

© [2008] IEEE. Reprinted, with permission, from [Debenham, John., Sierra, Carles Unifying Trust, Honour and Reliability, Second IEEE International Conference on Digital Ecosystems and Technologies, February 2008]. This material is posted here with permission of the IEEE. Such permission 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

Unifying Trust, Honour and Reliability

John Debenham

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

Carles Sierra

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

Abstract—An argumentation based negotiation model is supported by information theory. Argumentative dialogues change the models of agents with respect to ongoing relationships. Trust and Honour are key components. Trust measures expected deviations of behaviour in the execution of commitments. Honour measures the expected integrity of the arguments exchanged. We understand rhetorical moves in dialogues as actions to project the current relationships into the future.

I. INTRODUCTION

This paper is in the area labelled: *information-based agency* [13]. 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. Negotiation dialogues have been traditionally organised a-round 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 [6], [3]. Game theory [11], possibilistic logic and first-order logic [7] 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* [14]. However, no formal model of the meaning of these speech acts has been proposed yet. Expanded negotiation dialogues, including these 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. We will understand argumentation in this paper as an information exchange process between agents. Every illocution that an agent utters gives away valuable information. To evaluate each illocution exchanged we base our information-based agent architecture on information theory.

We distinguish between the *trust* that an agent displays through the enactment of its contracts, and the *honour* that an agent displays in sustaining its trading relationships. The enactment of a contract is uncertain to some extent, and trust, precisely, is a measure of how uncertain the enactment of a contract is. Trust is therefore a *measure of expected deviations of behaviour* along a dimension determined by the type of the contract [12]. In this sense, the higher the trust the lower the expectation that a (significant) deviation from what is signed occurs. Exchanged information may turn out to be partly false,

and promises may be partly broken. We introduce the concept of Honour to measure the *expected integrity of the information and promises exchanged*. This leads to a natural hierarchy: individual deals, (business) relationships and businesses.

We first introduce an argumentation language in Section II. Then we define an information based model in Section III that supports our model of Honour in Section IV. How the information is managed is explained in Section V.

II. INTERACTION LANGUAGE

In order for communication to be effective agents need to agree upon an ontology and a communication language.

To model the agent dialogues we define *ontology* to include a (minimum) repertoire of elements: a set of *concepts* organised in an *is-a hierarchy*, captured by a partial order relation, a set of *relations* defined over these concepts, and a set of axioms defined over the concepts and relations. We model ontologies following an algebraic approach and present ontologies as logical theories in the form of pairs $O = (S, A)$ where S is the *ontological signature*, representing the vocabulary, and A is a set of ontological axioms, specifying the interpretation of the vocabulary in a given context. The signature is usually represented as a mathematical structure and the set of axioms are usually written in a logical formalism. In our case, we define an ontology signature as a tuple $S = (C, R, \leq, \sigma)$ where C is a finite set of concept symbols (including basic data types); R is a finite set of relation symbols; \leq is a reflexive, transitive and anti-symmetric relation on C (a partial order); and, $\sigma : R \rightarrow C^+$ is the function assigning to each relation symbol its arity. Concepts play somehow the role of *type*, and the is-a hierarchy is the notion of subtype. Thus, type inference mechanisms can be used to type all symbols appearing in expressions. Once the basic ontology signature is fixed we need to define a language to express, in our case, contracts, promises, rewards and so on. We denote that language as L_S . The axioms in the ontology will be written in that language, $A \subseteq L_S$, a first order language defined over the ontological expressions.

In our dialogues an essential concept is that of *contract*. Contracts are usually thought of as having a number of issues (dimensions) associated with values (or regions). We define the *ontological context* of formula $\varphi \in L_S$, denoted $O(\varphi)$, as the set of concepts in C used by the literals of φ . We extend the subtype relationship, \leq , in the following way,

$$\varphi \leq \psi \iff \forall c_i \in O(\varphi) (\exists c_j \in O(\psi). c_i \leq c_j)$$

We base the definition of the *semantic distance* between two concepts on the path length over the \leq of the signature (more distance in the graph is more semantic distance), and the *depth* of the subsumer concept on the shortest path between the two concepts (the deeper in the hierarchy, the closer the meaning of the concepts) [8]:

$$Sim(c_1, c_2) = e^{-\kappa_1 l} \cdot \frac{e^{\kappa_2 h} - e^{-\kappa_2 h}}{e^{\kappa_2 h} + e^{-\kappa_2 h}}$$

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

From the semantic distance between two concepts we define the semantic distance between formulae as the maxmin distance between concepts in their ontological context.

$$Sim(\varphi, \psi) = \max_{c_i \in O(\varphi)} \min_{c_j \in O(\psi)} \{Sim(c_i, c_j)\}$$

III. AN INFORMATION-BASED MODEL

We ground our argumentation model on information-based concepts. *Entropy*, H , is a measure of uncertainty [9] in a probability distribution for a discrete random variable X : $H(X) \triangleq -\sum_i p(x_i) \log p(x_i)$ where $p(x_i) = P(X = x_i)$. Maximum entropy inference and minimum relative entropy inference are chosen partly because of their encapsulation of common sense reasoning [10].

Maximum entropy inference is used to derive sentence probabilities for that which is not known by constructing the “maximally noncommittal” [5] probability distribution, and minimum relative entropy inference is used to update these distributions. These forms of inference are criticised [4] for their dependence on the representation chosen — such as the way in which values for a continuous variable are represented as intervals. We argue to the contrary, that this choice enables the tailoring of the model in fine detail.

Given a prior probability distribution $\underline{q} = (q_i)_{i=1}^n$ and a set of constraints, the *principle of minimum relative entropy* chooses the posterior probability distribution $\underline{p} = (p_i)_{i=1}^n$ that has the least relative entropy with respect to \underline{q} , and that satisfies the constraints.

[1] describes the estimation of both $P(Acc(\alpha, \beta, \delta))$ and the estimation of $P(Acc(\beta, \alpha, \delta))$ which is α ’s estimate of β ’s willingness to accept δ . These estimates are derived by applying maximum entropy inference to the observed behaviour of the agents. In the subsequent subsection we’ll see how α updates its sentence probabilities for Honour(\cdot) by observation following receipt of the illocutionary particles described in Section II.

A. Updating honour from observations

The choice of trading partner is influenced by the observation of the fulfilment of informs, rewards and threats. Such illocutions are a *conditional commitment* to act. Agent α has the opportunity to observe the extent to which agent β sticks to his word in Reward(\cdot), Threat(\cdot) and Inform(\cdot) illocutions, and to observe how accurate his *info* is. We base our measure of honour as the negative entropy of the probability distribution of possible outcomes following such a given illocution — honour measures the relationship between commitment and execution of promises. More precisely, between issued illocutions and their *perceived* execution. In this way, a natural way to base our modelling of honour is on a conditional probability, P^t , between commitment and observation given a context e as:

$$P^t(\text{Observe}(\alpha, \varphi') \mid (\text{Reward}(\beta, \alpha, \delta, \varphi); \text{Accept}(\alpha, \beta, \delta)), e)$$

where every reward execution represents a point in that distribution. For threats and informs the distributions are:

$$\begin{aligned} &P^t(\text{Observe}(\alpha, \varphi') \mid (\text{Threat}(\beta, \alpha, \delta, \varphi); \neg \text{Accept}(\alpha, \beta, \delta)), e) \\ &P^t(\text{Observe}(\alpha, \varphi') \mid (\text{Inform}(\beta, \alpha, \varphi)), e) \end{aligned}$$

For simplicity, we denote these relationships between the commitment, φ , and the observation, φ' , as $P^t(\varphi'|\varphi, e)$.

An important aspect that we want to model is the fact that beliefs ‘evaporate’ as time goes by. If we don’t keep an ongoing relationship, we become unsure how *honourable* a trading partner is. If I stop buying meat from my butcher, I’m not sure anymore that he will commit to his promises. This decay is what justifies a continuous relationship between individuals. In our model, the conditional probabilities should tend to ignorance as represented by the *decay limit distribution* $\{d_i\}$. If we have the set of observations $\Phi = \{\varphi_1, \varphi_2, \dots, \varphi_n\}$ then complete ignorance of the opponent’s expected behaviour means that given the opponent commits to φ the conditional probability for each observable outcome φ' becomes $d_i = \frac{1}{n}$ — i.e. the unconstrained maximum entropy distribution — but α may be prepared to make assumptions about β ’s decay limit distribution so that its entropy is less than this. This natural decay of belief is offset by new observations.

We define the evolution of the probability distribution that supports the previous definition of decay using an equation inspired by pheromone like models [2]:

$$P^{t+1}(\phi'|\phi) = \kappa \cdot ((1 - \nu) \cdot d_i + \nu \cdot (P^t(\phi'|\phi) + \Delta^t P(\phi'|\phi))) \quad (1)$$

where κ is a normalisation constant and $\{d_i\}$ is the decay limit distribution for β . This equation models the passage of time for a conveniently large $\nu \in [0, 1]$ and where the term $\Delta^t P(\phi'|\phi)$ represents the increment in an instant of time according to the last experienced event as the following shows.

Similarity based. The question is how to use the observation φ' given a commitment φ (generated by a sequence of illocutions Reward($\alpha, \beta, \delta, \varphi$); Accept(β, α, δ) — similarly

for others) in the context of a deal $\delta \leq \rho$ in the update of the overall probability distribution over the set of all possible outcomes defined over L_C . Here we use the idea that given a particular deviation in a region of the space, *similar* deviations should be expected in other regions. The intuition behind the update is that if my butcher has not honoured commitments with respect to red meat, then I might expect similar deviations with respect to poultry.

This idea is built upon the previously defined *Sim* function (see Section II) to take into account the difference between acceptance probabilities and similarity between the perception of the execution x of a reward (or threat) y . Thus, after the observation of φ' the increment of probability distribution at time $t + 1$ is:

$$\Delta^t P(\phi'|\phi) = (1 - |Sim(\varphi', \varphi) - Sim(\phi', \phi)|) \quad (2)$$

Entropy based. Suppose that α has a business relationship ρ with agent β , that β promises ϕ , and this promise is sound. The material value of ϕ to ρ will depend on the future use that α makes of it, and that is unlikely to be known. So α estimates the value of ϕ to the honour he holds for β in ρ using a probability distribution (p_1, \dots, p_n) over a relationship evaluation space $E = (e_1, \dots, e_n)$ that could range from “that is what I expect from the perfect trading partner” to “it is totally useless” — E may contain hard or fuzzy values. $p_i = w_i(\rho, \phi)$ is the probability that e_i is the correct evaluation of the enactment ϕ in the context of relationship ρ , and $\underline{w} : L_S \times L_S \rightarrow [0, 1]^n$ is the promise *evaluation function*.

Let $L_S = (\phi_1, \dots, \phi_m)$ in some order. Then for a given ϕ_k , $(P_\beta^t(\phi_1|\phi_k), \dots, P_\beta^t(\phi_m|\phi_k))$ is the prior distribution of α 's estimate of what will actually occur if β promised that ϕ_k would occur and $\underline{w}(\rho, \phi_k) = (w_1(\rho, \phi_k), \dots, w_n(\rho, \phi_k))$ is α 's evaluation over E with respect to the relationship ρ of β 's promise ϕ_k . α 's expected evaluation of what will occur given that β has promised that ϕ_k will occur is: $\underline{w}^{\text{exp}}(\rho, \phi_k) =$

$$\left(\sum_{j=1}^m P_\beta^t(\phi_j|\phi_k) \cdot w_1(\rho, \phi_j), \dots, \sum_{j=1}^m P_\beta^t(\phi_j|\phi_k) \cdot w_n(\rho, \phi_j) \right).$$

Now suppose that α observes the event $(\varphi'|\varphi)$ in another relationship ρ' also with agent β . Eg: α may buy wine and cheese from the same supplier. α may wish to revise the prior estimate $\underline{w}^{\text{exp}}(\rho, \phi_k)$ in the light of the observation $(\varphi'|\varphi)$ to: $(\underline{w}^{\text{rev}}(\rho, \phi_k) | (\varphi'|\varphi)) =$

$$\underline{g}(\underline{w}^{\text{exp}}(\rho, \phi_k), \underline{w}(\rho', \varphi), \underline{w}(\rho', \varphi'), \rho, \rho', \phi, \varphi, \varphi'),$$

for some function \underline{g} — the idea being, for example, that if the promise, φ , concerning the purchase of cheese, ρ' , was not kept then α 's expectation that the promise, ϕ , concerning the purchase of wine, ρ , will not be kept should increase. The entropy based approach estimates $\Delta_\beta^t P(\phi'|\phi)$ by applying the principle of minimum relative entropy. Let:

$$(P_{\beta, C}^t(\phi_j|\phi))_{j=1}^m = \arg \min_{\underline{p}} \sum_{i=1}^m p_i \log \frac{p_i}{P_\beta^t(\phi_i|\phi)} \quad (3)$$

satisfying the n constraints C , and $\underline{p} = (p_j)_{j=1}^m$. Then:

$$\Delta^t P(\phi'|\phi) = P_{\beta, C}^t(\phi'|\phi) - P_\beta^t(\phi'|\phi) \quad (4)$$

Where the n constraints C are: $\sum_{j=1}^m p_j \cdot w_i(\rho, \phi_j) =$

$$g_i(\underline{w}^{\text{exp}}(\rho, \phi_k), \underline{w}(\rho', \varphi), \underline{w}(\rho', \varphi'), \rho, \rho', \phi, \varphi, \varphi')$$

for $i = 1, \dots, n$. This is a set of n linear equations in m unknowns, and so the calculation of the minimum relative entropy distribution may be impossible if $n > m$. In this case, we take only the m equations for which the change from the prior to the posterior value is greatest. That is, we attempt to select the most significant factors.

IV. AN HONOUR MODEL

Honour as expected behaviour. Consider a distribution of expected fulfilment of promises that represent α 's “ideal” for a relationship with β , in the sense that it is the best that α could reasonably expect β to do. This distribution will be a function of β , α 's history with β , anything else that α believes about β , and general environmental information including time — denote all of this by e , then we have $P_I^t(\varphi'|\varphi, e)$. For example, if α considers that it is unacceptable for the execution φ' to be less preferred than the promise φ then $P_I^t(\varphi'|\varphi, e)$ will only be non-zero for those φ' that α prefers to φ . The distribution $P_I^t(\cdot)$ represents what α expects, or hopes, β will do. Honour is the relative entropy between this ideal distribution, $P_I^t(\varphi'|\varphi, e)$, and the distribution of the observation of fulfilled promises, $P_\beta^t(\varphi'|\varphi)$. That is:

$$H(\alpha, \beta, \varphi) = 1 - \sum_{\varphi'} P_I^t(\varphi'|\varphi, e) \log \frac{P_I^t(\varphi'|\varphi, e)}{P_\beta^t(\varphi'|\varphi)} \quad (5)$$

where the “1” is an arbitrarily chosen constant being the maximum value that honour may have. This equation defines honour for one, single promise φ — for example, my honour in my butcher if he promises me a 10% discount for the rest of the year. It makes sense to aggregate these values over a class of promises, say over those φ that are subtypes of a particular relationship ρ , that is $\varphi \leq \rho$. In this way we measure the honour that I have in my butcher in relation to the promises he makes for red meat generally:

$$H(\alpha, \beta, \rho) = 1 - \frac{\sum_{\varphi: \varphi \leq \rho} P_\beta^t(\varphi) \left[\sum_{\varphi'} P_I^t(\varphi'|\varphi, e) \log \frac{P_I^t(\varphi'|\varphi, e)}{P_\beta^t(\varphi'|\varphi)} \right]}{\sum_{\varphi: \varphi \leq \rho} P_\beta^t(\varphi)}$$

where $P_\beta^t(\varphi)$ is a probability distribution over the space of promises that the next promise β will make to α is φ . Similarly, for an overall estimate of α 's honour in β :

$$H(\alpha, \beta) = 1 - \sum_{\varphi} P_\beta^t(\varphi) \left[\sum_{\varphi'} P_I^t(\varphi'|\varphi, e) \log \frac{P_I^t(\varphi'|\varphi, e)}{P_\beta^t(\varphi'|\varphi)} \right]$$

Honour as expected ‘preferability’. The previous notion of honour requires that an ideal distribution, $P_I^t(\varphi'|\varphi, e)$,

has to be specified for each φ . The idea behind honour as expected ‘preferability’ is that in a trading relationship, ρ , if the execution φ' of a promise is preferable to that that was promised φ then α will honour β more. The extent to which an enactment, φ' , can be ‘preferable’ to a promise, φ , is limited. For example, if my grocer promises me a 10% discount on tomatoes and gives me 15% then I may be delighted, but if he gives me a 95% discount then my desire to trade with him may diminish. Given a predicate $\text{Prefer}(c_1, c_2, e)$ meaning that α prefers c_1 to c_2 in environment e . An evaluation of $P^t(\text{Prefer}(c_1, c_2, e))$ may be defined using $\text{Sim}(\cdot)$ and the evaluation function \underline{w} — but we do not detail it here. Then if $\varphi \leq \rho$:

$$H(\alpha, \beta, \varphi) = \sum_{\varphi'} P^t(\text{Prefer}(\varphi', \varphi, \rho)) P_{\beta}^t(\varphi' | \varphi)$$

and:

$$H(\alpha, \beta, \rho) = \frac{\sum_{\varphi: \varphi \leq \rho} P_{\beta}^t(\varphi) \left[\sum_{\varphi'} P^t(\text{Prefer}(\varphi', \varphi, \rho)) P_{\beta}^t(\varphi' | \varphi) \right]}{\sum_{\varphi: \varphi \leq \rho} P_{\beta}^t(\varphi)}$$

Honour as certainty in promise execution. Honour is consistency in expected acceptable executions of promises, or “the lack of expected uncertainty in those possible executions that are better than the promise as specified”. Then if $\varphi \leq \rho$ let: $\Phi_+(\varphi, \rho, \kappa) = \{\varphi' | P^t(\text{Prefer}(\varphi', \varphi, \rho)) > \kappa\}$ for some constant κ . The idea here is that α will honour β more if variations, φ' , from expectation, φ , are not random. The Honour that an agent α has on agent β with respect to the execution of a promise φ is:

$$H(\alpha, \beta, \varphi) = 1 + \frac{1}{B^*} \cdot \sum_{\varphi' \in \Phi_+(\varphi, \rho, \kappa)} P_+^t(\varphi' | \varphi) \log P_+^t(\varphi' | \varphi)$$

where $P_+^t(\varphi' | \varphi)$ is the normalisation of $P_{\beta}^t(\varphi' | \varphi)$ for $\varphi' \in \Phi_+(\varphi, \rho, \kappa)$,

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

As above we aggregate this measure for those promises of a particular type ρ :

$$\begin{aligned} H(\alpha, \beta, \rho) &= 1 + \\ &\frac{\sum_{\varphi: \varphi \leq \rho} \left[P_{\beta}^t(\varphi) \cdot \sum_{\varphi' \in \Phi_+(\varphi, \rho, \kappa)} P_+^t(\varphi' | \varphi) \log P_+^t(\varphi' | \varphi) \right]}{B^* \cdot \sum_{\varphi: \varphi \leq \rho} P_{\beta}^t(\varphi)} \\ &= 1 + \frac{\sum_{\varphi: \varphi \leq \rho} \sum_{\varphi' \in \Phi_+(\varphi, \rho, \kappa)} [P_+^t(\varphi', \varphi) \log P_+^t(\varphi' | \varphi)]}{B^* \cdot \sum_{\varphi: \varphi \leq \rho} P_{\beta}^t(\varphi)} \end{aligned}$$

where $P_{\beta}^t(\varphi', \varphi)$ is the joint probability distribution, and $P_+^t(\varphi', \varphi)$ is the normalisation of it as above. And, similarly

as before:

$$\begin{aligned} H(\alpha, \beta) &= \\ 1 + &\frac{\sum_{\varphi} \left[P_{\beta}^t(\varphi) \cdot \sum_{\varphi' \in \Phi_+(\varphi, \rho, \kappa)} P_+^t(\varphi' | \varphi) \log P_+^t(\varphi' | \varphi) \right]}{B^*} \\ &= 1 + \frac{\sum_{\varphi} \sum_{\varphi' \in \Phi_+(\varphi, \rho, \kappa)} [P_+^t(\varphi', \varphi) \log P_+^t(\varphi' | \varphi)]}{B^*} \end{aligned}$$

Evaluating honour. The various expressions for honour given above involve extensive calculations. For example, Eqn. 5 contains $\sum_{\varphi'}$ that sums over *all* possible enactments φ' . To develop the above into a more computationally friendly model we appeal to the structure of the ontology described in Sec. ???. For example, Eqn. 5 may be approximated to:

$$H(\alpha, \beta, \varphi) = 1 - \sum_{\varphi': \text{Sim}(\varphi', \varphi) \geq \eta} P_{\eta, I}^t(\varphi' | \varphi, e) \log \frac{P_{\eta, I}^t(\varphi' | \varphi, e)}{P_{\eta, \beta}^t(\varphi' | \varphi)}$$

where $P_{\eta, I}^t(\varphi' | \varphi, e)$ is the normalisation of $P_I^t(\varphi' | \varphi, e)$ for $\text{Sim}(\varphi', \varphi) \geq \eta$, and similarly for $P_{\eta, \beta}^t(\varphi' | \varphi)$. 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 example, if β promises a 10% discount on tomatoes next month, we may choose to evaluate $H(\alpha, \beta, \varphi)$ for promises φ' that are *both* close to this promise, φ , and that concern tomatoes, ψ .

V. DEALING WITH [info].

In this section we discuss how to deal with the information that β communicates in the illocutions, $\text{info} \in L_C$. We will base this treatment on β ’s *reliability* as an estimate of the extent to which this *info* is correct. For example, β may send α the *info* that “the price of fish will go up by 10% next week”, and it may actually go up by 9%. α ’s argumentation and relationship building strategies are based on plans that are not described here. Those plans determine a set of probability distributions from which a world model is derived. *info* is represented as a set of linear constraints on one or more of those probability distributions. A chunk of *info* may not be directly related to one of those distributions, or may not be expressed naturally as constraints, and so some inference machinery is required to derive these constraints — this inference is performed by a set of model building function, J_s , for each active plan s chosen by α . If a plan s calls for the distribution D then J_s^D is the model building function for D , and $J_s^D(\text{info})$ is the set of constraints on D derived from *info*.

A. Updating the world model with info

If at time u , α receives a message containing *info* it is time-stamped and source-stamped $\text{info}_{(\beta, \alpha, u)}$, and placed in a repository \mathcal{Y}^t . If α has an active plan, s , then the model building function, J_s , is applied to $\text{info}_{(\beta, \alpha, u)}$ to derive constraints on some, or none, of α ’s distributions. The extent to which those constraints are permitted to effect the distributions is determined by a value for the *reliability* of β , $R^t(\alpha, \beta, O(\text{info}))$.

An agent may have models of integrity decay for some particular distributions, but general models of integrity decay for,

say, a chunk of information taken at random from the World Wide Web are generally unknown. However the values to which decaying integrity should tend in time *are* often known. For example, a prior value for the truth of the proposition that a “22 year-old male will default on credit card repayment” is well known to banks. As described in Sec. III-A, if α attaches such prior values to a distribution D they are called the *decay limit distribution* for D , $(d_i^D)_{i=1}^n$. No matter how integrity of *info* decays, in the absence of any other relevant information it should decay to the decay limit distribution.

In the absence of new *info* the integrity of distributions decays. If $D = (q_i)_{i=1}^n$ then we use a geometric model of decay:

$$q_i^{t+1} = (1 - \nu^D) \times d_i^D + \nu^D \times q_i^t, \text{ for } i = 1, \dots, n \quad (6)$$

where $\nu^D \in (0, 1)$ is the decay rate. This raises the question of how to determine ν^D . Just as an agent may know the decay limit distribution it may also know something about ν^D . In the case of an information-overfed agent there is no harm in conservatively setting ν^D “a bit on the low side” as the continually arriving *info* will sustain the estimate for D .

We now describe how new *info* is imported to the distributions. A single chunk of *info* may affect a number of distributions. Suppose that a chunk of *info* is received from β and that α attaches the epistemic belief probability $R^t(\alpha, \beta, O(\text{info}))$ to it. Each distribution models a facet of the world. Given a distribution $D = (q_i)_{i=1}^n$, q_i is the probability that the possible world ω_i for D is the true world for D . The effect that a chunk *info* has on distribution D is to enforce the set of linear constraints on D , $J_s^D(\text{info})$. If the constraints $J_s^D(\text{info})$ are taken by α as valid then α could update D to the posterior distribution $(p_i^{\text{info}})_{i=1}^n$ that is the distribution with least relative entropy with respect to $(q_i^t)_{i=1}^n$ satisfying the constraint:

$$\sum_i \{p_i^{\text{info}} : J_s^D(\text{info}) \text{ are all } \top \text{ in } \omega_i\} = 1. \quad (7)$$

But $R^t(\alpha, \beta, O(\text{info})) = r \in [0, 1]$ and α should only treat the $J_s^D(\text{info})$ as valid if $r = 1$. In general r determines the extent to which the effect of *info* on D makes it closer to $(p_i^{\text{info}})_{i=1}^n$ or to the prior $(q_i^t)_{i=1}^n$ distribution by:

$$p_i^t = r \times p_i^{\text{info}} + (1 - r) \times q_i^t \quad (8)$$

But, we should only permit a new chunk of *info* to influence D if doing so gives us new information. For example, if 5 minutes ago a trusted agent advises α that the interest rate will go up by 1%, and 1 minute ago a very unreliable agent advises α that the interest rate may go up by 0.5%, then the second unreliable chunk should not be permitted to ‘overwrite’ the first. We capture this by only permitting a new chunk of *info* to be imported if the resulting distribution has more information *relative* to the decay limit distribution than the existing distribution has. Precisely, this is measured using the Kullback-Leibler distance measure, and *info* is only used if:

$$\sum_{i=1}^n p_i^t \log \frac{p_i^t}{d_i^D} > \sum_{i=1}^n q_i^t \log \frac{q_i^t}{d_i^D} \quad (9)$$

In addition, we have described in Eqn. 6 how the integrity of each distribution D will decay in time. Combining these two into one result, distribution D is revised to:

$$q_i^{t+1} = \begin{cases} (1 - \nu^D) \times d_i^D + \nu^D \times p_i^t & \text{if usable } \text{info} \text{ is} \\ & \text{received at time } t \\ (1 - \nu^D) \times d_i^D + \nu^D \times q_i^t & \text{otherwise} \end{cases}$$

for $i = 1, \dots, n$, and decay rate ν^D as before. We have yet to estimate $R^t(\alpha, \beta, O(\text{info}))$ — that is described in Sec. V-B. We now illustrate the above by showing how an important class of *info* illocutions, namely preference information, is managed.

Preferences. Preference information is a statement by an agent that it prefers one class of deals to another where deals may be multi-issue. Preference illocutions may refer to particular issues within deals — e.g. “I prefer red to yellow”, or to combinations of issues — e.g. “I prefer a car with a five year warranty to the same car with a two year warranty than costs 15% less”. Here α receives preference information, *info*, from β through an *Inform*($\beta, \alpha, \text{info}$) illocution, or through one of the other argumentation illocutions — see Sec. II — and attaches $R^t(\alpha, \beta, O(\text{info}))$ to it. What happens next will depend on α ’s plans. Suppose that α has a plan s that constructs the probability distribution $P^t(\text{UPrefer}(\beta, \alpha, \delta))$ over all δ meaning that “ δ is the deal that β prefers most”. α may know the decay limit distribution for $\text{UPrefer}(\cdot)$. Suppose α has a prior distribution $(q_i^t)_{i=1}^n$ for $\text{UPrefer}(\cdot)$. Now suppose that α learns the *info* “x% of the time β prefers deals with property Q_1 to those with property Q_2 ”. J_s^{UPrefer} derives the following linear constraint on the $P^t(\text{UPrefer}(\beta, \alpha, \delta))$ distribution:

$$\frac{x}{100} = \frac{\sum_{\delta: Q_1(\delta)} p_\delta}{(\sum_{\delta: Q_1(\delta)} p_\delta) + (\sum_{\delta: Q_2(\delta)} p_\delta) - (\sum_{\delta: Q_1 \wedge Q_2(\delta)} p_\delta)}$$

and the procedure described above follows as described. Note that this manages two different probabilities: first, the probability within the *info* “x% of the time...”, and second α ’s estimate of β ’s reliability in providing information of this sort, $R^t(\alpha, \beta, O(\text{info}))$, that is determined by the ontological context $O(\text{info})$ of the *info*.

B. A reliability model

We estimate $R^t(\alpha, \beta, O(\text{info}))$ by measuring the error in information. α ’s plans will have constructed a set of distributions. We measure the ‘error’ in information as the error in the effect that information has on each of α ’s distributions. Suppose that a chunk of *info* is received from agent β at time u and is verified at some later time t . For example, a chunk of information could be “the interest rate will rise by 0.5% next week”, and suppose that the interest rate actually rises by 0.25% — call that correct information *fact*. What does all this tell agent α about agent β ’s reliability? Consider one of α ’s distributions D that is $\{q_i^u\}$ at time u . Let $(p_i^{\text{info}})_{i=1}^n$ be the minimum relative entropy distribution given that *info* has been received as calculated in Eqn. 7, and let $(p_i^{\text{fact}})_{i=1}^n$ be that

distribution if *fact* had been received instead. Suppose that the reliability estimate for distribution D was R_D^u . This section is concerned with what R_D^u should have been in the light of knowing *now*, at time t , that *info* should have been *fact*, and how that knowledge effects our current reliability estimate for D , $R^t(\alpha, \beta, O(\text{info}))$.

The idea of Eqn. 8, is that the current value of r should be such that, *on average*, $(p_i^u)_{i=1}^n$ will be seen to be “close to” $(p_i^{\text{fact}})_{i=1}^n$ when we eventually discover *fact* — no matter whether or not *info* was used to update D , as determined by the acceptability test in Eqn. 9 at time u . That is, given *info*, *fact* and the prior $(q_i^u)_{i=1}^n$, calculate $(p_i^{\text{info}})_{i=1}^n$ and $(p_i^{\text{fact}})_{i=1}^n$ using Eqn. 7. Then the *observed reliability* for distribution D , $R_D^{(\text{info}|\text{fact})}$, on the basis of the verification of *info* with *fact* is the value of r that minimises the Kullback-Leibler distance between $(p_i^u)_{i=1}^n$ and $(p_i^{\text{fact}})_{i=1}^n$:

$$\arg \min_r \sum_{i=1}^n (r \cdot p_i^{\text{info}} + (1-r) \cdot q_i^u) \log \frac{r \cdot p_i^{\text{info}} + (1-r) \cdot q_i^u}{p_i^{\text{fact}}}$$

If E^{info} is the set of distributions that *info* affects, then the overall *observed reliability* on the basis of the verification of *info* with *fact* is: $R^{(\text{info}|\text{fact})} = 1 - (\max_{D \in E^{\text{info}}} |1 - R_D^{(\text{info}|\text{fact})}|)$. Then for each ontological context o_j , at time t when, perhaps, a chunk of *info*, with $O(\text{info}) = o_k$, may have been verified with *fact*:

$$R^{t+1}(\alpha, \beta, o_j) = (1 - \nu) \times R^t(\alpha, \beta, o_j) + \nu \times R^{(\text{info}|\text{fact})} \times \text{Sim}(o_j, o_k)$$

where Sim measures the semantic distance between two sections of the ontology, and ν is the learning rate. Over time, α notes the ontological context of the various chunks of *info* received from β and over the various ontological contexts calculates the relative frequency, $P^t(o_j)$, of these contexts, $o_j = O(\text{info})$. This leads to an overall expectation of the *reliability* that agent α has for agent β :

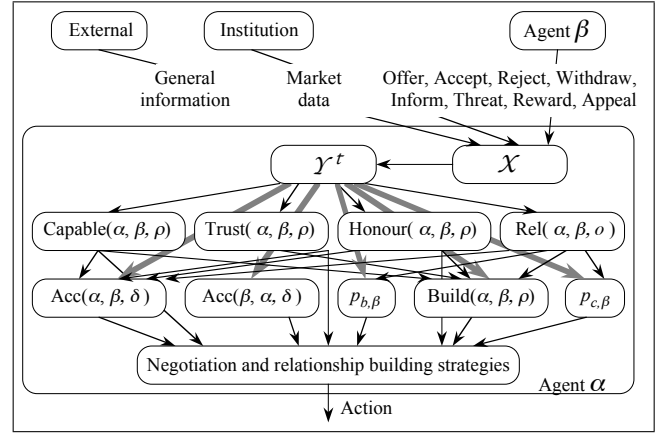
$$R^t(\alpha, \beta) = \sum_j P^t(o_j) \times R^t(\alpha, \beta, o_j)$$

VI. CONCLUSION

The agent architecture is summarised in Fig. 1. α 's actions are determined by its relationship-building strategy that identifies trading partners, and by its negotiation strategy that manages individual negotiations. We described two new concepts: *Honour*(\cdot) and *Reliability*(\cdot). *Honour*(\cdot) models the fulfilment of the promises derived from the use of argumentative particles such as *appeals*, *rewards* and *threats*. This concept plays a similar role to that that the concept of *Trust*(\cdot) plays with respect to the fulfilment of contract agreements, as developed in [12]. *Reliability*(\cdot) models the reliability of the information tabled by an opponent. We have proposed a semantics for these two concepts based on information theory.

Acknowledgements. Carles Sierra's research is partially supported by the OpenKnowledge STREP project, sponsored

Fig. 1. Trust, Honour and Reliability in the informed agent. The figure shows the social model and a sample world model. The social model consists of the four components: Capable(\cdot), Trust(\cdot), Honour(\cdot) and Rel(\cdot) standing for ‘reliability’. The sample world model contains just five distributions with the thick arcs representing the five model building functions J_s^D .



by the European Commission under contract number FP6-027253, and partially by the Spanish project “Agreement Technologies” (CONSOLIDER CSD2007-0022, INGENIO 2010).

REFERENCES

- [1] 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.
- [2] M. Dorigo and T. Stützle. *Ant Colony Optimization*. MIT Press, Cambridge, MA, 2004.
- [3] P. Faratin, C. Sierra, and N. Jennings. Using similarity criteria to make issue trade-offs in automated negotiation. *Journal of Artificial Intelligence*, 142(2):205–237, 2003.
- [4] J. Halpern. *Reasoning about Uncertainty*. MIT Press, 2003.
- [5] E. Jaynes. *Probability Theory — The Logic of Science*. Cambridge University Press, 2003.
- [6] N. Jennings, P. Faratin, A. Lomuscio, S. Parsons, C. Sierra, and M. Wooldridge. Automated negotiation: Prospects, methods and challenges. *International Journal of Group Decision and Negotiation*, 10(2):199–215, 2001.
- [7] S. Kraus. Negotiation and cooperation in multi-agent environments. *Artificial Intelligence*, 94(1–2):79–97, 1997.
- [8] Y. Li, Z. A. Bandar, and D. McLean. An approach for measuring semantic similarity between words using multiple information sources. *IEEE Transactions on Knowledge and Data Engineering*, 15(4):871 – 882, July / August 2003.
- [9] D. MacKay. *Information Theory, Inference and Learning Algorithms*. Cambridge University Press, 2003.
- [10] J. Paris. Common sense and maximum entropy. *Synthese*, 117(1):75 – 93, 1999.
- [11] J. S. Rosenschein and G. Zlotkin. *Rules of Encounter*. The MIT Press, Cambridge, USA, 1994.
- [12] 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.
- [13] 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.
- [14] C. Sierra, N. Jennings, P. Noriega, and S. Parsons. *Proceedings of the 4th International Workshop on Intelligent Agents IV, Agent Theories, Architectures, and Languages*, chapter A Framework for Argumentation-Based Negotiation, pages 177 – 192. Springer-Verlag, London, UK, 1997.