

Persuasive Negotiation for Autonomous Agents: A Rhetorical Approach

Sarvapali D. Ramchurn,
Nicholas R. Jennings

IAM group, ECS
University of Southampton, UK.
{sdr01r,nrj}@ecs.soton.ac.uk

Carles Sierra,
IIIA, CSIC

Bellaterra, Barcelona, Spain.
sierra@iiia.csic.es

Abstract

Persuasive negotiation occurs when autonomous agents exchange proposals that are backed up by rhetorical arguments (such as threats, rewards, or appeals). The role of such rhetorical arguments is to persuade the negotiation opponent to accept proposals more readily. To this end, this paper presents a rhetorical model of persuasion that defines the main types of rhetorical particles that are used and that provides a decision making model to enable an agent to determine what type of rhetorical argument to send in a given context and how to evaluate rhetorical arguments that are received. The model is empirically evaluated and we show that it is effective and efficient in reaching agreements.

1 Introduction

Automated negotiation is a fundamental interaction mechanism for managing inter-agent dependencies in multi-agent systems. It normally takes place via an exchange of proposals and counter proposals between a proponent and an opponent until either a mutually acceptable agreement is reached or one of the parties withdraws. In this paper we develop a model of persuasive negotiation whereby proposals¹ are supported by *rhetorical arguments* (such as threats, rewards or appeals). By ‘rhetorical’ we mean that the proponent believes that the enactment (in the case of rewards and threats) or validity (in the case of appeals) of the content of these arguments will, in some way, influence the opponent’s evaluation of the issues at stake so that the deal being proposed is more likely to be accepted (whether this is actually the case depends on many factors that we explore in this paper).

To date, several formalisations and some computational models of persuasion have been proposed (see section 7 for more details). However, most of them deal with logic-based argumentation (e.g. abductive [Sadri *et al.*, 2002] and deductive [Amgoud *et al.*, 2000a; Parsons *et al.*, 1998]). In these works, arguments are considered to be logical sentences and the dialogue is concerned with proving the truth of the sentences. If the argument is not believed, it can simply be rejected (or accepted until proof of the contrary is supplied). Also, they assume that agents completely trust each

¹Proposals, here, are not to be understood as logic propositions but as offers regarding a deal constituting of particular issues and values [Faratin *et al.*, 2002].

other [Parsons *et al.*, 1998], which is not the case in most competitive negotiation settings [Sabater and Sierra, 2002; Jennings *et al.*, 2002]. Moreover, these descriptions tend to focus on the formalization of the dialogues [Amgoud *et al.*, 2000b], and the protocols of interaction [McBurney *et al.*, 2002]. They do not deal with rhetorical arguments such as threats or appeals (e.g. to self-interest or common practice), nor with models for deciding which argument to send given some knowledge about the recipient’s beliefs, nor with how an agent should evaluate the economic utility obtainable by accepting a particular offer implied by an argument. A notable exception to this is the work by Kraus *et al.* [Kraus *et al.*, 1998] who do address such issues in a computational model. However, we believe their model offers an overly restricted view of persuasion that is not readily generalizable or efficient (see sections 6 and 7 for more details).

Against this background, this paper develops and empirically evaluates a general model of persuasive negotiation. In particular, we advance the state of the art in the following ways. Firstly, we provide a general characterisation (in terms of pre- and post- conditions) of the three main types of rhetorical arguments [Ury and Fisher, 1999; Karlins and Abelson, 1970] (i.e. threats, rewards, and appeals). Secondly, we provide a way of modelling the strength of a rhetorical argument in a given environment. Thirdly, we develop a decision making mechanism that enables an agent to select the most appropriate argument to send. Finally, we detail a mechanism an agent can use to evaluate the rhetorical arguments (i.e. the actions they imply) it receives.

The remainder of this paper is structured in the following manner. Section 2 defines what it means for an argument to be a threat, reward, or appeal. Section 3 describes how we model the strength of an argument and section 4 deals with the decision making about which arguments to send. The evaluation functions for the different types of arguments are given in section 5. The model is empirically evaluated in section 6. A comparison to related work is given in section 7 and section 8 concludes.

2 Defining Arguments

The complete formal description of the following language is given in [Ramchurn and Jennings, 2002]. Here, for reasons of space, we only present the key concepts. The world contains agents $\alpha, \beta, \dots \in A_g$ (in this paper we only consider a pair of them). Agents perform actions: (i) *illocutionary acts* ($i_1, i_2, \dots \in I$) are utterances or speech acts [Searle, 1969];

(ii) *environment actions* are those performed on the environment of the agents ($e_1, e_2, \dots \in EA$)². The complete set of actions available is $A = EA \cup I$. All actions $a \in A$ have pre-conditions that must be true before they can be executed ($pre(a)$) and post-conditions that follow from their execution ($post(a)$).

Each agent $\alpha \in A_g$ is characterised by its mental state B^α (noted as $b_1, b_2, \dots \in B^\alpha$) (Here we just focus on the beliefs an agent has). The set of possible mental states of all agents is $B^{A_g} : B^{\alpha_1} \times B^{\alpha_2} \times \dots \times B^{\alpha_{|A_g|}}$. Let W be the set of fully observable environmental states (noted as $\omega, \omega', \omega'', \dots \in W$). Now the set of world states, $S : B^{A_g} \times W$, is composed of pairs of tuples of agents' mental states and an environmental state (noted as $s = \langle b, \omega \rangle = \langle \langle b^{\alpha_1}, \dots, b^{\alpha_{|A_g|}} \rangle, \omega \rangle$). Each agent, $\alpha_n \in A_g$, has an evaluation function that indicates the desirability of a particular state $V^\alpha : S \rightarrow [0, 1]$. Actions cause transitions³ between world states; expressed as $\delta(s, a) = s' (a \in A)$. The expected value of an action(s) to an agent in a given state is expressed as $EV^\alpha : S \times A \rightarrow [0, 1]$.

When agents negotiate they exchange proposals. So, let P be a set of proposals to perform some actions, noted as $p_1, p_2, \dots \in P$, made by the agents. Proposals may consist of a number of environment actions⁴. Each proposal, $p_i \in P$, is a set of atomic actions to be performed by the proponent and the opponent, defined as $p = \langle a^\alpha, a^\beta \rangle$, where $p \in P$, $a^\alpha \subseteq A$, $a^\beta \subseteq A$, and $\alpha, \beta \in A_g$. The proposals are negotiated upon via illocutionary acts. As in [Sierra *et al.*, 1998], we distinguish between standard negotiation illocutions (I_{neg}) for a proposal, $p \in P$, such as *accept*(α, β, p), *reject*(α, β, p) and *propose*(α, β, p) and those that are specific to persuasive negotiation (I_{pers}) such as: *threaten*(α, β, p, th), *reward*(α, β, p, rw), and *appeal*(α, β, p, m), where $\alpha \in A_g$ is the sender and $\beta \in A_g$ is the hearer, $th, rw \in A$ and $m \in I$. A typical example of th is the removal of a privilege, of rw is a promise to give money, and m is *assert*(b) (where $b \in B^\alpha$)⁵.

Agents have various social relationships with one another. The key one in this work is that of trust (here viewed as the expectation that an exchange partner will behave benignly, based on the attribution of positive dispositions and intentions to the partner in a situation of uncertainty and risk [Marsh, 1994]). The trust one agent has in the other has a value between 0 (no trust) and 1 (absolute trust) and in this work we limit it to be a function of two agents as $T : A_g \times A_g \rightarrow [0, 1]$.

With this framework in place we can now define the persuasive illocutions more precisely. These are defined in a rhetorical sense as they imply that the sender anticipates what the hearer believes [Tindale, 1999; Ury and Fisher, 1999] rather than looking at the logical defeasibility or truth of the

²For example, 'update your beliefs with the fact that I am your boss' $\in I$, whereas 'move block one step' $\in EA$.

³Note that we do not focus on the formal semantics of composite actions but foresee using dynamic logic to do so (as in [Harel, 1984]). This is left as future work.

⁴We specialise this for our problem but expect that proposals could be illocutionary acts as well.

⁵*assert* is a function that attempts to update the beliefs of a recipient agent with a specified belief. Note that if the appeal is not supplemented with the belief b , then the appeal becomes a *proposal* itself, consisting of an unsupported proposal alone.

statements [Prakken and Vreeswijk, 2002]. However, the primary precondition for any sort of $\iota \in I_{pers}$ to be sent is that the proponent believes that the achievement of the proposal in ι is preferred to its current state. Thus for all $p \in P$, $\{EV^\alpha(s, p) > V^\alpha(s)\} \subseteq pre(propose(\alpha, \beta, p))$. The post condition is that the recipient β believes⁶ α prefers the proposal to be executed rather than staying in its current state s (i.e. $\{B^\beta(B^\alpha(EV^\alpha(s, p) > V^\alpha(s)))\} \subseteq post(propose(\alpha, \beta, p))$).

2.1 Threats

If $\iota = threaten(\alpha, \beta, p, th)$, then:

$$pre(\iota) = \left\{ \begin{array}{l} B^\alpha(V^\beta(s) > EV^\beta(s, p)), \\ B^\alpha(V^\beta(s) > EV^\beta(s, th)), \\ B^\alpha(V^\beta(\delta(s, p)) > V^\beta(\delta(s, th))) \end{array} \right\}$$

where $\alpha, \beta \in A_g$, $th \in A$, $s \in S$, and $p \in P$.

Thus, for an agent α to threaten another agent β , α must believe that β prefers staying in the current state than enacting the proposal and that β can be threatened (meaning α believes β will prefer to stay in the current state to having the threat effected). Also, α must believe that the state brought about by the threat is less preferred by β than the state brought about by the proposal (in the third condition) otherwise, we would not need threats. The following post-condition applies and it says that β knows something about what α believes of β 's evaluation function:

$$post(\iota) = \{ B^\beta(B^\alpha(V^\beta(s) > V^\beta(\delta(s, th)))) \}$$

2.2 Rewards

If $\iota = reward(\alpha, \beta, p, rw)$, then agent α must believe that enacting the proposal is less preferred by β to staying in the current state and that β can be rewarded with a more preferred alternative (rw) to the proposal. Moreover α should believe that the state brought about by rw is more preferred (by β) than the state brought about after the proposal is executed.

$$pre(\iota) = \left\{ \begin{array}{l} B^\alpha(EV^\beta(\delta(s, p), rw) > V^\beta(s) > EV^\beta(s, p)), \\ B^\alpha(V^\beta(\delta(s, p)) < V^\beta(\delta(\delta(s, p), rw))) \end{array} \right\}$$

where $\alpha, \beta \in A_g$, $rw \in A$, $p \in P$, and $s \in S$.

The post-condition of the reward is that β will come to know something about the mental model α has on it (i.e. it will know that α believes rw is more preferred by β to p and that p is less preferred to the current state (s) of the world by β).

$$post(\iota) = \left\{ \begin{array}{l} B^\beta(B^\alpha(V^\beta(\delta(s, p)) < V^\beta(\delta(\delta(s, p), rw))), \\ B^\beta(B^\alpha(V^\beta(s) > V^\beta(\delta(s, p)))) \end{array} \right\}$$

2.3 Appeals

We take appeals to be about either past promises, common practice, or self-interest (as per [Karlins and Abelson, 1970; Ury and Fisher, 1999]). Without being too specific about the individual type of appeals, we can generally say that the item being appealed to is a belief about the state of the world. Thus, one agent appeals to another agent's beliefs about a past promise, common practice in the domain, or the latter's possible preferences and goals (as far as these can be deduced).

⁶We abuse notation and use the set's name B^α as a predicate to denote the elements it should contain.

Appeals therefore contain the proposal and an illocution. This is different from rewards or threats which can be both illocutionary and environmental actions. Threats and rewards which are illocutionary actions are simply proposed illocutions that have the desired impact (e.g. a threat is “I will tell B that you are inefficient”, and a reward is “I will tell B to employ you”, while an appeal is “Remember i’m your boss”). In all these cases, the illocution is intended to influence the recipient’s current beliefs (hence its evaluation of the issues at stake) and ultimately change the state of the world depending on whether the argument is convincing enough.

If $\iota = \text{appeal}(\alpha, \beta, p, m)$ then the proposing agent α believes that the proposal (p) is less preferred by agent β to the current world state (s) and that the appeal is more preferred to that world state. Also, α believes that β would prefer to take up the appeal rather than only accepting the proposal:

$$\text{pre}(\iota) = \left\{ \begin{array}{l} B^\alpha(EV^\beta(s, m) > V^\beta(s) > EV^\beta(s, p)), \\ B^\alpha(V^\beta(\delta(s, p)) < V^\beta(\delta(s, m), p)) \end{array} \right\}$$

where $\alpha, \beta \in A_g$, $m \in I$, $p \in P$, and $s \in S$.

The post condition simply states that β will believe that α believes β prefers both the appeal and its current state to the proposal. Note that we do not force the recipient to actually believe that the proposal is more preferred to the current state since the agent could stay in its current environmental state and update its mental state without doing the proposal.

$$\text{post}(\iota) = \left\{ \begin{array}{l} B^\beta(B^\alpha(V^\beta(\delta(s, p)) < V^\beta(\delta(s, m)))) \\ B^\beta(B^\alpha(V^\beta(s) > V^\beta(\delta(s, p)))) \end{array} \right\}$$

3 The Rhetorical Strength of an Argument

Before an agent can decide what argument to send in a given context (see section 4), it needs to have a way of differentiating between the various rhetorical arguments at its disposal. The main way of doing this in our model is through the rhetorical strength of an argument; a strong argument is one that quickly persuades an opponent to do the proposal, while a weak argument is one that is likely to be less persuasive⁷. Naturally, these are viewed from the perspective of the proponent since the actual impact on the opponent will not be known.

It is easy to assume that the various rhetorical argument types can be ordered in a rigid domain-independent hierarchy of strength (as do [Kraus *et al.*, 1998]) (e.g. threats are more powerful than rewards). However, practical experience with humans indicates that this is not always so [Karlins and Abelson, 1970]. There are many factors that can help determine an argument’s rhetorical strength (see [Brembeck and Howell, 1976; Tindale, 1999; Toulmin, 1964] for a complete study), however, three of the most obvious ones are:

1. its success rate: the number of times it has caused a proposal to be accepted (i.e. its rhetorical adherence effect).
2. its effect on the trust the opponent has in the proponent (it may increase, decrease or remain the same). Increasing trust increases the opponent’s tendency to accept subsequent proposals from the proponent and so may be

⁷This takes a utility-maximising approach to defining strength which departs from the definitions taken in logic based argumentation in terms of defeasibility [Brewka, 2001; Prakken and Vreeswijk, 2002].

important in the long run (see sections 5 and 6). This is particularly important if future encounters have high payoffs [Axelrod, 1984].

3. its impact on the opponent’s evaluation function (its perceived payoffs). That is, how much it is likely to alter the opponent’s evaluation function (e.g. making a threat that leads the opponent to a very undesirable state means that even reasonably disadvantageous proposals may have to be accepted).

This meta-data about specific rhetorical arguments is taken into account in the proponent’s evaluation function and can be modified over time to reflect the agent’s experience in a given environment. However, the identification of the specific attributes that determine the strength of an argument in a given application will be up to the agent designer (the three above are only possibilities) since there is no universal solution. Thus for maximum generality, we allow the argument to have a dynamically varying strength during the lifetime of an agent depending on what happens in its environment and during the negotiations it undertakes.

To devise a value for the rhetorical strength of an argument, we need to consider the value of the state it brings the agent to. Since the meta-data attributes of the argument, as described above, are dynamic, the evaluation of the possible states that are led to by the argument will also vary accordingly. The specific way this is calculated is not central to our mechanism but in what follows we use the average over all possible states that it can result in (because this gives the overall impression of what the argument is worth given all possible states it can be executed in). Formally, we define the strength of a given argument a for an agent α , $SV^\alpha : A \rightarrow [-1, +1]$, as:

$$SV^\alpha(a) = \frac{\sum_{s \in S} (V^\alpha(\delta(s, a)) - V^\alpha(s))}{|S|} \quad (1)$$

where $V^\alpha : S \rightarrow [0, 1]$ is α ’s evaluation function for a state $s \in S$ of the world. The difference between $V(\delta(s, a))$ and $V(s)$ gives the actual value of enacting an argument a given a current state s .

4 Selecting Arguments

The key factors that determine what arguments to send are the desirability of the proposal to the proponent and the degree of trust (which describes how reliable or truthful the agent is)⁸ that exists between the two agents. These factors are then combined using a series of heuristics based on the believed motivations of the recipient (derived mainly from [Brembeck and Howell, 1976]) to determine what persuasive strength of argument should be chosen in the prevailing context. Since these heuristics involve significant degrees of uncertainty it was decided to exploit fuzzy reasoning techniques⁹. Thus

⁸We do not define a specific model of trust in this paper. However, models such as [Ramchurn *et al.*, 2003; Sabater and Sierra, 2002] can be plugged into our framework. Our only restriction is that the value of trust must be mapped into the range [0,1].

⁹The use of fuzzy techniques is not essential to our conceptual model of persuasive negotiation and any of the other techniques that exist for handling uncertainty could certainly be used instead. We chose fuzzy because it has a well proven record for providing computationally viable and robust solutions in a wide variety of domains [Jang, 1997].

rules of the following form are encoded into our agents:

RULE 1: if **trust** is **low** and **utility** of the proposal is **high**

(I need to do X and I don't trust you therefore I'll impose strong punishments for not doing X)

then send a **strong argument**

RULE 2: if **trust** is **high** and **utility** of the proposal is **low**

(I don't really need to do X and I trust you therefore I'll invite you to help me do X)

then send a **weak argument**

The rationale behind the rules is closely related to our concept of argument strength. The stronger an argument is, the more it is likely to lessen the opponent's trust in the proponent and the more it could coerce the opponent to change its preferences (e.g. by making a significant threat). However, this lowering of trust results in less cooperative behaviour which, in turn, makes it harder for the proponent to persuade the opponent to accept its future proposals. Thus strong arguments should only be sent, for example, when the negotiation needs to take place in the shortest time possible, when the proposal has high utility for the proponent or when it is known that the other partner cannot be trusted to reach effective agreements efficiently. Otherwise weaker arguments should be used.

For each of the fuzzy terms in the argument selection rules, a series of membership functions need to be defined¹⁰. Thus let the fuzzy set 'trust low' be represented by the membership function $T_{low} : [0, 1] \rightarrow [0, 1]$ (the exact details of the membership functions are domain specific but possible examples are $T_{low}(t) = \sin(\frac{t\pi}{2})$ or $T_{low}(t) = 1 - t$ [Jang, 1997]). We define the set of labels to represent trust as \mathcal{T} where $\mathcal{T} = \{T_{low}, T_{high}\}$ (for the sets 'trust high and low') and those to represent utility as \mathcal{U} where $\mathcal{U} = \{U_{low}, U_{high}\}$, where U_{high} and $U_{low} : [0, 1] \rightarrow [0, 1]$ (for the sets 'utility high' and 'utility low'). Similarly, assume that the remaining membership functions are Arg_{weak} and $Arg_{strong} : [-1, 1] \rightarrow [0, 1]$ for the sets (argument 'weak' and 'strong'). The rules can be represented as follows: $R_n = \mathcal{T} \times \mathcal{U} \rightarrow [-1, 1]$, where n corresponds to the n^{th} rule of the inference mechanism. The reasoning shown above can be summarized in the following fuzzy rules:

R_1 : if x is in T_{low} and y is in U_{high} then z is in Arg_{strong}

R_2 : if x is in T_{high} and y is in U_{low} then z is in Arg_{weak}

where x is the trust on the agent we are dealing with, $x = T(\alpha, \beta)$, y is the utility of the proposal, $y = EV^\alpha(s, p)$, and z is the rhetorical strength of the argument a to be sent, $z = SV^\alpha(a)$.

The rules choose the membership values θ_1 and θ_2 for the argument in each of the fuzzy sets 'argument strong' and 'argument weak' (i.e. $\theta_1 = T_{low}(x_0) \wedge U_{high}(y_0)$, and $\theta_2 = T_{high}(x_0) \wedge U_{low}(y_0)$ where x_0 and y_0 are the actual values of trust and utility in the prevailing situation). The fuzzy argument values derived from these sets (z_1 and z_2) can be calculated given that $\theta_1 = Arg_{strong}(z_1)$ and $\theta_2 = Arg_{strong}(z_2)$. From Tsukamoto's inference mechanism [Jang, 1997], we use the discrete Centre-of-Gravity

¹⁰The degree of membership in a fuzzy set varies between 0 (not a member of) and 1 (fully a member of). In between, the values of the membership can be defined as a percentage (e.g. 60% member of the set 'trust high').

method to calculate the overall crisp control action z_0 from the values of the inputs and outputs in the following equation:

$$z_0 = \frac{\theta_1 z_1 + \theta_2 z_2}{\theta_1 + \theta_2} \quad (2)$$

The value z_0 is the strength of the argument a that should ideally be chosen $SV^\alpha(a) = z_0$. However, if no argument has this exact value then the one that is closest is chosen. From this, the argument sent, a_{opt} , is then determined as:

$$a_{opt} = \arg \min_{a \in A} \{SV^\alpha(a) - z_0\} \quad (3)$$

This mechanism helps us balance the notions of trust and utility by calculating a representative average of the two in order to elicit the argument value. This is important because trust and utility are not directly related variables that can easily be aggregated to make a decision. Also, the extent that they each impact on the chosen argument cannot be precisely calculated. However, the fuzzy reasoning approach enables us to measure and manipulate the impact of both variables on the agent's decision-making. Also, the result obtained for the argument value is not overly sensitive to noise. By this we mean that small changes in the trust or utility value do not change the argument chosen by much. This is because of the averaging of results obtained over different fuzzy sets and the inherent properties of the membership functions.

Once a_{opt} has been sent, the recipient may decide to counter-propose and the proponent may decide to keep on arguing. This requires the mechanism to continue the dialogue with new arguments (sending the same argument repeatedly is unlikely to be effective). This means there needs to be a way for an agent to change its decision about which argument to send on reception of repeated counter proposals. The progression could be to move to ever stronger arguments as suggested by [Kraus *et al.*, 1998]. However, we believe this is likely to be inefficient (see section 6). Rather, we believe that the needed argument strength should be re-calculated since the trust may have been modified by the received counter proposal or there may have been changes in the environment (i.e. a revision of beliefs). In any case, a new argument is chosen according to the re-calculated value. If no, as yet un-sent, argument has a value close¹¹ to the needed value, then no argument is sent and the proponent may either choose to accept the counter proposal, withdraw, or make a more attractive proposal if it is still worth negotiating. To make the latter decision, the agent α needs to evaluate the counter proposal i to see whether the proposed world state is preferred to its initial proposal (i.e. $EV^\alpha(s, i) < EV^\alpha(s, p)$). If it is, then the agent chooses another argument, if available, (to back up its initial proposal p) to be sent as per the reasoning mechanism described earlier in this section, otherwise the agent accepts the received counter proposal.

5 Evaluating Proposals

We now focus on the evaluation processes involved in the dialogue. In particular, we propose specific evaluation functions

¹¹The range within which the argument strength should lie for a particular argument to be chosen is left to the agent designer. We believe the range should be determined by the number of arguments available and the need to negotiate further.

that agents can use whenever a proposal, whether or not supported by an argument, is received. These functions incorporate the notion of trust as the confidence in the opponent to fully carry out a proposed action (be it a proposal or an argument) as per section 2. In this way, the evaluation functions incorporate some rhetorical properties (based on trust) as well as some utility maximising properties (based on the utility of the proposals and arguments). Thus, trust guides the agent in believing the sender and affects (but does not define) the probability of the proposal getting accepted. Trust could, for example, be used to guess whether its opponent is actually going to give the reward it has promised.

5.1 Basic Proposals

First we consider proposals from α to β that have no accompanying argument:

$$EV^\beta(s, p) = \frac{V^\beta(\delta(s, p)) \cdot T(\beta, \alpha) + V^\beta(s) \cdot (1 - T(\beta, \alpha))}{V^\beta(s)} \quad (4)$$

where $p \in P$, $s \in S$, and $\alpha, \beta \in A_g$.

This says that agent β will evaluate the received proposal by calculating the expected utility of moving into the proposed state weighted by the trust in the sender α , added to the expected utility of remaining in the present state weighted by the amount of distrust in the other party. The rationale behind this function is that we evaluate proposed actions (mentioned in the proposal) by evaluating the state it brings about. However, there exists some uncertainty in the transition from the initial state to the proposed state and this is captured by the trust value which is equivalent to the probability of moving into the proposed state. The function computes the expected added value of the proposal relative to the current state. Thus, high values of trust will increase the weight of the value of the state to which the proposal aims, and decrease the weight of the value of the current state, hence giving more appeal to the proposal.

5.2 Threats

Let $i \in I_{pers}$ be *threaten*(α, β, p, th), then:

$$EV^\beta(s, i) = \begin{cases} \max(EV^\beta(s, p), EV^\beta(s, th)), & th \in P \\ \max(EV^\beta(s, p), EV^\beta(s, th) \cdot T(\beta, \alpha)), & th \in I \end{cases} \quad (5)$$

where $th \in P$ is a threat composed of a set of environment actions, $th \in I$ is an illocutionary threat, and $s \in S$ is the current state.

We assume that the rejection of the proposal directly entails the threat being executed (i.e. th subsumes the prior *reject* illocution). In this case, the agent compares the value of the threat being incurred against the value of the proposal being enacted. The evaluation function then chooses the maximum of the two. There are two types of threats to consider. One is where the threat is a set of environment actions (i.e. equivalent to a proposal) and the other is where it is an illocutionary act. In the former case, $EV^\beta(s, th)$ is evaluated as in equation 4 (because the threat explicitly describes the actions to be executed as in a proposal). In the latter case, th implicitly describes actions that will be performed (e.g. $th = propose(\alpha, \gamma, p')$ is a threat to propose another set of actions to a competitor agent γ). In this case, the proponent must be trusted to effect the illocution and then to effect the

content of the illocution (i.e. two levels of trust exist compared to the one level for the $th \in P$ case). Therefore the value obtained as a result of the initial evaluation that takes into account the trust in effecting the illocution, $EV^\beta(s, th)$, is also multiplied by the trust in agent α in doing the content of the illocution in th .

Having evaluated the threat, the agent must decide what to do: it will accept the proposal if the evaluation function results in the proposal having a higher value and it will reject the proposal otherwise.

5.3 Rewards

Rewards and threats can be considered to have symmetric semantics. Rewards are actions wanted and obtained after a proposal is enacted, while threats are not wanted and obtained when a proposal is rejected. This is reflected in our evaluation functions in that threats are evaluated relative to the rejection state while rewards are evaluated relative to the acceptance (and execution) state. Thus, let $i \in I_{pers}$ be *reward*(α, β, p, rw), then:

$$EV^\beta(s, i) = \begin{cases} \max(EV^\beta(\delta(s, p), rw) \cdot T(\beta, \alpha) + V^\beta(s) \cdot (1 - T(\beta, \alpha)), V^\beta(s)), & rw \in P \\ \max(EV^\beta(\delta(s, p), rw) \cdot T(\beta, \alpha) \cdot T(\beta, \alpha) + V^\beta(s) \cdot (1 - T(\beta, \alpha)), V^\beta(s)), & rw \in I \end{cases} \quad (6)$$

where $rw \in P$ is the reward as a set of environment actions, $rw \in I$ is an illocutionary reward, and $s \in S$ is the current state.

We apply the simplifying assumption that the rejection of the reward does not change the preferences of the agent. This enables us to calculate the value of the reward relative to the state in which the illocution is initially received. Thus, the value of the reward is compared to the state where the rejection is sent. The evaluation considers two cases. Firstly, if $rw \in P$, the reward is a set of environment actions (a proposal), it is evaluated as such (as per equation 4). In effect, this says that the reward will only be obtained if both agents do their part of the proposal. Secondly, if $rw \in I$, then β 's trust in α additionally weighs the value of receiving the reward and the value of rejecting the reward (i.e. staying in the initial state) by the trust in agent α . This adds the effect of the probability that agent α is making a truthful illocution (as described in section 5.2).

If the evaluation of the reward illocution selects the weighted proposal (including the reward rw), then the proposal is accepted and executed (else a rejection is sent).

5.4 Appeals

Appeals always contain an illocution (see section 2). The illocution may or may not be accepted depending on whether the agent attributes a high utility to that action. Let $i = appeal(\alpha, \beta, p, m)$, then:

$$EV^\beta(s, i) = \max(EV^\beta(s, m), EV^\beta(s, p), EV^\beta(\delta(s, m), p), V^\beta(s)) \quad (7)$$

where $m \in I$ is the appeal, and $s \in S$ is the current state.

The same simplifying assumptions as for threats and rewards apply for appeals. This means the evaluation of an

appeal selects the action(s) that will bring the agent to the best state (i.e. accepting the appeal and accepting the proposal, rejecting the appeal and accepting the proposal, accepting the appeal and rejecting the proposal, and rejecting both the proposal and the appeal so as to stay in the current state). We do not specify the evaluation function for an illocution $EV^\beta(s, m)$ in detail because this is likely to be highly domain dependent. However, in case m is another persuasive illocution, the evaluation functions for threats, rewards and appeals can be used. For other illocutions, attempting belief updates or removals, the evaluation function will reduce to $EV^\beta(s, m) = V(\delta(s, m))$ meaning that it is equivalent to the known value of the state in which the action has been executed.

Having evaluated the appeal, the agent has to decide what to do. In case the evaluation function selects the illocution only, then the proposal is rejected and the illocution is executed. If the proposal is chosen, then the illocution is rejected and the proposal is executed, while if both are chosen by the evaluation function, the agent will assimilate the content of the illocution and then execute the proposal. Otherwise, it stays in its current state.

6 Experimental Evaluation

The experiments aim to evaluate our argument selection mechanism. In particular, we postulate hypotheses about the efficiency (in terms of negotiation cycles and number of arguments exchanged) and effectiveness (utility gained by the agents) of our mechanism and test these against a variety of alternative mechanisms.

6.1 Setup

The environment is an abstract one involving just two agents. The agents are self-interested and negotiate via an exchange of proposals (backed up by persuasive arguments) that aim to change the state of the world to one that they prefer more than their current one. The states of the environment ($s, s', \dots \in S$) are denoted by a number ranging from $s = 1$ to $s = 1000$. Actions are denoted as a tuple composed of a start state and an end state. State changes only occur as a consequence of the proposals between the agents being enacted. The pre-conditions of a proposal p (as specified in section 2) must be satisfied before it is sent to an opponent. In each state, a restricted set of arguments can be uttered (between 0 and 10). These arguments (threats, rewards or appeals) define a transition to another state of the world similar to proposals and are hardwired into the agents. The argument’s illocutionary type determines the evaluation function used to assess its impact (described in section 5). This type also determines the effect the argument has on the opponent’s trust in the proponent and vice versa (as described in section 3). Upon receipt of proposals or illocutions an agent goes through the evaluation procedure presented in equations 4, 5, 6 and 7.

A single simulation run involves 1000 separate negotiation encounters between the two agents. They were paired against agents that used the same argument selection mechanism¹²; *Fuzzy* (using our mechanism), *Ramping* (using appeals, rewards and threats in sequence as in [Kraus *et al.*,

¹²This was mainly to test how effective the mechanism is for both proponent and opponent in the same environment. However, separate tests were also carried out with pairs of different agents as shown later.

Component	Range
$ S $	1000
$ A_g $	2 (α and β)
$ A $	1000^2
$V^\alpha(s)$	0 to 1000
$EV^\alpha(s, a)$	0 to 1000
$T(\alpha, \beta)$	0 to 1
No. of Arguments available per state	0 to 10
Effect of threat on trust	-0.0005
Effect of reward on trust	+0.0003
Effect of appeals on trust	+0.0002
Negotiation Cycles per proposal	1 to 10
Utility gained over 1000 proposals	0 to 10^6
No. of Agreements reached	0 to 1000

Table 1: Range of instantiated components.

1998]), *Random* (choosing any available argument at random) and *Non-arguing* (not using arguments). Non-arguing agents could only exchange proposals without arguments and they were not allowed to change their proposals during the course of an encounter. This was done to compare with the other agents which use different arguments to back up a similarly unchanged proposal. For each setting of the control variables (below), the simulation was repeated 20 times (making a total of 20000 encounters¹³). The control variables are: (i) the initial value of trust of the proponent in the opponent (uniformly distributed in the range 0 to 1), and (ii) the knowledge the participants have of each other’s argument evaluation function, that is:

1. The proponent has perfect knowledge of the opponent’s evaluation function. This means the agents know exactly what the values of the states brought about by the arguments are (e.g. a proposed high reward is indeed very good for the recipient).
2. The proponent has imperfect knowledge of the impact of its arguments. Thus an agent may send what it believes is a threat, and this may in fact turn out to be a reward. The degree of misalignment is normally distributed over the range $0 \leq EV^\alpha(s, i) \leq 1000$ (where s is the initial state and i is the argument).

One preliminary result of our simulations was that Ramping agents, when set against each other with imperfect knowledge, would often negotiate indefinitely because the argument generation mechanism tends to keep the trust value constant and decreases their tendency to reach agreements when trust is particularly low. Our tests were therefore made with the Ramping agent set against our Fuzzy agent. The experimental variables that were measured are: (i) the number of cycles taken on average to reach agreements, (ii) the number of agreements made, (iii) the number of arguments (threats, rewards, appeals) used, (iv) the final value of the trust of the opponent in the proponent and vice versa, and (v) the utility gained by interaction partners both individually ($V^\alpha(s') - V^\alpha(s)$) and overall ($(V^\alpha(s') - V^\alpha(s) + V^\beta(s') - V^\beta(s))/2$),

¹³We performed a regression analysis on samples of 50000 and 100000 encounters which resulted in P-values of 3.1×10^{-6} and 5.1×10^{-6} with standard error of 0.09 in the utility values obtained in both cases.

	Ramping	Random	Fuzzy
Agreements	735	745	710
Cycles/Agreement	2.2	2.9	3.4
Average Trust in Opponent	0.61	0.59	0.97

Table 2: Pairing the Fuzzy agent with other agents.

where s is the current state and s' is the end point of the negotiation.

To summarize, table 6.1 shows the range of values taken by the various simulation components. The trust model used was broadly based on [Ramchurn *et al.*, 2003; Marsh, 1994]. It should be noted that the effect of arguments on trust was arbitrarily chosen such that rewards and appeals have combined positive impact equating to the negative impact of threats. It was also deemed that rewards would bring about slightly greater trust than appeals (because rewards tend to show a more cooperative tendency than appeals [Karlins and Abelson, 1970]).

6.2 Experiments

We present our experiments as a series of hypotheses:

H1: *The Fuzzy mechanism fares better than the other mechanisms with perfect knowledge of argument values.* This can be deduced from the fact that the number of arguments used by the Fuzzy agent is significantly less than those used by the Ramping agent and slightly less than the Random agent (1373 vs 46922 vs 1405 in figure 2), while the gain in utility is higher for a Fuzzy agent than for the others (1371226 vs 861542 vs 1041130 in figure 2). The Fuzzy agent uses fewer cycles to reach agreements than other agents (1371 vs 1542 for a Ramping agent, and 1398 for a Random agent in figure 2). Ramping agents take a significantly larger number of cycles to reach agreements because this approach has no consideration for the actual strength of the argument itself (e.g. a threat might be known to have a low value for the opponent but its magnitude is not taken into consideration). Therefore, these agents are less persuasive and take more cycles to reach an agreement. The random approach does not consider the value of the argument, nor the type of the argument it wishes to send and this makes the arguments even less persuasive and leads to greater inefficiency.




Figure 1: Gain in utility, cycles and agreements made in imperfect knowledge case.

H2: *The Fuzzy mechanism fares better than the other mechanisms with imperfect knowledge of argument values.* This can

be supported by a number of factors; (i) Fuzzy and Random agents use fewer arguments on average than Ramping agents (2655 vs 2855 vs 4086, in figure 1), (ii) Fuzzy agents achieve higher utility than Ramping and Random agents (1410640 vs 791461 vs 568856 in figure 1), and (iii) Fuzzy agents use less cycles per agreement than the Ramping and Random agents (3.4 vs 8.8 vs 12.9 from figure 1). There is a strong negative correlation between the number of cycles required to reach an agreement and the utility gained (-0.98). This is due to the agent achieving more agreements and spending less time arguing. In particular, the Fuzzy agent has a more meaningful means of selecting its arguments, in terms of their value, than the other agents and hence does better. We also decided to test the Fuzzy agent's performance when set against the other agents (see table 2). In this case, the Fuzzy agent persuades the Ramping and Random agents more frequently than it does a Fuzzy opponent. The Ramping and Random agents cause the trust of the proponent in them to be relatively low compared to the Fuzzy opponent. Thus, the Fuzzy proponent uses more forceful arguments (hence convincing them) against Ramping and Random agents than it does with a Fuzzy opponent.

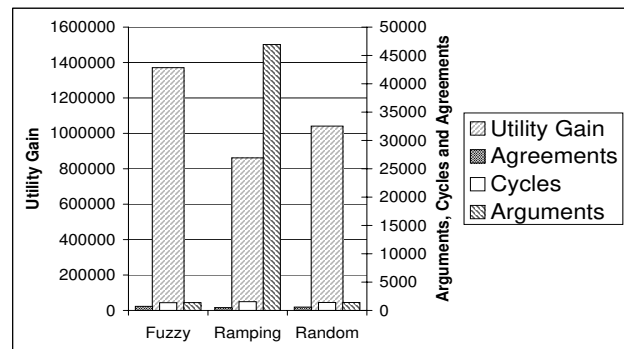


Figure 2: Gain in utility, cycles and agreements made in perfect knowledge case.

H3: *The Fuzzy mechanism fares better in the perfect knowledge case than in the imperfect knowledge one.* We expected the proponent to convince the opponent more easily than in the perfect knowledge case because it knows how the other agent evaluates its arguments. However, this hypothesis was **rejected** after analysing the results in figure 3 (note that the gain in utility for the proponent is slightly higher (by 39414) for fuzzy agents in the imperfect case than it is with perfect information). In the perfect knowledge case, the opponent can convince the proponent more easily since it knows its valuation of the arguments (hence reaching unprofitable agreements for the proponent in a fewer number of cycles!). On the other hand, in the imperfect information case, the counter proposals (backed up by arguments) do not convince the proponent so easily and allow it to continue negotiating (using more arguments, 2655 against 1373 in figure 3) until it reaches its intended agreement, rather than settling for a less profitable counter proposal (e.g. due to high trust).

H4: *Rhetorical arguments help achieve a larger number of agreements in a fewer number of cycles in both the perfect and imperfect cases.* Figures 1 and 2 show that all agents,

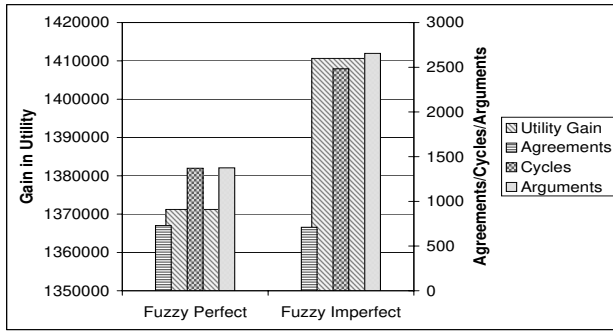


Figure 3: Comparing perfect and imperfect agents.

apart from the non-arguing ones, achieve a high number of agreements (> 300). Moreover, the difference is significant (of the order of 100 times bigger). This is because arguments can improve trust and can cause the agents to favour the proposed state relative to their current state (given that the argument is convincing enough). This is also partly due to the fact that agents are not allowed to change their proposals during the encounters to make them more attractive (see section 6.1) while arguments do change the trust values and can therefore give added value to the proposal without changing the proposal itself.




Figure 4: Individual and overall utility gain for Fuzzy agents in perfect knowledge.

From figures 1 and 2 we can calculate that the number of cycles per agreement reached is 3.47 (1.87 in the perfect case) for the Fuzzy agent, while it is 8.86 (3.11 in the perfect case) for Ramping agents, and 12.34 (2.34 for the perfect case) for Random agents. Non-arguing agents take a significantly larger number of cycles (> 100) to reach a negligible number of agreements. The low correlation between the number of arguments used and the number of cycles (+0.3) and the number of agreements reached (+0.3) signifies that the real factor determining the number of cycles per agreement is the quality of the argument chosen. In this respect, the Fuzzy agent does much better than the others because it is more precise in evaluating the strength of the argument to be sent.

H5: *Agents achieve greater individual and overall rewards as trust increases and trust nurtures trust.* Trust was found to be strongly positively correlated with the overall (+0.85) and individual rewards (+0.95) up to the point where trust is equal to 0.6. Beyond this value, trust has no effect on the overall or individual rewards (correlation of -0.35). This is due to the

Cycles	1468	1493	1439	1304	1215	1130	1001	1597	1477	1458	1494
Trust	0	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9	1.0

Table 3: Table of values for cycles (C) and trust values (T)

way our fuzzy reasoning works. When trust is above 0.6, the agents mainly use appeals and rewards as arguments. These tend to cause the trust in both agents to increase and hence the acceptance of proposals becomes much easier (meaning individual and overall gains settle at a stable value). Also, as the initial trust in the opponent was increased, the final trust at the end of the negotiations increased. From figure 4, the final trust in the opponent (i.e. the trust achieved at the end of 1000 negotiations) goes up as the initial trust increases. From this we deduce that agents that trust each other initially, will achieve higher levels of trust as they keep on negotiating.

H6: *High levels of trust help achieve agreements in fewer number of cycles.* The number of negotiation cycles for a pair of fuzzy agents is shown in table 6.2 for each value of trust. When trust is below 0.6, there is a high correlation between trust and the number of cycles to reach agreements (-0.96). After this threshold, the number of cycles go up and settle at a higher value. This trend follows from way the evaluation functions cause agreements to be reached more quickly as trust increases (see explanation in H5).

7 Related Work

Most of the literature on argumentation-based negotiation deals with protocols [McBurney *et al.*, 2002], abductive logic [Sadri *et al.*, 2001], and devising and formalizing rebutting and undercutting arguments [Parsons *et al.*, 1998; Amgoud *et al.*, 2000a]. Our approach is different from these models since they tend to focus on the product of argumentation (i.e. establishing the truth), while this work deals with the *process* of persuasion [Tindale, 1999; Perelman, 1982] whereby commitments to enact threats and rewards are used to convince another party (i.e. looking at the rhetorical aspect of the dialogue). Thus, these models can complement our work by handling the logical validity of appeals (i.e. arguments are rejected if they are not logically valid or can be defeated).

The closest implementation to our framework is that of [Kraus *et al.*, 1998]. In their work, agents are described as having a character depending on certain attributes (memoryless, knowledgeable, cooperative, etc.) such that only specific types of arguments can be used for each type of agent (e.g. appeals to past promise cannot be used for memoryless agents). The notion of argument generation and evaluation is based on the satisfaction of certain intentions (intention-to and intention-that). However, on the one hand, it would appear that such a model is too complex to be implemented in practical applications (it is based on a highly complex multi-modal logic), while on the other, the evaluation functions for the arguments are very specifically designed to work with agents that are explicitly coded in terms of intentions and desires. Also, their argumentation generation algorithm goes through a fixed order to select the argument to be sent. Threats are considered to be stronger than rewards which are, in turn, stronger than appeals. However, as demonstrated in section 6, this is not an effective strategy. The per-

suasive strength of an argument is partly dependent on the type of illocution used but there is no definite order established according to the persuasion literature [Tindale, 1999; Brembeck and Howell, 1976; Karlins and Abelson, 1970].

Sierra et al. [Sierra et al., 1998] also proposed a framework for persuasive negotiation. They exposed the basic prerequisites for designing a framework for persuasion using threats, rewards and appeals within a business process management scenario. However, their work does not give a thorough interpretation of arguments, nor an idea of how these can be chosen from a pool of available arguments. There is also a lack of a proper reasoning model which we provide for in this paper.

8 Conclusions and Future Work

We have presented a novel formal framework to describe persuasive negotiation based on rhetoric. In particular, we offered a general characterisation of the main types of persuasive illocutions and we developed a reasoning model that enables agents to determine which arguments to send and how to evaluate incoming arguments. Moreover, we have benchmarked our model against the only other computational model of *persuasive negotiation*¹⁴ and shown that it is more effective and efficient. Specifically, our agents achieved more agreements in fewer numbers of cycles than the Ramping agents while, at the same time, achieving greater utility.

Future work will concentrate on defining the evaluation functions for the states and illocutionary actions using multi attribute utility functions and showing how to develop particular negotiation strategies based on reinforcement learning and game theoretic techniques. We will develop the idea of social commitments with respect to the enactment of threats and rewards, and modify the mentalistic operational semantics given to rhetorical arguments. We will also extend our approach to cases in which there are more than two negotiating agents. We also aim to apply the model to pervasive computing environments in which agents manage the display of notifications on public and private devices by persuading the other agents to display or hide particular intrusions.

Acknowledgements

Sarvapali D. Ramchurn is supported by the FEEL project (IST-2000-26135) funded by the EU commission through the Disappearing Computer initiative. This work was also supported by the AgentCities mobility grant (IST-2000-28384).

References

[Amgoud et al., 2000a] L. Amgoud, N. Maudet, and S. Parsons. Modelling dialogues using argumentation. In Edmund Durfee, editor, *Proceedings of the Fourth International Conference on Multi-Agent Systems (ICMAS 1998)*, pages 31–38. IEEE, 2000.

[Amgoud et al., 2000b] L. Amgoud, S. Parsons, and N. Maudet. Arguments, dialogue, and negotiation. In W. Horn, editor, *Proceedings of ECAI2000*, pages 338–342. IOS Press, 2000.

[Axelrod, 1984] R. Axelrod. *The Evolution of Cooperation*. Basic Books, New York, USA, 1984.

[Brembeck and Howell, 1976] W. L. Brembeck and W. S. Howell. *Persuasion: A means of Social Influence*. Prentice Hall International, USA, 2nd edition, 1976.

[Brewka, 2001] G. Brewka. Dynamic argument systems: A formal model of argumentation processes based on situation calculus. *Journal of Logic and Computation*, 11(2):257–282, 2001.

[Faratin et al., 2002] P. Faratin, C. Sierra, and N. R. Jennings. Using similarity criteria to make trade-offs in automated negotiations. *Artificial Intelligence*, 142(2):205–237, 2002.

[Harel, 1984] D. Harel. Dynamic logic. In D. Gabbay and F. Guenther, editors, *Handbook of Philosophical Logic Volume II*, pages 497–604. D. Reidel Publishing Company, 1984.

[Jang, 1997] J.S.R Jang. *Neuro-Fuzzy and Soft Computing*. Prentice Hall, 1997.

[Jennings et al., 2002] N. R. Jennings, P. Faratin, A. R. Lomuscio, S. Parsons, C. Sierra, and M. Wooldridge. Automated negotiation: prospects, methods and challenges. *Int. J. of Group Decision and Negotiation*, 10(2):199–215, 2002.

[Karlins and Abelson, 1970] M. Karlins and H. I. Abelson. *Persuasion, how opinion and attitudes are changed*. Crosby Lockwood and Son, 2nd edition, 1970.

[Kraus et al., 1998] Sarit Kraus, Katia P. Sycara, and Amir Evenchik. Reaching agreements through argumentation: A logical model and implementation. *Artificial Intelligence*, 104(1-2):1–69, 1998.

[Marsh, 1994] S. Marsh. *Formalising Trust as a Computational Concept*. PhD thesis, Department of Mathematics and Computer Science, University of Stirling, 1994.

[McBurney et al., 2002] P. McBurney, S. Parsons, R. Van Eijk, and L. Amgoud. A dialogue-game protocol for agent purchase negotiations. *Journal of Autonomous Agents and Multi-Agent Systems. Special Issue on Argumentation in Inter-Agent Communication*, 2002.

[Parsons et al., 1998] S. Parsons, C. Sierra, and N. Jennings. Agents that reason and negotiate by arguing. *Journal of Logic and Computation*, 8(3):261–292, 1998.

[Perelman, 1982] C. Perelman. *The Realm of Rhetoric*. University of Notre Dame Press, 1st edition, 1982.

[Prakken and Vreeswijk, 2002] H. Prakken and Gerard Vreeswijk. Logics for defeasible argumentation. In D. Gabbay and F. Guenther, editors, *Handbook of Philosophical Logic*, volume 4, pages 219–318. Kluwer, Dordrecht, Netherlands, second edition, 2002.

[Ramchurn and Jennings, 2002] S.D Ramchurn and N. R. Jennings. Persuasive negotiation for autonomous agents. Technical report, University of Southampton, July 2002.

[Ramchurn et al., 2003] S. D. Ramchurn, C. Sierra, L. Godo, and N. R. Jennings. A computational trust model for multi-agent interactions based on confidence and reputation. In to appear, editor, *Workshop on Trust, Privacy, Deception and Fraud in Agent Societies, Melbourne, AAMAS, 2003*.

[Sabater and Sierra, 2002] Jordi Sabater and Carles Sierra. Regret: a reputation model for gregarious agents. In *Proc. of the 1st Int. Joint Conference on Autonomous Agents and Multi-Agent Systems*, pages 475–482, 2002.

[Sadri et al., 2001] F. Sadri, F. Toni, and P. Torroni. Dialogues for negotiation: agent varieties and dialogue sequences. In John-Jules Meyer and Milind Tambe, editors, *Intelligent Agent Series VIII: Proceedings of the 8th International Workshop on Agent Theories, Architectures, and Languages (ATAL 2001)*, volume 2333 of *Lecture Notes in Computer Science*, pages 69–84. Springer, 2001.

[Sadri et al., 2002] F. Sadri, F. Toni, and P. Torroni. Abductive logic programming architecture for negotiating agents. In *Proceedings of the 8th European Conference on Logics in Artificial Intelligence (JELIA 2002)*, volume 2424 of *Lecture Notes in Computer Science*, pages 419–431. Springer, 2002.

[Searle, 1969] J. Searle. *Speech Acts: An Essay in the Philosophy of Language*. Cambridge University Press, New York, 1969.

[Sierra et al., 1998] C. Sierra, N. R. Jennings, P. Noriega, and S. Parsons. A framework for argumentation-based negotiation. *Lecture Notes in Computer Science, Springer Verlag*, 1365:177–194, 1998.

¹⁴Other models of persuasion through logic-based argumentation exist (as seen in section 1) but these do not relate *specifically* to negotiation, hence their inapplicability to our context.

[Tindale, 1999] C. Tindale. *Acts of Arguing, A Rhetorical Model of Argument*. State University Press of New York, Albany, NY, USA, 1999.

[Toulmin, 1964] S. Toulmin. *The Uses of Argument*. Cambridge University Press, 1964.

[Ury and Fisher, 1999] W. Ury and R. Fisher. *Getting To Yes, Negotiating An Agreement Without Giving In*. Random House, 1999.