

A FOUNDATION FOR INFORMED NEGOTIATION

John Debenham and Simeon Simoff
University of Technology, Sydney
PO Box 123, Broadway, NSW 2007, Australia

Keywords: Intelligent Agents, Decision Support, Agents for Internet Computing.

Abstract: Approaches to the construction of agents that are to engage in competitive negotiation are often founded on game theory. In such an approach the agents are endowed with utility functions and assumed to be utility optimisers. In practice the utility function is derived in the context of massive uncertainties both in terms of the agent's priorities and of the raw data or information. To address this issue we propose an agent architecture that is founded on information theory, and that manages uncertainty with entropy-based inference. Our negotiating agent engages in multi-issue bilateral negotiation in a dynamic information-rich environment. The agent strives to make informed decisions. The agent may assume that the integrity of some of its information decays with time, and that a negotiation may break down under certain conditions. The agent makes no assumptions about the internals of its opponent — it focuses only on the signals that it receives. It constructs two probability distributions over the set of all deals. First the probability that its opponent will accept a deal, and second that a deal will prove to be acceptable to it in time.

1 INTRODUCTION

We propose an agent architecture that is founded on information theory, and that manages uncertainty with entropy-based inference. This architecture aims to address the problem of defining a utility function for a game-theoretic agent in the context of massive uncertainties both in terms of the agent's priorities and of the raw data or information. Our agent does not necessarily have a utility function and so is not necessarily a utility optimiser — it may have less selfish aims such as simply 'playing fair'. This *Negotiating Agent, NA*, engages in bilateral bargaining with an opponent, *OP*. It strives to make informed decisions in an information-rich environment that includes market data and general sources including the Internet. *NA* attempts to fuse the negotiation with the information generated both by and because of it. It reacts to information derived from its opponent and from the environment, and proactively seeks missing information that may be of value.

NA is founded on information theory. Game theory tells us what to do, and what outcome to expect, in many well-known negotiation situations, but these strategies and expectations are derived from assump-

tions about the internals of the opponent. Game theoretic analyses of bargaining are founded on the notion of agents as utility optimizers in the presence of complete and incomplete information about their opponents (Muthoo, 1999). Two probability distributions form the foundation of both the offer evaluation and the offer making processes. They are both over the set of all deals and are based on all information available to the agent. The first distribution is the probability that any deal is acceptable to *OP*. The second distribution is the probability that any deal will prove to be acceptable to *NA* — this distribution generalizes the notion of utility.

NA makes no assumptions about the internals of *OP* in particular whether it has a utility function. *NA* does make assumptions about: the way in which the integrity of information will decay, preferences that its opponent may have for some deals over others, and conditions that may lead to breakdown. It also assumes that unknown probabilities can be inferred using *maximum entropy probabilistic logic* (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). In the absence of knowledge about *OP*’s decision-making apparatus, *NA* assumes that the “maximally noncommittal” model is the correct model on which to base its reasoning.

A *preference relation* is an assumption that *NA* makes about *OP*’s preferences for some deals over others. For example, that she prefers to pay a lower price to a higher price. A *single-issue preference relation* assumes that she prefers deals on the basis of one issue alone, independent of the values of the other issues. A preference relation may be assumed prior to the negotiation, or during it based on the offers made. For example, the opponent may display a preference for items of a certain color; (Faratin et al., 2003) describes a basis for ordering colors. The preference relations illustrated here are single-issue orderings, but the agent’s reasoning operates equally well with any preference relation as long as it may be expressed in Horn clause logic.

Bilateral bargaining is known to be inherently inefficient (Myerson and Satterthwaite, 1983). (Bulow and Klemperer, 1996) shows that a seller is better off with an auction that attracts $n+1$ buyers than bargaining with n individuals, *no matter what* the bargaining protocol is. (Neeman and Vulkan, 2000) shows that the weaker bargaining types will fare better in exchanges leading to a gradual migration. These results hold for agents who aim to optimize their utility and do limit the work described here.

2 INFORMED AGENTS

NA operates in an information-rich environment. The integrity of its 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 is drawn. Little appears to be known about how the integrity of information, such as news-feeds, decays.

One source of *NA*’s information is the signals received from *OP*. These include offers to *NA*, and the acceptance or rejection of *NA*’s offers. If *OP* rejected *NA*’s offer of \$8 two days ago then what is *NA*’s belief now in the proposition that *OP* will accept another offer of \$8 now? Perhaps it is around 0.1. A linear model is used to model the integrity decay of these beliefs, and when the probability of a decaying belief approaches 0.5¹ the belief is discarded. This choice of a linear model is independent of the bargaining method. The model of decay could be exponential, quadratic or what ever.

¹A sentence probability of 0.5 represents “maybe, maybe not”.

2.1 Interaction Protocol

The agents communicate using sentences in a first-order language \mathcal{L} . This includes the exchange, acceptance and rejection of offers. \mathcal{L} contains the following predicates: *Offer*(δ), *Accept*(δ), *Reject*(δ) and *Quit*(.), where *Offer*(δ) means “the sender is offering you a deal δ ”, *Accept*(δ) means “the sender accepts your deal δ ”, *Reject*(δ) means “the sender rejects your deal δ ” and *Quit*(.) means “the sender quits — the negotiation ends”.

Two negotiation protocols are described. First, negotiation *without decay* in which all offers stand for the entire negotiation. Second, negotiation *with decay* in which offers stand only if accepted by return — *NA* represents *OP*’s offers as beliefs with sentence probabilities that decay in time.

NA and *OP* each exchange offers alternately at successive discrete times (Kraus, 2001). They enter into a commitment if one of them accepts a standing offer. The protocol has three stages:

1. Simultaneous, initial, binding offers from both agents;
2. A sequence of alternating offers, and
3. An agent quits and walks away from the negotiation.

The negotiation ceases *either* in the second round if one of the agents accepts a standing offer *or* in the final round if one agent quits and the negotiation breaks down.

In the first stage the agents simultaneously send *Offer*(.) messages to each other. These initial offers are taken as limits on the range of values that are considered possible. This is crucial to the method described in Sec. 3 where there are domains that would otherwise be unbounded. The exchange of initial offers “stakes out the turf” on which the subsequent negotiation will take place. In the second stage an *Offer*(.) message is interpreted as an implicit rejection, *Reject*(.), of the opponent’s offer on the table.

2.2 Agent Architecture

Incoming messages from all sources are time-stamped and placed in an “In Box”, \mathcal{X} , as they arrive. *NA* has a knowledge base \mathcal{K} and a belief set \mathcal{B} . Each of these two sets contains statements in \mathcal{L} . \mathcal{K} contains statements that are generally true, such as $\forall x(\text{Accept}(x) \leftrightarrow \neg \text{Reject}(x))$ — i.e. an agent does one thing or the other. The belief set $\mathcal{B} = \{\beta_i\}$ contains statements that are each qualified with a *given sentence probability*, $\mathbf{B}(\beta_i)$, that represents an agent’s belief in the truth of the statement. These sentence probabilities may decay in time.

The distinction between the knowledge base \mathcal{K} and the belief set \mathcal{B} is simply that \mathcal{K} contains unqualified

statements and \mathcal{B} contains statements that are qualified with sentence probabilities. \mathcal{K} and \mathcal{B} play different roles in the method described in Sec. 3.

NA 's actions are determined by its "strategy". A *strategy* is a function $\mathbf{S} : \mathcal{K} \times \mathcal{B} \rightarrow \mathcal{A}$ where \mathcal{A} is the set of actions. At certain distinct times the function \mathbf{S} is applied to \mathcal{K} and \mathcal{B} and the agent does something. The set of actions, \mathcal{A} , includes sending *Offer*(.), *Accept*(.), *Reject*(.) and *Quit*(.) messages to OP . The way in which \mathbf{S} works is described in Sec. 5. Momentarily before the \mathbf{S} function is activated, a "revision function" \mathbf{R} is activated:

$\mathbf{R} : (\mathcal{X} \times \mathcal{K} \times \mathcal{B}) \rightarrow (\mathcal{K} \times \mathcal{B})$

\mathbf{R} clears the "In Box", and stores the messages *either* in \mathcal{B} with a given sentence probability *or* in \mathcal{K} .

A *deal*, δ , is a commitment for the sender to do something, τ (the sender's "terms"), subject to the receiver committing to do something, ω (the receiver's "terms"): $\delta = (\tau, \omega)$. NA may have a real-valued *utility* function: $\mathbf{U} : \mathcal{T} \rightarrow \mathbb{R}$, where \mathcal{T} is the set of terms. If so, then for any deal $\delta = (\tau, \omega)$ the expression $\mathbf{U}(\omega) - \mathbf{U}(\tau)$ is called the *surplus* of δ . An agent may be unable to specify a utility function either precisely or with certainty.² Sec. 4 describes a predicate $NAAcc(\cdot)$ that represents the "acceptability" of a deal.

NA uses three things to make offers: an estimate of the likelihood that OP will accept any offer [Sec. 3], an estimate of the likelihood that NA will, in hindsight, feel comfortable accepting any particular offer [Sec. 4], and an estimate of when OP may quit and leave the negotiation [Sec. 5].

2.3 Random Worlds

Let \mathcal{G} be the set of all positive ground literals that can be constructed using the predicate, function and constant symbols in \mathcal{L} . A *possible world* is a valuation function $\mathbf{v} : \mathcal{G} \rightarrow \{\top, \perp\}$. \mathbf{V} denotes the set of all possible worlds, and $\mathbf{V}_{\mathcal{K}}$ denotes the set of possible worlds that are consistent with a knowledge base \mathcal{K} (Halpern, 2003).

A *random world* for \mathcal{K} is a probability distribution $\mathbf{W}_{\mathcal{K}} = \{p_i\}$ over $\mathbf{V}_{\mathcal{K}} = \{\mathbf{v}_i\}$, where $\mathbf{W}_{\mathcal{K}}$ expresses an agent's degree of belief that each of the possible worlds is the actual world. The *derived sentence probability* of any $\sigma \in \mathcal{L}$, with respect to a random world $\mathbf{W}_{\mathcal{K}}$ is:

$$\mathbf{P}_{\mathbf{W}_{\mathcal{K}}}(\sigma) \triangleq \sum_n \{p_n : \sigma \text{ is } \top \text{ in } \mathbf{v}_n\} \quad (1)$$

A random world $\mathbf{W}_{\mathcal{K}}$ is *consistent* with the agent's beliefs \mathcal{B} if: $(\forall \beta \in \mathcal{B})(\mathbf{B}(\beta) = \mathbf{P}_{\mathbf{W}_{\mathcal{K}}}(\beta))$. That

²The often-quoted oxymoron "I paid too much for it, but its worth it." attributed to Samuel Goldwyn, movie producer, illustrates that intelligent agents may choose to negotiate with uncertain utility.

is, for each belief its derived sentence probability as calculated using Eqn. 1 is equal to its given sentence probability.

The *entropy* of a discrete random variable X with probability mass function $\{p_i\}$ is (MacKay, 2003): $H(X) = -\sum_n p_n \log p_n$ where: $p_n \geq 0$ and $\sum_n p_n = 1$. Let $\mathbf{W}_{\{\mathcal{K}, \mathcal{B}\}}$ be the "maximum entropy probability distribution over $\mathbf{V}_{\mathcal{K}}$ that is consistent with \mathcal{B} ". Given an agent with \mathcal{K} and \mathcal{B} , its *derived sentence probability* for any sentence, $\sigma \in \mathcal{L}$, is:

$$(\forall \sigma \in \mathcal{L}) \mathbf{P}(\sigma) \triangleq \mathbf{P}_{\mathbf{W}_{\{\mathcal{K}, \mathcal{B}\}}}(\sigma) \quad (2)$$

Using Eqn. 2, the derived sentence probability for any belief, β_i , is equal to its given sentence probability. So the term *sentence probability* is used without ambiguity.

3 ESTIMATING $\mathbf{P}(OPAcc(\cdot))$

NA does two different things. First, it reacts to offers received from OP — that is described in Sec. 4. Second, it sends offers to OP . This section describes the estimation of $\mathbf{P}(OPAcc(\delta))$ where the predicate $OPAcc(\delta)$ means "the deal δ is acceptable to OP ".

When a negotiation commences NA may have no information about OP or about prior deals. If so then the initial offers may only be based on past experience or circumstantial information.³ So the opening offers are simply taken as given.

In the four sub-sections following, NA is attempting sell something to OP . In Secs. 3.1 and 3.2 NA 's terms τ are to supply a particular good, and OP 's terms ω are money — in those examples the amount of money ω is the subject of the negotiation. In Secs. 3.3 and 3.4 NA 's terms are to supply a particular good together with some negotiated warranty period, and OP 's terms are money — in those examples the amount of money p and the period of the warranty period w are the subject of the negotiation.

3.1 One Issue — Without Decay

The unary predicate $OPAcc(x)$ means "the amount of money $\$x$ is acceptable to OP ". NA is interested in whether the unary predicate $OPAcc(x)$ is true for various values of $\$x$. NA assumes the following preference relation on the $OPAcc$ predicate:

$$\kappa_1 : \forall x, y ((x > y) \rightarrow (OPAcc(x) \rightarrow OPAcc(y)))$$

³In rather dire circumstances King Richard III of England is reported to have initiated a negotiation with remarkably high stakes: "A horse! a horse! my kingdom for a horse!" [William Shakespeare]. Fortunately for Richard, a person named Catesby was nearby, and advised Richard to retract this rash offer "Withdraw, my lord", and so Richard's intention to honor his commitments was not put to the test.

Suppose that NA 's opening offer is $\bar{\omega}$, and OP 's opening offer is $\underline{\omega}$ where $\underline{\omega} < \bar{\omega}$. Then \mathcal{K} now contains two further sentences: $\kappa_2 : \neg OPAcc(\bar{\omega})$ and $\kappa_3 : OPAcc(\underline{\omega})$. There are now $\bar{\omega} - \underline{\omega}$ possible worlds, and the maximum entropy distribution is uniform.

Suppose that NA knows its true valuation for the good, u_{na} , and that NA has decided to make an "expected-utility-optimizing" offer: $x = \frac{\bar{\omega} + u_{na}}{2}$. This offer is calculated on the basis of the preference ordering κ_1 and the two signals that NA has received from OP . The response is in terms of only NA 's valuation u_{na} and the signal $Reject(\bar{\omega})$ — it is independent of the signal $Offer(\underline{\omega})$ which implies that $\underline{\omega}$ is acceptable.

In the standard game theoretic analysis of bargaining (Muthoo, 1999), NA assumes that OP has a utility, u_{op} , that it lies in some interval $[\underline{u}, \bar{u}]$, and that the expected value of u_{op} is uniformly distributed on that interval. On the basis of these assumptions NA then derives the expected-utility-optimizing offer: $\frac{\bar{u} + u_{na}}{2}$. These two offers differ by \bar{u} in the game-theoretic result and $\bar{\omega}$ in the maximum entropy result. The game theoretic approach relies on estimates for \underline{u} and \bar{u} :

$$\mathbf{E}([\underline{u}, \bar{u}] \mid Reject(\bar{\omega}) \wedge Accept(\underline{\omega}))$$

If OP has a utility, and it may not, then if OP is rational: $\underline{u} \leq \underline{\omega} \leq \bar{u}$. The inherent inefficiency of bilateral bargaining (Myerson and Satterthwaite, 1983) shows for an economically rational OP that u_{op} , and so consequently \bar{u} , may be greater than $\bar{\omega}$. There is no reason to suspect that \bar{u} and $\bar{\omega}$ will be equal.

3.2 One Issue — With Decay

As in the previous example, suppose that the opening offers at time t_0 are taken as given and are $\underline{\omega}$ and $\bar{\omega}$. Then \mathcal{K} contains κ_1 , κ_2 and κ_3 . Suppose \mathcal{L} contains n consecutive, integer constants in the interval $[\underline{\omega}, \bar{\omega}]$, where $n = \bar{\omega} - \underline{\omega} + 1$, that represent various amounts of money. κ_1 induces a total ordering on the sentence probabilities for $OPAcc(x)$ on the interval $[\underline{\omega}, \bar{\omega}]$, where the probabilities are ≈ 0 at $\bar{\omega}$, and ≈ 1 at $\underline{\omega}$.

Suppose that at time t_1 NA makes an offer ω_{na} which is rejected by OP , who has replied at time t_2 with an offer of ω_{op} where $\underline{\omega} \leq \omega_{op} \leq \omega_{na} \leq \bar{\omega}$. At time t_3 \mathcal{B} contains $\beta_1 : OPAcc(\omega_{na})$ and $\beta_2 : OPAcc(\omega_{op})$. Suppose that there is some level of integrity decay on these two beliefs: $0 < \mathbf{B}(\beta_1) < 0.5 < \mathbf{B}(\beta_2) < 1$. Then $\mathbf{V}_{\mathcal{K}}$ contains $n + 1$ possible worlds ranging from "all false" to "all true" each containing n literals. So a random world for \mathcal{K} will consist of $n + 1$ probabilities $\{p_i\}$, where, say, p_1 is the probability of "all true", and p_{n+1} is the probability of "all false". $\mathbf{P}_{\{\mathcal{K}, \mathcal{B}\}}$ will be the distribution that maximizes $-\sum_n p_n \log p_n$ subject to the constraints: $p_n \geq 0$, $\sum_n p_n = 1$, $\sum_{n=1}^{\bar{\omega} - \omega_{na} + 1} p_n = \mathbf{B}(\beta_1)$ and

$$\sum_{n=1}^{\bar{\omega} - \omega_{op} + 1} p_n = \mathbf{B}(\beta_2).$$

The optimization of entropy, H , subject to linear constraints is described in Sec. 3.2.1 below. $\mathbf{P}_{\{\mathcal{K}, \mathcal{B}\}}$ is:

$$p_n = \begin{cases} \frac{\mathbf{B}(\beta_1)}{\bar{\omega} - \omega_{na} + 1} & \text{if } 1 \leq n \leq \bar{\omega} - \omega_{na} + 1 \\ \frac{\mathbf{B}(\beta_2) - \mathbf{B}(\beta_1)}{\omega_{na} - \omega_{op}} & \text{if } \bar{\omega} - \omega_{na} + 1 < n < \bar{\omega} - \omega_{op} + 2 \\ \frac{1 - \mathbf{B}(\beta_2)}{\omega_{op} - \underline{\omega} + 1} & \text{if } \bar{\omega} - \omega_{op} + 2 \leq n \leq \bar{\omega} - \underline{\omega} + 2 \end{cases}$$

Using Eqn. 2, for $\omega_{op} \leq x \leq \omega_{na}$:

$$\mathbf{P}(OPAcc(x)) = \mathbf{B}(\beta_1) + \frac{\omega_{na} - x}{\omega_{na} - \omega_{op}} (\mathbf{B}(\beta_2) - \mathbf{B}(\beta_1)) \quad (3)$$

These probability estimates are used in Sec. 5 to calculate NA 's next offer.

The values for $\mathbf{P}(OPAcc(x))$ in the region $\omega_{op} \leq x \leq \omega_{na}$ are derived from only two pieces of information that are the two signals $Reject(\omega_{na})$ and $Offer(\omega_{op})$ each qualified with the time at which they arrived, and the decay rate on their integrity. The assumptions in the analysis given above are: the choice of values for $\underline{\omega}$ and $\bar{\omega}$ — which do not appear in Eqn. 3 in any case — and the choice of the "maximally non-committal" distribution.

If the agents continue to exchange offers then new beliefs will be acquired and the integrity of old beliefs will decay. If the next pair of offers lies within the interval $[\omega_{op}, \omega_{na}]$ and if the integrity of β_1 and β_2 decays then the sentence probabilities of β_1 and β_2 will be inconsistent with those of the two new beliefs due to the total ordering of sentence probabilities on $[\underline{\omega}, \bar{\omega}]$ induced by κ_1 . This inconsistency is resolved by the revision function \mathbf{R} that here discards inconsistent older beliefs, β_1 and β_2 , in favor of more recent beliefs. If the agents continue in this way then the sentence probabilities for the $OPAcc$ predicate are given simply by Eqn. 3 using the most recent values for ω_{na} and ω_{op} .

The analysis given above requires that values be specified for the opening offers $\underline{\omega}$ and $\bar{\omega}$. The only part of the probability distribution that depends on the values chosen for $\underline{\omega}$ and $\bar{\omega}$ are the two "tails" of the distribution. So the choice of values for these two opening offers is unlikely to effect the estimates. The two tails are necessary to "soak up" the otherwise unallocated probability.

3.2.1 Maximizing Entropy

If X is a discrete random variable taking a finite number of possible values $\{x_i\}$ with probabilities $\{p_i\}$ then the entropy is the average uncertainty removed by discovering the true value of X , and is given by $H = -\sum_n p_n \log p_n$. The direct optimization of

H subject to a number, θ , of linear constraints of the form $\sum_n p_n g_k(x_n) = \bar{g}_k$ for given constants \bar{g}_k , where $k = 1, \dots, \theta$, is a difficult problem. Fortunately this problem has the same unique solution as the *maximum likelihood problem* for the Gibbs distribution (Pietra et al., 1997). The solution to both problems is given by:

$$p_n = \frac{\exp(-\sum_{k=1}^{\theta} \lambda_k g_k(x_n))}{\sum_m \exp(-\sum_{k=1}^{\theta} \lambda_k g_k(x_m))} \quad (4)$$

for $n = 1, 2, \dots$, where the constants $\{\lambda_i\}$ may be calculated using Eqn. 4 together with the three sets of constraints: $p_n \geq 0$, $\sum_n p_n = 1$ and $\sum_n p_n g_k(x_n) = \bar{g}_k$. The distribution in Eqn. 4 is known as *Gibbs distribution*.

Calculating the expressions for the values of $\{p_n\}$ given in the example above in Sec. 3.2 does not require the full evaluation of the expressions in Eqn. 4. That equation shows that there are just three different values for the $\{p_n\}$. Applying simple algebra to that fact together with the constraints yields the expressions given.

3.3 Two Issues — Without Decay

The above approach to single-issue bargaining generalizes without modification to multi-issue bargaining, it is illustrated with two issues only for ease of presentation. The problem considered is the sale of an item with $0, \dots, 4$ years of warranty. The terms being negotiated specify an amount of money p and the number of years warranty w . The predicate $OPAcc(w, p)$ now means “ OP will accept the offer to purchase the good with w years warranty for $\$p$ ”.

NA assumes the following two preference orderings, and \mathcal{K} contains:

$$\kappa_{11} : \forall x, y, z((x > y) \rightarrow (OPAcc(y, z) \rightarrow OPAcc(x, z)))$$

$$\kappa_{12} : \forall x, y, z((x > y) \rightarrow (OPAcc(z, x) \rightarrow OPAcc(z, y)))$$

As in Sec. 3.1, these sentences conveniently reduce the number of possible worlds. The number of possible worlds will be finite as long as \mathcal{K} contains two statements of the form: $\neg OPAcc(4, a)$ and $OPAcc(0, b)$ for some a and b . Suppose that NA 's initial offer was “4 years warranty for \$21” and OP 's initial offer was “no warranty for \$10”. \mathcal{K} now contains:

$$\kappa_{13} : \neg OPAcc(4, 21) \quad \kappa_{14} : OPAcc(0, 10)$$

These two statements, together with the restriction to integers only, limit the possible values of w and p in $OPAcc(w, p)$ to a 5×10 matrix.

Suppose that NA knows its utility function for the good with $0, \dots, 4$ years warranty and that its values are: \$11.00, \$11.50, \$12.00, \$13.00 and \$14.50 respectively. Suppose that NA uses the strategy $\mathbf{S}^{(n)}$

which is described in Sec. 5 — the details of that strategy are not important now. If NA uses that strategy with $n = 2$, then NA offers $Offer(2, \$16)$ which suppose OP rejects and counters with $Offer(1, \$11)$. Then with $n = 2$ again, NA offers $Offer(2, \$14)$ which suppose OP rejects and counters with $Offer(3, \$13)$. $\mathbf{P}(OPAcc(w, p))$ now is:

	$w = 0$	$w = 1$	$w = 2$	$w = 3$	$w = 4$
$p = 20$	0.0000	0.0000	0.0000	0.0455	0.0909
$p = 19$	0.0000	0.0000	0.0000	0.0909	0.1818
$p = 18$	0.0000	0.0000	0.0000	0.1364	0.2727
$p = 17$	0.0000	0.0000	0.0000	0.1818	0.3636
$p = 16$	0.0000	0.0000	0.0000	0.2273	0.4545
$p = 15$	0.0000	0.0000	0.0000	0.2727	0.5454
$p = 14$	0.0000	0.0000	0.0000	0.3182	0.6364
$p = 13$	0.0455	0.0909	0.1364	1.0000	1.0000
$p = 12$	0.0909	0.1818	0.2727	1.0000	1.0000
$p = 11$	0.1364	1.0000	1.0000	1.0000	1.0000

and the expected-utility-optimizing offer is: $Offer(4, \$18)$. If NA makes that offer then the expected surplus is \$0.95. The matrix above contains the “maximally non-committal” values for $\mathbf{P}(OPAcc(w, p))$; those values are recalculated each time a signal arrives. The example demonstrates how the NA is able to conduct multi-issue bargaining in a focussed way without making assumptions about OP 's internals, in particular, whether OP is aware of a utility function (Osborne and Rubinstein, 1990).

3.4 Two Issues — With Decay

Following from the previous section, suppose that \mathcal{K} contains κ_{11} , κ_{12} , κ_{13} and κ_{14} . The two preference orderings κ_{11} and κ_{12} induce a partial ordering on the sentence probabilities in the $\mathbf{P}(OPAcc(w, p))$ array [as in Sec. 3.3] from the top-left where the probabilities are ≈ 0 , to the bottom-right where the probabilities are ≈ 1 . There are fifty-one possible worlds that are consistent with \mathcal{K} .

Suppose that \mathcal{B} contains: $\beta_{11} : OPAcc(2, 16)$, $\beta_{12} : OPAcc(2, 14)$, $\beta_{13} : OPAcc(1, 11)$ and $\beta_{14} : OPAcc(3, 13)$ — this is the same offer sequence as considered in Sec. 3.3 — and with a 10% decay in integrity for each time step: $\mathbf{P}(\beta_{11}) = 0.4$, $\mathbf{P}(\beta_{12}) = 0.2$, $\mathbf{P}(\beta_{13}) = 0.7$ and $\mathbf{P}(\beta_{14}) = 0.9$. Belief β_{11} is inconsistent with $\mathcal{K} \cup \{\beta_{12}\}$ as together they violate the sentence probability ordering induced by κ_{11} and κ_{12} . Resolving this issue is a job for the belief revision function \mathbf{R} which discards the older, and weaker, belief β_{11} .

Eqn. 4 is used to calculate the distribution $\mathbf{W}_{\{\mathcal{K}, \mathcal{B}\}}$ which has just five different probabilities in it. The resulting values for the three λ 's are: $\lambda_{12} = 2.8063$, $\lambda_{13} = -2.0573$ and $\lambda_{14} = -2.5763$.

$\mathbf{P}(OPAcc(w, p))$ now is:

	$w = 0$	$w = 1$	$w = 2$	$w = 3$	$w = 4$
$p = 20$	0.0134	0.0269	0.0286	0.0570	0.0591
$p = 19$	0.0269	0.0537	0.0571	0.1139	0.1183
$p = 18$	0.0403	0.0806	0.0857	0.1709	0.1774
$p = 17$	0.0537	0.1074	0.1143	0.2279	0.2365
$p = 16$	0.0671	0.1343	0.1429	0.2849	0.2957
$p = 15$	0.0806	0.1611	0.1714	0.3418	0.3548
$p = 14$	0.0940	0.1880	0.2000	0.3988	0.4139
$p = 13$	0.3162	0.6324	0.6728	0.9000	0.9173
$p = 12$	0.3331	0.6662	0.7088	0.9381	0.9576
$p = 11$	0.3500	0.7000	0.7447	0.9762	0.9978

In this array, the derived sentence probabilities for the three sentences in \mathcal{B} are shown in bold type; they are exactly their given values.

4 ESTIMATING $\mathbf{P}(NAAcc(\cdot))$

The proposition $NAAcc(\delta)$ means: “ δ is acceptable to NA ”. This section describes how NA attaches a conditional probability to the proposition: $\mathbf{P}(NAAcc(\delta) \mid \mathcal{I}_t)$ in the light of information \mathcal{I}_t . The meaning of “acceptable to NA ” is described below. This is intended to put NA in the position “looking back on it, I made the right decision at the time” — this is a vague notion but makes sense to the author. The idea is for NA to accept a deal δ when $\mathbf{P}(NAAcc(\delta) \mid \mathcal{I}_t) \geq \alpha$ for some threshold value α that is one of NA ’s mental states.

$\mathbf{P}(NAAcc(\delta) \mid \mathcal{I}_t)$ is derived from conditional probabilities attached to four other propositions:

$\mathbf{P}(Suited(\omega) \mid \mathcal{I}_t)$,

$\mathbf{P}(Good(OP) \mid \mathcal{I}_t)$,

$\mathbf{P}(Fair(\delta) \mid \mathcal{I}_t \cup \{Suited(\omega), Good(OP)\})$ and

$\mathbf{P}(Me(\delta) \mid \mathcal{I}_t \cup \{Suited(\omega), Good(OP)\})$.

meaning respectively: “terms ω are perfectly suited to my needs”, “ OP will be a good agent for me to be doing business with”, “ δ is generally considered to be a good deal for NA ”, and “on strictly subjective grounds, δ is acceptable to NA ”. The last two of these four probabilities factor out both the suitability of ω and the appropriateness of the opponent OP . The difference between the third and fourth is that the third captures the concept of “a good market deal” and the fourth a strictly subjective “what ω is worth to NA ”. The “ $Me(\cdot)$ ” proposition is related to the concept of a private valuation in game theory.

To determine $\mathbf{P}(Suited(\omega) \mid \mathcal{I}_t)$. If there are sufficiently strong preference relations to establish extrema for this distribution then they may be assigned extreme values ≈ 0.0 or 1.0 . NA is repeatedly asked to provide probability estimates for the offer ω that yields the greatest reduction in entropy for the resulting distribution (MacKay, 2003). This continues until NA considers the distribution to be “satisfactory”. This is tedious but the “preference acquisition bot-

tleneck” appears to be an inherently costly business (Castro-Schez et al., 2004).

To determine $\mathbf{P}(Good(OP) \mid \mathcal{I}_t)$ involves an assessment of the reliability of OP . For some retailers (sellers), information — of varying reliability — may be extracted from sites that rate them. For individuals, this may be done either through assessing their reputation established during prior trades (Ramchurn et al., 2003) (Sierra and Debenham, 2005), or by the inclusion of some third-party escrow service that is then rated for “reliability” instead.

$\mathbf{P}(Fair(\delta) \mid \mathcal{I}_t \cup \{Suited(\omega), Good(OP)\})$ is determined by market data. As for dealing with $Suited$, if the preference relations establish extrema for this distribution then extreme values may be assigned. Independently of this, real market data, qualified with given sentence probabilities, is fed into the distribution. The revision function \mathbf{R} identifies and removes inconsistencies, and missing values are estimated using the maximum entropy distribution.

Determining $\mathbf{P}(Me(\delta) \mid \mathcal{I}_t \cup \{Suited(\omega), Good(OP)\})$ is a subjective matter. It is specified using the same device as used for $Fair$ except that the data is fed in by hand “until the distribution appears satisfactory”. To start this process first identify those δ that “ NA would be never accept” — they are given a probability of ≈ 0.0 , and second those δ that “ NA would be delighted to accept” — they are given a probability of ≈ 1.0 . The Me proposition links the information-theory approach with “private valuations” in game-theory.

There is no causal relationship between the four probability distributions as they have been defined, with the possible exception of the third and fourth. To link the probabilities associated with the five propositions, the probabilities are treated as epistemic probabilities and the nodes form a simple Bayesian net. The weights on the four arcs of the Bayesian net are a subjective representation of what “acceptable” means to NA . The resulting net divides the problem of estimating $\mathbf{P}(NAAcc)$ into four simpler sub-problems.

The conditionals on the Bayesian network are subjective — they are easy to specify because twelve of them are zero — that is, for the cases in which NA believes that either Me or $Suited$ is “false”. For example, if the conditionals (set by NA) are:

$$\mathbf{P}(NAAcc \mid Me, Suited, Good, Fair) = 1.0$$

$$\mathbf{P}(NAAcc \mid Me, Suited, \neg Good, Fair) = 0.1$$

$$\mathbf{P}(NAAcc \mid Me, Suited, Good, \neg Fair) = 0.4$$

$$\mathbf{P}(NAAcc \mid Me, Suited, \neg Good, \neg Fair) = 0.05$$

then, with probabilities of 0.9 on each of the four evidence nodes, the probability $\mathbf{P}(NAAcc) = 0.75$. It then remains to manage the acquisition of information \mathcal{I}_t from the available sources to, if necessary, increase $\mathbf{P}(NAAcc(\delta) \mid \mathcal{I}_t)$ so that δ is acceptable. The conditional probabilities on the net represent an agent’s priorities for a deal, and so they are specified for each

class of deal.

The *NAAcc* predicate generalizes the notion of utility. Suppose that *NA* knows its utility function \mathbf{U} . If the conditionals on the Bayesian net are as in the previous paragraph and if either $\mathbf{P}(Me(\cdot))$ or $\mathbf{P}(Suited(\cdot))$ are zero then $\mathbf{P}(NAAcc(\cdot))$ will be zero. If the conditional probabilities on the Bayesian net are 1.0 when *Me* is true and are 0.0 otherwise then $\mathbf{P}(NAAcc) = \mathbf{P}(Me)$. Then define: $\mathbf{P}(Me(\tau, \omega)) = \frac{1}{2} \times (1 + \frac{\mathbf{U}(\omega) - \mathbf{U}(\tau)}{\mathbf{U}(\bar{\omega}) - \mathbf{U}(\tau)})$ for $\mathbf{U}(\omega) > \mathbf{U}(\tau)$ and zero otherwise, where $\bar{\omega} = \max_{\omega} \mathbf{U}(\omega)$.⁴ A bargaining threshold $\alpha > 0.5$ will then accept offers for which the surplus is positive. In this way *NAAcc* represents utility-based bargaining with a private valuation.

NAAcc also is intended to be able to represent apparently irrational bargaining situations (eg: “I’ve just got to have that hat”), as well as tricky multi-issue problems such as those typical in eProcurement. It enables an agent to balance the degree of suitability of the terms offered with the reliability of the opponent and with the fairness of the deal.

5 NEGOTIATION STRATEGIES

Sec. 3 estimated the probability distribution, $\mathbf{P}(OPAcc)$, that *OP* will accept an offer, and Sec. 4 estimated the probability distribution, $\mathbf{P}(NAAcc)$, that *NA* should be prepared to accept an offer. These two probability distributions represent the opposing interests of the two agents *NA* and *OP*. $\mathbf{P}(OPAcc)$ will change every time an offer is made, rejected or accepted. $\mathbf{P}(NAAcc)$ will change as the background information changes. This section discusses *NA*’s strategy \mathbf{S} .

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 infi-

⁴The introduction of $\bar{\omega}$ may be avoided by defining $\mathbf{P}(Me(\tau, \omega)) \triangleq \frac{1}{1 + \exp(-\beta \times (\mathbf{U}(\omega) - \mathbf{U}(\tau)))}$ for $\mathbf{U}(\omega) \geq \mathbf{U}(\tau)$ and zero otherwise, where β is some constant. This is the sigmoid transfer function used in some neural networks. This function is near-linear for $\mathbf{U}(\omega) \approx \mathbf{U}(\tau)$, and is concave, or “risk averse”, outside that region. The transition between these two behaviors is determined by the choice of β .

⁵This is reminiscent of Werner Heisenberg’s indeterminacy relation, or *unbestimmtheitsrelationen*: “you can’t measure one feature of an object without changing another” — with apologies.

nite 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.

An agent’s strategy \mathbf{S} is a function of the information \mathcal{I}_t that it has at time t . That information will be represented in the agent’s \mathcal{K} and \mathcal{B} , and will have been used to calculate $\mathbf{P}(OPAcc)$ and $\mathbf{P}(NAAcc)$. Simple strategies choose an offer only on the basis of $\mathbf{P}(OPAcc)$, $\mathbf{P}(NAAcc)$ and α . The greedy strategy \mathbf{S}^+ chooses $\arg \max_{\delta} \{ \mathbf{P}(NAAcc(\delta)) \mid \mathbf{P}(OPAcc(\delta)) \gg 0 \}$, it is appropriate for an agent that believes *OP* is desperate to trade. The *expected-acceptability-to-NA-optimizing strategy* \mathbf{S}^* chooses $\arg \max_{\delta} \{ \mathbf{P}(OPAcc(\delta)) \times \mathbf{P}(NAAcc(\delta)) \mid \mathbf{P}(NAAcc(\delta)) \geq \alpha \}$, it is appropriate for a confident agent that is not desperate to trade. The strategy \mathbf{S}^- chooses $\arg \max_{\delta} \{ \mathbf{P}(OPAcc(\delta)) \mid \mathbf{P}(NAAcc(\delta)) \geq \alpha \}$, 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 to make an acceptable offer by “walking round” the iso-curve of *NA*’s previous offer (that has, say, an acceptability of $\alpha_{na} \geq \alpha$) towards *OP*’s subsequent counter offer. In terms of the machinery described here, an analogue is to use the strategy \mathbf{S}^- : $\arg \max_{\delta} \{ \mathbf{P}(OPAcc(\delta)) \mid \mathbf{P}(NAAcc(\delta) \mid \mathcal{I}_t) \gtrsim \alpha_{na} \}$ 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 $\mathbf{P}(OPAcc)$ that contains foot prints of the prior offer sequence — see Sec. 3.4 — in that distribution more recent offers have stronger weights.

6 CONCLUSIONS

The negotiating agent achieves its goal of reaching informed decisions whilst making no assumptions about the internals of its opponent. The agent is

founded on information theory is ‘driven’ by real-time information flows derived from its environment that includes the Internet. Existing text and data mining bots have been used to feed information into *NA* in experiments including a negotiation between two agents in an attempt to swap a mobile phone for a digital camera with no cash involved.

NA has five ways of leading a negotiation towards a positive outcome. First, by making more attractive offers to *OP*. Second, by reducing its threshold α . Third, by acquiring information to hopefully increase the acceptability of offers received. Fourth, by encouraging *OP* to submit more attractive offers. Fifth, by encouraging *OP* to accept *NA*’s offers. The first two of these have been described. The third has been implemented but is not described here. The remaining two are the realm of argumentation-based negotiation which is the next step in this project. The integrated way in which *NA* manages both the negotiation and the information acquisition should provide a sound basis for an argumentation-based negotiator.

(Halpern, 2003) discusses problems with the random-worlds approach, and notes particularly representation and learning. Representation is particularly significant here — for example, the logical constants in the price domains could have been given other values, and, as long as they remained ordered, and as long as the input values remained unchanged, the probability distributions would be unaltered. Learning is not an issue now as the distributions are kept as simple as possible and are re-computed each time step. The assumptions of maximum entropy probabilistic logic exploit the agent’s limited rationality by attempting to assume “precisely no more than is known”. But, the computations involved will be substantial if the domains in the language \mathcal{L} are large, and will be infeasible if the domains are unbounded. If the domains are large then preference relations such as κ_1 can simplify the computations substantially.

Much has not been described here including: the data and text mining software, the use of the Bayesian net to prompt a search for information that may lead to *NA* raising — or perhaps lowering — its acceptability threshold, and the way in which the incoming information is structured to enable its orderly acquisition (Debenham, 2004). We have not described the belief revision and the identification of those random worlds that are consistent with \mathcal{K} .

The following issues are presently being investigated. The random worlds computations are performed each time the knowledge, \mathcal{K} , or beliefs, \mathcal{B} , alter — there is scope for using approximate updating techniques interspersed with the exact calculations. The offer accepting machinery operates independently from the offer making machinery — but not vice versa — this may mean that better deals could have been struck under some circumstances.

REFERENCES

- Bulow, J. and Klemperer, P. (1996). Auctions versus negotiations. *American Economic Review*, 86(1):180–194.
- Castro-Schez, J., Jennings, N., Luo, X., and Shadbolt, N. (2004). Acquiring domain knowledge for negotiating agents: a case study. *International Journal of Human-Computer Studies*, 61(1):3 – 31.
- Debenham, J. (2004). A bargaining agent aims to ‘play fair’. In *proceedings Twenty-fourth International Conference on Innovative Techniques and Applications of Artificial Intelligence*, pages 173–186. Springer-Verlag: Heidelberg, Germany.
- Faratin, P., Sierra, C., and Jennings, N. (2003). Using similarity criteria to make issue trade-offs in automated negotiation. *Journal of Artificial Intelligence*, 142(2):205–237.
- Halpern, J. (2003). *Reasoning about Uncertainty*. MIT Press.
- Jaynes, E. (1957). Information theory and statistical mechanics: Part I. *Physical Review*, 106:620 – 630.
- Kraus, S. (2001). *Strategic Negotiation in Multiagent Environments*. MIT Press.
- MacKay, D. (2003). *Information Theory, Inference and Learning Algorithms*. Cambridge University Press.
- Muthoo, A. (1999). *Bargaining Theory with Applications*. Cambridge UP.
- Myerson, R. and Satterthwaite, M. (1983). Efficient mechanisms for bilateral trading. *Journal of Economic Theory*, 29:1–21.
- Neeman, Z. and Vulkan, N. (2000). Markets versus negotiations. Technical report, Center for Rationality and Interactive Decision Theory, Hebrew University, Jerusalem.
- Osborne, M. J. and Rubinstein, A. (1990). *Bargaining and Markets*. Academic Press.
- Pietra, S. D., Pietra, V. D., and Lafferty, J. (1997). Inducing features of random fields. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 19(2):380–393.
- Ramchurn, S., Jennings, N., Sierra, C., and Godo, L. (2003). 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*.
- Sandholm, T. and Vulkan, N. (1999). Bargaining with deadlines. In *Proceedings of the National Conference on Artificial Intelligence (AAAI)*.
- Sierra, C. and Debenham, J. (2005). An information-based model for trust. In Dignum, F., Dignum, V., Koenig, S., Kraus, S., Singh, M., and Wooldridge, M., editors, *Proceedings Fourth International Conference on Autonomous Agents and Multi Agent Systems AAMAS-2005*, pages 497 – 504, Utrecht, The Netherlands. ACM Press, New York.