Information-Based Agents Model other Agents by Observing their Actions

John Debenham
Faculty of IT, University of Technology, Sydney, NSW, Australia,
debenham@it.uts.edu.au

Abstract

An 'information-based' agent is proposed for competitive multi-issue negotiation, where speculation about an opponent's motivation necessarily leads to an endless counter-speculation spiral of questionable value. Information-based agents model other agents by observing their behaviour and *not* by making assumptions concerning their motivations or internal reasoning. The integrity of these agents' information decays with time, and so these observations are represented as beliefs qualified with decaying epistemic probabilities. Entropy-based inference methods are applied to form expectations about the other agents' future actions.

1 Introduction

An agent, Π, attempts to fuse negotiation with the information that is generated both by and because of it. To achieve this, it draws on ideas from information theory rather than game theory. Π decides what to do — such as what deal to propose — on the basis of its information that may be qualified by expressions of degrees of belief. Π uses this information to calculate, and continually revise, probability distributions for that which it does not know. Two probability distributions form the foundation of competitive interaction — they are both over the set of all deals. The first distribution is the probability that any deal is acceptable to an opponent Ω_i. The second distribution is the probability that any deal will prove to be acceptable to Π — this is concerned with the integrity of the information about the deal as much as with the value of the deal itself. These distributions are calculated from Π's knowledge and beliefs using maximum entropy inference, *ME*. Π makes no assumptions about the internals of its opponents, including whether they have, or are even aware of the concept of, utility functions. Π is purely concerned with its opponents' behaviour — what they do — and not with assumptions about their motivations.

Maximum entropy inference is chosen because it enables inferences to be drawn from incomplete and uncertain information, and because of its encapsulation of common sense reasoning [Paris, 1999]. Unknown probability distributions are inferred using *maximum entropy inference* [MacKay, 2003] that is based on random worlds [Halpern, 2003]. The maximum entropy probability distribution is "the least biased estimate possible on the given information; i.e. it is maximally noncommittal with regard to missing information" [Jaynes, 1957].

The basic architecture of an "information-based" agent is presented in Sec. 2 — and its entropy-based inference machinery is described in Sec. 2.1. The integrity of the agent's information is in a permanent state of decay, Sec. 3 describes the agent's machinery for managing this decay leading to a characterisation of the "value" of information. An agent for bilateral bargaining is described in Sec. 4, and a market agent in Sec. 6. Sec. 7 describes a semi co-operative agent in a process management application. Each of these three agents derives its actions from its observations. Sec. 8 concludes.

2 Information-Based Agent Architecture

The essence of "information-based agency" is described following. An agent observes events in its environment including what other agents actually do. It chooses to represent some of those observations in its world model as beliefs. As time passes, an agent may not be prepared to accept such beliefs as being "true", and qualifies those representations with epistemic probabilities. Those qualified representations of prior observations are the agent's *information*. This information is primitive — it is the agent's representation of its beliefs about prior events in the environment and about the other agents prior actions. It is independent of what the agent is trying to achieve, or what the agent believes the other agents are trying to achieve. Given this information, an agent may then choose to adopt goals and strategies. Those strategies may be based on game theory, for example. To enable the agent's strategies to make good use of its information, tools from information theory are applied to summarise and process that information. Such an agent is called *information-based*.

An agent called Π is the subject of this discussion. Π engages in multi-issue negotiation with a set of other agents: $\{\Omega_1, \cdots, \Omega_o\}$. The foundation for Π's operation is the information that is generated both by and because of its negotiation exchanges — any message from one agent to another reveals information about the sender. Π also acquires information from the environment — including general information sources — to support its actions. Π uses ideas from information theory to process and summarise its information. Π's aim may not be "utility optimisation" — it may not be aware

Figure 1: Basic architecture of agent Π

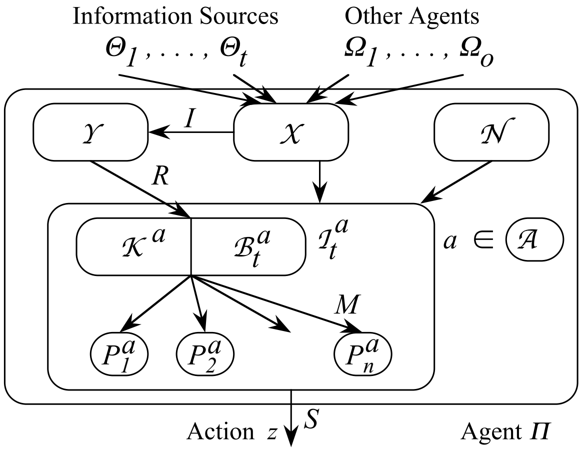

of a utility function. If Π *does* know its utility function *and* if it aims to optimise its utility *then* Π may apply the principles of game theory to achieve its aim. The information-based approach does not to reject utility optimisation — in general, the selection of a goal and strategy is secondary to the processing and summarising of the information.

In addition to the information derived from its opponents, Π has access to a set of information sources $\{\Theta_1, \cdots, \Theta_t\}$ that may include the marketplace in which trading takes place, and general information sources such as news-feeds accessed via the Internet. Together, Π, $\{\Omega_1, \cdots, \Omega_o\}$ and $\{\Theta_1, \cdots, \Theta_t\}$ make up a multi-agent system. The integrity of Π's information, including information extracted from the Internet, will decay in time. The way in which this decay occurs will depend on the type of information, and on the source from which it was drawn. Little appears to be known about how the integrity of real information, such as news-feeds, decays, although its validity can often be checked — "Is company X taking over company Y?" — by proactive action given a co-operative information source Θ_j. So Π has to consider how and when to refresh its decaying information.

Π has two languages: $\mathcal{C}$ and $\mathcal{L}$. $\mathcal{C}$ is an illocutionary-based language for communication and is not described here. $\mathcal{L}$ is a first-order language for internal representation — precisely it is a first-order language with sentence probabilities optionally attached to each sentence representing Π's epistemic belief in the truth of that sentence. Fig. 1 shows a high-level view of how Π operates. Messages expressed in $\mathcal{C}$ from $\{\Theta_i\}$ and $\{\Omega_i\}$ are received, time-stamped, source-stamped and placed in an *in-box* $\mathcal{X}$. The messages in $\mathcal{X}$ are then translated using an *import function* I into sentences expressed in $\mathcal{L}$ that have integrity decay functions (usually of time) attached to each sentence, they are stored in a *repository* $\mathcal{Y}$. And that is all that happens until Π triggers a goal.

Π triggers a goal in two ways: first in response to a message received from an opponent $\{\Omega_i\}$ "I'd like to purchase an apple from you", and second in response to some need, $\mathcal{N}$, "goodness, we've run out of coffee". Π's goals could be short-term such as obtaining some information "what is the time?", medium-term such as striking a deal with one of its opponents, or, rather longer-term such as building a (business) relationship with one of its opponents. For each goal that Π commits to, it has a mechanism for selecting a plan to achieve it. Π's plans reside in a plan library $\mathcal{A}$. Once a plan, a, has been activated, it extracts those sentences from the repository $\mathcal{Y}$ that are relevant to it, instantiates each of those sentences' integrity decay functions to the current time t, and selects a consistent sub-set of these sentences using its belief revision[1] function R. Those instantiated sentences that have no decay function are placed into the *knowledge base* $\mathcal{K}^a$, and those that have decay functions are placed along with their sentence probabilities into the *belief set* $\mathcal{B}_t^a$. $\mathcal{K}^a \cup \mathcal{B}_t^a = \mathcal{I}_t^a$ is the *information base* created by plan a at time t. Plan a then uses tools from information theory, including maximum entropy inference, M, to derive a set of probability distributions, $\{P_1^a, \cdots, P_n^a\}$, from $\mathcal{I}_t^a$. The way in which these derivations are performed are described in Sec. 2.1 following. Then plan a invokes some strategy S that uses the $\{P_1^a, \cdots, P_n^a\}$ to determine Π's action $z \in \mathcal{Z}$.

2.1 Π's Reasoning

Once Π has selected a plan $a \in \mathcal{A}$ it uses maximum entropy inference to derive the $\{P_i^a\}_{i=1}^n$ [see Fig. 1] and minimum relative entropy inference to update those distributions as new data becomes available. *Entropy*, $\mathbb{H}$, is a measure of uncertainty [MacKay, 2003] in a probability distribution for a discrete random variable X: $\mathbb{H}(X) \triangleq -\sum_i p(x_i) \log p(x_i)$ where $p(x_i) = \mathbb{P}(X = x_i)$. Maximum entropy inference is used to derive sentence probabilities for that which is not known by constructing the "maximally noncommittal" [Jaynes, 2003] probability distribution.

Let $\mathcal{G}$ be the set of all positive ground literals that can be constructed using Π's language $\mathcal{L}$. A *possible world*, v, is a valuation function: $\mathcal{G} \to \{\top, \bot\}$. $\mathcal{V}|\mathcal{K}^a = \{v_i\}$ is the set of all possible worlds that are consistent with Π's knowledge base $\mathcal{K}^a$ that contains statements which Π believes are true. A *random world* for $\mathcal{K}^a$, $W|\mathcal{K}^a = \{p_i\}$ is a probability distribution over $\mathcal{V}|\mathcal{K}^a = \{v_i\}$, where p_i expresses Π's degree of belief that each of the possible worlds, v_i, is the actual world. The *derived sentence probability* of any $\sigma \in \mathcal{L}$, *with respect to* a random world $W|\mathcal{K}^a$ is:

$$(\forall \sigma \in \mathcal{L}) \mathbb{P}_{\{W|\mathcal{K}^a\}}(\sigma) \triangleq \sum_n \{\, p_n \;:\; \sigma \; is \top in \; v_n \,\} \qquad (1)$$

The agent's *belief set* $\mathcal{B}_t^a = \{\beta_j\}_{j=1}^M$ contains statements to which Π attaches a *given sentence probability* $\mathbb{B}(.)$. A random world $W|\mathcal{K}^a$ is *consistent* with $\mathcal{B}_t^a$ if: $(\forall \beta \in \mathcal{B}_t^a)(\mathbb{B}(\beta) = \mathbb{P}_{\{W|\mathcal{K}^a\}}(\beta))$. Let $\{p_i\} = \{\overline{W}|\mathcal{K}^a, \mathcal{B}_t^a\}$ be the "maximum entropy probability distribution over $\mathcal{V}|\mathcal{K}^a$ that is consistent with $\mathcal{B}_t^a$". Given an agent with $\mathcal{K}^a$ and $\mathcal{B}_t^a$, *maximum entropy inference* states that the *derived sentence probability* for any sentence, $\sigma \in \mathcal{L}$, is:

$$(\forall \sigma \in \mathcal{L}) \mathbb{P}_{\{\overline{W}|\mathcal{K}^a, \mathcal{B}_t^a\}}(\sigma) \triangleq \sum_n \{\, p_n \;:\; \sigma \; is \top in \; v_n \,\} \qquad (2)$$

[1]This belief revision — consistency checking — exercise is non-trivial, and is not described here. For sake of illustration only, the strategy "discard the old in favour of the new" is sufficient.

From Eqn. 2, each belief imposes a linear constraint on the $\{p_i\}$. The maximum entropy distribution: $\arg\max_{\underline{p}} \mathbb{H}(\underline{p})$, $\underline{p} = (p_1, \ldots, p_N)$, subject to $M+1$ linear constraints:

$$g_j(\underline{p}) = \sum_{i=1}^{N} c_{ji} p_i - \mathbb{B}(\beta_j) = 0, \quad j = 1, \ldots, M.$$

$$g_0(\underline{p}) = \sum_{i=1}^{N} p_i - 1 = 0$$

where $c_{ji} = 1$ if β_j is $\top$ in v_i and 0 otherwise, and $p_i \geq 0, i = 1, \ldots, N$, is found by introducing Lagrange multipliers, and then obtaining a numerical solution using the multivariate Newton-Raphson method. In Sec. 3.1 we'll see how an agent updates its sentence probabilities as new information is received.

Given a prior probability distribution $\underline{q} = (q_i)_{i=1}^n$ and a set of constraints C, the *principle of minimum relative entropy* chooses the posterior probability distribution $\underline{p} = (p_i)_{i=1}^n$ that has the least *relative entropy*[2] with respect to $\underline{q}$:

$$\{\underline{W}|\underline{q}, C\} \triangleq \arg\min_{\underline{p}} \sum_{i=1}^{n} p_i \log \frac{p_i}{q_i}$$

and that satisfies the constraints C. This may be found by introducing Lagrange multipliers as above. Given a prior distribution $\underline{q}$ over $\{v_i\}$ — the set of all possible worlds, and a set of constraints C (that could have been derived as above from a set of new beliefs) *minimum relative entropy inference* states that the derived sentence probability for any sentence, $\sigma \in \mathcal{L}$, is:

$$(\forall \sigma \in \mathcal{L}) \mathbb{P}_{\{\underline{W}|\underline{q},C\}}(\sigma) \triangleq \sum_{n} \{\, p_n \,:\, \sigma \; is \top in \; v_n \,\} \tag{3}$$

where $\{p_i\} = \{\underline{W}|\underline{q}, C\}$. The principle of minimum relative entropy is a generalisation of the principle of maximum entropy. If the prior distribution $\underline{q}$ is uniform, then the relative entropy of $\underline{p}$ with respect to $\underline{q}$ differs from $-\mathbb{H}(\underline{p})$ only by a constant. So the principle of maximum entropy is equivalent to the principle of minimum relative entropy with a uniform prior distribution.

Discussion on entropy-based inference

Entropy-based inference is chosen because it enables complete probability distributions to be constructed from a small number of observations. In such a case the distributions could not be expected to be accurate. But the technique does derive the unique distribution that is least biased with respect to what is as yet unknown.

Maximum entropy inference presents four difficulties. First, it assumes that what the agent knows is "the sum total of the agent's knowledge, it is not a summary of the agent's knowledge, it is all there is" [Paris, 1999]. This assumption referred to as Watt's Assumption [Jaeger, 1996]. So if knowledge is absent then the agent may do strange things. Second, it may only be applied to a consistent set of beliefs — this may mean that valuable information is destroyed by the belief revision process that copes with the continuous arrival of new information. Third, its knowledge base is expressed in first-order logic. So issues that have unbounded domains — such as price — can only be dealt with either exactly as a large quantity of constants for each possible price, or approximately as price intervals. This decision will effect the inferences drawn and is referred to as representation dependence [Halpern, 2003]. Fourth, maximum entropy is not simple to calculate — numerical solutions are obtained by applying the Newton-Raphson method to as many non-linear, simultaneous equations as there are beliefs in the knowledge base.

Despite these four difficulties, maximum entropy inference is an elegant formulation of common sense reasoning [Paris, 1999]. Maximum entropy inference is also independent of any structure on the set of all possible deals. So it copes with single-issue and multi-issue negotiation without modification. It may also be applied to the probabilistic belief logic that is used here.

2.2 An Exemplar Application

An exemplar application is used following. Π is attempting to purchase of a second-hand motor vehicle, with some period of warranty, for cash. So the two issues in this negotiation are: the period of the warranty, and the cash consideration. A deal δ consists of this pair of issues, and the deal set has no natural ordering. Π observes its opponents' actions and uses these observations to estimate the probability that an offer of deal δ will be accepted for various δ. Suppose that the warranty period is simply $0, \cdots, 4$ years, and that the cash amount for this car will certainly be at least $5,000 with no warranty, and is unlikely to be more than $7,000 with four year's warranty. In what follows all price units are in thousands of dollars.

To represent the application in first-order logic, a finite set of intervals is chosen for the issue price. This choice will effect the results of the *ME* calculations, and *ME* is criticised [Halpern, 2003] because the way in which the knowledge is represented in $\mathcal{K}^a$ and $\mathcal{B}^a_t$ determines the values derived. This property is promoted here as a strength of the method because the correct representation, using the rich expressive power of first-order probabilistic logic, encapsulates features of the application at a fine level of detail. Suppose then that the set of intervals chosen for price contains the eleven intervals: D = { [5.0, 5.2), [5.2, 5.4), [5.4, 5.6), [5.6, 5.8), [5.8, 6.0), [6.0, 6.2), [6.2, 6.4), [6.4, 6.6), [6.6, 6.8), [6.8, 7.0), [7.0, ∞) }. Given some value for the warranty period w what is the least price the Π expects its opponent to accept? In the absence of any information, *ME* concludes that these eleven intervals are equi-probable with a probability of $\frac{1}{11}$. If any other complete and disjoint set of intervals had been chosen the answer would have been the same. Rather than condemn the method on the basis of this conclusion, we promote this as a strength. *Either* the intervals should be chosen so that they are *a priori* equi-probable, *or* knowledge should be included in $\mathcal{I}^a_t$ to express why that are not. The choice of intervals enables fine-grained modelling of the application domain.

Suppose then that the deal set chosen in this application consists of 55 individual deals in the form of pairs of warranty periods and price intervals (w, p) where $w = 0, \cdots, 4$ and $p \in$

[2] Otherwise called *cross entropy* or the *Kullback-Leibler* distance between the two probability distributions.

D. Suppose that Π has previously received two offers from Ω. The first was to offer 6.0 with no warranty, and the second to offer 6.9 with one year's warranty. Suppose Π believes that Ω still stands by these two offers with probability 0.8. Then this leads to two beliefs: $\beta_1 : \text{Acc}(\Omega, \Pi, (0, [6.0, 6.2]))$; $\mathbb{B}(\beta_1) = 0.8$, $\beta_2 : \text{Acc}(\Omega, \Pi, (1, [6.8, 7.0]))$; $\mathbb{B}(\beta_2) = 0.8$, where $\text{Acc}(\Omega, \Pi, \delta)$ means "Π believes that Ω will accept deal δ". Before "switching on" *ME*, Π should consider whether it believes that $\mathbb{P}(\text{Acc}(\Omega, \Pi, \delta))$ is uniform over δ. If it does then it includes both β_1 and β_2 in $\mathcal{B}$, and calculates $\{\overline{W} | \mathcal{K}^a, \mathcal{B}^a_t\}$ that yields estimates for $\mathbb{P}(\text{Acc}(\Omega, \Pi, \delta))$ for all δ. If it does not then it should include further knowledge in $\mathcal{K}^a$ and $\mathcal{B}^a_t$. For example, Π may believe that Ω prefers a greater warranty period the higher the price. If so, then this is a multi-issue constraint, that is represented in $\mathcal{B}^a_t$, and may be qualified with a sentence probability.

3 Managing Integrity Decay

Π's information base $\mathcal{I}^a_t$ contains sentences in first-order logic each with a sentence probability, $\mathbb{B}(\cdot)$, representing the agent's strength of belief in the truth of that statement. Maintaining the integrity of $\mathcal{I}^a_t$ is not just a matter of looking up information. Information may be temporarily unavailable, acquiring it may cost money, the information may be inherently unreliable, and its availability may be beyond the control of the agent. For example, if a chunk of information represents the action of another agent then that information can only be refreshed when the other agent acts.

3.1 Updating the probability distributions

Π's plans are partly driven by the probability distributions $\{P^a_1, \cdots, P^a_n\}$ in Fig. 1. These distributions are derived from the information in the repository $\mathcal{Y}$. The integrity of the information in $\mathcal{Y}$ will decay in accordance with the decay functions — unless it is refreshed. As the integrity decays, the entropy of the distributions $\{P^a_1, \cdots, P^a_n\}$ increases.

Suppose that $\mathcal{L}$ contains a unary predicate, $\text{A}(\cdot)$, whose domain is represented by the finite set of logical constants $\{c_i\}_{i=1}^d$. At time t let $p_i^t = \mathbb{B}(\text{A}(c_i))$ for $i = 1, \cdots, d$. As time t increases *either* no information is received and the entropy of (p_i^t) should increase, *or* information is received and the distribution (p_i^t) should be refreshed. If no information is received, a geometric decay to (maximum entropy) ignorance is achieved by:

$$p_i^{t+1} = \rho \cdot p_i^t + (1 - \rho) \cdot \frac{1}{d} \tag{4}$$

for a decay factor $\rho \in [0, 1]$ whose value depends on the meaning of $\text{A}(\cdot)$.

Now suppose that at time t, Π receives a message from source Θ that asserts the truth of $\text{A}(c_k)$, and suppose that Π decides to attach the given sentence probability $g(\Theta)$ to $\mathbb{B}(\text{A}(c_k))$, where the value $g(\Theta)$ is Π's confidence in the integrity of Θ's advice. Then the updated distribution is calculated by applying *minimum relative entropy inference*:

$$\left(p_j^t\right)_{j=1}^d = \arg\min_{\underline{b}} \sum_{i=1}^d b_i \log \frac{b_i}{p_i^{t-1}}$$

satisfying the constraint: $p_k^t = g(\Theta)$, and where $\underline{b} = (b_j)_{j=1}^d$. That is, the new distribution is the closest[3] to the previous one that satisfies the constraint. But, this depends on Π knowing $g(\Theta)$, ie: the strength of belief that Π allocates to Θ's information. The issue here is not, for example, forecasting the clearing price of a stock, it is estimating the confidence that Π has in the integrity of Θ's information. If Π can confirm the validity of Θ's advice *ex post* then a very simple way of estimating $g(\Theta)$ is by:

$$g_{\text{new}}(\Theta) = \begin{cases} \nu \cdot g_{\text{old}}(\Theta) + (1 - \nu) & \text{when } \Theta \text{ is correct} \\ \nu \cdot g_{\text{old}}(\Theta) & \text{when } \Theta \text{ is incorrect} \end{cases}$$

for a learning rate $\nu \in [0, 1]$.

3.2 Valuing Information

A chunk of information is valued only by the way that it enables Π to do something[4]. So information is valued in relation to the plans that Π is executing. A plan, a, is designed in the context of a particular representation, or environment, e. One way in which a chunk of information assists Π is by altering one of a's distributions $\{P^a_i\}$ — see Fig. 1. As a chunk of information could be "good" for one distribution and "bad" for another, the appropriate way to value information is by its effect on each distribution. For a plan a, the *value* to a of a message received at time t is the resulting decrease in entropy in a's distributions $\{P^a_i\}$ in Fig. 1. In general, suppose that a set of stamped messages $X = \{x_i\}$ is imported by plan a to the information base $\mathcal{I}^a_t$ where they are represented as the set of statements $D = \{d_i\} = R(I(X))$, where I is the import function and R the belief revision function. The *information* in D at time t with respect to a particular P^a_i, plan a and environment e is:

$$\mathbb{I}(D \mid P^a_i(\mathcal{I}^a_t), a, e) \triangleq \mathbb{H}(P^a_i(\mathcal{I}^a_t)) - \mathbb{H}(P^a_i(\mathcal{I}^a_t \cup D))$$

for $i = 1, \cdots, n$, where the argument of the $P^a_i(\cdot)$ is the state of Π's information base from which P^a_i was derived. And we define the information in the set of messages X (at time t with respect to a particular P^a_i, a and e) to be the information in $D = R(I(X))$. It is reasonable to aggregate the information in D over the distributions used by a, and to aggregate again over all plans to obtain the (potential) information in a statement. That is, the *potential information* in D with respect to environment e is:

$$\mathbb{I}(D \mid e) \triangleq \sum_{a \in \mathcal{A}} \sum_i \mathbb{I}(D \mid P^a_i(\mathcal{I}^a_t), a, e) \tag{5}$$

4 A Bargaining Agent

[Debenham, 2004b] describes a multi-issue, bilateral bargaining agent Π, with just one opponent Ω, whose strategies are all based on the three distributions: $\mathbb{P}(\text{Acc}(\Pi, \Omega, \delta))$ for all deals δ [ie: the probability that Π should accept deal δ from agent Ω], $\mathbb{P}(\text{Acc}(\Omega, \Pi, \delta))$ for all deals δ [ie: Π's estimate

[3]Precisely, it is the distribution that minimises the Kullback-Leibler distance from the prior distribution, (p_i^{t-1}), whilst satisfying the given constraint.

[4]That is we do not try to attach any intrinsic value to information.

of the probability that Ω would accept deal δ from agent Π], and $p_{b,\Omega}$ [ie: the probability of breakdown — the probability that Ω will "walk away" in the next negotiation round]. These three complete probability distributions are derived by observing the information in the signals received.

Π has do two different things. First, it must respond to offers received from Ω — this is described in Sec. 6. Second, it must send offers, and possibly information, to Ω. This section describes machinery for estimating the probabilities $\mathbb{P}(\mathrm{Acc}(\Omega, \Pi, \delta))$. Here Π is attempting to purchase of a particular second-hand motor vehicle, with some period of warranty, for cash from Ω as described in Sec. 2.2. So a deal δ will be represented by the pair (w, p) where w is the period of warranty in years and $\$p$ is the price.

Π assumes the following two preference relations for Ω, and $\mathcal{K}^a$ contains:

$\kappa_1 : \forall xyz((x < y) \rightarrow$
$(\mathrm{Acc}(\Omega, \Pi, (y, z)) \rightarrow \mathrm{Acc}(\Omega, \Pi, (x, z))))$

$\kappa_2 : \forall xyz((x < y) \rightarrow$
$(\mathrm{Acc}(\Omega, \Pi, (z, x)) \rightarrow \mathrm{Acc}(\Omega, \Pi, (z, y))))$

These sentences conveniently reduce the number of possible worlds. The two preference relations κ_1 and κ_2 induce a partial ordering on the sentence probabilities in the $\mathbb{P}(\mathrm{Acc}(\Omega, \Pi, w, p))$ array from the top-left where the probabilities are ≈ 1, to the bottom-right where the probabilities are ≈ 0. There are fifty-one possible worlds that are consistent with $\mathcal{K}^a$.

Suppose that the offer exchange has proceeded as follows: Ω asked for \$6,900 with one year warranty and Π refused, then Π offered \$5,000 with two years warranty and Ω refused, and then Ω asked for \$6,500 with three years warranty and Π refused. Then at the next time step $\mathcal{B}^a_t$ contains:

$\beta_3 : \mathrm{Acc}(\Omega, \Pi, (3, [6.8, 7.0)))$,
$\beta_4 : \mathrm{Acc}(\Omega, \Pi, (2, [5.0, 5.2)))$ and
$\beta_5 : \mathrm{Acc}(\Omega, \Pi, (1, [6.4, 6.6)))$,
and with a 10% decay in integrity for each time step: $\mathbb{P}(\beta_3) = 0.7$, $\mathbb{P}(\beta_4) = 0.2$ and $\mathbb{P}(\beta_5) = 0.9$

Eqn. 2 is used to calculate the distribution $\{\overline{W} | \mathcal{K}^a, \mathcal{B}^a_t\}$ which shows that there are just five different probabilities in it. The probability matrix for the proposition $\mathrm{Acc}(\Omega, \Pi, (w, p))$ is:

$p =$	$w = 0$	$w = 1$	$w = 2$	$w = 3$	$w = 4$
$[7.0, \infty)$	0.9967	0.9607	0.8428	0.7066	0.3533
$[6.8, 7.0)$	0.9803	0.9476	0.8330	**0.7000**	0.3500
$[6.6, 6.8)$	0.9533	0.9238	0.8125	0.6828	0.3414
$[6.4, 6.6)$	0.9262	**0.9000**	0.7920	0.6655	0.3328
$[6.2, 6.4)$	0.8249	0.8019	0.7074	0.5945	0.2972
$[6.0, 6.2)$	0.7235	0.7039	0.6228	0.5234	0.2617
$[5.8, 6.0)$	0.6222	0.6058	0.5383	0.4523	0.2262
$[5.6, 5.8)$	0.5208	0.5077	0.4537	0.3813	0.1906
$[5.4, 5.6)$	0.4195	0.4096	0.3691	0.3102	0.1551
$[5.2, 5.4)$	0.3181	0.3116	0.2846	0.2391	0.1196
$[5.0, 5.2)$	0.2168	0.2135	**0.2000**	0.1681	0.0840

In this array, the derived sentence probabilities for the three sentences in $\mathcal{B}^a_t$ are shown in bold type; they are exactly their given values. This matrix simplifies the expression of bargaining strategies such as [Faratin *et al.*, 2003].

5 Negotiation Strategies

Sec. 4 estimated the probability distribution, $\mathbb{P}(\mathrm{Acc}(\Omega, \Pi, w, p))$, that Ω will accept an offer, and Sec. 6 following estimates the probability distribution, $\mathbb{P}(\mathrm{Acc}(\Pi, \Omega, \delta))$, that Π should be prepared to accept an offer δ. These two probability distributions represent the opposing interests of the two agents Π and Ω. $\mathbb{P}(\mathrm{Acc}(\Omega, \Pi, w, p))$ will change every time an offer is made, rejected or accepted. $\mathbb{P}(\mathrm{Acc}(\Pi, \Omega, \delta))$ will change as the background information changes. This section discusses Π's strategy S. Sec. 5.2 considers the risk of breakdown.

Bargaining can be a game of bluff and counter-bluff in which an agent may even not intend to close the deal if one should be reached. A basic conundrum in any offer-exchange bargaining is: it is impossible to force your opponent to reveal information about their position without revealing information about your own position. Further, by revealing information about your own position you may change your opponents position — and so on.[5] This infinite regress, of speculation and counter-speculation, is avoided here by ignoring the internals of the opponent and by focussing on what is known for certain — that is: *what* information is contained in the signals received and *when* did those signals arrive.

A fundamental principle of competitive bargaining is "never reveal your best price", and another is "never reveal your deadline — if you have one" [Sandholm and Vulkan, 1999]. It is not possible to be prescriptive about what an agent *should* reveal. All that can be achieved is to provide strategies that an agent may choose to employ. The following are examples of such strategies.

5.1 Without Breakdown

An agent's strategy S is a function of the information $\mathcal{I}_t$ that is has at time t. That information will be represented in the agent's $\mathcal{K}^a$ and $\mathcal{B}^a_t$. Simple strategies choose an offer only on the basis of $\mathbb{P}(\mathrm{Acc}(\Pi, \Omega, \delta))$, $\mathbb{P}(\mathrm{Acc}(\Omega, \Pi, \delta))$ and α. The greedy strategy S^+ chooses:

$$\arg\max_\delta \{\mathbb{P}(\mathrm{Acc}(\Pi, \Omega, \delta)) \mid \mathbb{P}(\mathrm{Acc}(\Omega, \Pi, \delta)) \gg 0\},$$

it is appropriate for an agent that believes Ω is desperate to trade. The *expected-acceptability-to-Π-optimizing strategy* S^* chooses:

$$\arg\max_\delta \{\mathbb{P}(\mathrm{Acc}(\Omega, \Pi, \delta)) \times \mathbb{P}(\mathrm{Acc}(\Pi, \Omega, \delta)) \mid \mathbb{P}(\mathrm{Acc}(\Pi, \Omega, \delta)) \geq \alpha\}$$

it is appropriate for a confident agent that is not desperate to trade. The strategy S^- chooses:

$$\arg\max_\delta \{\mathbb{P}(\mathrm{Acc}(\Omega, \Pi, \delta)) \mid \mathbb{P}(\mathrm{Acc}(\Pi, \Omega, \delta)) \geq \alpha\}$$

it optimizes the likelihood of trade — it is a good strategy for an agent that is keen to trade without compromising its own standards of acceptability.

An approach to issue-tradeoffs is described in [Faratin *et al.*, 2003]. The bargaining strategy described there attempts

[5] This a reminiscent of Werner Heisenberg's indeterminacy relation, or *unbestimmtheitsrelationen*: "you can't measure one feature of an object without changing another" — with apologies.

to make an acceptable offer by "walking round" the iso-curve of Π's previous offer (that has, say, an acceptability of $\alpha_{na} \geq \alpha$) towards Ω's subsequent counter offer. In terms of the machinery described here, an analogue is to use the strategy S^{-}:

$$\arg\max_{\delta}\{\mathbb{P}(\mathrm{Acc}(\Pi,\Omega,\delta)) \mid \mathbb{P}(\mathrm{Acc}(\Pi,\Omega,\delta) \mid \mathcal{I}_t) \gtrsim \alpha_{na}\ \}$$

for $\alpha = \alpha_{na}$. This is reasonable for an agent that is attempting to be accommodating without compromising its own interests. Presumably such an agent will have a policy for reducing the value α_{na} if her deals fail to be accepted. The complexity of the strategy in [Faratin *et al.*, 2003] is linear with the number of issues. The strategy described here does not have that property, but it benefits from using $\mathbb{P}(\mathrm{Acc}(\Omega,\Pi,\delta))$ that contains foot prints of the prior offer sequence — see Sec. 4 — in that distribution more recent offers have stronger weights.

5.2 With Breakdown

A negotiation may break down because one agent is not prepared to continue for some reason. p_B is the probability that the opponent will quit the negotiation in the next round. There are three ways in which Π models the risk of breakdown. First, p_B is a constant determined exogenously to the negotiation, in which case at any stage in a continuing negotiation the expected number of rounds until breakdown occurs is $\frac{1}{p_B}$. Second, p_B is a monotonic increasing function of time — this attempts to model an impatient opponent. Third, p_B is a monotonic increasing function of $(1 - \mathbb{P}(\mathrm{Acc}(\Omega,\Pi,\delta)))$ — this attempts to model an opponent who will react to unattractive offers.

At any stage in a negotiation Π may be prepared to gamble on the expectation that Ω will remain in the game for the next n rounds. This would occur if there is a constant probability of breakdown $p_B = \frac{1}{n}$. Let $\mathcal{I}_t$ denote the information stored in Π's $\mathcal{K}^a$ and $\mathcal{B}^a_t$ at time t. S is Π's strategy. If Π offered to trade with Ω at $S(\mathcal{I}_1)$ then Ω may accept this offer, but may have also been prepared to settle for terms more favorable than this to Π. If Π offered to trade at $S(\mathcal{I}_1 \cup \{\mathrm{Acc}(\Omega,\Pi,S(\mathcal{I}_1))\})$ then Ω will either accept this offer or reject it. In the former case trade occurs at more favorable terms than $S(\mathcal{I}_1)$, and in the latter case a useful piece of information has been acquired: $\neg\mathrm{Acc}(S(\Omega,\Pi,\mathcal{I}_1))$ which is added to $\mathcal{I}_1$ before calculating the next offer. This process can be applied twice to generate the offer $S(\mathcal{I}_1 \cup \{\neg\mathrm{Acc}(\Omega,\Pi,S(\mathcal{I}_1 \cup \{\neg\mathrm{Acc}(\Omega,\Pi,S(\mathcal{I}_1))\}))\})$, or any number of times, optimistically working backwards on the assumption that Ω will remain in the game for n rounds. The strategy $S^{(n)}$, where $S^{(1)} = S^{*}$ the expected-acceptability-to-Π-optimizing strategy defined in Sec. 5.1. $S^{(n)}$ is the strategy of working back from $S^{(1)}$ $(n-1)$ times. At each stage $S^{(n)}$ will benefit also from the information in the intervening counter offers presented by Ω. The strategy $S^{(n)}$ is reasonable for a risk-taking, expected-acceptability-optimizing agent.

5.3 Information Revelation

Π's *negotiation strategy* is a function $S : \mathcal{K} \times \mathcal{B} \to \mathcal{A}$ where $\mathcal{A}$ is the set of actions that send *Offer*$(.)$, *Accept*$(.)$, *Reject*$(.)$ and *Quit*$(.)$ messages to Ω. If Π sends *Offer*$(.)$, *Accept*$(.)$ or *Reject*$(.)$ messages to Ω then she is giving Ω information about herself. In an infinite-horizon bargaining game where there is no incentive to trade now rather than later, a self-interested agent will "sit and wait", and do nothing except, perhaps, to ask for information. The well known bargaining response to an approach by an interested party "Well make me an offer" illustrates how a shrewd bargainer may behave in this situation.

An agent may be motivated to act for various reasons — three are mentioned. First, if there are costs involved in the bargaining process due *either* to changes in the value of the negotiation object with time *or* to the intrinsic cost of conducting the negotiation itself. Second, if there is a risk of breakdown caused by the opponent walking away from the bargaining table. Third, if the agent is concerned with establishing a sense of trust [Ramchurn *et al.*, 2003] with the opponent —this could be the case in the establishment of a business relationship. Of these three reasons the last two are addressed here. The risk of breakdown may be reduced, and a sense of trust may be established, if the agent appears to its opponent to be "approaching the negotiation in an even-handed manner". One dimension of "appearing to be even-handed" is to be equitable with the value of information given to the opponent. Various bargaining strategies, both with and without breakdown, are described in [Debenham, 2004b], but they do not address this issue. A bargaining strategy is described here that is founded on a principle of "equitable information gain". That is, Π attempts to respond to Ω's messages so that Ω's expected information gain similar to that which Π has received.

Π models Ω by observing her actions, and by representing beliefs about her future actions in the probability distribution $\mathbb{P}(\mathrm{Acc}(\Omega,\Pi,\delta))$. Π measures the value of information that it receives from Ω by the change in the entropy of this distribution. More generally, Π measures the value of information received in a message, μ, by the change in the entropy in its entire representation, $\mathcal{J}_t = \mathcal{K}_t \cup \mathcal{B}_t$, as a result of the receipt of that message; this is denoted by: $\Delta_\mu|\mathcal{J}^{\Pi}_t|$, where $|\mathcal{J}^{\Pi}_t|$ denotes the value (as negative entropy) of Π's information in $\mathcal{J}$ at time t — see Sec. 3.2. Although both Π and Ω will build their models of each other using the same data — the messages exchanged — the observed information gain will depend on the way in which each agent has represented this information. It is "not unreasonable to suggest" that these two representations should be similar. To support its attempts to achieve "equitable information gain" Π assumes that Ω's reasoning apparatus mirrors its own, and so is able to estimate the change in Ω's entropy as a result of sending a message μ to Ω: $\Delta_\mu|\mathcal{J}^{\Omega}_t|$. Suppose that Π receives a message μ = *Offer*$(.)$ from Ω and observes an information gain of $\Delta_\mu|\mathcal{J}^{\Pi}_t|$. Suppose that Π wishes to reject this offer by sending a counter-offer, *Offer*(δ), that will give Ω expected "equitable information gain".

$$\delta = \{\arg\max_{\delta}\mathbb{P}(\mathrm{Acc}(\Pi,\Omega,\delta) \mid \mathcal{I}_t) \geq \alpha \mid$$
$$(\Delta_{\mathit{Offer}(\delta)}|\mathcal{J}^{\Omega}_t| \ \approx \Delta_\mu|\mathcal{J}^{\Pi}_t|)\}.$$

That is Π chooses the most acceptable deal to herself that gives her opponent expected "equitable information gain" provided that there is such a deal. If there is not then Π chooses the best available compromise:

$$\delta = \{\arg\max_{\delta}(\Delta_{\mathit{Offer}(\delta)}|\mathcal{J}_t^{\Omega}|) \mid \mathbb{P}(\mathrm{Acc}(\Pi,\Omega,\delta) \mid \mathcal{I}_t) \geq \alpha\}$$

provided there is such a deal — this strategy is rather generous, it rates information gain ahead of personal acceptability. If there is not then Π does nothing.

The "equitable information gain" strategy generalizes the simple-minded alternating offers strategy. Suppose that Π is trying to buy something from Ω with bilateral bargaining in which all offers and responses stand — ie: there is no decay of offer integrity. Suppose that Π has offered \$1 and Ω has refused, and Ω has asked \$10 and Π has refused. If amounts are limited to whole dollars only then the deal set $\mathcal{D} = \{1, \cdots, 10\}$. Π knows that $\mathbb{P}(\mathrm{Acc}(\Omega,\Pi,1)) = 0$ and $\mathbb{P}(\mathrm{Acc}(\Omega,\Pi,10)) = 1$. The remaining eight values in this distribution are provided by Eqn. 2, and the entropy of the resulting distribution is 2.2020. To apply the "equitable information gain" strategy Π assumes that Ω's decision-making machinery mirrors its own. In which case Ω is assumed to have constructed a mirror-image distribution to model Π that will have the same entropy. At this stage, time $t = 0$, calibrate the amount of information held by each agent at zero — ie: $|\mathcal{J}_0^{\Pi}| = |\mathcal{J}_0^{\Omega}| = 0$. Now if, at time $t = 1$, Ω asks Π for \$9 then Ω gives information to Π and $|\mathcal{J}_1^{\Pi}| = 0.2548$. If Π rejects this offer then she gives information to Ω and $|\mathcal{J}_1^{\Omega}| = 0.2548$. Suppose that Π wishes to counter with an "equitable information gain" offer. If, at time $t = 2$, Π offers Ω \$2 then $|\mathcal{J}_2^{\Omega}| = 0.2548 + 0.2559$. Alternatively, if Π offers Ω \$3 then $|\mathcal{J}_2^{\Omega}| = 0.2548 + 0.5136$. And so \$2 is a near "equitable information gain" response by Π at time $t = 2$.

6 A Market Agent

We now consider how Π values a deal in the context of multi-issue auctions where, for protocols with a truth telling dominant strategy, Π only constructs $\mathbb{P}(\mathrm{Acc}(\Pi,\Omega,\delta))$ [ie: the probability that Π should bid deal δ]. For protocols with equilibrium solutions that are expressed in terms of the number of bidders, Π also requires a probability distribution over the various possible number of bidders. From the auctioneer's point of view, the distribution $\mathbb{P}(\mathrm{WinningBid}(\delta))$ [ie: the probability that δ will be the winning bid] may be expressed analytically in terms of the number of bidders, and the size of the domain chosen to represent the various possible deals.

The distribution $\mathbb{P}(\mathrm{Acc}(\Pi,\Omega,\delta))$ is Π's analogue of a game-theoretic agent's utility function, but it is quite different. $\mathbb{P}(\mathrm{Acc}(\Pi,\Omega,\delta))$ is Π's estimate of her certainty that deal δ is acceptable. That is, Π's estimate of the validity of her information, and not an estimate of the intrinsic value of δ. This estimate will be based on various factors one of which may be Π's estimate of whether δ is considered acceptable in the open market, $\mathrm{Fair}(\delta)$, that may be determined by reference to market data. Suppose that recently a similar vehicle sold with three year's warranty for \$6,500, and another less similar was sold for \$5,500 with one year's warranty. These are fed into $\mathcal{I}_t^a$ and are represented as two beliefs in $\mathcal{B}_t^a$: $\beta_6 : \mathrm{Fair}(3, [6.4, 6.6))$; $\mathbb{B}(\beta_6) = 0.9$, $\beta_7 : \mathrm{Fair}(3, [5.4, 5.6))$; $\mathbb{B}(\beta_7) = 0.8$. In an open-cry auction one source of market data is the bids made by other agents. The sentence probabilities that are attached to this data may be derived from knowing the identity, and so too the reputation, of the bidding agent. In this way the acceptability value is continually adjusted as information becomes available. In addition to β_6 and β_7, there are three chunks of knowledge in $\mathcal{K}^a$. First, $\kappa_3 : \mathrm{Fair}(4, 4999)$ that determines a base value for which $\mathbb{P}(\mathrm{Fair}(\cdot)) = 1$, and two other chunks that represent Π's preferences concerning price and warranty:

$$\kappa_4 : \forall x, y, z((x > y) \to (\mathrm{Fair}(z, x) \to \mathrm{Fair}(z, y)))$$

$$\kappa_5 : \forall x, y, z((x > y) \to (\mathrm{Fair}(y, z) \to \mathrm{Fair}(x, z)))$$

The deal set is a 5×11 matrix. The three statements in $\mathcal{K}^a$ mean that there are 56 possible worlds. The two beliefs are consistent with each other and with $\mathcal{K}^a$. A complete matrix for $\mathbb{P}(\mathrm{Fair}(\delta) \mid \mathcal{I}_t)$ is derived by solving two simultaneous equations of degree two using Eqn. 2:

$p =$	$w = 0$	$w = 1$	$w = 2$	$w = 3$	$w = 4$
$[7.0, \infty)$	0.0924	0.1849	0.2049	0.2250	0.2263
$[6.8, 7.0)$	0.1849	0.3697	0.4099	0.4500	0.4526
$[6.6, 6.8)$	0.2773	0.5546	0.6148	0.6750	0.6789
$[6.4, 6.6)$	0.3697	0.7394	0.8197	**0.9000**	0.9053
$[6.2, 6.4)$	0.3758	0.7516	0.8331	0.9147	0.9213
$[6.0, 6.2)$	0.3818	0.7637	0.8466	0.9295	0.9374
$[5.8, 6.0)$	0.3879	0.7758	0.8600	0.9442	0.9534
$[5.6, 5.8)$	0.3939	0.7879	0.8734	0.9590	0.9695
$[5.4, 5.6)$	0.4000	**0.8000**	0.8869	0.9737	0.9855
$[5.2, 5.4)$	0.4013	0.8026	0.8908	0.9790	0.9921
$[5.0, 5.2)$	0.4026	0.8053	0.8947	0.9842	0.9987

The two evidence values are shown above in bold face. As new evidence becomes available it is represented in $\mathcal{B}_t^a$, and the matrix is refreshed using Eqn. 3. If new evidence renders $\mathcal{B}_t^a$ inconsistent then this inconsistency will be detected by the failure of the process to yield values for the probabilities in $[0, 1]$ — if that occurs then the revision function R identifies and removes inconsistencies.

The examples discussed above for competitive interaction have been related to the basic offer exchange process. These information-theoretic tools may also be applied to capture deeper issues in negotiation such as trust [Ramchurn *et al.*, 2003]. For example, by defining *trust* as a measure of expected deviations in behaviour, the Kullback-Leibler distance[2] may be used to measure the deviation between a contract and its execution, and so to estimate trust.

7 A Co-operative Agent

Here information-based agents are applied to managing the collaboration in emergent process management. Although the agents in process management systems should attempt to co-operate, the management of emergent processes relies on good will. One operation is the delegation of responsibility for a process by one agent to another. Here Π estimates the probability that one of its collaborators will accept such responsibility — again this estimate is made solely on the basis of observations of past actions.

Π interacts with its collaborators $\{\Omega_i\}_{i=1}^n$. It is assumed that processes are initially triggered externally to the system. For example, Π's 'owner' may have an idea that she believes has value, and triggers an emergent process to explore the idea's worth. The interaction protocol is simple, if Π sends a *Delegate*$(\cdot)$ message to Ω_i then interaction continues until one agent sends an *Accept*$(\cdot)$ or a *Quit*$(\cdot)$ message. This assumes that agents respond in reasonable time which is fair in an essentially co-operative system.

To support the agreement-exchange process, Π has do two different things. First, it must respond to proposals received from Ω_i — that is not described here. Second, it must construct proposals, and possibly information, to send to Ω_i — that is described now. Maximum entropy inference is used to 'fill in' missing values with the "maximally noncommittal" probability distribution. To illustrate this suppose that Π proposes to delegate a sub-process to Ω_i. That sub-process involves Ω_i delivering — using an *Inform*$(\cdot)$ message — u chapters for a report in so-many days v. We describe the machinery for estimating the probabilities $\mathbb{P}(\mathrm{Acc}(\Omega_i, \Pi, (u, v)))$ where the predicate $\mathrm{Acc}(\Omega_i, \Pi, (u, v))$ means "Ω_i will accept Π's delegation proposal (u, v)".

Π assumes the following two preference relations for Ω_i, and $\mathcal{K}^a$ contains:

$$\kappa_6 : \forall xyz((x < y) \rightarrow (\mathrm{Acc}(\Omega_i, \Pi, (y, z)) \rightarrow \mathrm{Acc}(\Omega_i, \Pi, (x, z))))$$

$$\kappa_7 : \forall x, y, z((x < y) \rightarrow (\mathrm{Acc}(\Omega_i, \Pi, (z, x)) \rightarrow \mathrm{Acc}(\Omega_i, \Pi, (z, y))))$$

As noted in Sec. 4, these sentences conveniently reduce the number of possible worlds. The two preference relations κ_6 and κ_7 induce a partial ordering on the sentence probabilities in the $\mathbb{P}(\mathrm{Acc}(\Omega_i, \Pi, u, v))$ array. There are fifty-one possible worlds that are consistent with $\mathcal{K}^a$.

Suppose that Π has the following historical data on similar dealings with Ω_i. Two weeks ago Π asked Ω_i to deliver five chapters in three days — Ω_i refused, and offered to deliver four chapters in nine days. One week ago Π asked Ω_i to deliver two chapters in six days and Ω_i accepted this responsibility. $\mathcal{B}_t^a$ contains: $\beta_8 : \mathrm{Acc}(\Omega_i, \Pi, (5, 3))$; $\beta_9 : \mathrm{Acc}(\Omega_i, \Pi, (4, 9))$ and $\beta_{10} : \mathrm{Acc}(\Omega_i, \Pi, (2, 6))$, and assuming a 20% decay in integrity for each week: $\mathbb{P}(\beta_8) = 0.4$, $\mathbb{P}(\beta_9) = 0.6$ and $\mathbb{P}(\beta_{10}) = 0.8$

Eqn. 2 is used to calculate the distribution $\{\overline{W}|\mathcal{K}^a, \mathcal{B}_t^a\}$ which shows that there are just five different probabilities in it. The probability matrix for the proposition $\mathrm{Acc}(\Omega_i, \Pi, u, v)$ is:

$v \setminus u$	1	2	3	4	5
11	0.9937	0.9240	0.7747	0.6253	0.4852
10	0.9620	0.8987	0.7557	0.6127	0.4789
9	0.9304	0.8734	0.7367	**0.6000**	0.4726
8	0.8996	0.8489	0.7186	0.5882	0.4667
7	0.8688	0.8245	0.7004	0.5764	0.4608
6	0.8380	**0.8000**	0.6823	0.5646	0.4549
5	0.7558	0.7242	0.6261	0.5280	0.4366
4	0.6737	0.6484	0.5699	0.4914	0.4183
3	0.5916	0.5726	0.5137	0.4549	**0.4000**
2	0.3944	0.3817	0.3425	0.3032	0.2667
1	0.1972	0.1909	0.1712	0.1516	0.1333

In this array, the derived sentence probabilities for the three sentences in $\mathcal{B}_t^a$ are shown in bold type; they are exactly their given values. As described in Sec. 6, as new information becomes available this matrix may be refreshed using Eqn. 3.

8 Discussion

The object of this paper is to promote discussion on the value of this approach that treats the integrity of an agent's information as fundamental. This leads to the issue of modelling integrity decay. Entropy-based inference and information theory is used to model other agents on the basis of observations of their actions. This approach has been applied to: multi-issue bilateral bargaining, multi-issue auctions and multi-issue semi co-operative negotiation in process management. Current work is extending these ideas to argumentation including the estimation of the trust that one agent has for another, and to modelling inter-agent relationships.

References

[Debenham, 2004a] J. Debenham. Auctions and bidding with information. In P. Faratin and J.A. Rodriguez-Aguilar, editors, *Proceedings Agent-Mediated Electronic Commerce VI: AMEC*, pages 15 – 28, July 2004.

[Debenham, 2004b] J. Debenham. Bargaining with information. In N.R. 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, July 2004.

[Faratin *et al.*, 2003] 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.

[Halpern, 2003] J.Y. Halpern. *Reasoning about Uncertainty*. MIT Press, 2003.

[Jaeger, 1996] M. Jaeger. Representation independence of nonmonotonic inference relations. In *Proceedings of KR'96*, pages 461–472. Morgan Kaufmann, 1996.

[Jaynes, 1957] ET Jaynes. Information theory and statistical mechanics: Part I. *Physical Review*, 106:620 – 630, 1957.

[Jaynes, 2003] E.T. Jaynes. *Probability Theory — The Logic of Science*. Cambridge University Press, 2003.

[MacKay, 2003] D. MacKay. *Information Theory, Inference and Learning Algorithms*. Cambridge University Press, 2003.

[Paris, 1999] J. Paris. Common sense and maximum entropy. *Synthese*, 117(1):75 – 93, 1999.

[Ramchurn *et al.*, 2003] S Ramchurn, N Jennings, C Sierra, and L Godo. A computational trust model for multi-agent interactions based on confidence and reputation. In *Proceedings 5th Int. Workshop on Deception, Fraud and Trust in Agent Societies*, 2003.

[Sandholm and Vulkan, 1999] T. Sandholm and N. Vulkan. Bargaining with deadlines. In *Proceedings of the National Conference on Artificial Intelligence (AAAI)*, 1999.