $\Delta_1, \ldots, \Delta_n$ such that beliefs in Δ_i are all equally preferred and are preferred over elements in Δ_j where $i > j$. The *preference level* of a nonempty subset $S \subset \Delta$, where different elements $s \in S$ may belong to different layers Δ_i, is valued at the highest numbered layer which has a member in S and is referred to as $\textit{level}(S)$. In other words, S is only as strong as its weakest member. Note that the strength of a belief as used in this context is a separate concept from the notion of support discussed earlier.

DEFINITION 2.3. *Let* A_1 *and* A_2 *be arguments in* $\mathcal{A}(\Delta)$*.* A_1 *is* preferred *to* A_2 *according to* $\textit{Pref}$*,* $A_1 \gg^{\textit{Pref}} A_2$*, iff* $\textit{level}(\textit{Support}(A_1)) > \textit{level}(\textit{Support}(A_2))$*. If* A_1 *is preferred to* A_2*, we say that* A_1 *is* stronger *than* A_2*.*

We can now define the argumentation system we will use:

DEFINITION 2.4. *An* argumentation system *is a triple:*

$$\langle \mathcal{A}(\Delta), \textit{Undercut}, \textit{Pref} \rangle$$

such that:

- $\mathcal{A}(\Delta)$ *is a set of the arguments built from* Δ*,*
- *Undercut is a binary relation representing the defeat relationship between arguments,* $\textit{Undercut} \subseteq \mathcal{A}(\Delta) \times \mathcal{A}(\Delta)$*, and*
- *Pref is a pre-ordering on* $\mathcal{A}(\Delta) \times \mathcal{A}(\Delta)$*.*

The preference order makes it possible to distinguish different types of relations between arguments:

DEFINITION 2.5. *Let* A_1*,* A_2 *be two arguments of* $\mathcal{A}(\Delta)$*.*

- *If* A_2 *undercuts* A_1 *then* A_1 defends itself *against* A_2 *iff* $A_1 \gg^{\textit{Pref}} A_2$*. Otherwise,* A_1 does not defend itself.
- *A set of arguments* $\mathcal{A}$ defends A_1 *iff for every* A_2 *that undercuts* A_1*, where* A_1 *does not defend itself against* A_2*, then there is some* $A_3 \in \mathcal{A}$ *such that* A_3 *undercuts* A_2 *and* A_2 *does not defend itself against* A_3*.*

We write $\mathcal{A}_{\textit{Undercut},\textit{Pref}}$ to denote the set of all non-undercut arguments and arguments defending themselves against all their undercutting arguments. The set $\underline{\mathcal{A}}(\Delta)$ of acceptable arguments of the argumentation system

$$\langle \mathcal{A}(\Delta), \textit{Undercut}, \textit{Pref} \rangle$$

is [1] the least fixpoint of a function $\mathcal{F}$:

$$\begin{aligned} \mathcal{A} &\subseteq \mathcal{A}(\Delta) \\ \mathcal{F}(\mathcal{A}) &= \{(S, p) \in \mathcal{A}(\Delta) \mid (S, p) \text{ is defended by } \mathcal{A}\} \end{aligned}$$

DEFINITION 2.6. *The set of* acceptable *arguments for an argumentation system* $\langle \mathcal{A}(\Delta), \textit{Undercut}, \textit{Pref} \rangle$ *is recursively defined as:*

$$\begin{aligned} \underline{\mathcal{A}}(\Delta) &= \bigcup \mathcal{F}_{i \geq 0}(\emptyset) \\ &= A_{\textit{Undercut},\textit{Pref}} \cup \left[\bigcup \mathcal{F}_{i \geq 1}(A_{\textit{Undercut},\textit{Pref}}) \right] \end{aligned}$$

An argument is acceptable *if it is a member of the acceptable set, and a proposition is* acceptable *if it is the conclusion of an acceptable argument.*

An acceptable argument is one which is, in some sense, proven since all the arguments which might undermine it are themselves undermined.

DEFINITION 2.7. *If there is an acceptable argument for a proposition* p*, then the* status *of* p *is* accepted*, while if there is not an acceptable argument for* p*, the status of* p *is* not accepted.

Argument A is said to *affect the status* of another argument A' if changing the status of A will change the status of A'.

3. DIALOGUES

Systems like those described in [2, 18], lay down sets of *locutions* that agents can make to put forward propositions and the arguments that support them, and *protocols* that define precisely which locutions can be made at which points in the dialogue. We are not concerned with such a level of detail here. Instead we are interested in the interplay between arguments that agents put forth. As a result, we will consider only that agents are allowed to put forward arguments. We do not discuss the detail of the mechanism that is used to put these arguments forward — we just assume that arguments of the form (S, p) are inserted into an agent's commitment store where they are then visible to other agents.

We then have a typical definition of a dialogue:

DEFINITION 3.1. *A* dialogue D *is a sequence of* moves*:*

$$m_1, m_2, \ldots, m_n.$$

A given move m_i *is a pair* $\langle \alpha, A_i \rangle$ *where* A_i *is an argument that* α *places into its commitment store* CS_α*.*

Moves in an argumentation-based dialogue typically attack moves that have been made previously. While, in general, a dialogue can include moves that undercut several arguments, in the remainder of this paper, we will only consider dialogues that put forward moves that undercut at most one argument. For now we place no additional constraints on the moves that make up a dialogue. Later we will see how different restrictions on moves lead to different kinds of dialogue.

The sequence of arguments put forward in the dialogue is determined by the agents who are taking part in the dialogue, but they are usually not completely free to choose what arguments they make. As indicated earlier, their choice is typically limited by a protocol. If we write the sequence of n moves $m_1, m_2, \ldots, m_n$ as $\vec{m}_n$, and denote the empty sequence as $\vec{m}_0$, then we can define a profocol in the following way:

Definition 3.2. *A* protocol P *is a function on a sequence of moves* $\vec{m}_i$ *in a dialogue* D *that, for all* $i \geq 0$, *identifies a set of possible moves* $\mathbf{M}_{i+1}$ *from which the* m_{i+1}*th move may be drawn:*

$$P : \vec{m}_i \mapsto \mathbf{M}_{i+1}$$

In other words, for our purposes here, at every point in a dialogue, a protocol determines a set of possible moves that agents may make as part of the dialogue. If a dialogue D always picks its moves m from the set $\mathbf{M}$ identified by protocol P, then D is said to *conform* to P.

Even if a dialogue conforms to a protocol, it is typically the case that the agent engaging in the dialogue has to make a choice of move — it has to choose which of the moves in $\mathbf{M}$ to make. This excercise of choice is what we refer to as an agent's use of *rhetoric* (in its oratorical sense of "influencing the thought and conduct of an audience"). Some of our results will give a sense of how much scope an agent has to exercise rhetoric under different protocols.

As arguments are placed into commitment stores, and hence become public, agents can determine the relationships between them. In general, after several moves in a dialogue, some arguments will undercut others. We will denote the set of arguments $\{A_1, A_2, \ldots, A_j\}$ asserted after moves $m_1, m_2, \ldots, m_j$ of a dialogue to be $\mathcal{A}_j$ — the relationship of the arguments in $\mathcal{A}_j$ can be described as an argumentation graph, similar to those described in, for example, [3, 4, 9]:

Definition 3.3. *An* argumentation graph AG *over a set of arguments* $\mathcal{A}$ *is a directed graph* (V, E) *such that every vertex* $v, v' \in V$ *denotes one argument* $A \in \mathcal{A}$, *every argument* A *is denoted by one vertex* v, *and every directed edge* $e \in E$ *from* v *to* v' *denotes that* v *undercuts* v'.

We will use the term *argument graph* as a synonym for "argumentation graph".

Note that we do not require that the argumentation graph is connected. In other words the notion of an argumentation graph allows for the representation of arguments that do not relate, by undercutting or being undercut, to any other arguments (we will come back to this point very shortly).

We adapt some standard graph theoretic notions in order to describe various aspects of the argumentation graph. If there is an edge e from vertex v to vertex v', then v is said to be the *parent* of v' and v' is said to be the child of v. In a reversal of the usual notion, we define a root of an argumentation graph[1] as follows:

Definition 3.4. *A* root *of an argumentation graph* $AG = (V, E)$ *is a node* $v \in V$ *that has no children.*

Thus a root of a graph is a node to which directed edges may be connected, but from which no directed edges connect to other nodes. Thus a root is a node representing an argument that is undercut, but which itself does no undercutting. Similarly:

[1]Note that we talk of "a root" rather than "the root" — as defined, an argumentation graph need not be a tree.

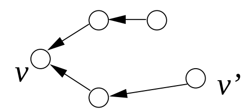

Figure 1: An example argument graph

Definition 3.5. *A* leaf *of an argumentation graph* $AG = (V, E)$ *is a node* $v \in V$ *that has no parents.*

Thus a leaf in an argumentation graph represents an argument that undercuts another argument, but does no undercutting. Thus in Figure 1, v is a root, and v' is a leaf. The reason for the reversal of the usual notions of root and leaf is that, as we shall see, we will consider dialogues to construct argumentation graphs from the roots (in our sense) to the leaves. The reversal of the terminology means that it matches the natural process of tree construction.

Since, as described above, argumentation graphs are allowed to be not connected (in the usual graph theory sense), it is helpful to distinguish nodes that are connected to other nodes, in particular to the root of the tree. We say that node v is *connected* to node v' if and only if there is a path from v to v'. Since edges represent undercut relations, the notion of connectedness between nodes captures the influence that one argument may have on another:

Proposition 3.1. *Given an argumentation graph* AG, *if there is any argument* A, *denoted by node* v *that affects the status of another argument* A', *denoted by* v', *then* v *is connected to* v'. *The converse does not hold.*

Proof. *Given Definitions 2.5 and 2.6, the only ways in which* A *can affect the status of* A' *is if* A *either undercuts* A', *or if* A *undercuts some argument* A'' *that undercuts* A', *or if* A *undercuts some* A''' *that undercuts some* A'' *that undercuts* A', *and so on. In all such cases, a sequence of undercut relations relates the two arguments, and if they are both in an argumentation graph, this means that they are connected.*

Since the notion of path ignores the direction of the directed arcs, nodes v *and* v' *are connected whether the edge between them runs from* v *to* v' *or vice versa. Since* A *only undercuts* A' *if the edge runs from* v *to* v', *we cannot infer that* A *will affect the status of* A' *from information about whether or not they are connected.* □

The reason that we need the concept of the argumentation graph is that the properties of the argumentation graph tell us something about the set of arguments $\mathcal{A}$ the graph represents. When that set of arguments is constructed through a dialogue, there is a relationship between the structure of the argumentation graph and the protocol that governs the dialogue. It is the extent of the relationship between structure and protocol that is the main subject of this paper. To study this relationship, we need to establish a correspondence between a dialogue and an argumentation graph. Given the definitions we have so far, this is simple:

Definition 3.6. *A dialogue* D, *consisting of a sequence of moves* $\vec{m}_n$, *and an argument graph* $AG = (V, E)$ correspond *to one another iff* $\forall m \in \vec{m}_n$, *the argument* A_i *that* *is advanced at move m_i is represented by exactly one node $v \in V$, and $\forall v \in V$, v represents exactly one argument A_i that has been advanced by a move $m \in \vec{m}_n$.*

Thus a dialogue corresponds to an argumentation graph if and only if every argument made in the dialogue corresponds to a node in the graph, and every node in the graph corresponds to an argument made in the dialogue. This one-to-one correspondence allows us to consider each node v in the graph to have an index i which is the index of the move in the dialogue that put forward the argument which that node represents. Thus we can, for example, refer to the "third node" in the argumentation graph, meaning the node that represents the argument put forward in the third move of the dialogue.

4. RELEVANCE

Most work on dialogues is concerned with what we might call *coherent* dialogues, that is dialogues in which the participants are, as in the work of Walton and Krabbe [27], focused on resolving some question through the dialogue[2] To capture this coherence, it seems we need a notion of *relevance* to constrain the statements made by agents. Here we study three notions of relevance:

DEFINITION 4.1. *Consider a dialogue D, consisting of a sequence of moves $\vec{m}_i$, with a corresponding argument graph AG. The move m_{i+1}, $i > 1$, is said to be* relevant *if one or more of the following hold:*

- *R1 Making m_{i+1} will change the status of the argument denoted by the first node of AG.*
- *R2 Making m_{i+1} will add a node v_{i+1} that is connected to the first node of AG.*
- *R3 Making m_{i+1} will add a node v_{i+1} that is connected to the last node to be added to AG.*

R2-relevance is the form of relevance defined by [3] in their study of strategic and tactical reasoning[3]. R1-relevance was suggested by the notion used in [15], and though it differs somewhat from that suggested there, we believe it captures the essence of its predecessor.

Note that we only define relevance for the second move of the dialogue onwards because the first move is taken to identify the *subject* of the dialogue, that is, the central question that the dialogue is intended to answer, and hence it must be relevant to the dialogue, no matter what it is. In assuming this, we focus our attention on the same kind of dialogues as [18].

We can think of relevance as enforcing a form of parsimony on a dialogue — it prevents agents from making statements that do not bear on the current state of the dialogue. This promotes efficiency, in the sense of limiting the number of moves in the dialogue, and, as in [15], prevents agents revealing information that they might better keep hidden. Another form of parsimony is to insist that agents are not allowed to put forward arguments that will be undercut by arguments that have already been made during the dialogue. We therefore distinguish such arguments.

[2]See [11, 12] for examples of dialogues where this is not the case.

[3]We consider such reasoning sub-types of rhetoric.

DEFINITION 4.2. *Consider a dialogue D, consisting of a sequence of moves $\vec{m}_i$, with a corresponding argument graph AG. The move m_{i+1} and the argument it puts forward, A_{i+1}, are both said to be* pre-empted*, if A_{i+1} is undercut by some $A \in \mathcal{A}_i$.*

We use the term "pre-empted" because if such an argument is put forward, it can seem as though another agent anticipated the argument being made, and already made an argument that would render it useless. In the rest of this paper, we will only deal with protocols that permit moves that are relevant, in any of the senses introduced above, and are not allowed to be pre-empted. We call such protocols *basic* protocols, and dialogues carried out under such protocols basic dialogues.

The argument graph of a basic dialogue is somewhat restricted.

PROPOSITION 4.1. *Consider a basic dialogue D. The argumentation graph AG that corresponds to D is a tree with a single root.*

PROOF. *Recall that Definition 3.3 requires only that AG be a directed graph. To show that it is a tree, we have to show that it is acyclic and connected.*

That the graph is connected follows from the construction of the graph under a protocol that enforces relevance. If the notion of relevance is R3, each move adds a node that is connected to the previous node. If the notion of relevance is R2, then every move adds a node that is connected to the root, and thus is connected to some node in the graph. If the notion of relevance is R1, then every move has to change the status of the argument denoted by the root. Proposition 3.1 tells us that to affect the status of an argument A', the node v representing the argument A that is effecting the change has to be connected to v', the node representing A', and so it follows that every new node added as a result of an R1-relevant move will be connected to the argumentation graph. Thus AG is connected.

Since a basic dialogue does not allow moves that are pre-empted, every edge that is added during construction is directed from the node that is added to one already in the graph (thus denoting that the argument A denoted by the added node, v, undercuts the argument A' denoted by the node to which the connection is made, v', rather than the other way around). Since every edge that is added is directed from the new node to the rest of the graph, there can be no cycles. Thus AG is a tree.

To show that AG has a single root, consider its construction from the initial node. After m_1 the graph has one node, v_1 that is both a root and a leaf. After m_2, the graph is two nodes connected by an edge, and v_1 is now a root and not a leaf. v_2 is a leaf and not a root. However the third node is added, the argument earlier in this proof demonstrates that there will be a directed edge from it to some other node, making it a leaf. Thus v_1 will always be the only root. The ruling out of pre-empted moves means that v_1 will never cease to be a root, and so the argumentation graph will always have one root. □

Since every argumentation graph constructed by a basic dialogue is a tree with a single root, this means that the first node of every argumentation graph is the root.

Although these results are straightforward to obtain, they allow us to show how the notions of relevance are related.

PROPOSITION 4.2. *Consider a basic dialogue D, consisting of a sequence of moves $\vec{m}_i$, with a corresponding argument graph AG.*

1. *Every move m_{i+1} that is R1-relevant is R2-relevant. The converse does not hold.*
2. *Every move m_{i+1} that is R3-relevant is R2-relevant. The converse does not hold.*
3. *Not every move m_{i+1} that is R1-relevant is R3-relevant, and not every move m_{i+1} that is R3-relevant is R1-relevant*

PROOF. *For 1, consider how move m_{i+1} can satisfy R1. Proposition 3.1 tells us that if A_{i+1} can change the status of the argument denoted by the root v_1 (which, as observed above, is the first node) of AG, then v_{i+1} must be connected to the root. This is precisely what is required to satisfy R2, and the relatiosnhip is proved to hold.*

To see that the converse does not hold, we have to consider what it takes to change the status of r (since Proposition 3.1 tells us that connectedness is not enough to ensure a change of status — if it did, R1 and R2 relevance would coincide). For m_{i+1} to change the status of the root, it will have to (1) make the argument A represented by r either unacceptable, if it were acceptable before the move, or (2) acceptable if it were unacceptable before the move. Given the definition of acceptability, it can achieve (1) either by directly undercutting the argument represented by r, in which case v_{i+1} will be directly connected to r by some edge, or by undercutting some argument A' that is part of the set of non-undercut arguments defending A. In the latter case, v_{i+1} will be directly connected to the node representing A' and by Proposition 4.1 to r. To achieve (2), v_{i+1} will have to undercut an argument A'' that is either currently undercutting A, or is undercutting an argument that would otherwise defend A. Now, further consider that m_{i+1} puts forward an argument A_{i+1} that undercuts the argument denoted by some node v', but this latter argument defends itself against A_{i+1}. In such a case, the set of acceptable arguments will not change, and so the status of A_r will not change. Thus a move that is R2-relevant need not be R1-relevant.

For 2, consider that m_{i+1} can satisfy R3 simply by adding a node that is connected to v_i, the last node to be added to AG. By Proposition 4.1, it is connected to r and so is R2-relevant.

To see that the converse does not hold, consider that an R2-relevant move can connect to any node in AG.

The first part of 3 follows by a similar argument to that we just used — an R1-relevant move does not have to connect to v_i, just to some v that is part of the graph — and the second part follows since a move that is R3-relevant may introduce an argument A_{i+1} that undercuts the argument A_i put forward by the previous move (and so v_{i+1} is connected to v_i), but finds that A_i defends itself against A_{i+1}, preventing a change of status at the root. □

What is most interesting is not so much the results but why they hold, since this reveals some aspects of the interplay between relevance and the structure of argument graphs. For example, to restate a case from the proof of Proposition 4.2, a move that is R3-relevant by definition has to add a node to the argument graph that is connected to the last node that was added. Since a move that is R2-relevant can add a node that connects anywhere on an argument graph, any move that is R3-relevant will be R2-relevant, but the converse does not hold.

It turns out that we can exploit the interplay between structure and relevance that Propositions 4.1 and 4.2 have started to illuminate to establish relationships between the protocols that govern dialogues and the argument graphs constructed during such dialogues. To do this we need to define protocols in such a way that they refer to the structure of the graph. We have:

DEFINITION 4.3. *A protocol is* single-path *if all dialogues that conform to it construct argument graphs that have only one branch.*

PROPOSITION 4.3. *A basic protocol P is single-path if, for all i, the set of permitted moves $\mathbf{M}_i$ at move i are all R3-relevant. The converse does not hold.*

PROOF. *R3-relevance requires that every node added to the argument graph be connected to the previous node. Starting from the first node this recursively constructs a tree with just one branch, and the relationship holds. The converse does not hold because even if one or more moves in the protocol are R1- or R2-relevant, it may be the case that, because of an agent's rhetorical choice or because of its knowledge, every argument that is chosen to be put forward will undercut the previous argument and so the argument graph is a one-branch tree.* □

Looking for more complex kinds of protocol that construct more complex kinds of argument graph, it is an obvious move to turn to:

DEFINITION 4.4. *A basic protocol is* multi-path *if all dialogues that conform to it can construct argument graphs that are trees.*

But, on reflection, since any graph with only one branch is also a tree:

PROPOSITION 4.4. *Any single-path protocol is an instance of a multi-path protocol.*

and, furthermore:

PROPOSITION 4.5. *Any basic protocol P is multi-path.*

PROOF. *Immediate from Proposition 4.1* □

So the notion of a multi-path protocol does not have much traction. As a result we distinguish multi-path protocols that permit dialogues that can construct trees that have more than one branch as *bushy* protocols. We then have:

PROPOSITION 4.6. *A basic protocol P is bushy if, for some i, the set of permitted moves $\mathbf{M}_i$ at move i are all R1- or R2-relevant.*

PROOF. *From Proposition 4.3 we know that if all moves are R3-relevant then we'll get a tree with one branch, and from Proposition 4.1 we know that all basic protocols will build an argument graph that is a tree, so providing we exclude R3-relevant moves, we will get protocols that can build multi-branch trees.* □

Of course, since, by Proposition 4.2, any move that is R3-relevant is R2-relevant and can quite possibly be R1-relevant (all that Proposition 4.2 tells us is that there is no guarantee that it will be), all that Proposition 4.6 tells us is that dialogues that conform to bushy protocols *may* have more than one branch. All we can do is to identify a bound on the number of branches:

PROPOSITION 4.7. *Consider a basic dialogue D that includes m moves that are not R3-relevant, and has a corresponding argumentation graph AG. The number of branches in AG is less than or equal to $m + 1$.*

PROOF. *Since it must connect a node to the last node added to AG, an R3-relevant move can only extend an existing branch. Since they do not have the same restriction, R1 and R2-relevant moves may create a new branch by connecting to a node that is not the last node added. Every such move could create a new branch, and if they do, we will have m branches. If there were R3-relevant moves before any of these new-branch-creating moves, then these m branches are in addition to the initial branch created by the R3-relevant moves, and we have a maximum of $m + 1$ possible branches.* □

We distinguish bushy protocols from multi-path protocols, and hence R1- and R2-relevance from R3-relevance, because of the kinds of dialogue that R3-relevance enforces. In a dialogue in which all moves must be R3-relevant, the argumentation graph has a single branch — the dialogue consists of a sequence of arguments each of which undercuts the previous one and the last move to be made is the one that settles the dialogue. This, as we will see next, means that such a dialogue only allows a subset of all the moves that would otherwise be possible.

5. COMPLETENESS

The above discussion of the difference between dialogues carried out under single-path and bushy protocols brings us to the consideration of what [18] called "predeterminism", but we now prefer to describe using the term "completeness". The idea of predeterminism, as described in [18], captures the notion that, under some circumstances, the result of a dialogue can be established without actually having the dialogue — the agents have sufficiently little room for rhetorical manoeuver that were one able to see the contents of all the Σ_i of all the $\alpha_i \in \mathbf{A}$, one would be able to identify the outcome of any dialogue on a given subject[4]. We develop this idea by considering how the argument graphs constructed by dialogues under different protocols compare to benchmark *complete* dialogues. We start by developing ideas of what "complete" might mean. One reasonable definition is that:

DEFINITION 5.1. *A basic dialogue D between the set of agents* $\mathbf{A}$ *with a corresponding argumentation graph AG is* topic-complete *if no agent can construct an argument A that undercuts any argument A' represented by a node in AG.*

The argumentation graph constructed by a topic-complete dialogue is called a *topic-complete* argumentation graph and is denoted $AG(D)_T$.

[4] Assuming that the Σ_i do not change during the dialogue, which is the usual assumption in this kind of dialogue.

A dialogue is topic-complete when no agent can add anything that is directly connected to the subject of the dialogue. Some protocols will prevent agents from making moves even though the dialogue is not topic-complete. To distinguish such cases we have:

DEFINITION 5.2. *A basic dialogue D between the set of agents* $\mathbf{A}$ *with a corresponding argumentation graph AG is protocol-complete under a protocol P if no agent can make a move that adds a node to the argumentation graph that is permitted by P.*

The argumentation graph constructed by a protocol-complete dialogue is called a *protocol-complete* argumentation graph and is denoted $AG(D)_P$. Clearly:

PROPOSITION 5.1. *Any dialogue D under a basic protocol P is protocol-complete if it is topic-complete. The converse does not hold in general.*

PROOF. *If D is topic-complete, no agent can make a move that will extend the argumentation graph. This means that no agent can make a move that is permitted by a basic protocol, and so D is also protocol complete.*

The converse does not hold since some basic dialogues (under a protocol that only permits R3-relevant moves, for example) will not permit certain moves (like the addition of a node that connects to the root of the argumentation graph after more than two moves) that would be allowed in a topic-complete dialogue. □

COROLLARY 5.1. *For a basic dialogue D, $AG(D)_P$ is a sub-graph of $AG(D)_T$.*

Obviously, from the definition of a sub-graph, the converse of Corollary 5.1 does not hold in general.

The important distinction between topic- and protocol-completeness is that the former is determined purely by the state of the dialogue — as captured by the argumentation graph — and is thus independent of the protocol, while the latter is determined entirely by the protocol. Any time that a dialogue ends in a state of protocol-completeness rather than topic completeness, it is ending when agents still have things to say but can't because the protocol won't allow them to.

With these definitions of completeness, our task is to relate topic-completeness — the property that ensures that agents can say everything that they have to say in a dialogue that is, in some sense, important — to the notions of relevance we have developed — which determine what agents are allowed to say. When we need very specific conditions to make protocol-complete dialogues topic-complete, it means that agents have lots of room for rhetorical maneouver when those conditions are not in force. That is there are many ways they can bring dialogues to a close before everything that can be said has been said. Where few conditions are required, or conditions are absent, then dialogues between agents with the same knowledge will always play out the same way, and rhetoric has no place. We have:

PROPOSITION 5.2. *A protocol-complete basic dialogue D under a protocol which only allows R3-relevant moves will be topic-complete only when $AG(D)_T$ has a single branch in which the nodes are labelled in increasing order from the root.*

PROOF. *Given what we know about R3-relevance, the condition on $AG(D)_P$ having a single branch is obvious. This is not a sufficient condition on its own because certain protocols may prevent — through additional restrictions, like strict turn-taking in a multi-party dialogue — all the nodes in $AG(D)_T$, which is not subject to such restrictions, being added to the graph. Only when $AG(D)_T$ includes the nodes in the exact order that the corresponding arguments are put forward is it necessary that a topic-complete argumentation graph be constructed.* □

Given Proposition 5.1, these are the conditions under which dialogues conducted under the notion of R3-relevance will always be predetermined, and given how restrictive the conditions are, such dialogues seem to have plenty of room for rhetoric to play a part.

To find similar conditions for dialogues composed of R1- and R2-relevant moves, we first need to distinguish between them. We can do this in terms of the structure of the argumentation graph:

PROPOSITION 5.3. *Consider a basic dialogue D, with argumentation graph AG which has root r denoting an argument A. If argument A', denoted by node v is an an R2-relevant move m, m is not R1-relevant if and only if:*

1. *there are two nodes v' and v'' on the path between v and r, and the argument denoted by v' defends itself against the argument denoted by v''; or*
2. *there is an argument A'', denoted by node v'', that affects the status of A, and the path from v'' to r has one or more nodes in common with the path from v to r.*

PROOF. *For the first condition, consider that since AG is a tree, v is connected to r. Thus there is a series of undercut relations between A and A', and this corrresponds to a path through AG. If this path is the only branch in the tree, then A* will *affect the status of A' unless the chain of "affect" is broken by an undercut that can't change the status of the undercut argument because the latter defends itself.*

For the second condition, as for the first, the only way that A' cannot affect the status of A is if something is blocking its influence. If this is not due to "defending against", it must be because there is some node u on the path that represents an argument whose status is fixed somehow, and that must mean that there is another chain of undercut relations, another branch of the tree, that is incident at u. Since this second branch denotes another chain of arguments, and these affect the status of the argument denoted by u, they must also affect the status of A. Any of these are the A'' in the condition. □

So an R2-relevant move m is not R1-relevant if either its effect is blocked because an argument upstream is not strong enough, or because there is another line of argument that is currently determining the status of the argument at the root. This, in turn, means that if the effect is not due to "defending against", then there is an alternative move that is R1-relevant — a move that undercuts A'' in the second condition above[5]. We can now show

[5]Though whether the agent in question can make such a move is another question.

PROPOSITION 5.4. *A protocol-complete basic dialogue D will always be topic-complete under a protocol which only includes R2-relevant moves and allows every R2-relevant move to be made.*

The restriction on R2-relevant rules is exactly that for topic-completeness, so a dialogue that has *only* R2-relevant moves will continue until every argument that any agent can make has been put forward. Given this, and what we revealed about R1-relevance in Proposition 5.3, we can see that:

PROPOSITION 5.5. *A protocol-complete basic dialogue D under a protocol which only includes R1-relevant moves will be topic-complete if $AG(D)_T$:*

1. *includes no path with adjacent nodes v, denoting A, and v', denoting A', such that A undercuts A' and A' is stronger that A; and*
2. *is such that the nodes in every branch have consecutive indices and no node with degree greater than two is an odd number of arcs from a leaf node.*

PROOF. *The first condition rules out the first condition in Proposition 5.3, and the second deals with the situation that leads to the second condition in Proposition 5.3. The second condition ensures that each branch is constructed in full before any new branch is added, and when a new branch is added, the argument that is undercut as part of the addition will be acceptable, and so the addition will change the status of the argument denoted by that node, and hence the root. With these conditions, every move required to construct $AG(D)_T$ will be permitted and so the dialogue will be topic-complete when every move has been completed.* □

The second part of this result only identifies one possible way to ensure that the second condition in Proposition 5.3 is met, so the converse of this result does not hold.

However, what we have is sufficient to answer the question about "predetermination" that we started with. For dialogues to be predetermined, every move that is R2-relevant must be made. In such cases every dialogue is topic complete. If we do not require that all R2-relevant moves are made, then there is some room for rhetoric — the way in which alternative lines of argument are presented becomes an issue. If moves are forced to be R3-relevant, then there is considerable room for rhetorical play.

6. SUMMARY

This paper has studied the different ideas of relevance in argumentation-based dialogue, identifying the relationship between these ideas, and showing how they can impact the extent to which the way that agents choose moves in a dialogue — what some authors have called the strategy and tactics of a dialogue. This extends existing work on relvance, such as [3, 15] by showing how different notions of relevance can have an effect on the outcome of a dialogue, in particular when they render the outcome predetermined. This connection extends the work of [18] which considered dialogue outcome, but stopped short of identifying the conditions under which it is predetermined.

There are two ways we are currently trying to extend this work, both of which will generalise the results and extend its applicability. First, we want to relax the restrictions that we have imposed, the exclusion of moves that attack several arguments (without which the argument graph can be mulitply-connected) and the exclusion of pre-empted moves, without which the argument graph can have cycles. Second, we want to extend the ideas of relevance to cope with moves that do not only add undercutting arguments, but also supporting arguments, thus taking account of *bipolar* argumentation frameworks [5].

Acknowledgments

The authors are grateful for financial support received from the EC, through project IST-FP6-002307, and from the NSF under grants REC-02-19347 and NSF IIS-0329037. They are also grateful to Peter Stone for a question, now several years old, which this paper has finally answered.

7. REFERENCES

[1] L. Amgoud and C. Cayrol. On the acceptability of arguments in preference-based argumentation framework. In *Proceedings of the 14th Conference on Uncertainty in Artificial Intelligence*, pages 1–7, 1998.

[2] L. Amgoud, S. Parsons, and N. Maudet. Arguments, dialogue, and negotiation. In W. Horn, editor, *Proceedings of the Fourteenth European Conference on Artificial Intelligence*, pages 338–342, Berlin, Germany, 2000. IOS Press.

[3] J. Bentahar, M. Mbarki, and B. Moulin. Strategic and tactic reasoning for communicating agents. In N. Maudet, I. Rahwan, and S. Parsons, editors, *Proceedings of the Third Workshop on Argumentation in Muliagent Systems*, Hakodate, Japan, 2006.

[4] P. Besnard and A. Hunter. A logic-based theory of deductive arguments. *Artificial Intelligence*, 128:203–235, 2001.

[5] C. Cayrol, C. Devred, and M.-C. Lagasquie-Schiex. Handling controversial arguments in bipolar argumentation frameworks. In P. E. Dunne and T. J. M. Bench-Capon, editors, *Computational Models of Argument: Proceedings of COMMA 2006*, pages 261–272. IOS Press, 2006.

[6] B. Chaib-Draa and F. Dignum. Trends in agent communication language. *Computational Intelligence*, 18(2):89–101, 2002.

[7] F. Dignum, B. Dunin-Kęplicz, and R. Verbrugge. Agent theory for team formation by dialogue. In C. Castelfranchi and Y. Lespérance, editors, *Seventh Workshop on Agent Theories, Architectures, and Languages*, pages 141–156, Boston, USA, 2000.

[8] P. M. Dung. On the acceptability of arguments and its fundamental role in nonmonotonic reasoning, logic programming and n-person games. *Artificial Intelligence*, 77:321–357, 1995.

[9] P. M. Dung, R. A. Kowalski, and F. Toni. Dialectic proof procedures for assumption-based, admissable argumentation. *Artificial Intelligence*, 170(2):114–159, 2006.

[10] R. A. Flores and R. C. Kremer. To commit or not to commit. *Computational Intelligence*, 18(2):120–173, 2002.

[11] D. M. Gabbay and J. Woods. More on non-cooperation in Dialogue Logic. *Logic Journal of the IGPL*, 9(2):321–339, 2001.

[12] D. M. Gabbay and J. Woods. Non-cooperation in Dialogue Logic. *Synthese*, 127(1-2):161–186, 2001.

[13] C. L. Hamblin. Mathematical models of dialogue. *Theoria*, 37:130–155, 1971.

[14] S. Kraus, K. Sycara, and A. Evenchik. Reaching agreements through argumentation: a logical model and implementation. *Artificial Intelligence*, 104(1–2):1–69, 1998.

[15] N. Oren, T. J. Norman, and A. Preece. Loose lips sink ships: A heuristic for argumentation. In N. Maudet, I. Rahwan, and S. Parsons, editors, *Proceedings of the Third Workshop on Argumentation in Muliagent Systems*, Hakodate, Japan, 2006.

[16] S. Parsons and N. R. Jennings. Negotiation through argumentation — a preliminary report. In *Proceedings of Second International Conference on Multi-Agent Systems*, pages 267–274, 1996.

[17] S. Parsons, M. Wooldridge, and L. Amgoud. An analysis of formal inter-agent dialogues. In *1st International Conference on Autonomous Agents and Multi-Agent Systems*. ACM Press, 2002.

[18] S. Parsons, M. Wooldridge, and L. Amgoud. On the outcomes of formal inter-agent dialogues. In *2nd International Conference on Autonomous Agents and Multi-Agent Systems*. ACM Press, 2003.

[19] H. Prakken. On dialogue systems with speech acts, arguments, and counterarguments. In *Proceedings of the Seventh European Workshop on Logic in Artificial Intelligence*, Berlin, Germany, 2000. Springer Verlag.

[20] H. Prakken. Relating protocols for dynamic dispute with logics for defeasible argumentation. *Synthese*, 127:187–219, 2001.

[21] H. Prakken and G. Sartor. Modelling reasoning with precedents in a formal dialogue game. *Artificial Intelligence and Law*, 6:231–287, 1998.

[22] I. Rahwan, P. McBurney, and E. Sonenberg. Towards a theory of negotiation strategy. In I. Rahwan, P. Moraitis, and C. Reed, editors, *Proceedings of the 1st International Workshop on Argumentation in Multiagent Systems*, New York, NY, 2004.

[23] C. Reed. Dialogue frames in agent communications. In Y. Demazeau, editor, *Proceedings of the Third International Conference on Multi-Agent Systems*, pages 246–253. IEEE Press, 1998.

[24] M. Rovatsos, I. Rahwan, F. Fisher, and G. Weiss. Adaptive strategies for practical argument-based negotiation. In I. Rahwan, P. Moraitis, and C. Reed, editors, *Proceedings of the 1st International Workshop on Argumentation in Multiagent Systems*, New York, NY, 2004.

[25] M. Schroeder, D. A. Plewe, and A. Raab. Ultima ratio: should Hamlet kill Claudius. In *Proceedings of the 2nd International Conference on Autonomous Agents*, pages 467–468, 1998.

[26] K. Sycara. Argumentation: Planning other agents' plans. In *Proceedings of the Eleventh Joint Conference on Artificial Intelligence*, pages 517–523, 1989.

[27] D. N. Walton and E. C. W. Krabbe. *Commitment in Dialogue: Basic Concepts of Interpersonal Reasoning*. State University of New York Press, Albany, NY, USA, 1995.