

Building Business Relationships with Negotiation

John Debenham¹ and Carles Sierra²

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

² Institut d'Investigacio en Intel.ligencia Artificial, Spanish Scientific Research Council, UAB
08193 Bellaterra, Catalonia, Spain
sierra@iia.csic.es

Abstract. Successful negotiators prepare by determining their position along five dimensions. We introduce a negotiation model based on these dimensions and on 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. 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.

1 Introduction

In this paper we propose a new negotiation model to deal with long term relationships that are founded on successive negotiation encounters. The model is grounded on results from business and psychological studies [1,2,3], and acknowledges that negotiation is an information exchange process as well as a utility exchange process [4,5]. We believe that if agents are to succeed in real application domains they have to reconcile both views: informational and game-theoretical. Before a negotiation starts human negotiators prepare the dialogic exchanges that can be made along the five LOGIC 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?

Our aim is to model trading scenarios where agents represent their 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. These two concepts, intimacy and balance are key in the model, and enable us to understand competitive and co-operative game theory as two particular theories of agent relationships (i.e. at different intimacy levels). These two theories are too specific and distinct to describe how a (business) relationship might grow because interactions have some aspects of these two extremes on a continuum in which, for example, agents reveal increasing amounts of private information as their intimacy grows. Negotiation strategies can then 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 2 we outline the aspects of human negotiation modelling that we cover in this work. Section 3 explains in outline the architecture and the concepts of intimacy and balance, and how they influence the negotiation. Section 4 contains a description of the different metrics used in the agent model including intimacy. Finally, Section 5 outlines how strategies and tactics use the LOGIC framework, intimacy and balance.

2 Human Negotiation

Negotiation dialogues 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 LOGIC 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.

2.1 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 [2,7] 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, *equality* being the allocation in equal amounts, and *need* being the allocation proportional to the need

for the resource. 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 LOGIC 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.

The perceived balance in a negotiation dialogue allows negotiators to infer information about their opponent, about its LOGIC 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 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.

3 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 ?? 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}^i . A generic *information-based* architecture is described in detail in [4].

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

¹ 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.

² 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.

dealt with by α 's reactive reasoning.³ Each plan prepares for the negotiation by assembling the contents of a 'LOGIC 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 each trading relationship this strategy generates a *relationship target* that is expressed in the LOGIC 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 LOGIC framework in terms of the *value* of Ψ^t to both α and β — see Section 4.2. 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 5. 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 [5,4]. $S_\alpha(\beta, a)$ 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 5. 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

³ 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 [5,4].

⁴ Empirical 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.

⁵ If α chooses to inflate her opening Options then this is achieved in Section 5 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.

h increases to 1. The basis for the acceptance criterion has thus developed from equity to equality, and then to need.

$\mathbb{P}^l(\text{acc}(\beta, \alpha, \delta))$ estimates the probability that β would accept δ , by observing β 's responses. For example, if β sends the message $\text{Offer}(\delta_1)$ then α derives the constraint: $\{\mathbb{P}^l(\text{acc}(\beta, \alpha, \delta_1)) = 1\}$ on the distribution $\mathbb{P}^l(\beta, \alpha, \delta)$, and if this is a counter offer to a former offer of α 's, δ_0 , then: $\{\mathbb{P}^l(\text{acc}(\beta, \alpha, \delta_0)) = 0\}$. 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" [9].

4 Summary Measures

A *dialogue*, Ψ^l , between agents α and β is a sequence of inter-related utterances in context. A *relationship*, Ψ^{*l} , 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 LOGIC 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.

4.1 Confidence

Confidence measures generalise what are commonly called *trust*, *reliability* and *reputation* measures into a single computational framework that spans the LOGIC categories. In Section 4.2 confidence measures are applied to valuing fulfilment of promises in the Legitimacy category — we formerly called this "honour" [5], to the execution of commitments — we formerly called this "trust" [10], and to valuing dialogues in the Goals category — we formerly called this "reliability" [5].

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 $\mathbb{P}_I^l(\varphi'|\varphi, e)$. Here we measure the relative entropy between this ideal distribution, $\mathbb{P}_I^l(\varphi'|\varphi, e)$, and the distribution of expected observations, $\mathbb{P}^l(\varphi'|\varphi)$. That is:

$$\mathbb{C}(\alpha, \beta, \varphi) = 1 - \sum_{\varphi'} \mathbb{P}_I^l(\varphi'|\varphi, e) \log \frac{\mathbb{P}_I^l(\varphi'|\varphi, e)}{\mathbb{P}^l(\varphi'|\varphi)} \quad (2)$$

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 $\varphi \leq o$:

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

where $\mathbb{P}'_{\beta}(\varphi)$ 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 :

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

Preferred observations. The previous measure requires, $\mathbb{P}'_I(\varphi'|\varphi, e)$, an ideal distribution 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 $\varphi \leq o$:

$$\mathbb{C}(\alpha, \beta, \varphi) = \sum_{\varphi'} \mathbb{P}'(\text{Prefer}(\varphi', \varphi, o)) \mathbb{P}'(\varphi'|\varphi)$$

and:

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

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”. Let: $\Phi_+(\varphi, o, \kappa) = \{\varphi' \mid \mathbb{P}'(\text{Prefer}(\varphi', \varphi, o)) > \kappa\}$ where $\varphi \leq o$ for some constant κ , and:

$$\mathbb{C}(\alpha, \beta, \varphi) = 1 + \frac{1}{B^*} \cdot \sum_{\varphi' \in \Phi_+(\varphi, o, \kappa)} \mathbb{P}'_+(\varphi'|\varphi) \log \mathbb{P}'_+(\varphi'|\varphi)$$

where $\mathbb{P}'_+(\varphi'|\varphi)$ is the normalisation of $\mathbb{P}'(\varphi'|\varphi)$ 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. 2 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. 2 may be approximated to:

$$1 - \sum_{\varphi': \text{Sim}(\varphi', \varphi) \geq \eta} \mathbb{P}'_{\eta, I}(\varphi'|\varphi, e) \log \frac{\mathbb{P}'_{\eta, I}(\varphi'|\varphi, e)}{\mathbb{P}'_{\eta}(\varphi'|\varphi)}$$

where $\mathbb{P}'_{\eta, I}(\varphi'|\varphi, e)$ is the normalisation of $\mathbb{P}'_I(\varphi'|\varphi, e)$ for $\text{Sim}(\varphi', \varphi) \geq \eta$, and similarly for $\mathbb{P}'_{\eta}(\varphi'|\varphi)$. 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 .

4.2 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 2×5 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.

We use the term “ 2×5 array” loosely to describe V_α in that the elements of the array are lists of measures that will be determined by the agent's requirements. Table 1 shows a sample measure for each of the ten categories, in it the dialogue commences at time s and terminates at time t . In that Table, $\mathbb{U}(\cdot)$ is a suitable utility evaluation function, $\text{needs}(\beta, \chi)$ means “agent β needs the need χ ”, $\text{cho}(\beta, \chi, \gamma)$ means “agent β satisfies need χ by choosing to negotiate with agent γ ”, N is the set of needs chosen from the ontology at some suitable level of abstraction, T^t is the set of offers on the table at time t , $\text{com}(\beta, \gamma, b)$ means “agent β has an outstanding commitment with agent γ to execute the commitment b ” where b is defined in the ontology at some suitable level of abstraction, B is the number of such commitments, and there are $n + 1$ agents in the system.

4.3 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 :

$$V_x^{*t+1} = v \times V_x(\Psi^t) + (1 - v) \times V_x^{*t} \quad (3)$$

Table 1. Sample measures for each category in $V_\alpha(\Psi^t)$. (Similarly for $V_\beta(\Psi^t)$.)

$$\begin{aligned}
I_\alpha^L(\Psi^t) &= \sum_{\varphi \in \Psi^t} \mathbb{C}^t(\alpha, \beta, \varphi) - \mathbb{C}^s(\alpha, \beta, \varphi) \\
U_\alpha^L(\Psi^t) &= \sum_{\varphi \in \Psi^t} \sum_{\varphi'} \mathbb{P}_\beta^t(\varphi' | \varphi) \times \mathbb{U}_\alpha(\varphi') \\
I_\alpha^O(\Psi^t) &= \frac{\sum_{\delta \in T^t} \mathbb{H}^s(\text{acc}(\beta, \alpha, \delta)) - \sum_{\delta \in T^t} \mathbb{H}^t(\text{acc}(\beta, \alpha, \delta))}{|T^t|} \\
U_\alpha^O(\Psi^t) &= \sum_{\delta \in T^t} \mathbb{P}^t(\text{acc}(\beta, \alpha, \delta)) \times \sum_{\delta'} \mathbb{P}^t(\delta' | \delta) \mathbb{U}_\alpha(\delta') \\
I_\alpha^G(\Psi^t) &= \frac{\sum_{\chi \in N} \mathbb{H}^s(\text{needs}(\beta, \chi)) - \mathbb{H}^t(\text{needs}(\beta, \chi))}{|N|} \\
U_\alpha^G(\Psi^t) &= \sum_{\chi \in N} \mathbb{P}^t(\text{needs}(\beta, \chi)) \times \mathbb{E}^t(\mathbb{U}_\alpha(\text{needs}(\beta, \chi))) \\
I_\alpha^I(\Psi^t) &= \frac{\sum_{i=1}^o \sum_{\chi \in N} \mathbb{H}^s(\text{cho}(\beta, \chi, \beta_i)) - \mathbb{H}^t(\text{cho}(\beta, \chi, \beta_i))}{n \times |N|} \\
U_\alpha^I(\Psi^t) &= \sum_{i=1}^o \sum_{\chi \in N} \mathbb{U}^t(\text{cho}(\beta, \chi, \beta_i)) - \mathbb{U}^s(\text{cho}(\beta, \chi, \beta_i)) \\
I_\alpha^C(\Psi^t) &= \frac{\sum_{i=1}^o \sum_{\delta \in B} \mathbb{H}^s(\text{com}(\beta, \beta_i, b)) - \mathbb{H}^t(\text{com}(\beta, \beta_i, b))}{n \times |B|} \\
U_\alpha^C(\Psi^t) &= \sum_{i=1}^o \sum_{\delta \in B} \mathbb{U}^t(\text{com}(\beta, \beta_i, b)) - \mathbb{U}^s(\text{com}(\beta, \beta_i, b))
\end{aligned}$$

where v 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 LOGIC 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 LOGIC framework.

5 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 way in which these agents go about building relationships with negotiation. 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 injects bounded random noise into the process, where the bound tightens as intimacy increases. The stance, $S'_{\alpha\beta}$, 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 $[\frac{1}{l(I'_{\alpha\beta}, x, y)}, l(I'_{\alpha\beta}, x, y)]$ where $l(I'_{\alpha\beta}, x, y)$ is the bound, and $I'_{\alpha\beta}$ 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 3) by deriving s from the $S'_{\alpha\beta}$ 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 3. Fix h, g, s and c , set the Options to the empty set, let $D'_s = \{\delta \mid \mathbb{P}'(\text{acc}(\alpha, \beta, \chi, \delta) > c)\}$, then:

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

By using $\mathbb{P}'(\text{acc}(\beta, \alpha, \delta))$ 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 3, α will want to achieve a *negotiation target*, $N_{\beta}(\Psi^t)$ such that: $v \cdot N_{\beta}(\Psi^t) + (1 - v) \cdot V_{\beta}^{*t}$ is "a bit on the T_{β}^{*t} side of" V_{β}^{*t} :

$$N_{\beta}(\Psi^t) = \frac{v - \kappa}{v} V_{\beta}^{*t} \oplus \frac{\kappa}{v} T_{\beta}^{*t} \quad (4)$$

for small $\kappa \in [0, v]$ 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 LOGIC 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 4. $N_{\alpha}(\Psi^t)$ 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 LOGIC dimensions. A cautious tactic could use the balance to bound the response μ to each utterance μ' from β by the constraint: $V_{\alpha}(\mu') \ominus V_{\beta}(\mu) \approx S'_{\alpha\beta} \otimes (N_{\alpha}(\Psi^t) \ominus N_{\beta}(\Psi^t))$, where \otimes is element-by-element matrix multiplication, and $S'_{\alpha\beta}$ 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 LOGIC category then α

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

6 Discussion

In this paper we have introduced a novel approach to negotiation 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. Each dialogical move produces a change in a 2×5 matrix that evaluates the dialogue along five information-based measures and five utility-based measures. The current Balance and intimacy levels and the desired, or target, levels are then used by the tactics to determine what to say next. 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.

References

1. Adams, J.S.: Inequity in social exchange. In: Berkowitz, L. (ed.) *Advances in experimental social psychology*, vol. 2, Academic Press, New York (1965)
2. 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)
3. Lewicki, R.J., Saunders, D.M., Minton, J.W.: *Essentials of Negotiation*. McGraw Hill, New York (2001)
4. Sierra, C., Debenham, J.: Information-based agency. In: *Proceedings of Twentieth International Joint Conference on Artificial Intelligence IJCAI-07*, Hyderabad, India, pp. 1513–1518 (2007)
5. Sierra, C., Debenham, J.: Trust and honour in information-based agency. In: Stone, P., Weiss, G. (eds.) *Proceedings Fifth International Conference on Autonomous Agents and Multi Agent Systems AAMAS-2006*, Hakodate, Japan, pp. 1225–1232. ACM Press, New York (2006)
6. Fischer, R., Ury, W., Patton, B.: *Getting to Yes: Negotiating agreements without giving in*. Penguin Books (1995)
7. 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, Greenwich (1995)
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. Debenham, J.: Bargaining with information. In: Jennings, N., Sierra, C., Sonenberg, L., Tambe, M. (eds.) *Proceedings Third International Conference on Autonomous Agents and Multi Agent Systems AAMAS-2004*, pp. 664–671. ACM Press, New York (2004)
10. Sierra, C., Debenham, J.: An information-based model for trust. In: Dignum, F., Dignum, V., Koenig, S., Kraus, S., Singh, M., Wooldridge, M. (eds.) *Proceedings Fourth International Conference on Autonomous Agents and Multi Agent Systems AAMAS-2005*, Utrecht, The Netherlands, pp. 497–504. ACM Press, New York (2005)