

© [2009] IEEE. Reprinted, with permission, from [John Debenham, Modelling Trust, Honour and Reliability in Business Relationships, 2009, Third IEEE International Conference on Digital Ecosystems and Technologies, 2009]. This material is posted here with permission of the IEEE. Such ermission of the IEEE does not in any way imply IEEE endorsement of any of the University of Technology, Sydney's products or services. Internal or personal use of this material is permitted. However, permission to reprint/republish this material for advertising or promotional purposes or for creating new collective works for resale or redistribution must be obtained from the IEEE by writing to pubs-permissions@ieee.org. By choosing to view this document, you agree to all provisions of the copyright laws protecting it

Modelling Trust, Honour and Reliability in Business Relationships

John Debenham
Faculty of Information Technology
University of Technology, Sydney
NSW 2007, Australia
Email: debenham@it.uts.edu.au

Abstract—Successful negotiators prepare by determining their position along five dimensions: Legitimacy, Options, Goals, Independence, and Commitment. We model business relationships in terms of these dimensions and two primitive concepts: *intimacy* (degree of closeness) and *balance* (degree of fairness). The *intimacy* is a pair of matrices that evaluate both an agent's contribution to the relationship and its opponent's contribution each from an information view and from a utilitarian view across the five dimensions. The *balance* is the difference between these matrices. A *relationship strategy* maintains a *target intimacy* for each relationship that an agent would like the relationship to move towards in future. The *negotiation strategy* maintains a set of Options that are in-line with the current intimacy level, and then *tactics* wrap the Options in argumentation with the aim of attaining a successful deal and manipulating the successive negotiation balances towards the target intimacy.

I. INTRODUCTION

In this paper we model long term relationships that are founded on successive negotiation encounters. The model is grounded on results from business and psychological studies [1], [12], [7], and acknowledges that negotiation is an information exchange process as well as a utility exchange process [11], [10]. We believe that if agents are to succeed in real application domains they have to reconcile both views: informational and game-theoretical.

Our aim is to model relationships between agents that represent human principals, and thus we want their behaviour to be comprehensible by humans and to respect usual human negotiation procedures, whilst being consistent with, and somehow extending, game theoretical and information theoretical results. In this sense, agents are not just utility maximisers, but aim at building long lasting relationships with progressing levels of *intimacy* that determine what *balance* in information and resource sharing is acceptable to them.

Negotiation strategies can naturally be seen as procedures that select tactics used to attain a successful deal *and* to reach a *target* intimacy level. It is common in human settings to use tactics that compensate for unbalances in one dimension of a negotiation with unbalances in another dimension. In this sense, humans aim at a *general sense of fairness* in an interaction.

In Section II we outline the aspects of human negotiation modelling that we cover in this work. Section III explains in outline the architecture and the concepts of intimacy and

balance, and how they influence the negotiation. Section IV contains a description of the different metrics used in the agent model including intimacy. Finally, Section V outlines how strategies and tactics use the relationship model, intimacy and balance.

II. HUMAN NEGOTIATION

Before a negotiation starts human negotiators prepare the dialogic exchanges that can be made along the five dimensions [6]:

- Legitimacy. What information is relevant to the negotiation process? What are the persuasive arguments about the fairness of the options?
- Options. What are the possible agreements we can accept?
- Goals. What are the underlying things we need or care about? What are our goals?
- Independence. What will we do if the negotiation fails? What alternatives have we got?
- Commitment. What outstanding commitments do we have?

Negotiation dialogues, in this context, exchange dialogical moves, i.e. messages, with the intention of getting information about the opponent or giving away information about us along these five dimensions: request for information, propose options, inform about interests, issue promises, appeal to standards A key part of any negotiation process is to build a model of our opponent(s) along these dimensions. All utterances agents make during a negotiation give away information about their current model, that is, about their legitimacy, options, goals, independence, and commitments. Also, several utterances can have a utilitarian interpretation in the sense that an agent can associate a preferential gain to them. For instance, an offer may inform our negotiation opponent about our willingness to sign a contract in the terms expressed in the offer, and at the same time the opponent can compute what is its associated expected utilitarian gain. These two views: information-based and utility-based, are central in the model proposed in this paper.

A. Intimacy and Balance in relationships

There is evidence from psychological studies that humans seek a *balance* in their negotiation relationships. The classical

view [1] 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 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.

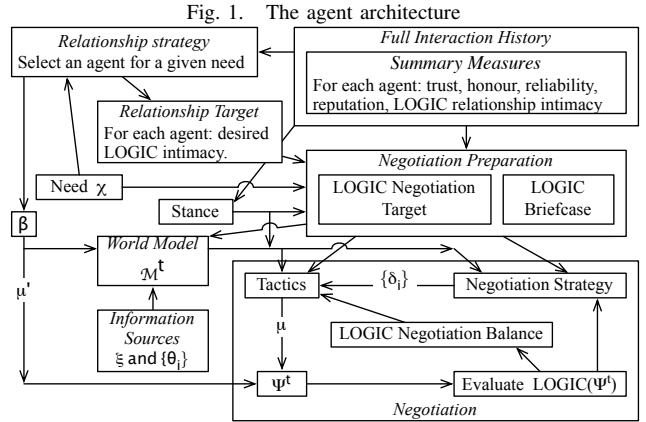
We believe that the perception of balance in dialogues (in negotiation or otherwise) is grounded on social relationships, and that every dimension of an interaction between humans can be correlated to the social closeness, or *intimacy*, between the parties involved. According to the previous studies, the more intimacy across the five dimensions 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 structure¹ to a shelter-based one² 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 [3] 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 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

¹In 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.

²In 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.



returned, or no complaints about not 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.

III. AGENT ARCHITECTURE

A multiagent system $\{\alpha, \beta_1, \dots, \beta_n, \xi, \theta_1, \dots, \theta_t\}$, 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 [2]. 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 III-A 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 [11].

The agent architecture is shown in Figure 1. Agent α acts in response to a *need* that is expressed in terms of the ontology. A need may be exogenous such as a need to trade profitably and may be triggered by another agent offering to trade, or endogenous such as α deciding that it owns more wine than it requires. Needs trigger α 's goal/plan proactive reasoning, while other messages are dealt with by α 's reactive reasoning.³ Each plan prepares for the negotiation by assembling the contents of a 'briefcase' that the agent 'carries' into the negotiation⁴. 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

³Each 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 [10], [11].

⁴Empirical evidence shows that in human negotiation, better outcomes are achieved by skewing the opening Options in favour of the proposer. We are unaware of any empirical investigation of this hypothesis for autonomous agents in real trading scenarios.

each trading relationship this strategy generates a *relationship target* that is expressed in the 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 framework in terms of the *value* of Ψ^t to both α and β — see Section IV-B. The *negotiation strategy* then determines the current set of Options $\{\delta_i\}$, and then the *tactics*, guided by the *negotiation target*, decide which, if any, of these Options to put forward and wraps them in argumentation dialogue — see Section V. 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 [10], [11]. $S_\alpha(\beta, a)$ is a random variable denoting α 's estimate of β 's subjective valuation of a over some finite, numerical evaluation space. 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*⁵ described in Section V. 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

⁵If α chooses to inflate her opening Options then this is achieved in Section V 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 Myerson and Satterthwaite in 1983.

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’ [5].

A. 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 .

B. Update of $\mathbb{P}^t(\varphi'|\varphi)$ given φ

First, if φ_k is observed then α may set $\mathbb{P}^{t+1}(\varphi_k|\varphi)$ to some value d where $\{\varphi_1, \varphi_2, \dots, \varphi_m\}$ is the set of all possible observations. We estimate the complete posterior distribution $\mathbb{P}^{t+1}(\varphi'|\varphi)$ by applying the principle of minimum

relative entropy⁶ as follows. Let $\vec{p}_{(\mu)}$ be the distribution: $\arg \min_{\vec{x}} \sum_j x_j \log \frac{x_j}{\mathbb{P}^t(\varphi'|\varphi)_j}$ that satisfies the constraint $\vec{p}_{(\mu)k} = d$. Then let $\vec{q}_{(\mu)}$ be the distribution:

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

and then let:

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

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

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

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

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 .

C. Update of $\mathbb{P}^t(\phi'|\phi)$ given φ

The sim method: Given as above $\mu = illoc(\alpha, \beta, \varphi, t)$ and the observation φ_k we define the vector \vec{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 $i = 1, \dots, p$. \vec{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 $\vec{r}_{(\mu)}$ defined to be the normalisation of \vec{t} .

The valuation method: For a given ϕ_k , $w^{\text{exp}}(\phi_k) = \sum_{j=1}^m \mathbb{P}^t(\phi_j|\phi_k) \cdot w(\phi_j)$ is α 's expectation of the value of what will be observed given that β has stated that ϕ_k will be observed, for some measure w . Now suppose that, as before, α observes φ_k after agent β has stated φ . α revises the prior estimate of the expected valuation $w^{\text{exp}}(\phi_k)$ in the light of the observation φ_k to:

$$(w^{\text{rev}}(\phi_k) | (\varphi_k|\varphi)) = g(w^{\text{exp}}(\phi_k), \text{Sim}(\phi_k, \varphi), w(\phi_k), w(\varphi), w_i(\varphi_k))$$

for some function g — the idea being, for example, that if the execution, φ_k , of the commitment, φ , to supply cheese was devalued then α 's expectation of the value of a commitment,

⁶Given a probability distribution \vec{q} , the *minimum relative entropy distribution* $\vec{p} = (p_1, \dots, p_I)$ subject to a set of J linear constraints $\vec{g} = \{g_j(\vec{p}) = \vec{a}_j \cdot \vec{p} - c_j = 0, j = 1, \dots, J$ (that must include the constraint $\sum_i p_i - 1 = 0$) is: $\vec{p} = \arg \min_{\vec{p}} \sum_j r_j \log \frac{r_j}{q_j}$. This may be calculated by introducing Lagrange multipliers $\vec{\lambda}$: $L(\vec{p}, \vec{\lambda}) = \sum_j p_j \log \frac{p_j}{q_j} + \vec{\lambda} \cdot \vec{g}$. Minimising L , $\{\frac{\partial L}{\partial \lambda_j} = g_j(\vec{p}) = 0, j = 1, \dots, J$ is the set of given constraints \vec{g} , and a solution to $\frac{\partial L}{\partial p_i} = 0, i = 1, \dots, I$ leads eventually to \vec{p} . Entropy-based inference is a form of Bayesian inference that is convenient when the data is sparse [4] and encapsulates common-sense reasoning [8].

ϕ , to supply wine should decrease. We estimate the posterior by applying the principle of minimum relative entropy as for Equation 3, where the distribution $\vec{p}_{(\mu)} = \vec{p}_{(\phi'|\phi)}$ satisfies the constraint:

$$\sum_{j=1}^p \mathcal{P}_{(\varphi', \varphi)_j} \cdot w_i(\phi_j) = g(w^{\text{exp}}(\phi_k), \text{Sim}(\phi_k, \varphi), w(\phi_k), w(\varphi), w_i(\varphi_k))$$

IV. SUMMARY MEASURES

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 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.

A. Confidence

Confidence measures generalise what are commonly called *trust*, *reliability* and *reputation* measures into a single computational framework that spans the five categories. In Section IV-B confidence measures are applied to valuing fulfilment of promises in the Legitimacy category — we formerly called this “honour” [10], to the execution of commitments — we formerly called this “trust” [9], and to valuing dialogues in the Goals category — we formerly called this “reliability” [10].

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:

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

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 $\varphi \leq o$:

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

where $\mathbb{P}_I^t(\varphi)$ 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 :

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

Preferred observations. The previous measure requires that an ideal distribution, $\mathbb{P}_1^t(\varphi'|\varphi, e)$, 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, Eqn. 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, and the right-hand side of Eqn. 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 .

B. 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 2×5 array $V_\alpha(\Psi^t)$ where:

$$V_x(\Psi^t) = \begin{pmatrix} I_x^L(\Psi^t) & I_x^O(\Psi^t) & I_x^G(\Psi^t) & I_x^I(\Psi^t) & I_x^C(\Psi^t) \\ U_x^L(\Psi^t) & U_x^O(\Psi^t) & U_x^G(\Psi^t) & U_x^I(\Psi^t) & U_x^C(\Psi^t) \end{pmatrix}$$

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. For example, if α receives the utterance “Today is Tuesday” then this may be translated into a constraint on a single distribution, and the resulting decrease in entropy is the information gain. Attaching a utilitarian measure to this utterance may not be so simple.

C. Intimacy and Balance

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 2×5 arrays V_α^{*t} and V_β^{*t} that are computed by an update function as each negotiation round terminates, $I_{\alpha\beta}^{*t} = (V_\alpha^{*t}, V_\beta^{*t})$. 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 2×5 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 eight 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 framework.

V. STRATEGIES AND TACTICS

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 2×5 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 $y = L, O, G, I, C$, 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 Options. If the set of options is empty then α will quit the negotiation. α perturbs the acceptance machinery (see Section III) by deriving s from the $S_{\alpha\beta}^t$ matrix such as the value at the (I, O) position. In line with the comment in Footnote 4, in the early stages of the negotiation α may decide to inflate her opening Options. This is achieved by increasing the value of s in Equation 1. The following strategy uses the machinery described in Section III. Fix h, g, s and c , set the Options 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 Options, 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 Options 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 2×5 matrix containing variations in the five 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 five dimensions. 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 category then

α will contribute nothing to it, and will leave this to the natural decay to the reputation D as described above.

VI. DISCUSSION

In this paper we have introduced a model of relationships that uses information and game-theoretical measures. 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. Negotiation is understood as a dialogue that affect five basic dimensions: Legitimacy, Options, Goals, Independence, and Commitment.

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 modeling of the negotiation interactions and has motivated some criticisms from its users about the lack of modeling of long-lasting relationships that our model could solve.

REFERENCES

- [1] J. S. Adams. Inequity in social exchange. In L. Berkowitz, editor, *Advances in experimental social psychology*, volume 2. New York: Academic Press, 1965.
- [2] J. L. Arcos, M. Esteva, P. Noriega, J. A. Rodríguez, and C. Sierra. Environment engineering for multiagent systems. *Journal on Engineering Applications of Artificial Intelligence*, 18, 2005.
- [3] M. H. Bazerman, G. F. Loewenstein, and S. B. White. Reversal of preference in allocation decisions: judging an alternative versus choosing among alternatives. *Administration Science Quarterly*, (37):220–240, 1992.
- [4] P. Cheeseman and J. Stutz. *Bayesian Inference and Maximum Entropy Methods in Science and Engineering*, chapter On The Relationship between Bayesian and Maximum Entropy Inference, pages 445 – 461. American Institute of Physics, Melville, NY, USA, 2004.
- [5] J. Debenham. Bargaining with information. In N. Jennings, C. Sierra, L. Sonenberg, and M. Tambe, editors, *Proceedings Third International Conference on Autonomous Agents and Multi Agent Systems AAMAS-2004*, pages 664 – 671. ACM Press, New York, July 2004.
- [6] R. Fischer, W. Ury, and B. Patton. *Getting to Yes: Negotiating agreements without giving in*. Penguin Books, 1995.
- [7] R. J. Lewicki, D. M. Saunders, and J. W. Minton. *Essentials of Negotiation*. McGraw Hill, 2001.
- [8] J. Paris. Common sense and maximum entropy. *Synthese*, 117(1):75 – 93, 1999.
- [9] C. Sierra and J. Debenham. An information-based model for trust. In F. Dignum, V. Dignum, S. Koenig, S. Kraus, M. Singh, and M. Wooldridge, editors, *Proceedings Fourth International Conference on Autonomous Agents and Multi Agent Systems AAMAS-2005*, pages 497 – 504, Utrecht, The Netherlands, July 2005. ACM Press, New York.
- [10] C. Sierra and J. Debenham. Trust and honour in information-based agency. In P. Stone and G. Weiss, editors, *Proceedings Fifth International Conference on Autonomous Agents and Multi Agent Systems AAMAS-2006*, pages 1225 – 1232, Hakodate, Japan, May 2006. ACM Press, New York.
- [11] C. Sierra and J. Debenham. Information-based agency. In *Proceedings of Twentieth International Joint Conference on Artificial Intelligence IJCAI-07*, pages 1513–1518, Hyderabad, India, January 2007.
- [12] H. Sondak, M. A. Neale, and R. Pinkley. The negotiated allocations of benefits and burdens: The impact of outcome valence, contribution, and relationship. *Organizational Behaviour and Human Decision Processes*, (3):249–260, December 1995.
- [13] K. L. Valley, M. A. Neale, and E. A. Mannix. Friends, lovers, colleagues, strangers: The effects of relationships on the process and outcome of negotiations. In R. Bies, R. Lewicki, and B. Sheppard, editors, *Research in Negotiation in Organizations*, volume 5, pages 65–94. JAI Press, 1995.