Agent Argumentation with Opinions and Advice

John Debenham & Carles Sierra

Abstract In argumentation-based negotiation the rhetorical illocutionary particles *Appeals*, *Rewards* and *Threats* have implications for the players that extend beyond a single negotiation and are concerned with building (business) relationships. This paper extends an agent's relationship-building argumentative repertoire with *Opinions* and *Advice*. A framework is described that enables agents to model their relationships and to use argumentative dialogue strategically both to achieve good negotiation outcomes and to build and sustain valuable relationships.

1 Introduction

The term *argumentation-based negotiation* has various meanings in multiagent systems [11]. *Classical argumentation* is the generation of arguments, usually as logical proofs, for and against a given course of action that support decision making processes. *Dialectical argumentation* is concerned with the argumentative process, and procedures by which argumentative dialogues are conducted. *Rhetorical argumentation* uses rhetorical illocutionary particles with the intention of modifying the beliefs of the listener. This paper is concerned with rhetorical argumentation.

Rhetorical argumentative dialogues have been traditionally organised around the rhetorical illocutionary particles *Offer*, *Accept* and *Reject* with the addition of particles such as *Appeals*, *Rewards* and *Threats* [17]. This form of argumentation has implications for the players that extend beyond a single negotiation and is concerned with building (business) relationships. When we reward or threaten we refer

John Debenham
QCIS, UTS, Broadway, NSW 2007, Australia, e-mail: debenham@it.uts.edu.au

Carles Sierra
IIIA, CSIC, Campus UAB, 08193 Bellaterra, Catalonia, Spain e-mail: sierra@iiia.csic.es

M. Bramer et al. (eds.), *Research and Development in Intelligent Systems XXVII*,
DOI 10.1007/978-0-85729-130-1_2,

to a future instant of time where the reward or threat will be effective, its scope goes beyond the current negotiation round.

This paper discusses the use of opinions and advice as rhetorical particles with the intention of building relationships. The intuition is that if an agent believes that another agent gives reliable and pertinent information, opinions and advice then this will strengthen their relationship. It is generally accepted that human agents rely heavily on their relationships in order to conduct business [19] [12]. An agent's *relationships* is a model of its interactions with other agents that enables it to exhibit a strategic social sense. The concept of trust [13] is one component of a relatlonship model — but there is far more to relationships than trust.

We will understand argumentation in this paper as an information exchange process between agents. Every illocution that an agent utters gives away valuable information. To evaluate each illocution exchanged we have built an information-based agent architecture on information theory. Information-based agents [15] have embedded tools from information theory that enable them to measure and manage strategic information — this makes them well-suited to measuring and managing the development of relationships through the exchange of rhetoric particles. The way in which this is achieved is described in detail in Section 6. As we will see, each agent has a world model that contains representations of every aspect of the world (including other agents) that the agent is interested in, as well as its strength of belief in the validity of the representation. When an information-based agent receives a rhetoric particle it updates its world model and calculates the *value* of that particle as information gain on the world model. Each particle is classified by identifying the components of the relationship model that it is relevant to, and its value is then used to update the strength of those components.

Section 2 describes the rhetoric particles: *informs*, *opinions*, and *advice*. The characteristics of relationships between agents are described and formalised in the LOGIC framework in Section 3. The formal model of relationships contains four component models that are all described in Section 4. Two further models, for trust and integrity, are presented in Section 5. The work then comes together in a discussion of strategies in Section 6, and Section 7 concludes.

2 Opinions and Advice

In this Section we describe the two rhetoric particles: `opinion` and `advise`:

- An `opinion` communicative act is a speaker's evaluation of a particular aspect of a thing in context, where the *context* is the set of all things that the thing is being, explicitly or implicitly, evaluated with or against.
- An `advise` communicative act is a speaker's evaluation of a particular aspect of a thing in the context of the speaker's beliefs of the listener's context. It is a directive in Searle's classification of speech acts.

Together with `inform`, these two acts are, in a sense, of increasing complexity. An `inform` communicates a verifiable fact, an `opinion` communicates the speaker's personal evaluation of something, and an `advice` communicates the speaker's opinion on what the listener's future states or actions should be.

We use the standard FIPA notation [1]. The rational effect (in the sense of the FIPA semantics) of these three particles is taken in two senses. First, the immediate effect that they will have on the listener's state or actions, and second, the effect that an act has on the relationship when the integrity of the information communicated has been verified. For the `inform` communicative act , agent i informs agent j of proposition p. The dual rational effects are:

$$\begin{array}{ll} \multicolumn{2}{l}{\langle i,\ \mathtt{inform}(j,\ p\ )\rangle} \\ \quad \mathtt{RE1:} & \mathrm{B}_j p \\ \quad \mathtt{RE2:} & \mathtt{Done}(\langle j,\ \mathtt{eval}(p,\ x)\rangle, \phi) \end{array}$$

where RE1 is as in FIPA standard, and $\langle j,\ \mathtt{eval}(p,\ x)\rangle$ is the action of agent j rating the integrity of proposition p as x, and proposition ϕ is true when the integrity of p is known to j. The evaluation is performed *ex post* at a time when opportunities to use the contents of the `inform` are well understood. It is over a fuzzy scale, $\mathtt{eval} \in [0,1]$, that must contain 0 (meaning "is of absolutely no use"), and must contain 1 (meaning "valued most highly").

The `opinion` and `advise` communicative acts are not part of the FIFA specification and are now defined using the FIPA notation.

The representation of the `opinion` communicative act contains:

- the thing that is the *subject* of the opinion,
- the *aspect*, or attribute, of the thing that is being evaluated,
- a distribution over some evaluation space representing the *rating* of the aspect of the thing in the context, and
- optionally the *context* in which the evaluation is made, and a *reason* supporting the opinion.

An `opinion` action indicates that the speaker:

- believes he knows that the listener holds a particular intention,
- believes his opinion of a thing is related to the listener's intention, and is more accurate than the listener's opinion of it

In the following, the speaker, i, informs the listener, j that his rating of an aspect, s, of a thing, t, is e in (the optional) context c for the (optional) reason, r. The two rational effects following represent the dual motives for uttering the illocution:

$$\begin{array}{ll} \multicolumn{2}{l}{\langle i,\ \mathtt{opinion}(j,s,t,e[,c,r])\rangle} \\ \quad \mathtt{FP:} & \mathrm{B}_j\ \mathtt{Rates}(j,s,t,e'[,c,r])\ [\wedge \mathrm{B}_i \mathrm{I}_j c \wedge \mathrm{B}_i r] \wedge \\ & \mathrm{B}_i \mathrm{I}_j\ \mathtt{Done}(\langle j,\ \mathtt{eval}(s,t,e,x[,c,r])\rangle, \phi) \\ \quad \mathtt{RE1:} & \mathrm{B}_j\ \mathtt{Rates}(j,s,t,e''[,c,r]) \wedge |e-e''| < |e-e'| \\ \quad \mathtt{RE2:} & \mathtt{Done}(\langle j,\ \mathtt{eval}(s,t,e,x[,c,r])\rangle, \phi) \end{array}$$

That is, i believes that as a result of expressing an opinion about t, j's rating of t is now closer to i's rating that it was prior to the opinion being uttered, where some suitable distance measure between distributions is assumed, and $\texttt{eval}(s,t,e,x[,c,r])$ is the action of evaluating the rating e in context, and ϕ is true when the evaluation is performed.

An `advise` *action* indicates that the speaker:

- believes he knows that the listener holds a particular intention,
- believes his knowledge of facts concerning the listener's intention is better than the listener's knowledge of them,
- intends the listener to believe that the advised action is in the listener's interests, and
- believes that the listener may act otherwise.

In the following, the speaker, i, advises the listener, j, that the speaker believes the listener should perform some action, a, if the listener's context includes the intention to achieve a goal, c. The two feasibility preconditions are alternative representations of i's beliefs of the superiority of his knowledge, and the two rational effects represent the dual motives for uttering the illocution:

$$
\begin{array}{ll}
\multicolumn{2}{l}{<i,\ \texttt{advise}(j,a,c)>} \\
\texttt{FP:} & \texttt{B}_i\texttt{I}_j c \ \wedge\ \texttt{B}_i(\ \texttt{W}_i(c) \rightarrow\ \texttt{W}_{j\backslash i}(c)\ )\wedge \neg\texttt{B}_i\texttt{I}_j\ \texttt{Done}(<j,a>)\ \wedge \\
 & \texttt{B}_i\texttt{I}_j\ \texttt{Done}(<j,\ \texttt{eval}(a,c,x)>,\phi) \\
\texttt{or:} & \texttt{B}_i\texttt{I}_j c \ \wedge\ \texttt{B}_i(\mathbb{H}(\texttt{W}_i(c)) < \mathbb{H}(\texttt{W}_{j\backslash i}(c)))\wedge\neg\texttt{B}_i\texttt{I}_j\ \texttt{Done}(a)\ \wedge \\
 & \texttt{B}_i\texttt{I}_j\ \texttt{Done}(<j,\ \texttt{eval}(a,c,x)>,\phi) \\
\texttt{RE1:} & \texttt{Done}(<j,a>) \\
\texttt{RE2:} & \texttt{Done}(<j,\ \texttt{eval}(a,c,x)>,\phi)
\end{array}
$$

where:
$\texttt{eval}(a,c,x)$ is the action of evaluating action a as x in context c, as above
$\texttt{W}_i(c)$ denotes all of i's beliefs concerning c — i.e. that part of i's world model
$\texttt{W}_{j\backslash i}(c)$ denotes i's beliefs concerning all of j's beliefs concerning c
$\texttt{W}_i(c) \rightarrow \texttt{W}_{j\backslash i}(c)$ denotes that everything in $\texttt{W}_{j\backslash i}(c)$ can be derived from a subset of $\texttt{W}_i(c)$
$\mathbb{H}(S)$ denotes the overall uncertainty of the set of beliefs S — possibly as entropy

3 Relationships

A *relationship* between two agents is somehow encapsulated in their *history* that is a complete record of their interactions. This potentially large amount of information is usually summarised by agents into various models. For example, the majority of agents construct a world model and a trust model [3]. There is evidence from psychological studies that humans seek a *balance* in their negotiation relationships. The classical view [2] is that people perceive resource allocations as being distributively fair (i.e. well balanced) if they are proportional to inputs or contributions (i.e. equitable). In the case of partners there is some evidence [4] that the allocations of goods and burdens (i.e. positive and negative utilities) are perceived as fair, or in balance, based on equity for burdens and equality for goods.

The LOGIC illocutionary framework for classifying argumentative interactions was first described in [16] where it was used to help agents to prepare for a negotiation in the *prelude stage* of an interaction. The work in this paper generalises that framework and uses it to define one of the two dimensions of the relationship model described in Section 4; the second dimension is provided by the structure of the ontology as specified by a partial order $\leq$ defined by the is-a hierarchy, and a distance measure between concepts such as Equation 1. The five LOGIC categories for information are quite general:

- Legitimacy contains *information* that may be part of, relevant to or in justification of contracts that have been signed.
- Options contains information about *contracts* that an agent may be prepared to sign.
- Goals contains information about the *objectives* of the agents.
- Independence contains information about the agent's *outside options* — i.e. the set of agents that are capable of satisfying each of the agent's needs.
- Commitments contains information about the *commitments* that an agent has.

and are used here to categorise all incoming communication that feeds into the agent's relationship model. As we will see this categorisation is not a one-to-one mapping and some illocutions fall into multiple categories. These categories are designed to provide a model of the agents' information that is relevant to their relationships, and are *not* intended to be a universal categorising framework for all utterances.

This paper is written from the point of view of an agent α is in a *multiagent system* with a finite number of other agents $\mathscr{B} = \{\beta_1, \beta_2, \dots\}$, and a finite number of *information providing agents* $\Theta = \{\theta_1, \theta_2, \dots\}$ that provide the *context* for all events in the system — Θ^t denotes the state of these agents at time t. α observes the actions of another agent β in the context Θ^t. The only thing that α 'knows for certain' is its *history* of past communication that is retains in the repository $\mathscr{H}_\alpha^t$. Each *utterance* in the history contains: an illocutionary statement, the sending agent, the receiving agent, the time that the utterance was sent or received.

Observations are of little value unless they can be verified. α may not posses a comprehensive range of reliable sensory input devices. Sensory inadequacy is dealt with invoking an *institution agent*, ξ, that truthfully, accurately and promptly reports what it sees.

All communication is recorded in α's history $\mathscr{H}_\alpha^t$ that in time may contain a large amount of data. The majority of agent architectures include models that summarise the contents of $\mathscr{H}^t$; for example, a *world model* and a *trust model*. In this paper we describe two models, a *relationship model* and an *integrity model* that are specifically designed to assist an agent to manage information asymmetries. To build the relationship model we will use the LOGIC framework to categorise the information in utterances received. That is, α requires a categorising function $v : U \rightarrow \mathscr{P}(\{\mathsf{L},\mathsf{O},\mathsf{G},\mathsf{I},\mathsf{C}\})$ where U is the set of utterances. The power set, $\mathscr{P}(\{\mathsf{L},\mathsf{O},\mathsf{G},\mathsf{I},\mathsf{C}\}$, is required as some utterances belong to multiple categories. For example, "I will not pay more for wine than the price that John charges" is categorised as both Option and Independence.

We assume an ontology that includes a (minimum) repertoire of elements: a set of *concepts* (e.g. quantity, quality, material) organised in a is-a hierarchy (e.g. platypus is a mammal, australian-dollar is a currency), and a set of relations over these concepts (e.g. price(beer,AUD)).[1]

We model ontologies following an algebraic approach [8]. An ontology is a tuple $\mathscr{O} = (C, R, \leq, \sigma)$ where:

1. C is a finite set of concept symbols (including basic data types);
2. R is a finite set of relation symbols;
3. $\leq$ is a reflexive, transitive and anti-symmetric relation on C (a partial order)
4. $\sigma : R \rightarrow C^+$ is the function assigning to each relation symbol its arity

where $\leq$ is a traditional *is-a* hierarchy, and R contains relations between the concepts in the hierarchy.

The concepts within an ontology are closer, semantically speaking, depending on how far away they are in the structure defined by the $\leq$ relation. Semantic distance plays a fundamental role in strategies for information-based agency. A measure [9] bases the *semantic similarity* between two concepts on the path length induced by $\leq$ (more distance in the $\leq$ graph means less semantic similarity), and the *depth* of the subsumer concept (common ancestor) in the shortest path between the two concepts (the deeper in the hierarchy, the closer the meaning of the concepts). Semantic similarity could then be defined as:

$$\mathrm{Sim}(c, c') = e^{-\kappa_1 l} \cdot \frac{e^{\kappa_2 h} - e^{-\kappa_2 h}}{e^{\kappa_2 h} + e^{-\kappa_2 h}} \tag{1}$$

where l is the length (i.e. number of hops) of the shortest path between the concepts, h is the depth of the deepest concept subsuming both concepts, and κ_1 and κ_2 are parameters scaling the contribution of shortest path length and depth respectively.

4 The Relationship Model $\mathscr{R}^t_{\alpha\beta}$

This Section describes how an agent's relationships are modelled using both the LOGIC framework (described in Section 3) and the structure of the ontology (described in Section 2). The relationship model is used in Section 6 to manage the argumentative discourse between two agents. Two models are described in Section 5

[1] Usually, a set of axioms defined over the concepts and relations is also required. We will omit this here.

that are used in Section 6 to select which agent to interact with in the context of a particular need.

All of α's models are summaries of its history $\mathscr{H}^t_\alpha$. The *relationship model* that α has of β consists of four component models. First, α's *intimacy model* of β's private information describes *how much* α knows about β, $I^t_{\alpha\beta}$ — this information will have been extracted from the dialogue including `inform`, `opinion` and `advise` utterances. Second, α's *reliability model* of *how reliable* the information summarised in $I^t_{\alpha\beta}$ is, $R^t_{\alpha\beta}$. Third, α's *reflection model* of β's model of α's private information, $J^t_{\alpha\beta}$. Fourth, a *balance model*, $B^t_{\alpha\beta}$, that measures the difference in the rate of growth of $I^t_{\alpha\beta}$ and $J^t_{\alpha\beta}$. Abusing notation we denote this by: $\frac{d}{dt}I^t_{\alpha\beta}$ and $\frac{d}{dt}J^t_{\alpha\beta}$ across the structure $\{\mathsf{L,O,G,I,C}\} \times \mathscr{O}$.

The remainder of this section details how these four component models are calculated. In addition to the models described in this Section, α is assumed to have a world model, $\mathscr{M}^t$, that represents everything in its world that it is interested in. The procedure for updating the world model relies on estimates of the reliability of all incoming utterances. $R^t_{\alpha\beta}$ is used for this purpose, and is used both to support the update process for $\mathscr{M}^t$ and to estimate the reliability of $I^t_{\alpha\beta}$. The description given employs the machinery to update the world model in our information-based agents [15]. However it can be adapted to the machinery used by any agent that represents uncertainty in its world model using probability distributions, that is: $\mathscr{M}^t = \{X_i\}_i$ where X_i are random variables. In addition to the world model and the models described in this paper an agent may construct other models such as an *honour model* [14].

Utterances are represented in the world model $\mathscr{M}^t_\alpha$ as probability distributions, (X_i), in first-order probabilistic logic $\mathscr{L}$. Representing an utterance in the world model requires its semantics. Semantics of utterances are specified as constraints on distributions in the world model. For example, in a simple multi-issue contract negotiation α may estimate $\mathbb{P}^t(\text{acc}(\beta,\alpha,\delta))$, the probability that β would accept contract δ, by observing β's responses. The distribution $\mathbb{P}^t(\text{acc}(\beta,\alpha,\delta)) \in \mathscr{M}^t_\alpha$ is classified as an Option in LOGIC. Using shorthand notation, if β sends the message Offer(δ_1) then α derives the constraint: $K_{\text{acc}(\beta,\alpha,\delta)}(\text{Offer}(\delta_1)) = \{\mathbb{P}^t(\text{acc}(\beta,\alpha,\delta_1)) = 1\}$, and if this is a counter offer to a former offer of α's, δ_0, then: $K_{\text{acc}(\beta,\alpha,\delta)}(\text{Offer}(\delta_1)) = \{\mathbb{P}^t(\text{acc}(\beta,\alpha,\delta_0)) = 0\}$.[2]

Updating $\mathscr{M}^t_\alpha$ is complicated if the reliability of utterances received is taken into account — it would certainly be foolish for α to believe that every utterance received from β was correct — whereas all utterances received from the institution agent ξ *are* assumed to be correct. The procedure for doing this, and for attaching reliability estimates to utterances is described below.

The idea of intimacy and balance is that intimacy summarises the degree of closeness, and *balance* is degree of fairness. Informally, *intimacy* measures how

[2] In the not-atypical special case of multi-issue bargaining where the agents' preferences over the individual issues *only* are known and are complementary to each other's, maximum entropy reasoning can be applied to estimate the probability that any multi-issue offer will be acceptable to β by enumerating the possible worlds that represent β's "limit of acceptability" [15].

much one agent knows about another agent's private information, and *balance* measures the extent to which information revelation process between the agents is 'fair'. The *intimacy* and *balance* models are structured using the LOGIC illocutionary framework and the ontology $\mathcal{O}$[3]. For example, an utterance meaning that agent β accepts agent α's previously offered deal δ is classified as an Option, and $< \alpha, \texttt{inform}(\beta, \mathit{info}) >$ meaning that agent β informs α about *info* and commits to the truth of it is classified as Legitimacy.

4.1 The Intimacy Model: $I^t_{\alpha\beta}$

The *intimacy* of α's relationship with β, $I^t_{\alpha\beta}$, models how much α knows about β's private information and is represented as real numeric values over $\{\mathsf{L,O,G,I,C}\} \times \mathcal{O}$. Suppose α receives an utterance u from β and that the LOGIC category $f \in v(u)$, where v is the categorising function described in Section 3. For any concept $c \in \mathcal{O}$, we extend the definition of Sim by defining $\text{Sim}(u,c) = \max_{c' \in u} \text{Sim}(c',c)$ where Sim is a semantic distance function such as that described in Equation 1. Denote the value of $I^t_{\alpha\beta}$ in position $(f,c) \in \{\mathsf{L,O,G,I,C}\} \times \mathcal{O}$ by $I^t_{\alpha\beta(f,c)}$ then:

$$I^t_{\alpha\beta(f,c)} = \begin{cases} \rho \times I^{t-1}_{\alpha\beta(f,c)} + (1-\rho) \times \mathbb{I}^t(u) \times \text{Sim}(u,c) & \text{if } u \text{ received,} \\ \mu \times I^{t-1}_{\alpha\beta(f,c)} & \text{otherwise.} \end{cases} \quad (2)$$

for any c, where $\mu < 1$ is the decay rate, ρ is the learning rate, and $\mathbb{I}^t(u)$ is Shannon information gain as given by Equation 7 that is described below. The method for estimating $\mathbb{I}^t(u)$ takes account of the reliability of u. The decay rate μ is a constant just less than 1 ensures the decay of $I^t_{\alpha\beta}$ towards a zero state if no utterances are received. α's estimate of β's intimacy on α, $J^t_{\alpha\beta}$, is constructed similarly by assuming that β's reasoning apparatus mirrors α's.

Equation 2 above requires an estimate of the information gain in an utterance, $\mathbb{I}^t(u)$. The calculation is fairly technical but as it is part of the procedure for updating the world model the marginal cost in building the relationship model is very low.

α's world model $\mathcal{M}^t_\alpha$ is a set of random variables, $\mathcal{M}^t = \{X_i, \ldots, X_n\}$ each representing an aspect of the world that α is interested in. In the absence of in-coming messages the integrity of $\mathcal{M}^t$ decays. α may have background knowledge concerning the expected integrity as $t \to \infty$. Such background knowledge is represented as a *decay limit distribution*. One possibility is to assume that the decay limit distribution has maximum entropy whilst being consistent with observations. Given a distribution, $\mathbb{P}(X_i)$, and a decay limit distribution $\mathbb{D}(X_i)$, $\mathbb{P}(X_i)$ decays by:

$$\mathbb{P}^{t+1}(X_i) = \Delta_i(\mathbb{D}(X_i), \mathbb{P}^t(X_i)) \quad (3)$$

[3] Only a subset of the ontology is required. The idea is simply to capture "How much has Carles told me about wine", or "how much do I know about Carles' commitments (possibly with other agents) concerning cheese".

where Δ_i is the *decay function* for the X_i satisfying the property that $\lim_{t\to\infty} \mathbb{P}^t(X_i) = \mathbb{D}(X_i)$. For example, Δ_i could be linear: $\mathbb{P}^{t+1}(X_i) = (1-\mu_i) \times \mathbb{D}(X_i) + \mu_i \times \mathbb{P}^t(X_i)$, where $\mu_i < 1$ is the decay rate for the i'th distribution. Either the decay function or the decay limit distribution could also be a function of time: Δ_i^t and $\mathbb{D}^t(X_i)$.

Suppose that α receives an utterance u from agent β at time t. This utterance could be an `inform`, an `opinion` or an `advise`. Suppose that this utterance's contents is qualified with probability z. α attaches an epistemic belief $R^t_{\alpha\beta}(u)$ to u — the reliability model $R^t_{\alpha\beta}$ is described below in Section 4.2. The semantics of utterance u is given by specifying constraints on those random variables in the world model that the receipt of u will effect. For $X_i \in \mathscr{M}^t$ we denote the constraint on X_i due to the receipt of u as $K_{X_i}(u)$ that are called *update functions*.

Given a prior distribution $\mathbf{p}_i = \mathbb{P}^t(X_i)$ let $\mathbf{p}_{i(u)}$ be the distribution with minimum relative entropy[4] with respect to $\mathbf{p}_i$: $\mathbf{p}_{i(u)} = \arg\min_{\mathbf{r}} \sum_j r_j \log \frac{r_j}{p_j}$ that satisfies the constraints $K_{X_i}(u)$. Then let $\mathbf{q}_{i(u)}$ be the distribution:

$$\mathbf{q}_{i(u)} = \begin{cases} R^t_{\alpha\beta}(u) \times \mathbf{p}_{i(u)} + (1 - R^t_{\alpha\beta}(u)) \times \mathbf{p}_i & \text{if } R^t_{\alpha\beta}(u) > 0.5, \\ \mathbf{p}_i & \text{otherwise.} \end{cases} \tag{4}$$

where $R^t_{\alpha\beta}(u)$ is determined by the reliability model below. The condition $R^t_{\alpha\beta}(u) > 0.5$ prevents information with an expected evaluation less than the ambivalence point (i.e. 0.5 as discussed in Section 4.2) from entering the process for updating $\mathscr{M}^t$. For example, $R^t_{\alpha\beta}(u) = 0$ means that u is certainly of no value. Then let:

$$\mathbb{P}^t(X_{i(u)}) = \begin{cases} \mathbf{q}_{i(u)} & \text{if } \mathbf{q}_{i(u)} \text{ is ``more interesting'' than } \mathbf{p}_i \\ \mathbf{p}_i & \text{otherwise} \end{cases} \tag{5}$$

A general measure of whether $\mathbf{q}_{i(u)}$ is *more interesting* than $\mathbf{p}_i$ is: $\mathbb{K}(\mathbf{q}_{i(u)} \| \mathbb{D}(X_i)) > \mathbb{K}(\mathbf{p}_i \| \mathbb{D}(X_i))$, where $\mathbb{K}(\mathbf{x}\|\mathbf{y}) = \sum_j x_j \ln \frac{x_j}{y_j}$ is the Kullback-Leibler distance between two probability distributions $\mathbf{x}$ and $\mathbf{y}$.

Finally merging Equation 5 and Equation 3 we obtain the method for updating a distribution X_i on receipt of a message u:

$$\mathbb{P}^{t+1}(X_i) = \Delta_i(\mathbb{D}(X_i), \mathbb{P}^t(X_{i(u)})) \tag{6}$$

This procedure deals with integrity decay, and with two probabilities: first, the probability z in the utterance u, and second the reliability $R^t_{\alpha\beta}(u)$ that α attached to u.

[4] Given a probability distribution $\mathbf{q}$, the *minimum relative entropy distribution* $\mathbf{p} = (p_1, \ldots, p_I)$ subject to a set of n linear constraints $\mathbf{g} = \{g_j(\mathbf{p}) = \mathbf{a_j} \cdot \mathbf{p} - c_j = 0\}, j = 1, \ldots, n$ (that must include the constraint $\sum_i p_i - 1 = 0$) is: $\mathbf{p} = \text{MRE}(\mathbf{q}, \mathbf{g}) = \arg\min_{\mathbf{r}} \sum_j r_j \log \frac{r_j}{q_j}$. This may be calculated by introducing Lagrange multipliers λ: $L(\mathbf{p}, \lambda) = \sum_j p_j \log \frac{p_j}{q_j} + \lambda \cdot \mathbf{g}$. Minimising L, $\{\frac{\partial L}{\partial \lambda_j} = g_j(\mathbf{p}) = 0\}, j = 1, \ldots, n$ is the set of given constraints $\mathbf{g}$, and a solution to $\frac{\partial L}{\partial p_i} = 0, i = 1, \ldots, I$ leads eventually to $\mathbf{p}$. Entropy-based inference is a form of Bayesian inference that is convenient when the data is sparse [5] and encapsulates common-sense reasoning [10].

The Shannon information gain in X_i is: $\mathbb{I}^t X_i = \mathbb{H}^t(X_i) - \mathbb{H}^{t-1}(X_i)$, and if the distributions in $\mathscr{M}^t$ are independent then the Shannon information gain for $\mathscr{M}^t$ following the receipt of utterance u is:

$$\mathbb{I}^t(u) = \sum_{X_i} \mathbb{I}^t X_i \tag{7}$$

4.2 The Reliability Model: $R^t_{\alpha\beta}$

Equation 4 above requires an estimate of the reliability of an utterance, $R^t_{\alpha\beta}(u)$, which is detailed in this Section. The reliability model is constructed by observing the difference between β's utterance u at time t and its subsequent evaluation[5] at time t'. This means that for β, building a strong $R^t_{\alpha\beta}$ will be a slow process. This is consistent with the observation that business relationships between human agents tend to build gradually over time.

We now consider how the estimates $R^t_{\alpha\beta(f,c)}$ develop in time. At each time step:

$$R^t_{\alpha\beta(f,c)} = \mu \times R^{t-1}_{\alpha\beta(f,c)} + (1-\mu) \times 0.5$$

representing the decay of the reliability towards the maximum entropy, or ambivalence point, value. Now suppose that u is received from agent β at some time and is evaluated, possibly with the assistance of the institution agent, ξ, at some later time t as $\texttt{eval}(u)$ as described in Section 2. This evaluation is on a fuzzy scale in $[0,1]$ that contains 0 and 1, i.e. $\texttt{eval}(u) \in [0,1]$. Suppose that the LOGIC category $f \in v(u)$, where v is the categorising function described in Section 3. For any category c, let $r = R^{t-1}_{\alpha\beta(f,c)}$ and:

$$\begin{aligned}
e &= (\rho \times 0.5) + (1-\rho) \times \texttt{eval}(u) \\
e' &= e \times \text{Sim}(u,c) \\
e'' &= (\text{Sim}(u,c) \times (e-1)) + 1
\end{aligned}$$

where ρ is the learning rate, then $R^t_{\alpha\beta(f,c)} = g(r, e', e'')$ where:

$$g(r,e',e'') = \begin{cases} \text{comb}(r,e') & \text{if } e > 0.5 \text{ and } e' > 0.5, \\ \text{comb}(r,e'') & \text{if } e < 0.5 \text{ and } e'' < 0.5, \\ r & \text{otherwise.} \end{cases} \tag{8}$$

where $\text{comb}(x,y) = \frac{x \times y}{(x \times y)+(1-x)\times(1-y)}$ is the combination of independent probabilities x and y. The assumption of independence is rather radical and the moderation of $\texttt{eval}(u)$ to e using the learning rate ρ is intended to compensate for this. The con-

[5] Evaluation is meant in the sense of the $\texttt{eval}$ functions that are part of the rational effect expressions in Section 2.

ditions in Equation 8 ensures that the update is only applied when Sim is reasonably large. When $\text{Sim} = 1$, $e' = e'' = e$. Those conditions limit the update to those values of e' and e'' that are "on the same side of" 0.5 as e.

5 Trust and Integrity

We now describe two measures that are attached to *complete dialogues* that are used in Section 6 to assist with the selection of negotiation partners for a particular need. The first of these is *trust* that measures the difference been commitments made during a dialogue and the eventual enactment of those commitments. The second is *integrity* that measures the difference between expectation and evaluation of the dialogue — the integrity measure is aggregated from values of the `eval` function.

The estimation of trust and integrity can be interpreted as a pattern mining exercise from the information in $\mathscr{C}^t_{\alpha\beta}$ to find the 'best' hypothesis that describes $\mathscr{C}^t_{\alpha\beta}$, where $\mathscr{C}^t_{\alpha\beta} \subset \mathscr{H}^t_{\alpha\beta}$ contains those utterances that contain evaluations of enactments, for trust, and of consumption, for integrity. One neat way to perform this induction is the *minimum description length principle* [7] that is founded on the minimisation of the cost of communicating a body of knowledge from one agent to another that thus has a fundamental affinity with distributed autonomous systems:

$$\mathscr{I}^t_{\alpha\beta} \triangleq \arg\min_M (\mathbb{L}(M) + \mathbb{L}(\mathscr{C}^t_{\alpha\beta} \mid M)) \tag{9}$$

where $\mathbb{L}(M)$ is the length of the shortest encoding of M, and $\mathbb{L}(\mathscr{C}^t_{\alpha\beta} \mid M)$ is the length of the shortest encoding of $\mathscr{C}^t_{\alpha\beta}$ given M. This definition is as neat as it is computationally expensive — it divides $\mathscr{C}^t_{\alpha\beta}$ into that which may be generalised and that which may not.

The definition of $\mathscr{I}^t_{\alpha\beta}$ in Equation 9 appears problematic for three reasons. First, if M can be any Turing computable model the definition is not computable, second a single language is required for representing M, and third the meaning of 'the length of the shortest encoding' is not clear. The second and third reason have been resolved [7]. The first, computability problem can be solved by restricting the models to some specific class. If the models are restricted to Bayesian decision graphs over finite spaces then Equation 9 is computable [18].

Equation 9 does not take time into account. To allow for varying strength of observations with time we construct instead $\mathscr{C}^{*t}_{\alpha\beta}$ that is the same as $\mathscr{C}^t_{\alpha\beta}$ except each evaluation, x, is replaced by a random variable X over evaluation space. These probability distributions are constructed by: $\lambda \times X + (1-\lambda) \times D_X$ where D_X is the *decay limit distribution*[6] for X — and X is a distribution with a '1' indicating the position of the evaluation and 0's elsewhere. Despite its elegance, Equation 9 is computationally expensive. [15] describes a computationally friendly method for evaluating trust that may also be used for integrity.

[6] If the decay limit distribution is unknown we use a maximum entropy distribution.

6 'Relationship-aware' Argumentation Strategies

Given a need ν in context Θ^t one way for agent α to select an interaction partner, β, on the basis of their past behaviour, is by reference to the trust model and the integrity model. Suppose that α uses model $M^t_{\alpha\beta}$, then the integrity of that model will decay in time, $M^t_{\alpha\beta} \rightarrow D$, by Equation 3 or similar, that is, the uncertainty $\mathbb{H}(M^t_{\alpha\beta})$ will increase in time, until the model is refreshed with new observations. So the rate of refreshment with new observations needs to be such that the uncertainty, $\mathbb{H}(M^t_{\alpha\beta})$, generally decreases in time. Given a need ν the set of partners that α considers are called the *pool* for ν. For each potential partner, β, we assume that α is able to estimate, $\mathbb{P}^t(\beta \gg)|\nu$, the probability that a negotiation with β to satisfy ν will lead to a better outcome than with any other in the pool. Then select a partner using the stochastic strategy: $\mathbb{P}^t(\text{Select } \beta_i) = \mathbb{P}^t(\beta_i \gg)|\nu$.

For each agent in the pool α has a view on the desired form of the relationship model, particularly the intimacy model, $I^t_{\alpha\beta}$ — that is the model that he would realistically wish to have. This is the *relationship target* $T^t_{\alpha\beta}$ for agent β. Using the $\{\mathsf{L,O,G,I,C}\}$ structure the target is expressed as a target for the pair of intimacy models in Section 4.1: $T^t_{\alpha\beta(f,c)} = \left(TI^t_{\alpha\beta(f,c)}, TJ^t_{\alpha\beta(f,c)}\right)$ where $TI_{\alpha\beta}$ and $TJ_{\alpha\beta}$ are respectively the targets for $I_{\alpha\beta}$ and $J_{\alpha\beta}$.

Having selected the interaction partner, and having set the relationship target, α now manages the interaction itself. α has a model of its intimacy with β, $I^t_{\alpha\beta(f,n)}$ where $\nu \in n \in \mathscr{O}$, and its target for β, $T^t_{\alpha\beta(f,n)}$. When the interaction with β is complete, the intimacy model, $I^{t'}_{\alpha\beta(f,n)}$ will have changed. Before the interaction commences, α may desire that $I^{t'}_{\alpha\beta(f,n)}$ will have changed "in the direction of" $T^t_{\alpha\beta(f,n)}$. This is formalised as the *negotiation target*, $N^t_{\alpha\beta(f,c)} = \left(NI^t_{\alpha\beta(f,c)}, NJ^t_{\alpha\beta(f,c)}\right)$, that is α's aspirations at time t for intimacy at time t'. Given the uncertainty in behaviour in any negotiation, the negotiation target is an approximate indication only of what should be achieved.

Any utterance that an agent makes gives away information if the receiving agent revises its world model as a result of receiving the utterance. In single-issue `offer`, `accept` and `reject` negotiation the *equitable information revelation* strategy is: α responds to β's offer with an offer o that gives β equivalent information gain as α has observed provided that o is acceptable to α[7]. Formally, if α receives an offer o from β at time t then α will observe information gain $\mathbb{H}(\mathscr{M}^{t-1}) - \mathbb{H}(\mathscr{M}^t)$ and so responds with an offer o' which is such that: $\mathbb{H}(\mathscr{M}^t_\beta) - \mathbb{H}(\mathscr{M}^t_\beta \oplus o') \approx \mathbb{H}(\mathscr{M}^{t-1}) - \mathbb{H}(\mathscr{M}^t)$ as long as o' is acceptable to α. If the negotiation is single-issue then this strategy determines a unique offer and yields a sequence of alternating offer exchanges that is almost "classic market haggling".

[7] This assumes, not unreasonably, that α and β model each other's limit price with a random variable in their respective world models.

For multi-issue `offer`, `accept` and `reject` negotiation we assume that α estimates the probability that any proposed deal, δ, is acceptable, $\mathbb{P}^t(Acc_\alpha(\delta|\mathscr{H}^t_\alpha))$, that is accompanied by a threshold value τ meaning that if $\mathbb{P}^t(Acc_\alpha(\delta|\mathscr{H}^t_\alpha)) > \tau$ then δ is acceptable. We also assume that α estimates the probability that the deal will be acceptable to β, $\mathbb{P}^t(Acc_\beta(\delta|\mathscr{H}^t_\alpha))$; an estimate for this may be derived from the offers that β has both made and rejected using maximum entropy inference [14]. Given these two estimates then an analogue of the issue-tradeoffs strategy in [6] is for α to offer: $\delta^* = \arg\max_\delta\{\mathbb{P}^t(Acc_\beta(\delta|\mathscr{H}^t_\alpha)) \mid \mathbb{P}^t(Acc_\alpha(\delta|\mathscr{H}^t_\alpha)) > \tau\}$.

The issue-tradeoffs strategy described above does not take into account the expected information gain from making such a proposal: $\mathbb{H}(\mathscr{M}^t_\beta) - \mathbb{H}(\mathscr{M}^t_\beta \oplus \delta)$. The consideration of information gain adds an interesting dimension. Consider the set of deals of similar acceptability to β as δ^*: $\Delta = \{\delta \mid \mathbb{P}^t(Acc_\beta(\delta|\mathscr{H}^t_\alpha)) \approx \delta^* \wedge \mathbb{P}^t(Acc_\alpha(\delta|\mathscr{H}^t_\alpha)) > \tau\}$. Each $\delta \in \Delta$ are similarly acceptable to each agent but are of potentially different information gain to β: $\mathbb{H}(\mathscr{M}^t_\beta) - \mathbb{H}(\mathscr{M}^t_\beta \oplus \delta)$. α is now is a position to decide how to manage the revelation of information in the proposals it makes, and may decide to do so equitably or otherwise.

The term *tactics* is used to refer to the strategy that wraps a possibly empty proposal in argumentation to form a complete utterance. The equitable information revelation strategy extends without modification to argumentation across the full structure of $\{\mathsf{L,O,G,I,C}\}$. If α receives an utterance u from β at time t then α responds with u' which is such that: $\mathbb{H}(\mathscr{M}^t_\beta) - \mathbb{H}(\mathscr{M}^t_\beta \oplus u') \approx \mathbb{H}(\mathscr{M}^{t-1}) - \mathbb{H}(\mathscr{M}^t)$ as long as any contractual commitment in u' is acceptable to α. The idea is that α uses the negotiation target as a guide to go above or below an equitable information revelation response.

The negotiation literature consistently advises that an agent's behaviour should not be predictable even in close, intimate relationships. This variation of behaviour is normally described as varying the negotiation *stance* that informally varies from "friendly guy" to "tough guy". The stance injects bounded random noise into the process, where the bound tightens as intimacy increases. For software agents, the role of stance is to prevent an observer from decrypting an agent's strategies.

7 Discussion

The prospect of automating the negotiation process in electronic business is a powerful motivation for research into robust negotiation strategies. A considerable effort is being made by game theorists to build strategies on a utilitarian basis. The work described in this paper is concerned with aspects of negotiation that are difficult, if not impossible, to capture within the utilitarian framework. Specifically the work is concerned with building relationships with the intention that they will provide agents with some degree of protection against the exploitation of information asymmetries in the marketplace. The strategic use of opinions and advice as argumentative illocutionary particles is one step on a long road to build reliable agents for business negotiation.

References

1. FIPA Communicative Act Library Specification. Tech. Rep. SC00037J, Foundation for Intelligent Physical Agents, Geneva, Switzerland (2002)
2. Adams, J.S.: Inequity in social exchange. In: L. Berkowitz (ed.) Advances in experimental social psychology, vol. 2. New York: Academic Press (1965)
3. Artz, D., Gil, Y.: A survey of trust in computer science and the semantic web. Web Semantics: Science, Services and Agents on the World Wide Web **5**(2), 58–71 (2007)
4. Bazerman, M.H., Loewenstein, G.F., White, S.B.: Reversal of preference in allocation decisions: judging an alternative versus choosing among alternatives. Administration Science Quarterly (37), 220–240 (1992)
5. Cheeseman, P., Stutz, J.: Bayesian Inference and Maximum Entropy Methods in Science and Engineering, chap. On The Relationship between Bayesian and Maximum Entropy Inference, pp. 445 – 461. American Institute of Physics, Melville, NY, USA (2004)
6. Faratin, P., Sierra, C., Jennings, N.: Using similarity criteria to make issue trade-offs in automated negotiation. Journal of Artificial Intelligence **142**(2), 205–237 (2003)
7. Grünwald, P.D.: The Minimum Description Length Principle. MIT Press, Cambridge, MA (2007)
8. Kalfoglou, Y., Schorlemmer, M.: IF-Map: An ontology-mapping method based on information-flow theory. In: S. Spaccapietra, S. March, K. Aberer (eds.) Journal on Data Semantics I, *Lecture Notes in Computer Science*, vol. 2800, pp. 98–127. Springer-Verlag: Heidelberg, Germany (2003)
9. Li, Y., Bandar, Z.A., McLean, D.: An approach for measuring semantic similarity between words using multiple information sources. IEEE Transactions on Knowledge and Data Engineering **15**(4), 871 – 882 (2003)
10. Paris, J.: Common sense and maximum entropy. Synthese **117**(1), 75 – 93 (1999)
11. Rahwan, I., Ramchurn, S., Jennings, N., McBurney, P., Parsons, S., Sonenberg, E.: Argumentation-based negotiation. Knowledge Engineering Review **18**(4), 343–375 (2003)
12. Rauyruena, P., Miller, K.E.: Relationship quality as a predictor of B2B customer loyalty. Journal of Business Research **60**(1), 21–31 (2007)
13. Sabater, J., Sierra, C.: Review on computational trust and reputation models. Artificial Intelligence Review **24**(1), 33–60 (2005)
14. Sierra, C., Debenham, J.: Trust and honour in information-based agency. In: P. Stone, G. Weiss (eds.) Proceedings Fifth International Conference on Autonomous Agents and Multi Agent Systems AAMAS-2006, pp. 1225 – 1232. ACM Press, New York, Hakodate, Japan (2006)
15. Sierra, C., Debenham, J.: Information-based agency. In: Proceedings of Twentieth International Joint Conference on Artificial Intelligence IJCAI-07, pp. 1513–1518. Hyderabad, India (2007)
16. Sierra, C., Debenham, J.: The LOGIC Negotiation Model. In: Proceedings Sixth International Conference on Autonomous Agents and Multi Agent Systems AAMAS-2007, pp. 1026–1033. Honolulu, Hawai'i (2007)
17. Sierra, C., Jennings, N., Noriega, P., Parsons, S.: Proceedings of the 4th International Workshop on Intelligent Agents IV, Agent Theories, Architectures, and Languages, chap. A Framework for Argumentation-Based Negotiation, pp. 177 – 192. Springer-Verlag, London, UK (1997)
18. Suzuki, J.: Learning bayesian belief networks based on the MDL principle: An efficient algorithm using the branch and bound technique. IEICE TRANSACTIONS on Information and Systems **E81-D**(12), 356–367 (1998)
19. Ulaga, W., Eggert, A.: Relationship value in business markets: The construct and its dimensions. Journal of Business To Business Marketing **12**(1), 73 – 99 (2005)