Chapter 29
Building Relationships with Trust

Carles Sierra and John Debenham

Abstract In this chapter we put together two major threads of work: trust in the enactment of contracts and the modelling of relationships between agents. We depart from previous work where trust is defined as the relationship between commitment of action and instant observation of the actual actions being performed. Here we generalise the approach by assuming a time delay between the observation of the actions and their valuation. The fundamental new idea being that commitment for action has a social dimension as the commitment of an agent should mean 'attempting to act' in the interest of the contractual partner, and that attempt has a time dimension that cannot be ignored.

29.1 Introduction

In this section trust is presented as the foundation for a rich sense of friendship between agents in a multiagent system. When agents interact their growing history of illocutionary dialogues is their *relationship*. An agent understands its relationships using various measures that summarise its dialogue history. These summary measures, of which trust is fundamental, enable relationships to be understood in the context of a multifaceted continuum rather than the simplistic cooperative / competitive divide. On the basis of this understanding an agent may choose: to form speculative beliefs concerning the future behaviour of other agents, to decide who to interact with under given circumstances, and to determine how to interact with them. This opens the way for an agent to proactively influence its dialogues with

Carles Sierra
IIIA - CSIC, Barcelona, e-mail: `sierra@iiia.csic.es`

John Debenham
QCIS, Univ. of Technology, Sydney, Australia, e-mail: `john.debenham@uts.edu.au`

the aim of shaping its relationships so that they provide some degree of protection against future unknowns in an uncertain world.

Section 29.2 introduces the framework within which the work is developed; in particular, the term trust is defined in the context of the signing, enactment and evaluation of contracts. Section 29.3 describes the components of the trust model: the ontology, the core trust mechanism, the representation of prior knowledge, and the context. Then in Section 29.4 the relationship model is introduced — this models the relationships between agents. Section 29.5 draws the previous ideas together by discussing negotiation.

29.2 Trust

The informal meaning of the statement "agent α trusts agent β" is that α expects β to act in a way that is somehow preferred by α. Human agents seldom trust another for *any* action that they may take — it is more usual to develop a trusted expectation with respect to a particular set of actions. For example, "I trust John to deliver fresh vegetables" whilst the quality of John's advice on investments may be terrible. This section describes trust when the set of actions is restricted to negotiating, signing and enacting contracts that are expressed using some particular ontology.

A multiagent system $\{\alpha, \beta_1, \ldots, \beta_o, \xi, \theta_1, \ldots, \theta_t\}$, contains an agent α that interacts with negotiating agents, $\mathscr{X} = \{\beta_i\}$, information providing agents, $\mathscr{I} = \{\theta_j\}$, and an *institutional agent*, ξ, that represents the institution where the interactions are assumed to happen [3]. Institutions give a normative context to interactions that simplify matters (e.g an agent can't make an offer, have it accepted, and then renege on it). The institutional agent ξ may form opinions on the actors and activities in the institution and may publish reputation estimates on behalf of the institution. The agent ξ also fulfils a vital role to compensate for any lack of sensory ability in the other agents by promptly and accurately reporting observations as events occur. For example, without such reporting an agent may have no way of knowing whether it is a fine day or not.

Our agents are information-based [10], they are endowed with machinery for valuing the information that they have, and that they receive. They were inspired by the observation that "everything an agent says gives away information", even if the utterances are not truthful. They model how much they know about other agents, how much they believe other agents know about them, and the extent to which they believe other agents are telling the truth. Everything in their world, including their information, is uncertain; their only means of reducing uncertainty is acquiring fresh information. To model this uncertainty, their world model, $\mathscr{M}^t$, consists of random variables each representing a point of interest in the world. Distributions are then derived for these variables on the basis of information received. Over time agents acquire large amounts of information that are distilled into convenient measures including trust. By classifying private information into functional classes, and by drawing on the structure of the ontology, information-based agents develop other measures including a map of the 'intimacy' [11] of their relationships with other agents.

In this section agent interaction is limited to dealing with contracts. The scenario is: two agents α and β negotiate with the intention of leading to a *signed contract* that is a pair of commitments, (a,b), where a is α's and b is β's. A contract is signed by both agents at some particular time t. At some later time, t', both agents will have enacted their commitments[1] in some way, as say (a',b'). At some later time again, t'', α will consume b' and will then be in a position to evaluate the extent to which β's enactment of (a,b), b', was in α's interests. See Figure 29.1.

α's trust of agent β is expressed as an expectation of β's future actions. We consider how α forms these expectations, how α will compare those expectations with observations, and how α then determines whether β's actions are preferred to α's expectations of them.

α forms expectations of β's future actions on the basis of all that it has: its full interaction history $H_\alpha \in \mathscr{H}_\alpha$ where $\mathscr{H}_\alpha$ is the set of all possible interaction histories that may be expressed in α's ontology[2]. H_α is a record of all interactions with each negotiating agent in $\mathscr{X}$ and with each information providing agent in $\mathscr{I}$. Let $\mathscr{B} = (b_1, b_2, \dots)$ denote that space of all enactments that β may make and $\mathscr{A}$ the space of α's enactments. α's expectations of β's behaviour will be represented as probability distributions over $\mathscr{B}$. Assuming that the space of contracts and enactments are the same, the space of all contracts and enactments is: $\mathscr{C} = \mathscr{A} \times \mathscr{B}$.

This raises the strategic question of given an expectation of some particular future requirements how should α strategically shape its interaction history to enable it to build a reliable expectation of β's future actions concerning the satisfaction of those particular requirements. At time t'' α compares b' with α's expectations of β's actions, β having committed at time t to enact b at time t'. That is:

$$\text{compare}^{t''}_\alpha(\mathbb{E}^t_\alpha(\text{Enact}^{t'}_\beta(b)|\text{sign}^t_{\alpha,\beta}((a,b)), H^t_\alpha), b')$$

where $\text{sign}^t_{\alpha,\beta}((a,b))$ is a predicate meaning that the joint action by α and β of signing the contract (a,b) was performed at time t, and $\text{Enact}^{t'}_\beta(b)$ is a random variable over $\mathscr{B}$ representing α's expectations over β's enactment action at time t', $\mathbb{E}^t_\alpha(\cdot)$ is α's expectation, and $\text{compare}(\cdot,\cdot)$ somehow describes the result of the comparison.

Expectations over β's enactment actions:

$$\mathbb{E}^t_\alpha(\text{Enact}^{t'}_\beta(b)|\text{sign}^t_{\alpha,\beta}((a,b)), H^t_\alpha)$$

could form the basis for trust. In practice, developing a sense on expectation over β's actions is tricky except possibly in the case that there is a history of contracts with a high degree of similarly to (a,b). Given such an expectation an agent may be prepared to use the structure of the ontology to propagate these expectations.

[1] For convenience it is assumed that both agents are presumed to have completed their enactments by the same time, t'.

[2] The ontology is not made explicit to avoid overburdening the notation.

Fig. 29.1: Contract signing, execution and evaluation.

Action	Sign	Enact	Evaluate
Object	*(a,b)*	*(a',b')*	*b'*
Time	*t*	*t'*	*t''*

For example, if α has a history of observing β's 'trusted' executions of orders for cow's cheese then it may be prepared to partially propagate this expectation to goat's cheese — perhaps on the basis that cow's cheese and goat's cheese are semantically close concepts in the ontology.

The discussion above is based on expectations of *what action* β will do. It makes more practical sense to develop a sense of expectation over *the evaluation of* β's actions. Let $\mathscr{V} = (v_1, v_2, \ldots, v_V)$ be the valuation space. Then α's expectation of the evaluation of a particular action that β may make is represented as a probability distribution over $\mathscr{V}$: $(f_1, f_2, \ldots, f_V)$. For example, a simple valuation space could be $(\text{good}, \text{ok}, \text{bad})$. The sequence $\mathscr{V}$ will generally be smaller than the sequence $\mathscr{B}$, and so developing a sense of expectation for the value of β's actions should be easier than for the actions themselves. That is, it is simpler to form the expectation:

$$\mathbb{E}^t_\alpha(\text{Value}^{t''}_\beta(b)|\text{sign}^t_{\alpha,\beta}((a,b)), H^t_\alpha)$$

where $\text{Value}^{t''}(b)$ is a random variable over $\mathscr{V}$ representing α's expectations of the value of β's enactment action given that he signed (a,b) and given H^t_α. At time t'' it then remains to compare expectation, $\mathbb{E}^t_\alpha(\text{Value}^{t''}_\beta(b)|\text{sign}^t_{\alpha,\beta}((a,b)), H^t_\alpha)$, with observation, $\text{val}_\alpha(b')$, where $\text{val}(\cdot)$ represents α's preferences — i.e. it is α's utility function[3].

We are now in a position to define 'trust'. *Trust*, $\tau_{\alpha\beta}(b)$, is a computable[4] estimate of the distribution: $\mathbb{E}^t_\alpha(\text{Value}^{t''}_\beta(b)|\text{sign}^t_{\alpha,\beta}((a,b)), H^t_\alpha)$. τ is a summarising function that distils the trust-related aspects of the (probably very large) set H_α into a probability distribution that may be computed. $\tau_{\alpha\beta}(b)$ summarises the large set H_α. The set of contracts $\mathscr{C}$ is also large. It is practically unfeasible to estimate trust for each individual contract. The structure of the ontology is used to deal with this problem by aggregating estimates into suitable classes such as John's trustworthiness in supplying Australian red wine.

In real world situations the interaction history may not reliably predict future action, in which case the notion of trust is fragile. No matter how trust is defined trusted relationships are expected to develop slowly over time. On the other hand

[3] It is arguably more correct to consider: $\text{Value}((a,b)) = \text{Value}(b) - \text{Value}(a)$, as β's actions may be influenced by his expectations of α's enactment of a — this additional complication is ignored.

[4] *Computable* in the sense that it is finitely computable, and hopefully not computationally complex.

they can be destroyed quickly by an agent whose actions unexpectedly fall below expectation. This highlights the importance of being able to foreshadow the possibility of untrustworthy behaviour.

$\tau_{\alpha\beta}(b)$ is predicated on α's ability to form an expectation of the value of β's future actions. This is related to the famous question posed by Laplace "what is the probability that the sun will rise tomorrow?". Assume that it has always previously been observed to do so and that there have been n prior observations. Then if the observer is in complete ignorance of the process he will assume that the probability distribution of a random variable representing the prior probability that the sun will rise tomorrow is the maximum entropy, uniform distribution on $[0,1]$. Further, using Bayes' theorem he will derive the posterior estimate $\frac{n+1}{n+2}$; the key assumption is that the observer is "in complete ignorance of the process". There may be many reasons why the sun may not rise such as the existence of a large comet on a collision trajectory with earth. These all important reasons are the *context* of the problem.

Laplace's naïve analysis above forms the basis of a very crude measure of trust. Suppose that the valuation space is: $\mathscr{V} = (\text{bad}, \text{good})$, and that α is considering signing contract (a,b) with β. Let the random variable B denote the value of β's next action. Then assume that nothing is known about the contract or about β except that this contract has been enacted by β on n prior occasions and that the valuation was "good" on s of those occasions. Using the maximum entropy prior distribution for B, $[0.5, 0.5]$, Bayes' theorem gives us a posterior distribution $[\frac{n-s+1}{n+2}, \frac{s+1}{n+2}]$. If at time t α signs the contract under consideration then the expected probability of a "good" valuation at time t'' is: $\frac{s+1}{n+2}$. This crude measure has little practical value although it readily extends to general discrete valuation spaces, and to continuous valuation spaces. The zero-information, maximum entropy distribution is the *trivial trust measure*. The crude Laplacian trust measure is in a sense the simplest non-trivial measure.

The weaknesses of the crude trust measure above show the way to building a reliable measure of trust. A reliable trust measure will include:

Prior knowledge. The use of the maximum entropy prior[5] is justified when there is absolutely no prior knowledge or belief of an agent's behaviour. In practical scenarios prior observations, reputation measures or the opinions of other agents are expected to be available and to be reflected in the prior.

Time. There is no representation of time. In the crude trust measure all prior observations have the same significance, and so an agent that used to perform well and is deteriorating may have the same trust measure as one that used to perform badly and is now performing well.

Context. There is no model of general events in the world or of *how* those events may affect an agent's behaviour. This includes modelling causality, namely *why* an agent might behave as it does.

[5] The maximum entropy prior expresses total uncertainty about what the prior distribution is.

29.3 Trust Model

The previous section defines trust as an optimistic[6] estimator of the expected value of future enactments, and concluded with three features of a reliable measure of trust. This section describes such a measure that uses the computational methods of information-based agents [10] particularly their information evaluation, acquisition and revelation strategies that ideally suits them to this purpose. Section 29.2 also described the fundamental role that the structure of the ontology plays in the trust model. This is described next followed by the core trust mechanism and then a reliable measure of trust.

29.3.1 Ontology

The structure of the ontology plays a central role in maintaining the trust model. Observations are propagated across the model moderated by their "semantic distance" from the concepts in the observation to nearby concepts.

Our agent communication language, U, is founded on three fundamental primitives: $\text{Commit}(\alpha,\beta,\varphi)$ to represent, in φ, the world that α aims at bringing about and that β has the right to verify, complain about or claim compensation for any deviations from, $\text{Observe}(\alpha,\varphi)$ to represent that a certain state of the world, φ, is observed, and $\text{Done}(u)$ to represent the event that a certain action u[7] has taken place. In our language, norms, contracts, and information chunks are represented as instances of $\text{Commit}(\cdot)$ where α and β can be individual agents or institutions, U is the set of expressions. $u \in U$ is defined as:

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

where φ_v is a formula with free variable v, $illoc$ is a predicate defining any appropriate set of illocutionary particles, ';' means sequencing, and *context* represents either previous agreements, previous illocutions, or code that aligns the ontological differences between the speakers needed to interpret an action u, and *term* represents logical predicates. t represents a point in time.[8] We denote by Φ the set of expressions φ used as the propositional content of illocutions.

---

[6] *Optimistic* in the sense that the estimation can be performed on the basis of the agent's interaction history.

[7] All actions are assumed to be dialogical.

[8] Usually omitted to simplify notation.

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

Offer(α, β,`spent`(β, α, `October`, X) ∧ X ≥ € 100 →
∀ y. Done(Inform(ξ, α, pay(β, α, y), November)) →
Commit(α, β, discount(y,10%)))

or, "If I tell you who I buy my tomatoes from then would you keep that information confidential?" as:

Offer(α, β, ∃δ. (Commit(α,β,Done(Inform(α,β,provider(δ,α,tomato)))) ∧
∀γ. ∀ t. Commit(β,α,¬Done(Inform(β,γ,provider(δ,α,tomato), t))))

In order to define the *terms* of the language introduced above (e.g. $\text{pay}(\beta,\alpha,y)$ or $\text{discount}(y,10\%)$) an ontology is required that includes a (minimum) repertoire of elements: a set of *concepts* (e.g. quantity, quality, material) organised in a is-a hierarchy (e.g. platypus is a mammal, australian-dollar is a currency), and a set of relations over these concepts (e.g. price(beer,AUD)).[9]

We model ontologies following an algebraic approach [6] as: an *ontology* is a tuple $\mathscr{O} = (C, R, \leq, \sigma)$ where:

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

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

The semantic distance between concepts plays a fundamental role in the estimation of trust. The concepts within an ontology are closer, semantically speaking, depending on how far away they are in the structure defined by the $\leq$ relation. Semantic distance plays a fundamental role in strategies for information-based agency. How signed contracts, $\text{Commit}(\cdot)$ about objects in a particular semantic region, and their execution $\text{Observe}(\cdot)$, *affect* our decision making process about signing future contracts on nearby semantic regions is crucial to modelling the common sense that human beings apply in managing trading relationships.

A measure [7] bases the *semantic similarity* between two concepts on the path length induced by $\leq$ (more distance in the $\leq$ graph means less semantic similarity), and the *depth* of the subsumer concept (common ancestor) in the shortest path between the two concepts (the deeper in the hierarchy, the closer the meaning of the concepts). [7] defines semantic similarity as:

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

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

where e is Euler's number (≈ 2.71828), l is the length (i.e. number of hops) of the shortest path between the concepts, h is the depth of the deepest concept subsuming both concepts, and κ_1 and κ_2 are parameters scaling the contribution of shortest path length and depth respectively. If $l = h = 0$ then $\mathrm{Sim}(c, c') = 1$; in general $\mathrm{Sim}(c, c') \in [0, 1]$.

29.3.2 The core trust mechanism

Section 29.2 ends with three essential components of a reliable trust model. Those three components will be dealt with in due course. This section describes the core trust estimation mechanism. In subsequent sections the core is enhanced with the three essential components. The final component, context, is unresolved as it relies on the solution to hard problems, such as modelling rare but significant contextual events, that are beyond the scope of this discussion.

The general idea is that whenever α evaluates $\mathrm{val}_{\alpha}^{t''}(b')$ for the enactment (a', b') of some previously signed contract (a, b) the trust estimates are updated. The contract space is typically very large and so estimates are not maintained for individual contracts; instead they are maintained for selected abstractions based on the ontology. Abstractions are denoted by the 'hat' symbol: e.g. $\hat{a}$. For example, "red wine orders for more that 24 bottles" or "supply of locally produced cheese". Whenever an evaluation $\mathrm{val}_{\alpha}^{t''}(b')$ is performed the trust estimates, $\tau_{\alpha\beta}(\hat{b})$, for certain selected nearby abstractions, $\hat{b}$, are updated.

In the absence of incoming information the integrity of an information-based agent's beliefs decays in time. In the case of the agent's beliefs concerning trust, incoming information is in the form of valuation observations $\mathrm{val}_{\alpha}^{t''}(b')$ for each enacted contract. If there are no such observations in an area of the ontology then the integrity of the estimate for that area should decay.

In the absence of valuation observations in the region of $\hat{b}$, $\tau_{\alpha\beta}(\hat{b})$ decays to a *decay limit distribution* $\overline{\tau_{\alpha\beta}(\hat{b})}$ (denoted throughout this section by 'overline'). The decay limit distribution is the zero-data distribution, but not the zero-information distribution because it takes account of reputation estimates and the opinions of other agents [13]. We assume that the decay limit distribution is known for each abstraction $\hat{b}$. At time s, given a distribution for random variable $\tau_{\alpha\beta}(\hat{b})^s$, and a decay limit distribution, $\overline{\tau_{\alpha\beta}(\hat{b})^s}$, $\tau_{\alpha\beta}(\hat{b})$ decays by:

$$\tau_{\alpha\beta}(\hat{b})^{s+1} = \Delta(\overline{\tau_{\alpha\beta}(\hat{b})^s}, \tau_{\alpha\beta}(\hat{b})^s)$$

where s is time and Δ is the *decay function* for the X satisfying the property that $\lim_{s\to\infty} \tau_{\alpha\beta}(\hat{b})^s = \overline{\tau_{\alpha\beta}(\hat{b})}$. For example, Δ could be linear:

$$\tau_{\alpha\beta}(\hat{b})^{s+1} = (1-\mu) \times \overline{\tau_{\alpha\beta}(\hat{b})^s} + \mu \times \tau_{\alpha\beta}(\hat{b})^s$$

where $0 < \mu < 1$ is the decay rate.

We now consider what happens when valuation observations are made. Suppose that at time s, α evaluates β's enactment b' of commitment b, $\mathrm{val}^s_\alpha(b') = v_k \in \mathscr{V}$. The update procedure updates the probability distributions for $\tau_{\alpha\beta}(\hat{b})^s$ for each $\hat{b}$ that is "moderately close to" b. Given such a $\hat{b}$, let $\mathbb{P}^s(\tau_{\alpha\beta}(\hat{b}) = v_k)$ denote the prior probability that v_k would be observed. The update procedure is in two steps. First, estimate the posterior probability that v_k would be observed, $\mathbb{P}^{s+1}(\tau_{\alpha\beta}(\hat{b}) = v_k)$ for the particular value v_k. Second, update the entire posterior distribution for $\tau_{\alpha\beta}(\hat{b})$ to accommodate this revised value.

Given a $\hat{b}$, to revise the probability that v_k would be observed three things are used: the observation: $\mathrm{val}^s_\alpha(b')$, the prior: $\mathbb{P}^s(\tau_{\alpha\beta}(\hat{b}) = v_k)$, and the decay limit value: $\mathbb{P}^s(\overline{\tau_{\alpha\beta}(\hat{b})} = v_k)$. The observation $\mathrm{val}^s_\alpha(b')$ may be represented as a probability distribution with a '1' in the k'th place and zero elsewhere, $\mathbf{u}_k$. To combine it with the prior its significance is discounted for two reasons:

- b may not be semantically close to $\hat{b}$, and
- $\mathrm{val}^s_\alpha(b') = v_k$ is a single observation whereas the prior distribution represents the accumulated history of previous observations.

to discount the significance of the observation $\mathrm{val}^s_\alpha(b') = v_k$ a value is determined in the range between '1' and the zero-data, decay limit value $\mathbb{P}^s(\overline{\tau_{\alpha\beta}(\hat{b})} = v_k)$ by:

$$\delta = \mathrm{Sim}(b,\hat{b}) \times \kappa + (1 - \mathrm{Sim}(b,\hat{b}) \times \kappa) \times \mathbb{P}^s(\overline{\tau_{\alpha\beta}(\hat{b})} = v_k)$$

where $0 < \kappa < 1$ is the learning rate, and $\mathrm{Sim}(\cdot,\cdot)$ is a semantic similarity function such as that shown in Equation 29.3.1. Then the posterior estimate $\mathbb{P}^{s+1}(\tau_{\alpha\beta}(\hat{b}) = v_k)$ is given by:

$$\mathbb{P}^{s+1}(\tau_{\alpha\beta}(\hat{b}) = v_k) = \frac{\rho\delta(1-\omega)}{\rho\delta(1-\omega) + (1-\rho)(1-\delta)\omega} = \nu$$

where $\rho = \mathbb{P}^s(\tau_{\alpha\beta}(\hat{b}) = v_k)$ is the prior value, and $\omega = \mathbb{P}^s(\overline{\tau_{\alpha\beta}(\hat{b})} = v_k)$ is the decay limit value.

It remains to update the entire posterior distribution for $\tau_{\alpha\beta}(\hat{b})$ to accommodate the constraint $\mathbb{P}^{s+1}(\tau_{\alpha\beta}(\hat{b}) = v_k) = \nu$. Information-based agents [10] employ a standard procedure for updating distributions, $\mathbb{P}^t(X = x)$ subject to a set of linear constraints on X, $c(X)$, using:

$$\mathbb{P}^{t+1}(X = x | c(X)) = \mathrm{MRE}(\mathbb{P}^t(X = x), c(X))$$

where the function MRE is defined by: $\mathrm{MRE}(\mathbf{q},\mathbf{g}) = \arg\min_{\mathbf{r}} \sum_j r_j \log \frac{r_j}{q_j}$ such that $\mathbf{r}$ satisfies $\mathbf{g}$, $\mathbf{q}$ is a probability distribution, and $\mathbf{g}$ is a set of n linear constraints $\mathbf{g} = \{g_j(\mathbf{p}) = \mathbf{a_j} \cdot \mathbf{p} - c_j = 0\}, j = 1,\ldots,n$ (including the constraint $\sum_i p_i - 1 = 0$).

The resulting $\mathbf{r}$ is the *minimum relative entropy distribution*[10] [8]. Applying this procedure to $\tau_{\alpha\beta}(\hat{b})$:

$$\mathbb{P}^{s+1}(\tau_{\alpha\beta}(\hat{b}) = v) = \mathrm{MRE}(\mathbb{P}^{s}(\tau_{\alpha\beta}(\hat{b}) = v), \mathbb{P}^{s+1}(\tau_{\alpha\beta}(\hat{b}) = v_k) = \nu)$$

where ν is the value given by Equation 29.3.2.

Whenever α evaluates an enactment $\mathrm{val}^{s}_{\alpha}(b')$ of some commitment b, the above procedure is applied to update the distributions for $\mathbb{P}(\tau_{\alpha\beta}(\hat{b}) = v)$. It makes sense to limit the use of this procedure to those distributions for which $\mathrm{Sim}(b, \hat{b}) > y$ for some threshold value y.

29.3.3 Prior knowledge

The decay-limit distribution plays a key role in the estimation of trust. It is not directly based on any observations and in that sense it is a "zero data" trust estimate. It is however not "zero information" as it takes account of opinions and reputations communicated by other agents [13]. The starting point for constructing the decay-limit distribution is the maximum entropy (zero-data, zero-information) distribution. This gives a two layer structure to the estimation of trust: opinions and reputations shape the decay-limit distribution that in turn plays a role in forming the trust estimate that takes account of observed data. Communications from other agents may not be reliable. α needs a means of estimating the reliability of other agents before they can be incorporated into the decay-limit distribution — reliability is discussed at the end of this section.

Reputation is the opinion (more technically, a social evaluation) of a group about something. So a group's reputation about a thing will be related in some way to the opinions that the individual group members hold towards that thing. An *opinion* is an assessment, judgement or evaluation of something. Opinions are represented in this section as probability distributions on a suitable ontology that for convenience is identified with the *evaluation space* $\mathscr{V}$. That is, opinions communicated by β concerning another agent's trustworthiness are assumed to be expressed as predicates using the same valuation space as $\mathscr{V}$ over which α represents its trust estimates.

An opinion is an evaluation of an *aspect* of a thing. A rainy day may be evaluated as being "bad" from the aspect of being suitable for a picnic, and "good" from the aspect of watering the plants in the garden. An aspect is the "point of view" that an agent has when forming his opinion. An opinion is evaluated in context. The *context* is everything that the thing is being, explicitly or implicitly, evaluated with or against. The set of valuations of all things in the context calibrates the valuation space; for example, "this is the best paper in the conference". The context can be

[10] This may be calculated by introducing Lagrange multipliers λ: $L(\mathbf{p}, \lambda) = \sum_j p_j \log \frac{p_j}{q_j} + \lambda \cdot \mathbf{g}$. Minimising L, $\{\frac{\partial L}{\partial \lambda_j} = g_j(\mathbf{p}) = 0\}, j = 1, \ldots, n$ is the set of given constraints $\mathbf{g}$, and a solution to $\frac{\partial L}{\partial p_i} = 0, i = 1, \ldots, I$ leads eventually to $\mathbf{p}$.

vague: "of all the presents you could have given me, this is the best". If agents are to discuss opinions then they must have some understanding of each other's context.

Summarising the above, an *opinion* is an agent's evaluation of a particular aspect of a thing in context. A representation of an opinion will contain: the thing, its aspect, its context, and a distribution on $\mathscr{V}$ representing the evaluation of the thing. α acquires opinions and reputations through communication with other agents. α estimates the reliability of those communicating agents before incorporating that information into the decay-limit distributions. The basic process is the same for opinions and reputations; the following describes the incorporation of opinions only.

Suppose agent β' informs agent α of his opinion of the trustworthiness of another agent β using an utterance of the form: $u = \texttt{inform}(\beta', \alpha, \tau_{\beta'\beta}(b))$, where conveniently b is in α's ontology. This information may not be useful to α for at least two reasons: β' may not be telling the truth, or β' may have a utility function that differs from α's. We will shortly estimate β''s "reliability", $R^t_\alpha(\beta')$ that measures the extent to which β' is telling the truth and that α and β' "are on the same page" or "think alike"[11]. Precisely, $0 < R^t_\alpha(\beta') < 1$; its value is used to moderate the effect of the utterance on α's decay-limit distributions. The estimation of $R^t_\alpha(\beta')$ is described below.

Suppose that α maintains the decay limit distribution $\overline{\tau_{\alpha\beta}(\hat{b})^s}$ for a chosen $\hat{b}$. In the absence of utterances informing opinions of trustworthiness, $\overline{\tau_{\alpha\beta}(\hat{b})^s}$ decays to the distribution with maximum entropy. As previously this decay could be linear:

$$\overline{\tau_{\alpha\beta}(\hat{b})^{s+1}} = (1-\mu) \times \mathrm{MAX} + \mu \times \overline{\tau_{\alpha\beta}(\hat{b})^s}$$

where $\mu < 1$ is the decay rate, and MAX is the maximum entropy, uniform distribution.

When α receives an utterance of the form u above, the decay limit distribution is updated by:

$$\begin{aligned}\overline{\tau_{\alpha\beta}(\hat{b})^{s+1}} \mid \texttt{inform}(\beta', \alpha, \tau_{\beta'\beta}(b)) = \\ \left(1 - \kappa \times \mathrm{Sim}(\hat{b}, b) \times R^s_\alpha(\beta')\right) \times \overline{\tau_{\alpha\beta}(\hat{b})^s} \\ +\kappa \times \mathrm{Sim}(\hat{b}, b) \times R^s_\alpha(\beta') \times \tau_{\beta'\beta}(b)\end{aligned}$$

where $0 < \kappa < 1$ is the learning rate and $R^s_\alpha(\beta')$ is α estimate of β''s reliability. It remains to estimate $R^s_\alpha(\beta')$.

Estimating $R^s_\alpha(\beta')$ is complicated by its time dependency. First, in the absence of input of the form described following, $R^s_\alpha(\beta')$ decays to zero by: $R^{s+1}_\alpha(\beta') = \mu \times R^s_\alpha(\beta')$. Second, describe how $R^s_\alpha(\beta')$ is increased by comparing the efficacy of $\overline{\tau_{\alpha\beta}(\hat{b})^s}$ and $\tau_{\beta'\beta}(b)^s$ in the following interaction scenario. Suppose at a time s, α is considering signing the contract (a, b) with β. α requests β''s opinion of β with respect to b, to which β may respond $\texttt{inform}(\beta', \alpha, \tau_{\beta'\beta}(b))$. α now has two

[11] The reliability estimate should perhaps also be a function of the commitment, $R^t_\alpha(\beta', b)$, but that complication is ignored.

estimates of β's trustworthiness: $\overline{\tau_{\alpha\beta}(\hat{b})^s}$ and $\tau_{\beta'\beta}(b)^s$; $\overline{\tau_{\alpha\beta}(\hat{b})^s}$ and $\tau_{\beta'\beta}(b)^s$ are both probability distributions that each provide an estimate of $\mathbb{P}^s(\text{Value}_\beta(b) = v_i)$ for each valuation v_i. α increases its reliability estimate of β if the trust estimate in β's `inform` is 'better' than α's current decay limit value. Suppose that α signs the contract (a,b) at time t, and at some later time t'' evaluates β's enactment $\text{val}_\alpha^{t''}(b') = v_k$, say. Then:

$$\mathbb{P}(\tau_{\beta'\beta}(b)^s = v_k) > \mathbb{P}(\overline{\tau_{\alpha\beta}(\hat{b})^s} = v_k)$$

and β''s trust estimate is better than α's; α increases $R_\alpha^s(\beta')$ using:

$$R_\alpha^{s+1}(\beta') = \kappa + (1-\kappa) \times R_\alpha^s(\beta')$$

where $0 < \kappa < 1$ is the learning rate.

29.3.4 Time

The core trust mechanism and the prior knowledge both give greater weight to recent observations than to historic data. This may be a reasonable default assumption but has no general validity. Trust, $\tau_{\alpha\beta}(\hat{b})^s$, estimates *how* β is expected to act. If an agent is considering repeated interaction with β then he may also be interested in how β's actions are expected to *change* in time.

The way in which the trust estimate is evolving is significant in understanding which agents to interact with. For example, an agent for whom $\tau_{\alpha\beta}^s(\hat{b})$ is fairly constant in time may be of less interest than an agent who is slightly less trustworthy but whose trust is consistently improving. To capture this information something like the finite derivative is required: $\frac{\delta}{\delta s}\tau_{\alpha\beta}^s(\hat{b})$. The sum of the elements in such a vector will be zero, and in the absence of any data it will decay to the zero vector.

Estimating the rate of change of $\tau_{\alpha\beta}^s(\hat{b})$ is complicated by the way it evolves that combines continual integrity decay with periodic updates. Evolution due to decay tells us nothing about the rate of change of an agent's behaviour. Evolution caused by an update is performed following a period of prior decay, and may result in compensating for it. Further, update effects will be very slight in the case that the commitment b is semantically distant from $\hat{b}$. In other words, the evolution of $\tau_{\alpha\beta}^s(\hat{b})$ itself is not directly suited to capturing the rate of change of agent behaviour.

The idea for an indirect way to estimate how β's actions are evolving comes from the observation that $\tau_{\alpha\beta}(\hat{b})^s$ is influenced more strongly by more recent observations, and the extent to which this is so depends on the decay rate. For example. if the decay rate is zero then $\tau_{\alpha\beta}(\hat{b})^s$ is a time-weighted "average" of prior observations. Suppose that $\tau_{\alpha\beta}(\hat{b})^s$ has been evaluated. We perform a parallel evaluation using a lower decay rate to obtain $\tau_{\alpha\beta}^-(\hat{b})^s$, then $\tau_{\alpha\beta}(\hat{b})^s - \tau_{\alpha\beta}^-(\hat{b})^s$ is a vector the sum of whose elements is zero, and in which a positive element indicates a value that is presently "on the increase" compared to the historic average.

The preceding method for estimating change effectively does so by calculating a first difference. If another first difference is calculated using an even lower decay rate then calculate a second difference to estimate the *rate of* change. This may be stretching the idea too far!

29.3.5 Trust in Context

The informal meaning of context is information concerning everything in the environment that could affect decision making *together with* rules that link that information to the deliberative process. That is, *context* consists of facts about the environment, including rare but significant events, *and* rules that link those facts to the agent's reasoning. Those rules typically rely on common sense reasoning. Dealing with context is a hard problem for intelligent agents generally and for their management of trust estimates in particular.

From an artificial intelligence point of view, artificial agents lack the skills of their human counterparts for dealing with context. Humans then rely on common sense and experience to learn how to key contextual information to their deliberation, and to identify incompleteness in their knowledge. For artificial agents; identifying and dealing with inconsistency and incompleteness is a hard problem, and so is keying general information to their own deliberative apparatus.

Even if 'trust in context' is narrowed to just one issue "is there any reason to distrust our trust estimate due to a *change* in context?" the problems remain hard. Supposing that α is considering signing a contract (a,b) at time t, to address this issue the following are required:

1. knowledge of the context of previous observations of behaviour. Their *context* is the state of each of the observables in the environment and of the states of the other agents when those previous observations of behaviour were made.
2. founded beliefs concerning the context that will pertain at the future time of the evaluation of the presumed future behaviour — i.e. at time t'' in Figure 29.1.
3. some reasoning apparatus that enables us to decide whether *differences* between the believed future context and the observed previous contexts cause us to modify our experience-based trust estimate.

The information-based architecture makes a modest contribution to trust in context in the following sense. An agent builds up a sense of trust on the basis of its own past experience and statements of opinion and reputation from other agents. In a sense those statements of opinions and reputation are contextual information for the business of estimating trust. It also moderates its trust estimates through the persistent decay of contextual information integrity by Equation 29.3.2. Beyond that no 'magic bullet' solutions are given to the contextual problems described above and the discussion is left as a pointer to the work that is required to increase the reliability of trust estimation in dynamic environments.

29.4 Relationship Model

The trust model described in Section 29.3 is a summary of the history of interaction between α and β, $H_{\alpha\beta}$, augmented by reputation estimates. Reputation estimates *per se* are outside the α's direct experience and are therefore part of the *context* of α's trust. Trust is not the only way in which the interaction history may be usefully summarised. The *relationship model* contains summary estimates that include trust. Before describing these measures human relationships are examined particularly ways in which they are summarised. This leads to a discussion of the formal representation of relationships using the LOGIC framework.

29.4.1 Relationships

A *relationship* between two human or artificial agents is their *interaction history* that is a complete record of their interactions evaluated *in context*. There is evidence from psychological studies that humans seek a *balance* in their negotiation relationships. The classical view [2] is that people perceive resource allocations as being distributively fair (i.e. well balanced) if they are proportional to inputs or contributions (i.e. equitable). However, more recent studies [14, 15] show that humans follow a richer set of norms of distributive justice depending on their *intimacy* level: equity, equality, and need. Here *equity* is allocation proportionally 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).

We believe that the perception of balance in dialogues, especially in negotiation, 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. The more intimacy 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 structure[12] to a shelter-based one[13] the probability of survival increased [14].

In this context, for example, 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. For example, a mother does not expect reciprocity when explaining everything to her children, or buying toys for them. In the case of partners there is

[12] 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.

[13] 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.

some evidence [4] that the allocations of goods and burdens (i.e. positive and negative utilities) are perceived as fair, or in balance, based on equity for burdens and equality for goods.

The 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 every time requested information 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, issues that are causing a burden to our opponent can be identified by observing an imbalance in their information or utilitarian utterances on that issue.

We assume that the interactions between agents can be organised into dialogues, where a *dialogue* is a set of related utterances. This section is concerned with *commitment dialogues* that contain at least one commitment, where a commitment may simply be the truth of a statement or may be a contractual commitment. We assume that all commitment dialogues take place in some or all of the following five stages:

1. the *prelude* during which agents prepare for the interaction
2. the *negotiation* that may lead to
3. *signing* a contract at time t
4. the *enactment* of the commitments in the contract at time t'
5. the *evaluation* at time t'' of the complete interaction process that is made when the goods or services acquired by enactment of the contract have been consumed

The notation of a commitment dialogue is broad in that a dialogue that does not contain any sort of commitment is arguably of little interest.

A major issue in building models of dialogues and relationships is dealing with the reliability of the utterances made. For an information-based agent the *reliability* of an utterance is an epistemic probability estimate of the utterance's veracity. For example, if the utterance is an `inform` containing a proposition then its reliability is an estimate of the probability that the proposition is correct. If the utterance is an `opinion` then its reliability is an estimate of the probability that the `opinion` will in time be judged to be sound. The difficulty with estimating reliability is that it may take months or years for an agent to be able to say: "Ah, that was good advice". Reliability is a measure attached to an utterance, and integrity is a measure attached to a complete dialogue. A blanket estimation of the reliability of an agent was described in Section 29.3.3.

29.4.2 The LOGIC Framework

The LOGIC illocutionary framework for classifying argumentative interactions was first described in [11] where it was used to help agents to prepare for a negotiation in the *prelude stage* of an interaction as described above. This section generalises that framework and uses it to define one of the two dimensions of the relationship model described below, the second dimension is provided by the structure of the ontology as specified by a partial order $\leq$ defined by the is-a hierarchy, and a distance measure between concepts such as Equation 29.3.1. The five LOGIC categories for information are quite general:

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

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

Taking a more formal view, the LOGIC framework categorises information in an utterance by its relationship to:

L $= \{B(\alpha, \varphi)\}$, that is a set of *beliefs*, communicated by: `inform`.
O $= \{\text{Accept}(\beta, \alpha, c)\}$, that is a set of *acceptable contracts*, communicated by: `offer`, `reject` and `accept`.
G $= \{D(\alpha, \varphi)\}$, that is a set of *needs* or *desires*, communicated by: `Ineed`.
I $= \{\text{Can}(\alpha, \text{Do}(p))\}$, that is a set of *capabilities*, communicated by: `canDo`.
C $= \{I(\alpha, \text{Do}(p))\} \cup \{\text{Commit}(\alpha, \text{Do}(p))\}$, that is a set of *commitments* and *intentions*, communicated by: `commit` (for future commitments), and `intend` (commitments being enacted).

Four predicates L, O, G, I and C recognise the category of an utterance. Information in an `inform` utterance is categorised as Goals, Independence and Commitments if the `inform` contains the illocutions listed above: `Ineed`, `canDo`, `commit` and `intend`. Otherwise it is categorised as Legitimacy.

Given a need ν and an agent β the variables $L^t_{\nu\beta}$, $O^t_{\nu\beta}$, $G^t_{\nu\beta}$, $I^t_{\nu\beta}$ and $C^t_{\nu\beta}$ are aggregated from observations of how forthcoming β was during prior dialogues. They are then used to form α's expectation of β's future readiness to reveal private information across the five LOGIC categories. They are updated at the end of each dialogue[14] using a linear form that is consistent with [5] for the human brain in a volatile environment.

In the following a dialogue Γ commences at time $t-s$ and terminates at time t when the five variables are updated. $t-d$ denotes the time at which these variables were previously updated. For convenience assume that $d \geq s$. Γ aims to satisfy need

[14] This is for efficiency. Updating the model following each utterance could expend resources to little effect.

ν. All the estimates given below are for the effect of Γ on variables for a nearby need ν' for which $\eta' = \eta \times \text{Sim}(\nu, \nu')$, η is the learning rate, and μ the decay rate.

$L^t_{\nu\beta}$ measures the amount of information in β's Legitimacy `inform` utterances. The procedure by which `inform` utterances update $\mathscr{M}^t$ is described in [10]. The Shannon information in a single `inform` statement, u, is: $\mathbb{I}(u) = \mathbb{H}(\mathscr{M}^{t-1}) - \mathbb{H}(\mathscr{M}^t|u)$. It is defined in terms of the contents of $\mathscr{M}^t$, and so the valuation is restricted to 'just those things of interest' to α. During Γ observe: $l = \sum_{u \in \Gamma, \mathsf{L}(u)} \mathbb{I}(u)$. Then update $L^t_{\nu'\beta}$ with:

$$L^t_{\nu'\beta} = \eta' \sum_{u \in \Gamma, \mathsf{L}(u)} \mathbb{I}(u) + (1 - \eta')\mu^d L^{t-d}_{\nu'\beta}$$

$O^t_{\nu\beta}$ measures the amount of information β reveals about the deals he will accept. β's limit contracts were modelled on the basis of observed behaviour in [10]. Let random variable Y over contract space $\mathscr{C}$ denote α's beliefs that a contract is a limit contract for β. The information gain in Y during Γ is: $\mathbb{H}^{t-s}(Y) - \mathbb{H}^t(Y)$, and $O^t_{\nu'\beta}$ is updated by:

$$O^t_{\nu'\beta} = \eta' \left(\mathbb{H}^{t-s}(Y) - \mathbb{H}^t(Y)\right) + (1 - \eta')\mu^d O^{t-d}_{\nu\beta}$$

$G^t_{\nu\beta}$ measures the information β reveals about his goals, and $I^t_{\nu\beta}$ about his suggested capabilities. $G^t_{\nu\beta}$ and $I^t_{\nu\beta}$ are similar in that both `Ineed` and `canDo` preempt the terms of a contract. Suppose β informs α that: $\texttt{Ineed}(\nu)$ and $\texttt{canDo}(\delta)$. If β is being forthcoming then this suggests that he has in mind an eventual contract (a, b) in which $a \le \nu$ and $b \le \delta$ (using $\le$ from the ontology). Suppose that Γ leads to the signing of the contract (a, b) then observe: $g = \text{Sim}(a, \nu)$ and $i = \max_\delta \text{Sim}(b, \delta)$; $\max_\delta$ is in case β utters more than one `canDo`. $G^t_{\nu'\beta}$ is aggregated by:

$$G^t_{\nu'\beta} = \eta' \text{Sim}(a, \nu) + (1 - \eta')\mu^d G^{t-d}_{\nu'\beta}$$

Similarly: $I^t_{\nu'\beta} = \eta' \max_\delta \text{Sim}(b, \delta) + (1 - \eta')\mu^d I^{t-d}_{\nu'\beta}$.

$C^t_{\nu\beta}$ measures the amount of information β reveals about his commitments and intentions. These are measured just as for $L^t_{\nu\beta}$ by aggregating the observation: $c = \sum_{u \in \Gamma, \mathsf{C}(u)} \mathbb{I}(u)$, and $C^t_{\nu'\beta}$ is updated by:

$$C^t_{\nu'\beta} = \eta' \sum_{u \in \Gamma, \mathsf{C}(u)} \mathbb{I}(u) + (1 - \eta')\mu^d C^{t-d}_{\nu'\beta}$$

The measures described above are based on what β says. In negotiation what was *not* said but *could have* been said may be equally significant. A *confidentiality* measure described in [12] addresses this issue.

In addition, if $\text{val}_\alpha(\cdot)$ is α's utilitarian evaluation function that is used to evaluate both the contract and the enactment *in context* then the observations

$$v^t(\Gamma) = \mathrm{val}_\alpha((a',b')|H^t) - \mathrm{val}_\alpha((a,b)|H^{t-s})$$

update the variable $U^t_{\nu\beta}$ that estimates *utility gain* during Γ:

$$U^t_{\nu'\beta} = \eta' \left(\mathrm{val}_\alpha((a',b')|H^t) - \mathrm{val}_\alpha((a,b)|H^{t-s})\right) + (1-\eta')\mu^d U^{t-d}_{\nu'\beta}$$

Finally the LOGIC evaluation of a complete dialogue is assembled. Putting the six measures together define α's evaluation function, $\mathrm{logic}^t(\Gamma)$, for a complete dialogue Γ in which the contract (a,b) is signed. With notation as above:

$$\mathrm{logic}^t(\Gamma) = \Bigg(\sum_{u\in\Gamma,\mathsf{L}(u)} \mathbb{I}(u), \mathbb{H}^{t-s}(Y) - \mathbb{H}^t(Y), \mathrm{Sim}(a,\nu), \max_\delta \mathrm{Sim}(b,\delta), \\ \sum_{u\in\Gamma,\mathsf{C}(u)} \mathbb{I}(u), \mathrm{val}_\alpha((a',b')|H^t) - \mathrm{val}_\alpha((a,b)|H^{t-s}) \Bigg)$$

We model our expectation of observing any particular value $\mathrm{logic}^t(\Gamma)$ with the six-dimensional random variable $\mathbf{E}^t_{\nu\beta}$ where $(E^t_{\nu\beta})_k$ is the expectation for $L^t_{\nu'\beta}$, $O^t_{\nu'\beta}$, $G^t_{\nu'\beta}$, $I^t_{\nu'\beta}$, $C^t_{\nu'\beta}$, $U^t_{\nu'\beta}$ respectively, $k = 1,\ldots,6$.

29.5 Negotiation

If α prefers to deal with trusted partners then because trust is established by interaction α needs to determine the pool of agents to interact with who are then potential negotiation partners for each generic need. If the pool is large then the integrity of the trust estimates will be low, and if the pool is small then α may deny itself access to new partners. The *pool selection* problem is to manage the size and composition of the pool of partners for each generic need so as to balance these conflicting values. Pool selection is addressed followed by the offer strategy, and finally the strategic use of argumentation to build strong and trusted relationships.

29.5.1 Pool Selection

The aim of the pool *selection* phase is to select a strategically *diverse* pool of agents, $\mathscr{P}_\nu$, for each of α's needs ν. Let $\mathscr{B}_n$ be the set of n-element subsets of $\{\beta_1,\ldots,\beta_o\}$, then

$$\mathscr{P}_\nu = \arg\max_n \{B \in \mathscr{B}_m \mid \\ \forall bb' \in B : \mathbb{P}((\mathbf{E}^t_{\nu b})_k > e_k) > c_k, \mathbb{H}((\mathbf{E}^t_{\nu b})_k) < h_k, \mathrm{div}(b,b') > d\}$$

where: **e**, **h**, **c** and d are selected constants, $k = 1, \ldots, 6$, and $\text{div}(\beta_i, \beta_j)$ is a measure of: geographic, political, economic and/or functional *agent diversity*. Suppose that α's *needs model* is such that the probability that need ν is triggered at any time is ε_ν.

The *uniform selection strategy* selects an agent from $\mathscr{P}_\nu$ when ν is triggered by: $\mathbb{P}^t(\text{Select } \beta|\nu) = \frac{1}{n}$, and each $\beta \in \mathscr{P}_\nu$ expects to be selected each $m = \frac{1}{n\varepsilon_\nu}$ time steps. If β is selected at time $t - s$ and if the value $\text{logic}^t(\Gamma)$ is observed for the resulting negotiation dialogue then:

$$\mathbb{P}(\mathbf{E}^t_{\nu\beta} = \text{logic}^t(\Gamma)) = \eta + (1-\eta) \times \mathbb{P}(\mathbf{E}^{t-1}_{\nu\beta} = \text{logic}^t(\Gamma))$$

The full distribution for $\mathbf{E}^t_{\nu\beta}$ is then calculated using the MRE (minimum relative entropy) process described in Section 29.3.2 using Equation 29.5.1 as the constraint. By time $t+m$ full distribution for $\mathbf{E}^t_{\nu\beta}$ will have decayed in line with Equation 29.3.2 and:

$$\begin{aligned}&\mathbb{P}(\mathbf{E}^{t+m}_{\nu\beta} = \text{logic}^t(\Gamma)) = \\ &\qquad (1-\mu^m)\mathbb{P}(\overline{\mathbf{E}^t_{\nu\beta}} = \text{logic}^t(\Gamma)) + \mu^m(\eta + (1-\eta)\mathbb{P}(\mathbf{E}^{t-1}_{\nu\beta} = \text{logic}^t(\Gamma)))\end{aligned}$$

To ensure decreasing entropy: $\mathbb{P}(\mathbf{E}^{t+m}_{\nu\beta} = \text{logic}^t(\Gamma)) > \mathbb{P}(\mathbf{E}^{t-1}_{\nu\beta} = \text{logic}^t(\Gamma))$. Suppose $\mathbf{p} = \mathbb{P}(\overline{\mathbf{E}^t_{\nu\beta}} = \text{logic}^t(\Gamma))$ and $\mathbb{P}(\mathbf{E}^{t-1}_{\nu\beta} = \text{logic}^t(\Gamma)) = \kappa \cdot \mathbf{p}$; i.e. expect $\kappa_k > 1$ and $\kappa_k \mathbf{p}_k < 1$ for $k = 1, \ldots, 6$. Let $\kappa = \kappa_k$ and $p = \mathbf{p}_k$ for some value of k. Then the expected least value of m to prevent integrity decay is such that: $\kappa p = (1-\mu^m)p + \mu^m(\eta + (1-\eta)\kappa p)$, and so:

$$m = \frac{\log(p(\kappa-1)) - \log(\eta - p + (1-\eta)\kappa p)}{\log \mu}$$

E.g. suppose $p = 0.2$, $\kappa = 3$, $\eta = 0.7$, $\mu = 0.98$ then $m = 26$. Alternatively, solving for η: $\eta = \frac{(\kappa-1)\left(1-\mu^{-m}\right)p}{\kappa p - 1}$. The lower limit for m in Equation 29.5.1 and a value for ε_ν gives an upper limit for n the size of the pool $\mathscr{P}_\nu$.

A *stochastic selection strategy* selects an agent from $\mathscr{P}_\nu$ when ν triggers by:

$$\mathbb{P}^t(\text{Select } \beta|\nu) = \mathbb{P}(\mathbf{E}^t_{\nu\beta} \gg)$$

where $\mathbb{P}^t(\mathbf{E}^t_{\nu\beta} \gg)$ denotes the probability that β is better than for all the others in the following sense. If the six-dimensional sample space for $E^t_{\nu\beta}$ is linearly ordered in increasing degree of satisfaction, then the probability that the evaluation of a dialogue for ν with β will be more satisfactory that, β' $\mathbb{P}^t(\mathbf{E}^t_{\nu\beta} > \mathbf{E}^t_{\nu\beta'})$, may be estimated. To prevent integrity decay of $\mathscr{P}_\nu$ for this strategy repeat the calculation above for the worst choice for ν that will expect to be selected every: $m' = \frac{1}{\varepsilon_\nu}\mathbb{P}(\mathbf{E}^t_{\nu\beta} \ll)$ time steps. For any stochastic strategy denote $\mathbb{P}^t(\text{Select } \beta|\nu)$ by $\mathbb{P}^t(S_\nu = s_{\nu,i})$ for random variable S_ν then $\mathbb{H}(S_\nu)$, $i = 1, \ldots, n$, measures *selection strategy diversity*, or normalised as: $\frac{1}{\log n}\mathbb{H}(S_\nu)$ [1].

29.5.2 Offer Strategy

The previous section analysed the intuition if an agent maintains too great a choice of trading partners then its certainty in their behaviour will decay — no matter whether their behaviour is good or bad. Having determined which negotiation partners to interact with the *offer strategy* that determines what offers to make is considered, and so this section is concerned with the options component of the LOGIC model, $O^t_{v\beta}$. In the following Section 29.5.3 considers *argumentation*, that 'wraps' utterances with rhetorical argumentation, and will address the remaining LOGIC components.

α is assumed to have a utilitarian negotiation strategy [9] that the following ideas are intended to embellish. That strategy may reference the estimate that β will accept the contract (a,b): $\mathbb{P}^t(\text{Accept}(\beta,\alpha,(a,b)))$ — an estimate[15] is derived in [10]. This leads to a variation of the issue-tradeoff strategy where α makes the offer that is acceptable to her that β is most likely to accept. If $\text{Accept}^t(\alpha,\beta,c)$ denotes that c is acceptable to α then offer c^* where:

$$c^* = \arg\max_c \{\mathbb{P}^t(\text{Accept}(\beta,\alpha,c)) \mid \text{Accept}^t(\alpha,\beta,c)\}$$

Setting utilitarian considerations aside for a moment estimate which offer to make for which β's response, accept or reject, gives α greatest information gain. If β was prepared to answer repeated questions of the form then "Is contract y acceptable to you?" then the expected shortest question sequence has a Shannon encoding that is optimum with respect to the prior expectation of offer acceptance.

We show that if there is one issue and if the prior is the maximum entropy distribution then the sequence with greatest information gain will select the 'mid-value' at each stage. Denote β's expected limit contract by random variable Y. Suppose Y's sample space is $(0,\dots,n)$ and β's preferences are known to be monotonic increasing over this space with n known to be acceptable and 0 known to be unacceptable. The prior for Y is the maximum entropy distribution over $(1,\dots,n)$ with $\mathbb{H}(Y) = \log_2 n$, and $\mathbb{P}(\text{Accept}(\beta,\alpha,y) = \frac{y}{n}$. If β reports that y is acceptable then $\mathbb{H}(Y|y \text{ acceptable}) = \log_2 y$, and the information gain is $\log_2[\frac{n}{y}]$. Likewise $\mathbb{H}(Y|y \text{ unacceptable}) = \log_2(n-y)$. Solving the continuous model for maximal expected information gain:

$$\frac{d}{dy}\left(y\log_2\frac{n}{y} + (n-y)\log_2\frac{n}{n-y}\right) = 1 - \log_2\frac{n}{n-y} = 0, \text{ and } y = \frac{n}{2}$$

Consideration of the offer with maximal expected information gain is more interesting in multi-issue negotiation where α may have a set of potential offers D_v all with similar material value v, and may then wish to prioritise them on the basis

[15] If α assumes the each dimension of the contract space may be ordered to reflect β's preferences and interprets β's illocutionary actions of `offer` as willingness to *accept* whilst *rejecting* α's previous offers then a probabilistic model of β's limit contracts is derived using maximum entropy inference.

of expected information gain. Given an estimate for $\mathbb{P}(\text{Accept}(\beta,\alpha,y)), y \in D_\nu$ (see [10]) the preceding ideas may be used to enumerate the expected information gain for each $y \in D_\nu$ and so to make the maximal offer.

This section takes both utilitarian gain and information gain into account in managing the offer sequence. Within a single negotiation dialogue utilitarian gain is what matters most. Information gain on the other hand is concerned with strengthening the relationship and trust models and so underpins the agent's long-term strategies to build secure trading relationships for the future. Information-based agents aim to strike a balance between short term gains and long term security.

29.5.3 Argumentation and Relationship Building

This section is concerned with trust and relationships between agents. Relationships are built through dialogical interaction that is modelled using the LOGIC framework. Argumentation strategies take account of bluff and counter-bluff in the cut and thrust of competitive interaction, and contribute to relationships in each of the five LOGIC categories. We discuss what an argumentation strategy should aim to achieve from the LOGIC model point of view — the construction of the illocutionary sequences to achieve this aim is beyond the scope of this discussion.

Rhetoric argumentation aims to alter the beliefs of the recipient; it is also an information acquisition and revelation process as measured using the LOGIC framework. Equation 29.4.2 is α's evaluation function that applies to both contract enactment and argumentation, it is *also* the basis for α's *relationship-building strategies* that aim to influence the strength of β's relationship through argumentation and offer acceptance.

For each generic need ν α maintains a pool of potential partners (Section 29.5), and for each negotiation partner β, α has a model of their relationship summarised as: $L^t_{\nu\beta}$, $O^t_{\nu\beta}$, $G^t_{\nu\beta}$, $I^t_{\nu\beta}$, $C^t_{\nu\beta}$ and $U^t_{\nu\beta}$. The idea is that for each agent in a pool α has a *target intimacy* that is its desired LOGIC model for that agent: $TL^t_{\nu\beta}$, $TO^t_{\nu\beta}$, $TG^t_{\nu\beta}$, $TI^t_{\nu\beta}$, $TC^t_{\nu\beta}$ and $TU^t_{\nu\beta}$. The prior to commencing an interaction dialogue Γ, α constructs a target LOGIC model for that dialogue: $DL^t_{\nu\beta}$, $DO^t_{\nu\beta}$, $DG^t_{\nu\beta}$, $DI^t_{\nu\beta}$, $DC^t_{\nu\beta}$ and $DU^t_{\nu\beta}$. The dialogue target then becomes a constraint on the argumentation strategy.

α does not give private information away freely, and seeks a level of *balance* in information revelation. This is achieved by building a speculative model of β's model of α — after all, α should have a fairly good idea of what β knows about α — this is the *reflection model*. As the dialogue proceeds information in the five logic categories is exchanged (or, 'traded') and whilst attempting to maintain a reasonable level of balance α aims to achieve its dialogue target. Conversely, α may deliberatively diverge from a balanced information exchange to send a (positive or negative) signal to β.

Contract enactment is α's final opportunity to adjust the balance of an interaction dialogue by enacting minor variations of the signed commitment or by further information revelation. This mechanism is used widely by human agents who may "add a little extra" to impress their partner, or may otherwise diverge from their commitment to signal their intent.

The preceding discussion is at a high level but given the detailed measurement of information exchange in the LOGIC framework it tightly constrains α's utterances possibly to the extent of making α's behaviour appear to be predictable. *Stance* is common device used by human agents to conceal their interaction strategies. Stance randomly varies along the axis 'tough guy' / 'nice guy'[16] and is applied as a filter on outgoing utterances to add strategic noise that aims to prevent its underlying interaction strategies from being decrypted by β.

29.6 Conclusions

This section has drawn together two major threads: trust in the enactment of contracts and the relationships between agents. Trust has been defined in terms of the expected *value* derived from signing a contract — this is in contrast to defining trust as the expected variation between commitment and enactment. The definition chosen is more general in that it assumes some time delay between the enactment and the valuation, and that the valuation reflects the personal preferences of the agent. For example, if a car is purchased it may be delivered exactly as specified but after driving the car for some time the agent may come to value the purchase as being imperfect in some way. This notion of trust treats commitment as not simply "acting as specified" but as attempting to act in the interests of the contractual partner. This is achieved with a model of the relationships between agents that enables agents to build relationships with trust.

References

1. Acar, W., Troutt, M.D.: A methodological analysis of the normalisation, calibration and linearity of continuous diversity measures. International Journal of Operational Research **3**(1–2), 52 – 76 (2008)
2. Adams, J.S.: Inequity in social exchange. In: L. Berkowitz (ed.) Advances in experimental social psychology, vol. 2. New York: Academic Press (1965)
3. 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)
4. Bazerman, M.H., Loewenstein, G.F., White, S.B.: Reversal of preference in allocation decisions: judging an alternative versus choosing among alternatives. Administration Science Quarterly (37), 220–240 (1992)

[16] When questioning suspects the police may have two officers present each with a deliberately different stance.

5. Behrens, T.E.J., Woolrich, M.W., Walton, M.E., Rushworth, M.F.S.: Learning the value of information in an uncertain world. Nature Neuroscience **10**, 1214 – 1221 (2007)
6. Kalfoglou, Y., Schorlemmer, M.: IF-Map: An ontology-mapping method based on information-flow theory. In: S. Spaccapietra, S. March, K. Aberer (eds.) Journal on Data Semantics I, *Lecture Notes in Computer Science*, vol. 2800, pp. 98–127. Springer-Verlag: Heidelberg, Germany (2003)
7. Li, Y., Bandar, Z.A., McLean, D.: An approach for measuring semantic similarity between words using multiple information sources. IEEE Transactions on Knowledge and Data Engineering **15**(4), 871 – 882 (2003)
8. MacKay, D.: Information Theory, Inference and Learning Algorithms. Cambridge University Press (2003)
9. Osborne, M., Rubinstein, A.: A Course in Game Theory. The MIT Press (1994)
10. Sierra, C., Debenham, J.: Information-based agency. In: M. Veloso (ed.) Twentieth International Joint Conference on AI, pp. 1513–1518 (2007)
11. Sierra, C., Debenham, J.: The LOGIC Negotiation Model. In: Proceedings Sixth International Conference on Autonomous Agents and Multi Agent Systems AAMAS-2007, pp. 1026–1033. Honolulu, Hawai'i (2007)
12. Sierra, C., Debenham, J.: Information-based deliberation. In: L. Padgham, D. Parkes, J. Müller, S. Parsons (eds.) Proceedings Seventh International Conference on Autonomous Agents and Multi Agent Systems AAMAS-2008. ACM Press, New York, Estoril, Portugal (2008)
13. Sierra, C., Debenham, J.: Information-based reputation. In: M. Paolucci (ed.) First International Conference on Reputation: Theory and Technology (ICORE'09), pp. 5–19. Gargonza, Italy (2009)
14. 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 (3), 249–260 (1995)
15. 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: R. Bies, R. Lewicki, B. Sheppard (eds.) Research in Negotiation in Organizations, vol. 5, pp. 65–94. JAI Press (1995)