Proceedings of the 24th IASTED International Multi-Conference
ARTIFICIAL INTELLIGENCE AND APPLICATIONS
February 13-16, 2006, Innsbruck, Austria

INTELLIGENT AGENTS FOR MULTI-ISSUE AUCTIONS AND BIDDING

John Debenham & Simeon Simoff
Faculty of IT
University of Technology, Sydney
PO Box 123 Broadway
Sydney, NSW, Australia
email: {debenham, simeon}@it.uts.edu.au

ABSTRACT
An approach to auctions and bidding is founded on observations and expectations of the opponents' behavior and not on assumptions concerning the opponents' motivations or internal reasoning. The approach draws ideas from information theory. A bidding agent employs maximum entropy inference to determine its actions on the basis of this uncertain data. Maximum entropy inference may be applied both to multi-issue and to single-issue negotiation. Multi-issue variants of the four common auction mechanisms are discussed.

KEY WORDS
Intelligent Agents. E-commerce.

1 Introduction

Game theory, dating back to the work of John von Neumann and Oscar Morgenstern, provides the basis for the analysis of auctions and bidding. There is a wealth of material in this analysis [1] originating with the work of William Vickrey. Fundamental to this analysis is the central role of the utility function, and the notion of rational behavior by which an agent aims to optimize its utility, when it is able to do so, and to optimize its expected utility otherwise. Analyses that are so founded on game theory are collectively referred to as *game theoretic*, or *GT*.

The application of *GT* to the design of auction mechanisms has been both fruitful and impressive — rational behavior provides a theoretical framework in which mechanism performance may be analyzed. A notable example being the supremely elegant *Generalized Vickrey* mechanism [2]. *GT* also leads to prescriptive results concerning agent behavior, such as the behavior of agents in the presence of hard deadlines [3]. The general value of *GT* as a foundation for a prescriptive theory of agent behavior is limited both by the extent to which an agent knows its own utility function, and by its certainty in the probability distributions of the utility functions (or, *types*) of its opponents.

In some negotiations — such as when an agent buys a hat, a car, a house or a company — she may not know her utility with certainty. Nor may she be aiming to optimize anything — she may simply want to buy it. Further, she may be even less certain of her opponents' types, or whether her opponents are even aware of the concept of utility. In such negotiations, an agent may be more driven towards establishing a feeling of personal "comfort" through a process of information acquisition, than by a desire to optimize an uncertain personal utility function.

A negotiation agent, Π, attempts to fuse the 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 whether to bid in an auction — on the basis of information that may be qualified by expressions of degrees of belief. Π uses this information to calculate, and continually re-calculate, probability distributions for that which it does not know. One such distribution, over the set of all possible deals, expresses Π's belief in the acceptability of a deal. Other distributions attempt to predict the behavior of its opponents — such as what they might bid in an auction. These distribution are calculated from Π's knowledge and beliefs using maximum entropy inference. 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' behavior — 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 [4]. Unknown probability distributions are inferred using *maximum entropy inference* [5] that is based on random worlds [6]. 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" [7]. As applied to the analysis of auctions, 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" [4]. This assumption referred to as Watt's Assumption [8]. So if knowledge is absent it 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 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 [6]. Fourth, maximum entropy can be tricky to calculate — although here the equivalent maximum likelihood problem for the Gibbs distribution [9] was solved numerically without incident 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. Maximum entropy inference is also independent of any structure on the set of all possible deals. So it copes with single-issue and multiple-issue negotiation without modification. It may also be applied to probabilistic belief logic. These properties are particularly useful in analyzing auctions and bidding.

The information-theory oriented analysis described here, which employs maximum entropy inference, is referred to as *ME* in contrast to *GT*.

2 Bidding Agent Π

The form of negotiation considered is between bidding agents and an auctioneer Υ in an information rich environment. The agent described here is called the *Bidding Agent*, or Π, it engages in auctions with a set of S opponents $\{\Omega_1, \cdots, \Omega_S\}$. General information is extracted from the World Wide Web using special purpose bots that import and continually confirm information that is then represented in pre-specified predicates. Π receives information by observing its opponents $\{\Omega_i\}$ and from these bots.

The integrity of information decays in time. Little appears to be known about how the integrity of information, such as news-feeds, decays. One source of information is the signals received by observing the behavior of the opponent agents both prior to a negotiation and during it. For example, if an opponent bid \$8 in an auction for an identical good two days ago then my belief that she will bid \$8 now could be 0.8. When the probability of a decaying belief approaches 0.5 the belief is discarded.

2.1 Agent Architecture

The agents communicate using the following predicate: *Bid*(.), where *Bid*(δ) means "the sender bids a deal δ". A *deal* is a pair of commitments $\delta_{\Pi:\Omega}(\pi, \omega)$ between an agent Π and an opponent agent Ω, where π is Π's commitment and ω is Ω's commitment. $\mathcal{D} = \{\delta_i\}_{i=1}^{D}$ is the deal set — ie: the set of all possible deals. If the discussion is from the point of view of a particular agent then the subscript "Π:" may be omitted, and if that agent has just one opponent that the "Ω" may be omitted as well. These commitments may involve multiple issues and not simply a single issue such as trading price. The set of *terms*, $\mathcal{T}$, is the set of all possible commitments that could occur in deals in the deal set. An agent may have a real-valued *utility* function: $\mathbf{U} : \mathcal{T} \to \Re$, that induces a total ordering on $\mathcal{T}$. For any deal $\delta = (\pi, \omega)$ the expression $\mathbf{U}(\omega) - \mathbf{U}(\pi)$ is called the *surplus* of δ, and is denoted by $\mathbf{L}(\delta)$ where $\mathbf{L} : \mathcal{T} \times \mathcal{T} \to \Re$. For example, the values of the function $\mathbf{U}$ may expressed in units of money. It may not be possible to specify the utility function either precisely or with certainty. This is addressed in Sec. 4 where a predicate *Accept*(.) represents the acceptability of a deal.

Π has a knowledge base $\mathcal{K}$ and a belief set $\mathcal{B}$. Each of these two sets contains statements in a first-order language $\mathcal{L}$. $\mathcal{K}$ contains statements that are generally true. The belief set $\mathcal{B} = \{\beta_i\}$ contains statements, β_i, that are each qualified with a *given sentence probability*, $\mathbf{B}(\beta_i)$ that represents the agent's belief in the truth of the statement. The integrity of the statements in $\mathbf{B}$ 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. This apparently odd distinction is made because $\mathcal{K}$ and $\mathcal{B}$ play different roles in the method described in Sec. 2.2.

Π's actions are determined by its "strategy". A *strategy* is a function $\mathbf{S} : \mathcal{K} \times \mathcal{B} \to \mathcal{A}$ where $\mathcal{A}$ is the set of actions. The idea is that 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, $\mathcal{A}$, is limited to sending *Bid*(.) messages to the auctioneer Υ. The way in which $\mathbf{S}$ works is described in Sec. 5. In between the discrete times at which $\mathbf{S}$ is activated, information may arrive. Incoming information from all sources is time-stamped and placed in an "In Box", $\mathcal{X}$, as it arrives. Then, momentarily before the $\mathbf{S}$ function is activated, a "revision function" $\mathbf{R}$ is activated: $\mathbf{R} : (\mathcal{X} \times \mathcal{K} \times \mathcal{B}) \to (\mathcal{K} \times \mathcal{B})$. $\mathbf{R}$ clears the "In Box", and updates $\mathcal{K}$ and $\mathcal{B}$ to ensure consistency. It is not described here.

2.2 Maximum Entropy Inference

Π uses maximum entropy inference. Let $\mathcal{G}$ be the set of all positive ground literals that can be constructed using the predicate and function symbols in $\mathcal{L}$. A *possible world* is a valuation function $\mathbf{v} : \mathcal{G} \to \{\top, \bot\}$. That is, a possible world assigns either true ($\top$) or false ($\bot$) to each ground literal in $\mathcal{G}$. $\mathbf{V}$ denotes the set of all possible worlds, and $\mathbf{V}_{\mathcal{K}}$ denotes the set of possible worlds that are consistent with the agent's knowledge base $\mathcal{K}$ [6].

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 sentence σ in $\mathcal{L}$, *with respect to* a random world $\mathbf{W}_{\mathcal{K}}$ is:

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

That is, we only admit those possible worlds in which σ is true. A random world $\mathbf{W}_{\mathcal{K}}$ is *consistent* with the agent's beliefs $\mathcal{B}$ if: $(\forall \beta \in \mathcal{B})(\mathbf{B}(\beta) = \mathbb{P}_{\mathbf{W}_{\mathcal{K}}}(\beta))$. That 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 defined in the usual way [5]: $H(X) = -\sum_n p_n \log p_n$ 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, σ, in $\mathcal{L}$, is:

$$\mathbf{P}(\sigma) = \mathbf{P}_{\mathbf{W}_{\{\mathcal{K},\mathcal{B}\}}}(\sigma) \quad (2)$$

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 from here on without ambiguity. Π uses *maximum entropy inference* which attaches the derived sentence probability to any given sentence σ.

3 Representation Dependence

ME is criticized [6] because the way in which the knowledge is formulated in $\mathcal{K}$ and $\mathcal{B}$ determines the values derived. This property is promoted here as a strength of the method because the correct formulation of the knowledge base, using the rich expressive power of first-order probabilistic logic, encapsulates features of the application at a fine level of detail.

Price is a common issue in auction and market applications. Two ways of representing price in logic are: to establish a logical constant for each possible price, and to work instead with price intervals. Admitting the possibility of an interval containing just one value, the second generalizes the first. To represent price using price intervals, we have to specify the "width" of each interval. Suppose in an application an item will be sold in excess of \$100. Suppose the predicate *TopBid*(Ω, δ) means "δ is the highest price that agent Ω is prepared to bid". This predicate will satisfy: $\forall xy((\mathit{TopBid}(\Omega, x) \wedge \mathit{TopBid}(\Omega, y)) \rightarrow (x = y))$. A crude representation of the set of possible bids is as two logical constants in $\mathcal{L}$: $[100, 200)$ and $[200, \infty)$. Following the development in Sec. 2.2, there are two positive ground literals in $\mathcal{G}$: *TopBid*(Ω, $[100, 200)$) and *TopBid*(Ω, $[200, \infty)$), and there are three possible worlds: $\{(\bot, \bot), (\top, \bot), (\bot, \top)\}$. In the absence of any further information, the maximum entropy distribution is uniform, and, for example, the probability that Ω's highest bid $\geq$ \$200 is $\frac{1}{3}$. Now if the set of possible bids had been represented as *three* logical constants: $[100, 150)$, $[150, 200)$ and $[200, \infty]$, then the same probability is $\frac{1}{4}$. Which is correct: $\frac{1}{3}$ or $\frac{1}{4}$? That depends on Π's beliefs about Ω. In both of these examples, by using *ME*, and by specifying no further knowledge about *TopBid*(.), we have implicitly asserted that the probability of each possible world being the true world is the same. In the first example all three are $\frac{1}{3}$, and in the second all four are $\frac{1}{4}$. This is what happens when the "maximally noncommittal" distribution is chosen. Conversely, if believe that: $\forall x, y(\mathbb{P}(\mathit{TopBid}(\Omega, x)) = \mathbb{P}(\mathit{TopBid}(\Omega, y)))$ then it is not necessary to include this in $\mathcal{K}$ — it is implicitly present and we should appreciate that it is so. Sec. 1 mentioned Watt's Assumption, that assumption says more than it might at first appear to.

Following from the previous paragraph with just two logical constants, suppose the predicate *MayBid*(Ω, δ) means "Ω is prepared to make a bid of δ". Assuming the Ω will prefer to pay less than more, this predicate will satisfy: $\kappa_1 : \forall x, y((\mathit{MayBid}(\Omega, x) \wedge (x \geq y)) \rightarrow \mathit{MayBid}(\Omega, y))$, where x and y are intervals and the meaning of "$\geq$" is obvious. With just κ_1 in $\mathcal{K}$ there are three possible worlds: $\{(\bot, \bot), (\top, \bot), (\top, \top)\}$. The maximum entropy distribution is uniform, and, $\mathbb{P}(\mathit{MayBid}(\Omega, [100, 200))) = \frac{2}{3}$, and $\mathbb{P}(\mathit{MayBid}(\Omega, [200, \infty])) = \frac{1}{3}$. With no additional information, $\mathbb{P}(\mathit{TopBid}(\Omega, x))$ will be uniform and $\mathbb{P}(\mathit{MayBid}(\Omega, x))$ will be linear decreasing in x.

An exemplar application is used following. It concerns the purchase of a particular 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. The meaning of the predicate *MayBid*(Ω, δ) is unchanged but δ now consists of a pair of issues and the deal set has no natural ordering. Suppose that Π wishes to apply *ME* to estimate values for: $\mathbb{P}(\mathit{MayBid}(\Omega, \delta))$ 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. Suppose then that the deal set in this application consists of 55 individual deals in the form of pairs of warranty periods and price intervals: { (w, [5.0, 5.2)), (w, [5.2, 5.4)), (w, [5.4, 5.6)), (w, [5.6, 5.8), (w, [5.8, 6.0)), (w, [6.0, 6.2)), (w, [6.2, 6.4)), (w, [6.4, 6.6)), (w, [6.6, 6.8)), (w, [6.8, 7.0)), (w, [7.0, ∞)) }, where $w = 0, \cdots, 4$. Suppose that Π has received intelligence that agent Ω is prepared to bid 6.0 with no warranty, and to bid 6.9 with one year's warranty, and Π believes this with probability 0.8. Then this leads to two beliefs: β_1 : $\mathit{TopBid}(0, [6.0, 6.2))$; $\mathbb{B}(\beta_1) = 0.8$, β_2 : $\mathit{TopBid}(1, [6.8, 7.0))$; $\mathbb{B}(\beta_2) = 0.8$. Following the discussion above, before "switching on" *ME*, Π should consider whether it believes that $\mathbb{P}(\mathit{MayBid}(\Omega, \delta))$ is uniform over δ. If it does then it includes both β_1 and β_2 in $\mathcal{B}$, and calculates $\mathbf{W}_{\{\mathcal{K},\mathcal{B}\}}$ that yields estimates for $\mathbb{P}(\mathit{MayBid}(\Omega, \delta))$ for all δ. If it does not then it should include further knowledge in $\mathcal{K}$ and $\mathcal{B}$. For example, Π may believe that Ω is more likely to bid for a greater warranty period the higher her bid price. If so, then this is a multi-issue constraint, that is represented in $\mathcal{B}$, and is qualified with some sentence probability.

4 From Utility to Acceptability

One aim of this discussion is lay the foundations for a normative theory of auctions and bidding that does not rely on knowledge of an agent's utility, and does not require an agent to make assumptions about her opponents' utilities or types, including whether they are aware of their utility [10]. Such a theory must provide some mechanism that determines the *acceptability* of a deal; ie: the probability that the deal is acceptable to an agent. Agent, Π, is attempting to buy or bid for a second-hand motor vehicle with a specific period of warranty as described in Sec. 3. Here, Π is bidding in a multi-issue auction for a vehicle, where the two issues are price and warranty period. Possible rules for this auction are described in Sec. 5.

The proposition $(Accept(\delta) \mid \mathcal{I}_t)$ means: "Π will be comfortable accepting the deal δ given that Π knows information $\mathcal{I}_t$ at time t". In an auction for terms ω, Π's strategy, S, may bid one or more π for which $\mathbb{P}(Accept((\pi, \omega)) \mid \mathcal{I}_t) \geq \alpha$ for some threshold constant α. This section describes how Π estimates: $\mathbb{P}(Accept(\delta) \mid \mathcal{I}_t)$. The meaning of $Accept(\delta)$ is described below, it is intended to put Π in the position "looking back on it, I made the right decision at the time" — this is a vague notion but makes good sense to the author.

With the motor vehicle application in mind, $\mathbb{P}(Accept(\delta) \mid \mathcal{I}_t)$ is derived from conditional probabilities attached to four other propositions: $Suited(\omega)$, $Good(\Omega)$, $Fair(\delta)$, and $Me(\delta)$. meaning respectively: "terms ω are perfectly suited to Π's needs", "Ω will be a good agent for Π to be doing business with", "δ is generally considered to be a fair deal at least", and "on strictly subjective grounds, the deal δ is acceptable to Π". These four probabilities are: $\mathbb{P}(Suited(\omega) \mid \mathcal{I}_t)$, $\mathbb{P}(Good(\Omega) \mid \mathcal{I}_t)$, $\mathbb{P}(Fair(\delta) \mid \mathcal{I}_t \cup \{Suited(\omega), Good(\Omega)\})$ and $\mathbb{P}(Me(\delta) \mid \mathcal{I}_t \cup \{Suited(\omega), Good(\Omega)\})$.
The last two of these four probabilities factor out both the suitability of ω and the appropriateness of the opponent Ω. The third captures the concept of "a fair market deal" and the fourth a strictly subjective "what ω is worth to Π". The "$Me(.)$" proposition is closely related to the concept of a private valuation in game theory. This derivation of $\mathbb{P}(Accept(\delta) \mid \mathcal{I}_t)$ from these four probabilities may not be suitable for assessing other types of deal.

$\mathbb{P}(Fair(\delta) \mid \mathcal{I}_t \cup \{Suited(\omega), Good(\Omega)\})$ is determined by reference to market data [11]. 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 $\mathcal{I}_t$ and are represented as two beliefs in $\mathcal{B}$: β_3 : $Fair(3, [6.4, 6.6))$; $\mathbb{B}(\beta_3) = 0.9$, β_4 : $Fair(3, [5.4, 5.6))$; $\mathbb{B}(\beta_4) = 0.8$. 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_3 and β_4, there are three chunks of knowledge in $\mathcal{K}$. First, κ_2 : $Fair(4, 4999)$ that determines a base value for which $\mathbb{P}(Fair) = 1$, and two other chunks that represent Π's preferences concerning price and warranty:

Figure 1. Acceptability of a deal

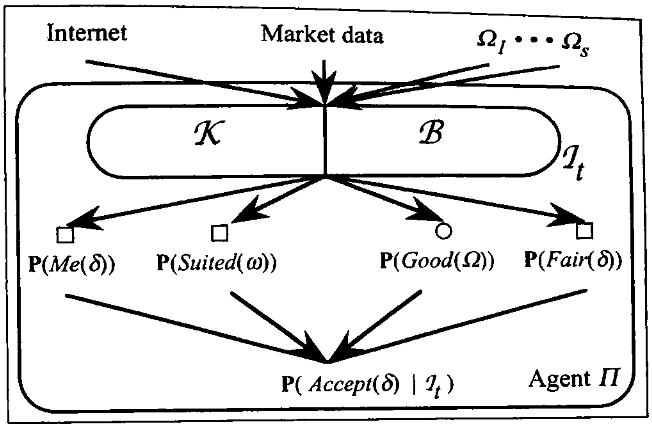

$\kappa_3 : \forall x, y, z((x > y) \rightarrow (Fair(z, x) \rightarrow Fair(z, y)))$
$\kappa_4 : \forall x, y, z((x > y) \rightarrow (Fair(y, z) \rightarrow Fair(x, z)))$
The deal set is a 5×11 matrix with highest interval $[7.0, \infty)$. The three statements in $\mathcal{K}$ mean that there are 56 possible worlds. The two beliefs are consistent with each other and with $\mathcal{K}$. A complete matrix for $\mathbf{P}(Fair(\delta) \mid \mathcal{I}_t)$ is derived by solving two simultaneous equations of degree two. As new evidence becomes available it is represented in $\mathcal{B}$, and the inference process is re-activated. If new evidence renders $\mathcal{B}$ inconsistent then this inconsistency will be detected by the failure of the process to yield values for the probabilities in $[0, 1]$. If $\mathcal{B}$ becomes inconsistent then the revision function $\mathbf{R}$ identifies and removes inconsistencies from $\mathcal{B}$ prior to re-calculating the probability distribution.

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

The two evidence values are shown above in bold face.

The whole "accept an offer" apparatus is illustrated in Fig. 1. The in-flow of information from the Internet, the market and from the opponent agents is represented as $\mathcal{I}_t$ and is stored in the knowledge base $\mathcal{K}$ and belief set $\mathcal{B}$. In that Figure the □ symbols denote probability distributions as described above, and the ∘ symbol denotes a single value. The probability distributions for $Me(\delta)$, $Suited(\omega)$ and $Fair(\delta)$ are derived as described above. *ME* inference is then used to derive the sentence probability of the $\mathbf{P}(Accept(\delta) \mid \mathcal{I}_t)$ predicate from the sentence probabilities attached to the *Me*, *Suited*, *Good* and *Fair* predicates. This derivation is achieved by two chunks of knowledge and two beliefs.

5 Auctions

The *ME* analysis of auctions focuses on what agents actually do rather than their reasons for doing what they do. The four common auction mechanisms are considered for an auctioneer, Υ, a single item and multi-issue bids each consisting of a set of deals. In the Dutch auction the auctioneer calls out successive sets of deals until one bidding agent shouts "mine". In the first- and second-price, sealed-bid mechanisms, bidding agents submit any number of multi-issue bids. The "Australian" mechanism is a variant of the common English mechanism in which agents alternately bid successive sets of deals until no further bids are received — as each set of deals is received the auctioneer identifies the current winning bid. So, unlike in the multi-issue English mechanism, in the Australian mechanism the auctioneer is not required to publicize fully her winner determination criterion in advance, and the bidders are not required to submit successive bids that are increasing with respect to that criterion. In the two sealed-bid mechanisms and the Australian mechanism the auctioneer determines the winner — and the runner up in the second-price mechanism — using a preference ordering on the set of all possible deals that may be made known to the bidding agents. The bids in these auctions may contain a large number of deals which is rather impractical.

Consider what happens from the auctioneer's point of view. Υ's expectation of what might happen will rely on both an understanding of the motivations and strategies of the agents taking part, and the rules of the auction mechanism. These two matters will effect Υ's choice of deal set, but otherwise the analysis is the same for the four common auction mechanisms. Suppose that there are S agents, $\{\Omega_i\}_{i=1}^{S}$, bidding in the auction, and the value set, $\mathcal{D} = \{\delta_i\}_{i=1}^{D}$, contains D elements. Suppose that Υ has a total preference ordering, $\gtrsim_\Upsilon$, on the deal set, $\mathcal{D}$, and that $\mathcal{D}$ is labeled such that if $i > j$ then $\delta_i \gtrsim_\Upsilon \delta_j$. Let the predicate $\mathit{TopBid}(\Omega, \delta)$ now mean "deal δ is the highest bid that Ω will make with respect to the order $\gtrsim_\Upsilon$". There are $S \times D$ ground literals in terms of this predicate. This predicate will satisfy: $\kappa_7 : \forall ixy((\mathit{TopBid}(\Omega_i, x) \wedge \mathit{TopBid}(\Omega_i, y)) \rightarrow (x = y))$. Suppose that the deal set, $\mathcal{D}$, has been chosen [see Sec. 3] so that Υ expects each of the $(D+1)^S$ possible worlds that are consistent with κ_7 to be equally probable for $\mathit{TopBid}(.)$ for each Ω_i for $i = 1, \ldots, S$.[1] The maximum entropy distribution is uniform and $\forall ij \mathbb{P}(\mathit{TopBid}(\Omega_i, \delta_j)) = \frac{1}{D+1}$. Let the predicate $\mathit{WinningBid}(\delta)$ mean "deal δ is the highest bid that the $\{\Omega_i\}_{i=1}^{S}$ will make with respect to the order $\gtrsim_\Upsilon$". Then: $\kappa_8 : \forall i(\mathit{WinningBid}(\delta_i) \leftrightarrow (\neg\exists jk \mathit{TopBid}(\Omega_j, \delta_k) \wedge (k > i)) \wedge (\exists n \mathit{TopBid}(\Omega_n, \delta_i))$. There are now $(S \times D) + D$ ground literals in terms of these two predicates, but still only $(D+1)^S$ possible worlds. So:

$$\mathbb{P}(\mathit{WinningBid}(\delta_i)) = \left(1 - \frac{D-i}{D+1}\right)^S \times \left(1 - \left(\frac{i}{i+1}\right)^S\right)$$

For example, if $S = 2$ and $D = 3$ then the probability of the highest of the three possible deals being bid by at least one of the two agents is $\frac{7}{16}$. If the total ordering $\gtrsim_\Upsilon$ is established by a utility function then this result enables the estimation of the expected utility.[2] The analysis completed so far may be applied to any sealed-bid auction, or to any open-cry auction prior to any bids being placed. Once the bidding starts in an open-cry auction, information about what agents are, or are not, prepared to bid is available. This information may alter a bidding agent's assessment of the acceptability of a deal by feeding into the *Fair*(.) predicate — see Sec. 4. It also alters the assessments of the probabilities of what the various opponents will bid, and of any deal being the winning bid. Bids made in an Australian auction provide lower limits, and bids not made in a Dutch auction provide upper limits, to what the opponents will bid.[3] As these limits change the assessment of these probabilities are revised. A formula for $\mathbb{P}(\mathit{WinningBid}(\delta_i))$ in terms of these limits is rather messy.[4] The value derived for $\mathbb{P}(\mathit{WinningBid}(\delta_i))$ relies on κ_7 and κ_8 in $\mathcal{K}$, together with expressions of the observed limits and the assumed expectation that each possible world is equally probable for *TopBid*(.).

Now consider the four auctions from a bidding agent's point of view. Two strategies, **S**, for bidding agents are described for illustration only. First, a *keen agent* who prefers to trade on any acceptable deal to missing out — they are not primarily trying to optimize anything — although in the Australian auction they may choose to bid strategically, and may attempt to reach the most acceptable deal possible. Second, a *discerning agent* who attempts to optimize expected acceptability, and is prepared to miss out on a deal as a result.

First, consider keen agents. In a first-price, sealed-bid auction these agents will bid the entire set $\{\delta \mid \mathbb{P}(\mathit{Accept}(\delta) \mid \mathcal{I}_t) \geq \alpha\}$. In an Australian, open-cry auction these agents agent may attempt to submit bids that are just "superior" to the bids already submitted by other

[1] This is the *symmetric* case when the expected performance of each of the S bidding agents is indistinguishable.

[2] In the continuous *GT* analysis, if X_i is a random variable representing the amount bid by Ω_i, and if the distributions for the X_i are uniform on [0, 1] then the expected value of the winning bid is given by the expected value of the Sth order statistic $\mathbb{E}(X_{(S)}) = \frac{S}{S+1}$.

[3] This is the *asymmetric* case.

[4] In the continuous *GT* analysis, given a sample of S non-identical, independent random variables $\{X_i\}_{i=1}^{S}$ where X_i is uniform on $[C_i, 1]$. For each sample, $p_i = \mathbb{P}(X_i \geq X) = \frac{1-X}{1-C_i}$ if $C_i \leq X \leq 1$ and zero otherwise. So the probability that *none* of the X_i exceed $Y \geq \max\{C_i\}$ is $\mathbb{P}(Y) = \prod_{j=1}^{S}(1 - p_i) = \prod_{j=1}^{S} \frac{Y-C_j}{1-C_j}$ which is the probability distribution function for the largest Y. Then $\mathbb{E}(Y) = \int_{Y=\max\{C_i\}}^{Y=1} Y \times f(Y) \times dY$ where $f(Y) = (\prod_{j=1}^{S} \frac{Y-C_j}{1-C_j}) \times (\sum_{i=1}^{S} \frac{1}{Y-C_i})$. For example, for $S = 2$, $C_1 = c$, $C_2 = d$, $0 \leq c \leq d \leq 1$ then $\mathbb{E}(Y) = \frac{4-(3\times d)-d^3+(3\times c\times(d^2-1))}{6\times(1-c)\times(1-d)}$, and if $c = d = 0$ then $\mathbb{E}(Y) = \frac{2}{3}$ as we expect.

agents. The meaning of "superior" is determined by $\gtrsim_\Upsilon$ and may be private information. If a bidding agent does not know $\gtrsim_\Upsilon$ then it will have to guess and assume it. Suppose that Δ is the set of bids submitted so far by the opponents in an Australian auction. First define the set of bids that are just superior to Δ: $\Delta^+ = \{\delta \in \mathcal{D} \mid \delta \notin \Delta, \exists \delta_1 \in \Delta, \delta \gtrsim_\Upsilon \delta_1, \forall \delta_2((\delta \gtrsim_\Upsilon \delta_2 \gtrsim_\Upsilon \delta_1) \rightarrow ((\delta_2 = \delta) \vee (\delta_2 = \delta_1)))\}$. Now bid $\{\arg\max_\delta\{\mathbb{P}(Accept(\delta) \mid \mathcal{I}_t) \mid (\mathbb{P}(Accept(\delta) \mid \mathcal{I}_t) \geq \alpha) \wedge (\delta \in \Delta^+)\}\}$. To avoid bidding against itself in a Vickrey auction an agent will bid a set of deals that forms a *shell*, Σ, with respect to $\gtrsim_\Upsilon$ [ie: $\forall \delta_i \delta_j \in \Sigma(\neg(\delta_i \gtrsim_\Upsilon \delta_j))$]. An agent will only bid in a Vickrey auction if $\gtrsim_\Upsilon$ is known, because that ordering will determine the "highest" non-winning bid. This uncertainty makes the Vickrey auction less attractive to keen agents than the other three forms. If keen agents do *not* feed bidding information into their acceptability mechanism in the open-cry cases, then the expected revenue will be greatest in the first-price, sealed-bid, followed by the Dutch and then by the Australian — it is not clear how the Vickrey auction fares due to the uncertainty in it. Feeding bidding information into the acceptability mechanisms of keen agents may have an inflationary effect on expected revenue in an Australian auction, and bidding non-information may have a deflationary effect in the Dutch auction. The extent to which these effects may change the expected-revenue ordering will be strategy-specific.

Second, consider discerning agents. A similar analysis to the above may be used by a discerning agent to optimize expected acceptability in the symmetric case. This analysis follows the general pattern of the standard *GT* analysis for utility optimizing agents — see for example [12] — it is not developed here. For a discerning agent, the Vickrey mechanism has a dominant strategy to bid at, and the Australian mechanism right up to, the acceptability margin. For the Dutch and first-price mechanisms, the acceptability of the deals bid will be shaded-down from the margin. In both the Dutch and the Australian mechanisms, the margin of acceptability may move as bidding information becomes available.

6 Conclusion

Auctions have been considered from the point of view of agents that bid because they feel comfortable as a result of knowledge acquisition, rather that being motivated by expected utility optimization. Information is derived generally from the World Wide Web, from market data and from observing the behavior of other agents in the market. The agents described do not make assumptions about the internals of their opponents. In competitive negotiation, an agent's motivations should be kept secret from its opponents. So speculation about an opponent's motivations necessarily leads to an endless counter-speculation spiral of questionable value. Maximum entropy inference is eminently suited to this requirement, and has the additional bonus of operating with logical constants and variables that represent individual deals. So the deals may be multi-issue. Four simple multi-issue auction mechanisms have been analyzed for two classes of agent: keen agents that are primarily motivated to trade, and discerning agents that are primarily motivated by the optimization of their expected acceptability. The acceptability mechanism generalizes game theoretic utility in that acceptability is expressed in terms of probabilities that are dynamically revised during a negotiation in response to both changes in the background information and the opponents' actions.

References

[1] Klemperer, P.: The Economic Theory of Auctions : Vols I and II. Edward Elgar (2000)

[2] Varian, H.: Mechanism design for computerized agents. In: Proceedings Usenix Workshop on Electronic Commerce. (1995) 13–21

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

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

[5] MacKay, D.: Information Theory, Inference and Learning Algorithms. Cambridge University Press (2003)

[6] Halpern, J.: Reasoning about Uncertainty. MIT Press (2003)

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

[8] Jaeger, M.: Representation independence of nonmonotonic inference relations. In: Proceedings of KR'96, Morgan Kaufmann (1996) 461–472

[9] Pietra, S.D., Pietra, V.D., Lafferty, J.: Inducing features of random fields. IEEE Transactions on Pattern Analysis and Machine Intelligence **19** (1997) 380–393

[10] Debenham, J.: Bargaining with information. In Jennings, N., Sierra, C., Sonenberg, L., Tambe, M., eds.: Proceedings Third International Conference on Autonomous Agents and Multi Agent Systems AAMAS-2004, ACM (2004) 664 – 671

[11] Castro-Schez, J., Jennings, N., Luo, X., Shadbolt, N.: Acquiring domain knowledge for negotiating agents: a case study. International Journal of Human-Computer Studies **61** (2004) 3 – 31

[12] Wolfstetter, E.: Topics in Microeconomics. Cambridge UP (1999)