

An Agent Model of Business Relationships

John Debenham¹ and Carles Sierra²

¹ QCIS, University of Technology, Sydney, Australia
debenham@it.uts.edu.au

² Institut d'Investigació en Intel·ligència Artificial - IIIA,
Spanish Scientific Research Council, CSIC
08193 Bellaterra, Catalonia, Spain
sierra@iia.csic.es

Abstract. Relationships are fundamental to all but the most impersonal forms of interaction in business. An agent aims to secure projected needs by attempting to build a set of (business) relationships with other agents. A relationship is built by exchanging private information, and is characterised by its intimacy — degree of closeness — and balance — degree of fairness. Each argumentative interaction between two agents then has two goals: to satisfy some immediate need, and to do so in a way that develops the relationship in a desired direction. An agent's desire to develop each relationship in a particular way then places constraints on the argumentative utterances. This paper describes argumentative interaction constrained by a desire to develop such relationships.

1 Introduction

Modelling long-term (business) relationships underpins the evolution of trust relationships. A basis for agent interaction is presented that manages the (business) relationships that an agent has with each agent that it interacts with. Our agent summarises its relationships using 'intimacy' and 'balance' measures. Its actions are then shaped by its desired values for these two measures that represent its foreseeable social aspirations, and are called the 'target intimacy' and 'target balance'. Given all of this, a particular interaction with another agent is approached both with the goal of negotiating towards a satisfactory conclusion, and as an opportunity to do so in a way that gradually develops the relationship towards its target. In this way the agent's target aspirations constrain and shape its argumentative behaviour in *relationship-based argumentation*.

Negotiation dialogues are traditionally organised around the basic illocutionary particles: *Offer*, *Accept* and *Reject*. Previous work has been centred on the design of negotiation strategies and on proposing agent architectures able to deal with the exchange of offers [1,2]. Game theory [3], possibilistic logic [4] and first-order logic [5] have been used for this purpose. Some initial steps in proposing rhetoric particles have been made, especially around the idea of *appeals*, *rewards* and *threats* [6]. Expanded negotiation dialogues, including these and other rhetoric moves, are known as *argumentation-based negotiations*. *Argumentation* in this sense is mainly to do with building (business) *relationships*. When we reward or threaten we refer to a future instant of time where the reward or threat will be effective, its scope goes beyond the current negotiation round.

This paper is in the area labelled: *information-based agency* [7]. An information-based agent has an identity, values, needs, plans and strategies all of which are expressed using a fixed ontology in probabilistic logic for internal representation and in an illocutionary language for communication. We assume that such an agent resides in an electronic institution [8] and is aware of the prevailing norms and interaction protocols. An information-based agent makes no *a priori* assumptions about the states of the world or the other agents in it — these are represented in a world model, \mathcal{M}^t , that is inferred solely from the messages that it receives. The intuition behind information-based agency is that all illocutionary acts give away (valuable) information.

An agent's world model, \mathcal{M}^t , is a set of probability distributions for a set of random variables each of which represents the agent's expectations about some point of interest about the world or the other agents in it. Each incoming utterance is translated into a set of (linear) constraints on one or more of these distributions, and then the posterior state of the world model is estimated using entropy-based inference. These distributions are the foundation for the agent's reasoning.

A pair of agents interact by passing messages to each other. We assume that they share a common ontology and that their interactions are organised into dialogues, where a *dialogue* is a finite sequence of inter-related utterances. A *commitment* is a consequence of an utterance by an agent that contains a promise that the world will be in some state in the future. A *contract* is a pair of commitments exchanged between a pair¹ of agents. The set of all dialogues between two agents up to the present is their *relationship*. This discussion is from the point of view of an information-based agent α in a multiagent system where α interacts with negotiating agents, β_i , information providing agents, θ_j , and an *institutional agent*, ξ , that represents the institution where we assume that all interactions happen [8]: $\{\alpha, \beta_1, \dots, \beta_o, \xi, \theta_1, \dots, \theta_t\}$.

Our communication model is described in Section 2. Relationships are formalised in Section 3, and the agent architecture in Section 4. Section 6 describes an elaborate means of measuring the intimacy — degree of closeness — and balance – degree of fairness — that are based on measures of the information in any utterance. Section 7 describes the argumentation framework, and Section 8 concludes.

2 Communication Model

The communication language is detailed below; we assume that utterances in the communication language may be classified into unique illocutionary categories² $\mathcal{L} = \{l_i\}_{i=1}^L$. In order to structure agent dialogues we also need 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)).³ We model ontologies following an algebraic approach [9] and an ontology is a tuple $\mathcal{O} = (C, R, \leq, \sigma)$ where:

¹ Sets of commitments between more than two agents are not considered here.

² In a simple bargaining scenario these utterances could be: “propose”, “accept” and “reject”.

³ Usually, a set of axioms defined over the concepts and relations is also required. We will omit this here.

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 the traditional *is-a* hierarchy. To simplify computations in the computing of probability distributions we assume that there is a number of disjoint *is-a* trees covering different ontological spaces (e.g. a tree for types of fabric, a tree for shapes of clothing, and so on). R contains relations between the concepts in the hierarchy, this is needed to define ‘objects’ (e.g. deals) that are defined as a tuple of issues. We then analyse dialogues in terms of the *dialogical framework* $\mathcal{L} \times \mathcal{O}$.

The semantic distance between concepts within an ontology depends 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. How signed contracts, $Commit(\cdot)$, about objects in a particular semantic region, and their execution, $Done(\cdot)$, affect our decision making process about signing future contracts in nearby semantic regions is crucial to modelling the common sense that human beings apply in managing trading relationships. A measure [10] 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 is then defined as:

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

where $e = 2.71828$, 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 contributions of the shortest path length and the depth respectively.

The shape of the language that α uses to represent the information received and the content of its dialogues depends on two fundamental notions. First, when agents interact within an overarching institution they explicitly or implicitly accept the *norms* that will constrain their behaviour, and accept the established sanctions and penalties whenever norms are violated. Second, the dialogues in which α engages are built around two fundamental actions: (i) passing information, and (ii) exchanging proposals and contracts. A contract $\delta = (a, b)$ between agents α and β is a pair where a and b represent the actions that agents α and β are responsible for respectively. *Contracts* signed by agents and *information* passed by agents, are similar to norms in the sense that they oblige agents to behave in a particular way, so as to satisfy the conditions of the contract, or to make the world consistent with the information passed. Contracts and Information can thus be thought of as normative statements that restrict an agent’s behaviour.

Norms, contracts, and information have an obvious temporal dimension. Thus, an agent has to abide by a norm while it is inside an institution, a contract has a validity period, and a piece of information is true only during an interval in time. The set of norms affecting the behaviour of an agent defines the *context* that the agent has to take into account.

α 's communication language has two fundamental primitives: $\text{Commit}(\alpha, \beta, \varphi)$ to represent, in φ , the world that α aims at bringing about and that β has the right to verify, complain about or claim compensation for any deviations from, and $\text{Done}(\mu)$ to represent the event that a certain action μ^4 has taken place. In this way, norms, contracts, and information chunks will be represented as instances of $\text{Commit}(\cdot)$ where α and β can be individual agents or institutions. \mathcal{C} is:

$$\begin{aligned} \mu &::= \text{illoc}(\alpha, \beta, \varphi, t) \mid \mu; \mu \mid \\ &\quad \mathbf{Let\ context\ In\ } \mu \mathbf{\ End} \\ \varphi &::= \text{term} \mid \text{Done}(\mu) \mid \text{Commit}(\alpha, \beta, \varphi) \mid \varphi \wedge \varphi \mid \\ &\quad \varphi \vee \varphi \mid \neg\varphi \mid \forall v. \varphi_v \mid \exists v. \varphi_v \\ \text{context} &::= \varphi \mid \text{id} = \varphi \mid \text{prolog_clause} \mid \text{context}; \text{context} \end{aligned}$$

where φ_v is a formula with free variable v , *illoc* is any appropriate set of illocutionary particles, ‘;’ means sequencing, and *context* represents either previous agreements, previous illocutions, the ontological working context, that is a projection of the ontological trees that represent the focus of the conversation, or code that aligns the ontological differences between the speakers needed to interpret an action a . Representing an ontology as a set predicates in Prolog is simple. The set *term* contains instances of the ontology concepts and relations.⁵

For example, we can represent the following offer: “If you spend a total of more than €100 in my shop during October then I will give you a 10% discount on all goods in November”, as:

$$\begin{aligned} \text{Offer}(\alpha, \beta, \text{spent}(\beta, \alpha, \text{October}, X) \wedge X \geq \text{€}100 \rightarrow \\ \forall y. \text{Done}(\text{Inform}(\xi, \alpha, \text{pay}(\beta, \alpha, y), \text{November})) \rightarrow \\ \text{Commit}(\alpha, \beta, \text{discount}(y, 10\%))) \end{aligned}$$

ξ is an institution agent that reports the payment.

3 Relationships

There is evidence from psychological studies that humans seek a *balance* in their negotiation relationships. The classical view [11] 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). However, more recent studies [12,13] show that humans follow a richer set of norms of distributive justice depending on their *intimacy* level: *equity*, *equality*, and *need*. *Equity* being the allocation proportional to the effort (e.g. the profit of a company goes to the stock holders proportional to their investment), *equality* being the allocation in equal amounts (e.g. two friends eat the same amount of a cake cooked by one of them), and *need* being the allocation proportional to the need for the

⁴ Without loss of generality we will assume that all actions are dialogical.

⁵ We assume the convention that $C(c)$ means that c is an instance of concept C and $r(c_1, \dots, c_n)$ implicitly determines that c_i is an instance of the concept in the i -th position of the relation r .

resource (e.g. in case of food scarcity, a mother gives all food to her baby). For instance, if we are in a purely economic setting (low intimacy) we might request equity for the Options dimension but could accept equality in the Goals dimension.

The perception of a relation being in balance (i.e. fair) depends strongly on the nature of the social relationships between individuals (i.e. the intimacy level). In purely economical relationships (e.g., business), equity is perceived as more fair; in relations where joint action or fostering of social relationships are the goal (e.g. friends), equality is perceived as more fair; and in situations where personal development or personal welfare are the goal (e.g. family), allocations are usually based on need.

We believe that the perception of balance in dialogues (in negotiation or otherwise) is grounded on social relationships, and that every dimension of an interaction between humans can be correlated to the social closeness, or *intimacy*, between the parties involved. According to the previous studies, the more intimacy across the illocutionary categories the more the need norm is used, and the less intimacy the more the equity norm is used. This might be part of our social evolution. There is ample evidence that when human societies evolved from a hunter-gatherer structure⁶ to a shelter-based one⁷ the probability of survival increased when food was scarce.

In this context, we can clearly see that, for instance, families exchange not only goods but also information and knowledge based on need, and that few families would consider their relationships as being unbalanced, and thus unfair, when there is a strong asymmetry in the exchanges (a mother explaining everything to her children, or buying toys, does not expect reciprocity). In the case of partners there is some evidence [14] 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. See Table 1 for some examples of desired balances along five illocutionary categories.

Table 1. Some desired balances (sense of fairness) for five illocutionary categories

Illoc. Category	A new trading partner	my butcher	my boss	my partner	my children
Legitimacy	equity	equity	equity	equality	need
Options	equity	equity	equity	mixed ^a	need
Goals	equity	need	equity	need	need
Independence	equity	equity	equality	need	need
Commitment	equity	equity	equity	mixed	need

^a equity on burden, equality on good

The perceived balance in a negotiation dialogue allows negotiators to infer information about their opponent, about its stance, and to compare their relationships with all negotiators. For instance, if we perceive that every time we request information it is provided, and that no significant questions are returned, or no complaints about not

⁶ In its purest form, individuals in these societies collect food and consume it when and where it is found. This is a pure equity sharing of the resources, the gain is proportional to the effort.

⁷ In these societies there are family units, around a shelter, that represent the basic food sharing structure. Usually, food is accumulated at the shelter for future use. Then the food intake depends more on the need of the members.

receiving information are given, then that probably means that our opponent perceives our social relationship to be very close. Alternatively, we can detect what issues are causing a burden to our opponent by observing an imbalance in the information or utilitarian senses on that issue.

4 Agent Architecture

A multiagent system $\{\alpha, \beta_1, \dots, \beta_n, \xi, \theta_1, \dots, \theta_i\}$, contains an agent α that interacts with other *argumentation agents*, β_i , *information providing agents*, θ_j , and an *institutional agent*, ξ , that represents the institution where we assume the interactions happen [8]. The institutional agent reports promptly and honestly on what actually occurs after an agent signs a contract, or makes some other form of commitment. In Section 5 this enables us to measure the difference between an utterance and a subsequent observation. Agents have a probabilistic first-order *internal language* \mathcal{L} used to represent a *world model*, \mathcal{M}^t . A generic *information-based* architecture is described in detail in [7].

The agent architecture is shown in Figure 1. Agent α acts in response to a *need* that is expressed in terms of the ontology. Needs trigger α 's goal/plan proactive reasoning, while other messages are dealt with by α 's reactive reasoning.⁸ Each plan prepares for the negotiation by assembling the contents of a 'briefcase' that the agent 'carries' into the negotiation⁹. The *relationship strategy* determines which agent to negotiate with for a given need; it uses risk management analysis to preserve a strategic set of trading relationships for each mission-critical need — this is not detailed here. For each trading relationship this strategy generates a *relationship target* that is expressed in the dialogical framework as a desired level of *intimacy* to be achieved in the long term.

Each negotiation consists of a dialogue, Ψ^t , between two agents with agent α contributing utterance μ and the partner β contributing μ' . Each dialogue, Ψ^t , is evaluated using the dialogical framework in terms of the *value* of Ψ^t to both α and β — see Section 6.2. The *negotiation strategy* then determines the current set of offers $\{\delta_i\}$, and then the *tactics*, guided by the *negotiation target*, decide which, if any, of these offers to put forward and wraps them in argumentation dialogue — see Section 7. We now describe two of the distributions in \mathcal{M}^t that support offer exchange.

$\mathbb{P}^t(\text{acc}(\alpha, \beta, \chi, \delta))$ estimates the probability that α should accept proposal δ in satisfaction of her need χ , where $\delta = (a, b)$ is a pair of commitments, a for α and b for β . α will accept δ if: $\mathbb{P}^t(\text{acc}(\alpha, \beta, \chi, \delta)) > c$, for level of certainty c . This estimate is compounded from subjective and objective views of acceptability. The *subjective estimate* takes account of: the extent to which the enactment of δ will satisfy α 's need χ , how much δ is 'worth' to α , and the extent to which α believes that she will be in a position to execute her commitment a [7]. $S_\alpha(\beta, a)$ is a random variable denoting α 's estimate of β 's subjective valuation of a over some finite, numerical evaluation space.

⁸ Each of α 's plans and reactions contain constructors for an initial *world model* \mathcal{M}^t . \mathcal{M}^t is then maintained from percepts received using *update functions* that transform percepts into constraints on \mathcal{M}^t — for details, see [7].

⁹ Empirical evidence shows that in human negotiation, better outcomes are achieved by skewing the opening offer in favour of the proposer. We are unaware of any empirical investigation of this hypothesis for autonomous agents in real trading scenarios.




Fig. 1. The agent architecture

The *objective estimate* captures whether δ is acceptable on the open market, and variable $U_\alpha(b)$ denotes α 's open-market valuation of the enactment of commitment b , again taken over some finite numerical valuation space. We also consider needs, the variable $T_\alpha(\beta, a)$ denotes α 's estimate of the *strength* of β 's motivating *need* for the enactment of commitment a over a valuation space. Then for $\delta = (a, b)$: $\mathbb{P}^t(\text{acc}(\alpha, \beta, \chi, \delta)) =$

$$\mathbb{P}^t \left(\left(\frac{T_\alpha(\beta, a)}{T_\alpha(\alpha, b)} \right)^h \times \left(\frac{S_\alpha(\alpha, b)}{S_\alpha(\beta, a)} \right)^g \times \frac{U_\alpha(b)}{U_\alpha(a)} \geq s \right) \quad (1)$$

where $g \in [0, 1]$ is α 's *greed*, $h \in [0, 1]$ is α 's degree of *altruism*, and $s \approx 1$ is derived from the *stance*¹⁰ described in Section 7. The parameters g and h are independent. We can imagine a relationship that begins with $g = 1$ and $h = 0$. Then as the agents share increasing amounts of their information about their open market valuations g gradually reduces to 0, and then as they share increasing amounts of information about their needs h increases to 1. The basis for the acceptance criterion has thus developed from equity to equality, and then to need.

$\mathbb{P}^t(\text{acc}(\beta, \alpha, \delta))$ estimates the probability that β would accept δ , by observing β 's responses. For example, if β sends the message Offer(δ_1) then α derives the constraint: $\{\mathbb{P}^t(\text{acc}(\beta, \alpha, \delta_1)) = 1\}$ on the distribution $\mathbb{P}^t(\beta, \alpha, \delta)$, and if this is a counter offer to a former offer of α 's, δ_0 , then: $\{\mathbb{P}^t(\text{acc}(\beta, \alpha, \delta_0)) = 0\}$. 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 δ will be acceptable to β by enumerating the possible worlds that represent β 's "limit of acceptability" [16].

¹⁰ If α chooses to inflate her opening offer then this is achieved in Section 7 by increasing the value of s . If $s \gg 1$ then a deal may not be possible. This illustrates the well-known inefficiency of bilateral bargaining established analytically by [15].

5 Updating the World Model \mathcal{M}^t

α 's world model consists of probability distributions that represent its uncertainty in the world state. α is interested in the degree to which an utterance accurately describes what will subsequently be observed. All observations about the world are received as utterances from an all-truthful institution agent ξ . For example, if β communicates the goal "I am hungry" and the subsequent negotiation terminates with β purchasing a book from α (by ξ advising α that a certain amount of money has been credited to α 's account) then α may conclude that the goal that β chose to satisfy was something other than hunger. So, α 's world model contains probability distributions that represent its uncertain expectations of what will be observed on the basis of utterances received.

We represent the relationship between *utterance*, φ , and subsequent *observation*, φ' , by $\mathbb{P}^t(\varphi'|\varphi) \in \mathcal{M}^t$, where φ' and φ may be ontological categories in the interest of computational feasibility. For example, if φ is "I will deliver a bucket of fish to you tomorrow" then the distribution $\mathbb{P}(\varphi'|\varphi)$ need not be over *all* possible things that β might do, but could be over ontological categories that summarise β 's possible actions.

In the absence of in-coming utterances, the conditional probabilities, $\mathbb{P}^t(\varphi'|\varphi)$, should tend to ignorance as represented by a *decay limit distribution* $\mathbb{D}(\varphi'|\varphi)$. α may have background knowledge concerning $\mathbb{D}(\varphi'|\varphi)$ as $t \rightarrow \infty$, otherwise α may assume that it has maximum entropy whilst being consistent with the data. In general, given a distribution, $\mathbb{P}^t(X_i)$, and a decay limit distribution $\mathbb{D}(X_i)$, $\mathbb{P}^t(X_i)$ decays by:

$$\mathbb{P}^{t+1}(X_i) = \Delta_i(\mathbb{D}(X_i), \mathbb{P}^t(X_i)) \quad (2)$$

where Δ_i is the *decay function* for the X_i satisfying the property that $\lim_{t \rightarrow \infty} \mathbb{P}^t(X_i) = \mathbb{D}(X_i)$. For example, Δ_i could be linear: $\mathbb{P}^{t+1}(X_i) = (1 - \nu_i) \times \mathbb{D}(X_i) + \nu_i \times \mathbb{P}^t(X_i)$, where $\nu_i < 1$ 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 $\mu = \text{illoc}(\alpha, \beta, \varphi, t)$ from agent β at time t . Suppose that α attaches an epistemic belief $\mathbb{R}^t(\alpha, \beta, \mu)$ to μ — this probability takes account of α 's level of personal *caution*. We model the update of $\mathbb{P}^t(\varphi'|\varphi)$ in two cases, one for observations given φ , second for observations given ϕ in the semantic neighbourhood of φ .

First, if $\{\varphi_1, \varphi_2, \dots, \varphi_m\}$ is the set of all possible observations and φ_k is observed then α may set $\mathbb{P}^{t+1}(\varphi_k|\varphi)$ to some value d . We estimate the complete posterior distribution $\mathbb{P}^{t+1}(\varphi'|\varphi)$ by applying the principle of minimum relative entropy as follows. Let $\mathbf{p}_{(\mu)}$ be the distribution: $\arg \min_{\mathbf{x}} \sum_j x_j \log \frac{x_j}{\mathbb{P}^t(\varphi'|\varphi)_j}$ that satisfies the constraint $\mathbf{p}_{(\mu)_k} = d$. Then let $\mathbf{q}_{(\mu)}$ be the distribution:

$$\mathbf{q}_{(\mu)} = \mathbb{R}^t(\alpha, \beta, \mu) \times \mathbf{p}_{(\mu)} + (1 - \mathbb{R}^t(\alpha, \beta, \mu)) \times \mathbb{P}^t(\varphi'|\varphi)$$

and then let:

$$\mathbf{r}_{(\mu)} = \begin{cases} \mathbf{q}_{(\mu)} & \text{if } \mathbf{q}_{(\mu)} \text{ is more interesting than } \mathbb{P}^t(\varphi'|\varphi) \\ \mathbb{P}^t(\varphi'|\varphi) & \text{otherwise} \end{cases}$$

A measure of whether $\mathbf{q}_{(\mu)}$ is more interesting than $\mathbb{P}^t(\varphi'|\varphi)$ is: $\mathbb{K}(\mathbf{q}_{(\mu)} \parallel \mathbb{D}(\varphi'|\varphi)) > \mathbb{K}(\mathbb{P}^t(\varphi'|\varphi) \parallel \mathbb{D}(\varphi'|\varphi))$, where $\mathbb{K}(\mathbf{x} \parallel \mathbf{y}) = \sum_j x_j \ln \frac{x_j}{y_j}$ is the Kullback-Leibler distance between two probability distributions \mathbf{x} and \mathbf{y} .

Finally incorporating Equation 2 we obtain the method for updating a distribution $\mathbb{P}^t(\varphi'|\varphi)$ on receipt of a message μ :

$$\mathbb{P}^{t+1}(\varphi'|\varphi) = \Delta_i(\mathbb{D}(\varphi'|\varphi), \mathbf{r}_{(\mu)}) \quad (3)$$

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

Second we consider the update of $\mathbb{P}^t(\phi'|\phi)$ given φ . Given $\mu = illoc(\alpha, \beta, \varphi, t)$ and the observation φ_k we define the vector \mathbf{t} by

$$t_i = \mathbb{P}^t(\phi_i|\phi) + (1 - |\text{Sim}(\varphi_k, \varphi) - \text{Sim}(\phi_i, \phi)|) \cdot \text{Sim}(\varphi_k, \phi)$$

with $\{\phi_1, \phi_2, \dots, \phi_p\}$ the set of all possible observations in the context of ϕ and $i = 1, \dots, p$. \mathbf{t} is not a probability distribution. The multiplying factor $\text{Sim}(\varphi', \phi)$ limits the variation of probability to those formulae whose ontological context is not too far away from the observation. The posterior $\mathbb{P}^{t+1}(\phi'|\phi)$ is obtained with Equation 3 with $\mathbf{r}_{(\mu)}$ defined to be the normalisation of \mathbf{t} .

6 Measuring the Confidence in a Relationship

A *dialogue*, Ψ^t , between agents α and β is a sequence of inter-related utterances in context. A *relationship*, Ψ^{*t} , is a sequence of dialogues. We first measure the *confidence* that an agent has for another by observing, for each utterance, the difference between what is said (the utterance) and what subsequently occurs (the observation). Second we *evaluate* each dialogue as it progresses in terms of the dialogical framework — this evaluation employs the confidence measures. Finally we define the *intimacy* of a relationship as an aggregation of the value of its component dialogues.

6.1 Confidence

Confidence measures generalise what are commonly called *trust*, *reliability* and *reputation* measures [17] into a single computational framework that spans the illocutionary categories \mathcal{C} . In Section 6.2 confidence measures are applied to valuing fulfilment of promises, to the execution of commitments, and to valuing dialogues.

Ideal observations. Consider a distribution of observations that represent α 's “ideal” in the sense that it is the best that α could reasonably expect to observe. This distribution will be a function of α 's *context* with β denoted by e , and is $\mathbb{P}_I^t(\varphi'|\varphi, e)$. Here we measure the relative entropy between this ideal distribution, $\mathbb{P}_I^t(\varphi'|\varphi, e)$, and the distribution of expected observations, $\mathbb{P}^t(\varphi'|\varphi)$. That is:

$$\mathbb{C}(\alpha, \beta, \varphi) = 1 - \sum_{\varphi'} \mathbb{P}^t(\varphi'|\varphi, e) \log \frac{\mathbb{P}_I^t(\varphi'|\varphi, e)}{\mathbb{P}^t(\varphi'|\varphi)} \quad (4)$$

where the “1” is an arbitrarily chosen constant being the maximum value that this measure may have. This equation measures confidence for a single statement φ . It makes

sense to aggregate these values over a class of statements, say over those φ that are in the ontological context o , that is $\varphi \leq o$:

$$\mathbb{C}(\alpha, \beta, o) = 1 - \frac{\sum_{\varphi: \varphi \leq o} \mathbb{P}_{\beta}^t(\varphi) [1 - \mathbb{C}(\alpha, \beta, \varphi)]}{\sum_{\varphi: \varphi \leq o} \mathbb{P}_{\beta}^t(\varphi)}$$

where $\mathbb{P}_{\beta}^t(\varphi)$ is a probability distribution over the space of statements that the next statement β will make to α is φ . Similarly, for an overall estimate of β 's *confidence* in α :

$$\mathbb{C}(\alpha, \beta) = 1 - \sum_{\varphi} \mathbb{P}_{\beta}^t(\varphi) [1 - \mathbb{C}(\alpha, \beta, \varphi)]$$

Preferred observations. The previous measure requires that: $\mathbb{P}_I^t(\varphi'|\varphi, e)$, has to be specified for each φ . Here we measure the extent to which the observation φ' is preferable to the original statement φ . Given a predicate $\text{Prefer}(c_1, c_2, e)$ meaning that α prefers c_1 to c_2 in environment e . Then if $\varphi \leq o$:

$$\mathbb{C}(\alpha, \beta, \varphi) = \sum_{\varphi'} \mathbb{P}^t(\text{Prefer}(\varphi', \varphi, o)) \mathbb{P}^t(\varphi'|\varphi)$$

and:

$$\mathbb{C}(\alpha, \beta, o) = \frac{\sum_{\varphi: \varphi \leq o} \mathbb{P}_{\beta}^t(\varphi) \mathbb{C}(\alpha, \beta, \varphi)}{\sum_{\varphi: \varphi \leq o} \mathbb{P}_{\beta}^t(\varphi)}$$

Certainty in observation. Here we measure the consistency in expected acceptable observations, or “the lack of expected uncertainty in those possible observations that are better than the original statement”.

If $\varphi \leq o$ let: $\Phi_+(\varphi, o, \kappa) = \{\varphi' \mid \mathbb{P}^t(\text{Prefer}(\varphi', \varphi, o)) > \kappa\}$ for some constant κ , and:

$$\mathbb{C}(\alpha, \beta, \varphi) = 1 + \frac{1}{B^*} \cdot \sum_{\varphi' \in \Phi_+(\varphi, o, \kappa)} \mathbb{P}_+^t(\varphi'|\varphi) \log \mathbb{P}_+^t(\varphi'|\varphi)$$

where $\mathbb{P}_+^t(\varphi'|\varphi)$ is the normalisation of $\mathbb{P}^t(\varphi'|\varphi)$ for $\varphi' \in \Phi_+(\varphi, o, \kappa)$,

$$B^* = \begin{cases} 1 & \text{if } |\Phi_+(\varphi, o, \kappa)| = 1 \\ \log |\Phi_+(\varphi, o, \kappa)| & \text{otherwise} \end{cases}$$

As above we aggregate this measure for observations in a particular context o , and measure confidence as before.

Computational Note. The various measures given above involve extensive calculations. For example, Equation 4 contains $\sum_{\varphi'}$ that sums over *all* possible observations φ' . We obtain a more computationally friendly measure by appealing to the structure of the ontology described, and the right-hand side of Equation 4 may be approximated to:

$$1 - \sum_{\varphi': \text{Sim}(\varphi', \varphi) \geq \eta} \mathbb{P}_{\eta, I}^t(\varphi'|\varphi, e) \log \frac{\mathbb{P}_{\eta, I}^t(\varphi'|\varphi, e)}{\mathbb{P}_{\eta}^t(\varphi'|\varphi)}$$

where $\mathbb{P}_{\eta, I}^t(\varphi'|\varphi, e)$ is the normalisation of $\mathbb{P}_I^t(\varphi'|\varphi, e)$ for $\text{Sim}(\varphi', \varphi) \geq \eta$, and similarly for $\mathbb{P}_{\eta}^t(\varphi'|\varphi)$. The extent of this calculation is controlled by the parameter η . An even tighter restriction may be obtained with: $\text{Sim}(\varphi', \varphi) \geq \eta$ and $\varphi' \leq \psi$ for some ψ .

6.2 Valuing Negotiation Dialogues

Suppose that a negotiation commences at time s , and by time t a string of utterances, $\Phi^t = \langle \mu_1, \dots, \mu_n \rangle$ has been exchanged between agent α and agent β . This negotiation dialogue is evaluated by α in the context of α 's world model at time s , \mathcal{M}^s , and the environment e that includes utterances that may have been received from other agents in the system including the information sources $\{\theta_i\}$. Let $\Psi^t = (\Phi^t, \mathcal{M}^s, e)$, then α estimates the *value* of this dialogue to itself in the context of \mathcal{M}^s and e as a $2 \times L$ array $V_\alpha(\Psi^t)$ where:

$$V_x(\Psi^t) = \begin{pmatrix} I_x^{l_1}(\Psi^t) & \dots & I_x^{l_L}(\Psi^t) \\ U_x^{l_1}(\Psi^t) & \dots & U_x^{l_L}(\Psi^t) \end{pmatrix}$$

where the $I(\cdot)$ and $U(\cdot)$ functions are information-based and utility-based measures respectively as we now describe. α estimates the *value* of this dialogue to β as $V_\beta(\Psi^t)$ by assuming that β 's reasoning apparatus mirrors its own.

In general terms, the information-based valuations measure the reduction in uncertainty, or information gain, that the dialogue gives to each agent, they are expressed in terms of decrease in entropy that can always be calculated. The utility-based valuations measure utility gain are expressed in terms of ‘‘some suitable’’ utility evaluation function $\mathbb{U}(\cdot)$ that can be difficult to define. This is one reason why the utilitarian approach has no natural extension to the management of argumentation that is achieved here by our information-based approach.

The *balance* in a negotiation dialogue, Ψ^t , is defined as: $B_{\alpha\beta}(\Psi^t) = V_\alpha(\Psi^t) \ominus V_\beta(\Psi^t)$ for an element-by-element difference operator \ominus that respects the structure of $V(\Psi^t)$. The *intimacy* between agents α and β , $I_{\alpha\beta}^{*t}$, is the pattern of the two $2 \times L$ arrays V_α^{*t} and V_β^{*t} that are computed by an update function as each negotiation round terminates, $I_{\alpha\beta}^{*t} = \begin{pmatrix} V_\alpha^{*t} \\ V_\beta^{*t} \end{pmatrix}$. If Ψ^t terminates at time t :

$$V_x^{*t+1} = \nu \times V_x(\Psi^t) + (1 - \nu) \times V_x^{*t} \quad (5)$$

where ν is the learning rate, and $x = \alpha, \beta$. Additionally, V_x^{*t} continually decays by: $V_x^{*t+1} = \tau \times V_x^{*t} + (1 - \tau) \times D_x$, where $x = \alpha, \beta$; τ is the decay rate, and D_x is a $2 \times L$ array being the decay limit distribution for the value to agent x of the intimacy of the relationship in the absence of any interaction. D_x is the *reputation* of agent x . The *relationship balance* between agents α and β is: $B_{\alpha\beta}^{*t} = V_\alpha^{*t} \ominus V_\beta^{*t}$. In particular, the intimacy determines values for the parameters g and h in Equation 1. As a simple example, if both $I_\alpha^O(\Psi^{*t})$ and $I_\beta^O(\Psi^{*t})$ increase then g decreases, and as the remaining information-based components increase, h increases.

The notion of balance may be applied to pairs of utterances by treating them as degenerate dialogues. In simple multi-issue bargaining the *equitable information revelation* strategy generalises the tit-for-tat strategy in single-issue bargaining, and extends to a tit-for-tat argumentation strategy by applying the same principle across the dialogical framework.

7 Strategies and Tactics for Building Relationships

Each negotiation has to achieve two goals. First it may be intended to achieve some contractual outcome. Second it will aim to contribute to the growth, or decline, of the relationship intimacy.

We now describe in greater detail the contents of the “*Negotiation*” box in Figure 1. The negotiation literature consistently advises that an agent’s behaviour should not be predictable even in close, intimate relationships. The required variation of behaviour is normally described as varying the negotiation *stance* that informally varies from “friendly guy” to “tough guy”. The stance is shown in Figure 1, it injects bounded random noise into the process, where the bound tightens as intimacy increases. The stance, $S_{\alpha\beta}^t$, is a $2 \times L$ matrix of randomly chosen multipliers, each ≈ 1 , that perturbs α ’s actions. The value in the (x, y) position in the matrix, where $x = I, U$ and $y \in \mathcal{L}$, is chosen at random from $[\frac{1}{l(I_{\alpha\beta}^{*t}, x, y)}, l(I_{\alpha\beta}^{*t}, x, y)]$ where $l(I_{\alpha\beta}^{*t}, x, y)$ is the bound, and $I_{\alpha\beta}^{*t}$ is the intimacy.

The negotiation *strategy* is concerned with maintaining a working set of proposals. If the set of proposals is empty then α will quit the negotiation. α perturbs the acceptance machinery (see Section 4) by deriving s from the $S_{\alpha\beta}^t$ matrix. In line with the comment in Footnote 9, in the early stages of the negotiation α may decide to inflate her opening offer. This is achieved by increasing the value of s in Equation 1. The following strategy uses the machinery described in Section 4. Fix h, g, s and c , set the Proposals to the empty set, let $D_s^t = \{\delta \mid \mathbb{P}^t(\text{acc}(\alpha, \beta, \chi, \delta) > c)\}$, then:

- repeat the following as many times as desired: add $\delta = \arg \max_x \{\mathbb{P}^t(\text{acc}(\beta, \alpha, x)) \mid x \in D_s^t\}$ to Proposals, remove $\{y \in D_s^t \mid \text{Sim}(y, \delta) < k\}$ for some k from D_s^t

By using $\mathbb{P}^t(\text{acc}(\beta, \alpha, \delta))$ this strategy reacts to β ’s history of Propose and Reject utterances.

Negotiation *tactics* are concerned with selecting some offers and wrapping them in argumentation. Prior interactions with agent β will have produced an intimacy pattern expressed in the form of $(V_{\alpha}^{*t}, V_{\beta}^{*t})$. Suppose that the relationship target is $(T_{\alpha}^{*t}, T_{\beta}^{*t})$. Following from Equation 5, α will want to achieve a *negotiation target*, $N_{\beta}(\Psi^t)$ such that: $\nu \cdot N_{\beta}(\Psi^t) + (1 - \nu) \cdot V_{\beta}^{*t}$ is “a bit on the T_{β}^{*t} side of” V_{β}^{*t} :

$$N_{\beta}(\Psi^t) = \frac{\nu - \kappa}{\nu} V_{\beta}^{*t} \oplus \frac{\kappa}{\nu} T_{\beta}^{*t} \quad (6)$$

for small $\kappa \in [0, \nu]$ that represents α ’s desired *rate of development* for her relationship with β . $N_{\beta}(\Psi^t)$ is a $2 \times L$ matrix containing variations in the dialogical framework’s dimensions that α would like to reveal to β during Ψ^t (e.g. I’ll pass a bit more information on options than usual, I’ll be stronger in concessions on options, etc.). It is reasonable to expect β to progress towards her target at the same rate and $N_{\alpha}(\Psi^t)$ is calculated by replacing β by α in Equation 6. $N_{\alpha}(\Psi^t)$ is what α hopes to receive from β during Ψ^t . This gives a *negotiation balance target* of: $N_{\alpha}(\Psi^t) \ominus N_{\beta}(\Psi^t)$ that can be used as the foundation for reactive tactics by striving to maintain this balance across the dialogical framework. A cautious tactic could use the balance to bound the response μ to each

utterance μ' from β by the constraint: $V_\alpha(\mu') \ominus V_\beta(\mu) \approx S_{\alpha\beta}^t \otimes (N_\alpha(\Psi^t) \ominus N_\beta(\Psi^t))$, where \otimes is element-by-element matrix multiplication, and $S_{\alpha\beta}^t$ is the stance. A less neurotic tactic could attempt to achieve the target negotiation balance over the anticipated complete dialogue. If a balance bound requires negative information revelation in one dialogical framework category then α will contribute nothing to it, and will leave this to the natural decay to the reputation D as described above.

The following are a list of components that we have described that could be combined into an agent's negotiation strategy. These components all constrain the agent's actions. We assume that they are all soft constraints and that they operate together with a hard constraint $C^t(\alpha, \beta, x^t)$ on the message x^t that α may send to β at time t .

Information-based strategies. Every communication gives away information and so has the potential to contribute to the intimacy and balance of a relationship. Information-based strategies manage the information revelation process. Let $M_{\alpha\beta}^t$ be the set of time-stamped messages that α has sent to β , and $M_{\beta\alpha}^t$ likewise both at time t . \mathcal{M}^t is α 's world model at time t and consists of a set of probability distributions. x^t denotes a message received at time t . $\mathbb{I}^t(\alpha, \beta, x^t)$ is the information gain — measured as the reduction of the entropy of \mathcal{M}^t — observed by α after receiving message x^t . $\mathbb{I}^t(\beta, \alpha, x^t)$ is α 's estimate of β 's information gain after receiving message x^t from α .

The complete *information history* of both the observed and the estimated information gain, $G^t(\alpha, \beta)$, is:

$$G^t(\alpha, \beta) = \{(x^s, \mathbb{I}^s(\alpha, \beta, x^s)) \mid x^s \in M_{\beta\alpha}^t\} \cup \{(x^s, \mathbb{I}^s(\beta, \alpha, x^s)) \mid x^s \in M_{\alpha\beta}^t\}$$

respectively.

In [7] we described to the model that α constructs of β . In general α can not be expected to guess β 's world model, \mathcal{M}_β^t , unless α knows what β 's needs are — even then, α would only know \mathcal{M}_β^t with certainty if it knew what plans β had chosen. However, α always knows the private information that it has sent to β — for example, in Propose(\cdot) and Reject(\cdot) messages. Such private information could be used by β to estimate α 's probability of accepting a proposal: $\mathbb{P}_\beta^t(\text{acc}(\alpha, \beta, \chi', z))$, where χ' is the need that β believes α to have.

α 's information-based strategies constrain its actions, x^t , on the basis of $\mathbb{I}^t(\beta, \alpha, x^t)$ and its relation to $G^t(\alpha, \beta)$. For example, the strategy that gives β greatest expected information gain:

$$\arg \max_z \{ \mathbb{I}_\beta^s(\beta, \alpha, z) \mid C^t(\alpha, \beta, z) \}$$

More generally, for some function f :

$$\arg \max_z \{ f(\mathbb{I}_\beta^s(\beta, \alpha, z), G^t(\alpha, \beta)) \mid C^t(\alpha, \beta, z) \}$$

the idea being that the f 'optimises' in some sense the information gain taking account of the interaction history.

Ontology-based strategies. The structure of the ontology may be used to manage the information revelation process in particular strategic areas. For example, α may prefer to build a relationship with β in the context of the supply of particular goods only [9]. The structure of the ontology is provided by the $\text{Sim}(\cdot)$ function. Given two contracts δ and δ' containing concepts $\{o_1, \dots, o_i\}$ and $\{o'_1, \dots, o'_j\}$ respectively, the (non-symmetric) distance of δ' from δ is the vector

$$\Gamma(\delta, \delta') = (d_k : o''_k)_{k=1}^i$$

where $d_k = \min_x \{\text{Sim}(o_k, o'_x) \mid x = 1, \dots, j\}$, $o''_k = \sup(\arg \min_x \{\text{Sim}(o_k, x) \mid x = o'_1, \dots, o'_j\}, o_k)$ and the function $\sup(\cdot, \cdot)$ is the supremum of two concepts in the ontology. $\Gamma(\delta, \delta')$ quantifies how different δ' is to δ and enables α to “work around” or “move away from” a contract under consideration. In general for some function g ;

$$\arg \max_z \{ g(\Gamma(z, x^s)) \mid x^s \in M_{\alpha\beta}^t \cup M_{\beta\alpha}^t \wedge C^t(\alpha, \beta, z) \}$$

the idea being that the g ‘optimises’ in some sense the ontological relationship with the interaction history.

8 Discussion

The ability of agents to conduct business relies on their ability to build business relationships with each other. In this paper we have introduced a novel approach to negotiation that incorporates a rich model of relationships that is dimensioned by the structure of the ontology and a set of illocutionary categories. It is grounded on business and psychological studies and introduces the concepts of *intimacy* and *balance* as key elements in understanding what is a negotiation strategy and tactic. Relationships are strengthened by managing the agent’s dialogical moves. Each dialogical move produces a change in an array structure. The current balance and intimacy levels and the desired, or target, levels are then used by the tactics to determine what to say next. The architecture is simple and the implementation of the agents straightforward using tools from information theory.

We are currently exploring the use of this model as an extension of a currently widespread eProcurement software commercialised by a spin-off company of the laboratory of one of the authors. This tool has only a utilitarian modelling of the negotiation interactions and has motivated some criticisms from its users about the lack of modelling of long-lasting relationships that our model could solve.

References

1. Jennings, N., Faratin, P., Lomuscio, A., Parsons, S., Sierra, C., Wooldridge, M.: Automated negotiation: Prospects, methods and challenges. *International Journal of Group Decision and Negotiation* 10, 199–215 (2001)
2. Faratin, P., Sierra, C., Jennings, N.: Using similarity criteria to make issue trade-offs in automated negotiation. *Journal of Artificial Intelligence* 142, 205–237 (2003)
3. Rosenschein, J.S., Zlotkin, G.: *Rules of Encounter*. The MIT Press, Cambridge (1994)

4. Giménez, E., Godo, L., Rodríguez-Aguilar, J.A., Garcia, P.: Designing bidding strategies for trading agents in electronic auctions. In: Proceedings of the Third International Conference on Multi-Agent Systems (ICMAS 1998), pp. 136–143 (1998)
5. Kraus, S.: Negotiation and cooperation in multi-agent environments. *Artificial Intelligence* 94, 79–97 (1997)
6. Sierra, C., Jennings, N., Noriega, P., Parsons, S.: A Framework for Argumentation-Based Negotiation. In: Rao, A., Singh, M.P., Wooldridge, M.J. (eds.) *ATAL 1997*. LNCS, vol. 1365, pp. 177–192. Springer, Heidelberg (1998)
7. Sierra, C., Debenham, J.: Information-based agency. In: Proceedings of Twentieth International Joint Conference on Artificial Intelligence, IJCAI 2007, Hyderabad, India, pp. 1513–1518 (2007)
8. Arcos, J.L., Esteva, M., Noriega, P., Rodríguez, J.A., Sierra, C.: Environment engineering for multiagent systems. *Journal on Engineering Applications of Artificial Intelligence* 18 (2005)
9. Kalfoglou, Y., Schorlemmer, M.: IF-Map: An ontology-mapping method based on information-flow theory. In: Spaccapietra, S., March, S., Aberer, K. (eds.) *Journal on Data Semantics I*. LNCS, vol. 2800, pp. 98–127. Springer, Heidelberg (2003)
10. 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, 871–882 (2003)
11. Adams, J.S.: Inequity in social exchange. In: Berkowitz, L. (ed.) *Advances in experimental social psychology*, vol. 2. Academic Press, New York (1965)
12. Sondak, H., Neale, M.A., Pinkley, R.: The negotiated allocations of benefits and burdens: The impact of outcome valence, contribution, and relationship. *Organizational Behaviour and Human Decision Processes*, 249–260 (1995)
13. Valley, K.L., Neale, M.A., Mannix, E.A.: Friends, lovers, colleagues, strangers: The effects of relationships on the process and outcome of negotiations. In: Bies, R., Lewicki, R., Sheppard, B. (eds.) *Research in Negotiation in Organizations*, vol. 5, pp. 65–94. JAI Press (1995)
14. 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*, 220–240 (1992)
15. Myerson, R., Satterthwaite, M.: Efficient mechanisms for bilateral trading. *Journal of Economic Theory* 29, 1–21 (1983)
16. Debenham, J.: Bargaining with information. In: Jennings, N., Sierra, C., Sonenberg, L., Tambe, M. (eds.) *Proceedings of Third International Conference on Autonomous Agents and Multi Agent Systems, AAMAS-2004*, pp. 664–671. ACM Press, New York (2004)
17. Sierra, C., Debenham, J.: Information-based reputation. In: Paolucci, M. (ed.) *First International Conference on Reputation: Theory and Technology (ICORE 2009)*, Gargonza, Italy, pp. 5–19 (2009)