FROM LOCAL ASSESSMENTS TO GLOBAL RATIONALITY

LOVE EKENBERG
MATS DANIELSON
MAGNUS BOMAN

The DECIDE Research Group

Department of Computer and Systems Sciences
Stockholm University and Royal Institute of Technology
Electrum 230, S-164 40 KISTA, SWEDEN
Phone: +46 8 16 1679
Fax: +46 8 703 9025
email: {lovek,mad,mab}@dsv.su.se
WWW: http://www.dsv.su.se/~mab/DECIDE.html

ABSTRACT

We present a theory and a tool for the treatment of problems arising when a decision making agent faces a situation involving a choice between a finite set of strategies, having access to a finite set of autonomous agents reporting their opinions. Each of these agents may itself be a decision making agent, and the theory is independent of whether there is a specific coordinating agent or not. Any decision making agent is allowed to assign different credibilities to the statements made by the other autonomous agents. The theory admits the representation of vague and numerically imprecise information, and the evaluation results in a set of admissible strategies by using criteria conforming to classical statistical decision theory. The admissible strategies can be further investigated with respect to strength and also with respect to the range of values that makes them admissible.

Keywords: Multi-agent systems, decision analysis, multi-criteria decision aid, uncertain reasoning

1. Introduction

Distributed AI emerged as a research field in its own right around 1980 [1] and a partition is often made into distributed problem solving systems (DPSs) and multi-agent systems (MASs). In the former, there is a global task that needs to be solved and usually a global notion of utility that can constrain the actions of the intelligent agents. In MASs, by contrast, there is no such global notion of utility (cf. [2]). Both parts of distributed AI are important to information systems. The DPS part covers the case when a coordinating agent controls a set of agents in order to accomplish some task in a distributed way. The MAS part covers the case when a set of agents must act on their own without the immediate aid from a coordinator.

The development of autonomous agent intelligence in recent years has affected the design of information systems in several ways, not least in that issues of computational complexity have come into sharper focus. One main issue of distributed artificial intelligence is how a group of agents can cooperate in order to solve different kinds of tasks and how such a system of agents can be coordinated. Not much work has been done based on results from the area of multi-criteria decision aid (MCDA) (cf. [3]) for handling such problems. Other aspects of decision theory, however, have influenced the area of MASs (cf. [4]), partly as a result of philosophical aspects of agent rationality [5], and partly because of interest in extending the principle of maximising the expected utility in efficient real-life applications [6]. The viewpoint taken in this paper is that MCDA is well suited for the treatment of systems of cooperating agents. A typical multi-criteria decision model is depicted in Fig. 1.

	C_1	C_2	…	C_k
S_1	R_{11}	R_{12}	…	R_{1k}
S_2	R_{21}	R_{22}	…	R_{2k}
	.	.	…	.
	.	.	…	.
	.	.	…	.
S_m	R_{m1}	R_{m2}	…	R_{mk}

Fig. 1. An MCDA decision model

In the figure, the C_j:s, S_i:s, and R_{ij}:s respectively denote the *criteria*, *strategies*, and *results* under consideration. If, for example, the decision situation consists of buying a car, some relevant criteria could be speed, safety, and durability. In the model, if the strategy S_1 is adopted (i.e. buying a particular car), the results obtained will be R_{11}, R_{12}, …, and R_{1k}. Each R_{1j} then denotes the outcome with respect to criteria C_j. The decision situation could then be evaluated relative to the weigths of the criteria according to the decision maker and his preferences among the R_{ij}:s.

The strategies may, for instance, be evaluated by a decision rule based on the principle of maximising the expected utility, where the expected utility of the strategy S_i, $E(S_i)$, by $c_1u_{i1} + \ldots + c_ku_{ik}$. c_j denotes the weight of the criteria C_j, and u_{ij} denotes the utility of R_{ij}.

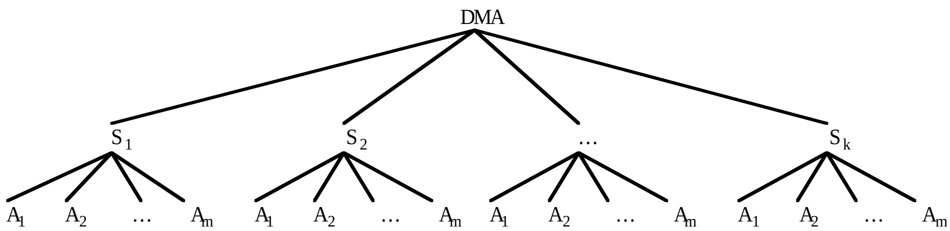

Fig. 2. A multi-agent decision model

Our work on distributed AI treats scenarios where a decision making agent (DMA), which in cooperative information systems may be a human coordinator as well as another agent process, faces a situation involving a choice between a finite set of strategies $\{S_i\}$ having access to a finite set of autonomous agents $\{A_i\}$ reporting their opinions on the strategies to the DMA, see Fig. 2.

In a situation modelled as in Fig. 2, some agents may be more reliable than others when evaluating the strategies involved, since different agents may have different capabilities to determine the utilities. The DMA may also have access to assessments expressing the credibility of the different agents.

This situation maps conveniently into the multi-criteria situation. The different attitudes of the agents correspond to the different criteria in the MCDA model in Fig. 1, and the decision can be evaluated with similar techniques. However, for the agents to carry out their tasks and to acquire sufficient and reliable knowledge *en route*, it is fundamental that they are able to evaluate information gathered from different sources, some unreliable and some noisy. The dynamic adaptation taking place over time as the agents interact with their environment, and with other agents, is affected by the means available to assess and evaluate imprecise information.

In the model in Fig. 2, the DMA is set on choosing the most preferred strategy given the agents' individual reports and their relative credibility. The statements are assumed to be assigned and revised, typically with incomplete background information, and the method we propose allows for vague and numerically imprecise assessments. Thus, the DMA may rank the credibilities of the different autonomous agents as well as quantifying them in imprecise terms. The autonomous agents have a similar expressibility regarding their respective opinions about the strategies under consideration.

Example: Assume a very simplified scenario where a set consisting of the agents A_1, A_2, A_3, and A_4 report to a decision making agent DMA on their respective preferences concerning the strategies for managing a mutual fund. Further assume that the DMA has to decide whether to make additional investments in real estate, in the stock market, or in bonds. Call these strategies S_1, S_2, and S_3, respectively. Moreover, assume that the agents A_1 through A_4 have reported to the DMA the following utility statements.[a] The utilities involved could, for example, be monetary values. In that case, they are linearly transformed to real values in the interval [0,1].

Statements according to agent A_1:

- The utility of strategy S_1 is between 0.50 and 0.70.
- The utility of strategy S_2 is between 0.10 and 0.70.
- The utility of strategy S_3 is at least 0.30.

Statements according to agent A_2:

- The utility of strategy S_1 is between 0.10 and 0.50.
- The utility of strategy S_2 is between 0.40 and 0.70.
- I have no opinion about the utility of strategy S_3.

[a] The agents may have evaluated the prospective strategies using any number of well-established financial models. In this paper, we only discuss the evaluation of the total investment situation.

Statements according to agent A_3:

- The utility of strategy S_1 is not less than that of S_2.
- The utility of strategy S_3 is between 0.50 and 0.70.

Statements according to agent A_4:

- The utility of strategy S_2 is not less than that of S_3.
- The utility of strategy S_1 is between 0.50 and 0.70.
- The utility of strategy S_2 is at most 0.70

Moreover, the DMA has estimated the credibility of A_1 through A_4 as numbers in the interval [0,1]. The number 0 denotes the lowest credibility, and 1 the highest:

- The credibility of agent A_1 is between 0.20 and 0.90.
- The credibility of agent A_2 is between 0.10 and 0.30.
- The credibility of agent A_3 is between 0.20 and 0.70.
- The credibility of agent A_4 is at most 0.50. ■

We will in this paper describe how the DMA may use the proposed method to evaluate problems such as the one above. The following section presents some earlier attempts at the treatment of vague and numerically imprecise statements. Section 3 gives an overview of a general theory for decision analysis using autonomous agents and summarises our results from the areas of decision analyses and MASs. The theory was originally developed for investigating problems of making rational decisions represented on a standard alternative–consequence form [7–10]. A significant feature of the method is that it does encourage agents not to present report statements with an unrealistic degree of precision. It provides algorithms for the efficient computation of admissible strategies and for sensitivity analyses, and an implementation of the method has been used in a number of real-life decision situations [9]. Finally, section 4 offers some concluding remarks.

2. Handling Imprecise Information

Investigations into relaxing the pointwise quantitative nature of estimates were made quite early in the modern history of probability, e.g., [11–16]. Some aspects on the relation between our work and earlier approaches to represent estimates are treated in [10], but our main concern has been how problems involving numerically imprecise information can be evaluated. A general approach to this problem has been investigated in [17]. However, Levi’s theory has some counter-intuitive implications, and is also problematic when confronted with empirical results [18]. Another suggestion that extend classical analysis includes the application of fuzzy measurement theory [19, 20] and, in particular, belief measures used in approaches related to the Dempster-Shafer theory [21]. By contrast, our approach attempts to conform to traditional statistical reasoning by using the concept of admissibility [22]. One reason is that the Dempster-Shafer theory seems to be unnecessarily strong [23]. Moreover, the possibility to state, for example, that one report is more credible than another without quantifying the credibilities is useful, particularly when handling qualitative information. Therefore, in addition to using interval statements, our approach uses inequality statements to express comparisons between credibilities and comparisons between utilities. We have generalised this work into a theory for distributed AI, applicable to decision problems involving autonomous agents [24, 25].

3. Theoretical Background

3.1. Credibility Bases

Consider the example above. Essentially, the model consists of a set of agents, a set of strategies, and two systems of statements concerning the credibilities and utilities involved. The sets of credibility and utility statements are transformed into linear systems of equations that are checked for consistency. The sets of statements constitute the *strategy base* and the *credibility base* respectively. A credibility base with k agents and m strategies is expressed in the credibility variables $\{c_1,\ldots,c_k\}$, stating the relative credibility of the different agents. The term c_j denotes the credibility of agent A_j. For instance, *the credibility of agent A_1 is between 0.20 and 0.60* can be represented by $c_1 \geq 0.20$ and $0.60 \geq c_1$.

A credibility base contains expressions about the credibility of each agent. To make the qualitative statements of credibility computable, we translate them into expressions of a certain form. Here we will handle four types of possible credibility statements.

1. Agent A_i is credible, the opinion of agent A_i is worth considering, agent A_i is not credible, etc.
2. The credibility of A_i equals a real number m, is at least m, is at most m.
3. The credibility of A_i is between the real numbers a_i and b_i.
4. The credibility of A_i is equal to the credibility of A_j, is approximately equal to that of A_j, is not less than that of A_j, etc.

The following interpretations are suggested for the above sentences. Each k_i below is a positive real number used for expressing the different statements above. Hence k_1 could be, e.g., 0.5 if we desire that for A_i to be credible, it must have a credibility of at least 0.5.

1.1 $(c_i \geq k_1, k_2 \geq c_i)$.

1.2 $(c_i \geq k_3, k_4 \geq c_i)$.

1.3 $(c_i \geq k_5, k_6 \geq c_i)$.

2.1 $c_i = m$.

2.2 $(c_i \geq m, k_7 \geq c_i)$.

2.3 $(m \geq c_i, c_i \geq k_8)$.

3.1 $(c_i \geq a_i, b_i \geq c_i)$.

4.1 $c_i = c_j$.

4.2 $c_j + k_9 \geq c_i, c_i + k_{10} \geq c_j$.

4.3 $c_i \geq c_j$.

In order for the credibility statements to be normalised, we also add the constraint $\sum c_j=1$ to the statements above. We call the conjunction of expressions of the four types above, together with $\sum c_j = 1$, the *credibility base* K(c). We will also refer to the expressions within parentheses as *intervals*.

Example (cont.): In the example above, the DMA has estimated the credibility of A_1 through A_4 as numbers in the interval [0,1]. The number 0 denotes the lowest credibility and 1 the highest. Thus the translation of the statements into a credibility base results in the following expressions. Below, we denote the expression $(x \geq k_i, k_j \geq x)$ by $x \in [k_i, k_j]$.

- $c_1 \in [0.20, 0.90]$
- $c_2 \in [0.10, 0.30]$
- $c_3 \in [0.20, 0.70]$
- $c_4 \leq 0.50$ ■

3.2. Solution Sets and Regular Consistency

A list of numbers $[n_1,\ldots,n_s]$ is a *solution vector* to a credibility base K(c) containing variables c_i, where $i = 1,\ldots,s$, if n_i can be consistently substituted for c_i in K(c). The set of solution vectors to a credibility base constitutes a *solution set*. This solution set is a convex polytope, i.e. an intersection of a finite number of closed half spaces. It seems a reasonable requirement that this polytope should not be empty; it must be possible to consistently substitute numbers for the credibility variables. (A number n_i satisfies an interval statement $(c_i \geq a_i, b_i \geq c_i)$ iff n_i, substituted for c_i, simultaneously satisfies $n_i \geq a_i$ and $b_i \geq n_i$.)

Since the DMA works with numerically imprecise information, it may be misleading only to check whether the credibility base is consistent or not. For example, when a DMA has stated that a certain credibility is within quite large bounds, implying great ignorance, it seems peculiar if the credibility base were only to be consistent with one solution vector. Therefore, the method also supports the possibility of checking if the base is consistent within certain regions. The DMA may require that there exists a minimal interval I_i for each interval statement in the credibility base with the following property: Let Y be the solution set to a credibility base K(c), let c_i be a credibility variable in K(c), let the i:th position in the solution vectors correspond to this variable, and let J_i denote the interval expressed by the interval statement containing c_i. For every number $n_i \in I_i \subseteq J_i$, there must exist numbers $n_1,\ldots,n_{i-1},n_{i+1},\ldots,n_s$, such that $[n_1,\ldots,n_{i-1},n_i,n_{i+1},\ldots,n_s] \in Y$.

When a credibility base satisfies this condition for all c_i in K(c), it is said to be *regularly consistent*. The concept of regular consistency can be described by a sequence of definitions.

Definition 1: Numbers in the interval [0,1] are associated with each interval $(c_i \geq a_i, b_i \geq c_i)$ in a credibility base K(c). The list of these numbers is denoted by $[r_i]^{K(c)}$.

Definition 2: Given a credibility base K(c) containing the intervals $(c_1 \geq a_1, b_1 \geq c_1),\ldots, (c_n \geq a_n, b_n \geq c_n)$, $\Delta(K(c))$ is the list $[l_1,\ldots,l_n]$, where $l_i = b_i - a_i$.

Definition 3: Given a credibility base K(c) containing the intervals ($c_1 \geq a_1$, $b_1 \geq c_1$),…, ($c_n \geq a_n$, $b_n \geq c_n$), *$\Delta(K(c)) \geq [r_i]^{K(c)}$* means that for each l_i in $\Delta(K(c))$ and r_i in $[r_i]^{K(c)}$, $l_i \geq r_i$.

In the definition below, we use the set union operation to denote list concatenation and the set difference operation to denote deletion of elements in a list.

Definition 4: Given a base X, containing a list Z of all interval statements [($c_1 \geq a_1$, $b_1 \geq c_1$),…, ($c_n \geq a_n$, $b_n \geq c_n$)] in X, *$\Omega(X)$* is generated by the following procedure.

- ***a.*** Let $X_0 = X$ and $Z_0 = Z$.
- ***b.*** For all interval statements ($c_i \geq a_i$, $b_i \geq c_i$) in the list Z:
  Substitute the statement ($c_i \geq \inf(\{n_i : n_i$ is at the i:th position in a solution vector to $X_{i-1}\})$, $\sup(\{n_i: n_i$ is at the i:th position in a solution vector to $X_{i-1}\}) \geq c_i$) for the interval statement ($c_i \geq a_i$, $b_i \geq c_i$) in the list Z_{i-1}, and call the result Z_i. Let $X_i = (X_{i-1} - Z_{i-1}) \cup Z_i$.
- ***c.*** Let $\Omega(X) = X_n$.

Informally, $\Omega(X)$ is like X, but contains the (in some sense) maximal intervals, where each value in the intervals is a solution relative to X.

Example (cont.): The credibility base in the example above contains four credibility variables with the properties that $c_1 \in [0.20, 0.90]$, $c_2 \in [0.10, 0.30]$, $c_3 \in [0.20, 0.70]$, and $c_4 \leq 0.50$. As we mentioned above, the credibilities are also subject to the normalisation constraint $\sum c_j = 1$. Consequently, the greatest value that can consistently be assigned to c_1 is 0.7 (the minimum value that $c_2 + c_3 + c_4$ can have is 0.3, since $c_1 + c_2 + c_3 + c_4$ should be 1). Therefore, an equivalent representation of the solution set to this base would be $\Omega(K(c)) = \{c_1 \in [0.20, 0.70], c_2 \in [0.10, 0.30], c_3 \in [0.20, 0.70], c_4 \in [0, 0.50], \sum c_j = 1\}$. Before making any calculations at all, we reduce the intervals in the base to give the (in some sense) maximal intervals, where each value in the intervals is a solution relative to the base. ■

The above definition has a computational meaning since the intervals are determined when supremum and infimum are determined. The sufficiency of computing supremum and infimum follows from two convexity observations:

(i) If Γ is a convex set, and f(x) is a convex function on Γ, then the set $\{x \in \Gamma : f(x) \geq k\}$ is convex for any real number k;

(ii) The intersection of two convex sets is also a convex set.

It is now possible to determine whether the lengths of the intervals represented by $\Omega(K(c))$ are greater than the numbers in $[r_i]^{K(c)}$ and the concept of regular consistency can be defined.

Definition 5: K(c) is *regularly consistent* given $[r_i]^{K(c)}$ iff $\Delta(\Omega(K(c))) \geq [r_i]^{K(c)}$.

Every new expression in the credibility base can be checked for regular consistency. If an earlier regularly consistent credibility base becomes inconsistent by adding a new interval ($c_i \geq a_i$, $b_i \geq c_i$), the above algorithms can be modified to give the agent information on consistent estimates by calculating the extremes of the maximal admissible interval of c_i.

Example (cont.): In the original formulation of the example, no r-values were considered, but assume that $[r_i]^{K(c)}$ is the list [0.20,0.30,0.20,0.20]. Then K(c) is not regularly consistent given $[r_i]^{K(c)}$ since the distance between the endpoints for c_2 is only 0.20 while the r-value is 0.30. ■

3.3. Strategy Bases

Strategy bases contain statements about individual agents' opinions of the utilities of different strategies, i.e., it consists of a number of expressions that express various strategy assessments. A strategy base is expressed in strategy variables $\{u_{11},\ldots,u_{1k},\ldots,u_{m1},\ldots,u_{mk}\}$ stating the utility of the strategies according to the different agents. The term u_{ij} denotes the utility of strategy S_i according to the report of agent A_j. Thus the statement *the utility of strategy S_2 is not less than that of S_3 according to the agent A_4* can be represented by the expression $u_{34} \leq u_{24}$.

We will handle five types of possible strategy statements.

Given an autonomous agent A_j:

1. The strategy S_i is desirable, S_i is fairly desirable, S_i is undesirable, etc.
2. The utility of the strategy S_i equals *m*, is at least *m*.
3. The utility of strategy S_i is between a_i and b_i.
4. The strategy S_i is as desirable (or undesirable) as strategy S_k, more desirable than S_k, the utility of S_i is approximately equal to the utility of S_k.
5. The difference in utility between S_i and S_k is not less than the difference in utility between S_m and S_n. (For simplicity, we assume that the utility of S_i is greater than the utility of S_k, and that the utility of S_m is greater than the utility of S_n.)

Statements of type 1 and 4 are qualitative, and they have to be transformed into quantitative statements. The following interpretations are suggested for the statements in 1 to 5:

1.1 $(u_{ij} \geq k_1, k_2 \geq u_{ij})$.
1.2 $(u_{ij} \geq k_3, k_4 \geq u_{ij})$.
1.3 $(u_{ij} \geq k_5, k_6 \geq u_{ij})$.
2.1 $u_{ij} = m$.
2.2 $(u_{ij} \geq m, k_7 \geq u_{ij})$.
3.1 $(u_{ij} \geq a_i, b_i \geq u_{ij})$.
4.1 $u_{ij} = u_{kj}$.
4.2 $u_{ij} \geq u_{kj}$.
4.3 $u_{kj} + k_8 \geq u_{ij}, u_{ij} + k_9 \geq u_{kj}$.
5.1 $u_{ij} + u_{nj} \geq u_{kj} + u_{mj}$.

Example (cont.): The reports provided by the agent system are translated to the following expressions.

$\mathbf{u_{11}} \in [0.50, 0.70]$	$\mathbf{u_{33}} \in [0.50, 0.70]$
$\mathbf{u_{21}} \in [0.10, 0.70]$	$\mathbf{u_{14}} \in [0.50, 0.70]$
$\mathbf{u_{31}} \geq 0.30$	$\mathbf{u_{24}} \leq 0.70$
$\mathbf{u_{12}} \in [0.10, 0.50]$	$\mathbf{u_{13}} \geq \mathbf{u_{23}}$
$\mathbf{u_{22}} \in [0.40, 0.70]$	$\mathbf{u_{24}} \geq \mathbf{u_{34}}$ ■

The system *Ω(S(u))* is a linear system of equations in the strategy variables, constructed from S(u) by a procedure similar to the construction of Ω(K(c)). Note that the treatment of the credibility and strategy bases are similar.

Definition 6: S(u) is *regularly consistent* given $[r_i]^{S(u)}$ iff $\Delta(\Omega(S(u))) \geq [r_i]^{S(u)}$.

3.4. Information Frames

A strategy base together with a credibility base constitute an *information frame*.

Definition 7: An *information frame I(c,u) representing k agents and m strategies* is a structure (K(c), $[r_i]^{K(c)}$, S(u), $[r_i]^{S(u)}$). K(c) is a credibility base with expressions in credibility variables $\{c_1,\ldots,c_k\}$, expressing the credibility of the different agents. S(u) is a strategy base with expressions in strategy variables $\{u_{11},\ldots,u_{1k},\ldots,u_{m1},\ldots,u_{mk}\}$ expressing the utility of the strategies according to the different agents. u_{ij} denotes agent A_j's opinion of strategy S_i. It is assumed that the variables' respective ranges are real numbers in the interval [0,1]. $[r_i]^{K(c)}$ and $[r_i]^{S(u)}$ are lists of real numbers in the interval [0,1].

3.5. Comparing Strategies

Relative to a particular information frame, which strategy should be chosen? To be able to answer that question, we introduce the concept of t-admissibility. The concept of a t-admissible strategy is crucial in understanding the central parts of the evaluative phase.

We define the expected utility of the strategy S_i, $E(S_i)$ or E_i in brief, by $c_1u_{i1} + \ldots + c_ku_{ik}$, and define S_i to be an *admissible strategy* if there are instances of the credibility and strategy variables that constitute a solution to $E(S_i) > E(S_j)$ for each $j \neq i$ with respect to the information frame. If only two strategies are considered, the strategy S_1 is admissible if the strategy S_2 is not strictly better for all credibility values and utilities that are consistent with the information frame. The definition of an admissible strategy is significant since it seems to be quite uncontroversial not to take strategies that are clearly worse than the other strategies involved, into serious account. The definition also conforms to the usual one in statistical decision theory and the problem of finding an admissible strategy can be formulated as an ordinary quadratic programming (QP) problem.

Definition 8: As before, let the decision situation involve the credibility variables $c_1,\ldots,c_k$, and let u_{ij} denote agent A_j's opinion of strategy S_i, then the expected utility $E_i =_{def} E(S_i) =_{def} c_1u_{i1} + \ldots + c_ku_{ik}$.

Definition 9: Given an information frame I(c,u) and a real number $t \in [-1,1]$, *S_i is at least as t-good as S_j* iff $E(S_i) - E(S_j) - t \geq 0$ for all instances of the credibility and strategy variables that are solutions to $K(c) \cup S(u)$.

Example (cont.): When comparing the strategies S_2 and S_1, the expression $E(S_2) - E(S_1) \geq t$ is considered. The expression could then for instance be tested for $t = 0$. If S_2 is at least as 0-good as S_1, we may increase *t* to 0.1 and investigate whether S_2 is at least as 0.1-good as S_1. The more *t* is increased, the harder it is for S_2 to be at least as t-good as S_1. In this sense the t-values express the ‘‘strengths’’ of the different strategies. Needless to say, we may also find the greatest value of *t* such that a strategy is as least as t-good as another strategy. It is not possible for S_2 to be at least as 0.4-good as S_1, since $E(S_2) - E(S_1) = c_1(u_{21}-u_{11}) + c_2(u_{22}-u_{12}) + c_3(u_{23}-u_{13}) + c_4(u_{24}-u_{14}) \leq c_1(0.2) + c_2(0.6) + c_4(0.5) \leq 0.04 + 0.18 + 0.15 < 0.4$. ■

Definition 10: Given an information frame I(c,u) and a real number $t \in [-1,1]$, *S_i is t-better than S_j* iff S_i is at least as t-good as S_j and there exists at least one instance of the credibility and strategy variables that is a solution to $\{E(S_i) - E(S_j) - t > 0\} \cup K(c) \cup S(u)$.

Definition 11: Given an information frame I(c,u) and a real number $t \in [-1,1]$, *S_i is a t-admissible$_1$ strategy* iff no other strategy is t-better.

The definition of a t-admissible strategy, with t = 0, conforms to the usual one in statistical decision theory [22]. However, for computational reasons, we will state the definition of t-admissibility somewhat differently.

Definition 12: Given an information frame I(c,u) and a real number $t \in [-1,1]$, S_i is a *t-admissible$_2$ strategy* iff for each $j \neq i$:

(i) there are instances of the credibility and strategy variables that constitute a solution to $\{E(S_i) - E(S_j) + t > 0\} \cup K(c) \cup S(u)$, or

(ii) there is no instance of the credibility and strategy variables that constitutes a solution to $\{E(S_j) - E(S_i) - t > 0\} \cup K(c) \cup S(u)$.

Theorem 1: Given an information frame I(c,u) and a real number $t \in [-1,1]$, S_i is a t-admissible$_1$ strategy iff S_i is a t-admissible$_2$ strategy.

To reduce the computational complexity in solving the QP problems we use methods from [8] and [26]. It is possible to determine if a strategy is admissible by reducing the problem to linear systems in the credibility variables, and then applying ordinary linear programming algorithms.

To simplify checks for consistency used in the definitions below, we use the constructive proof of a theorem from [8].

Theorem 2: Given an information frame I(c,u). Let $E(S_i) - E(S_j) + t > 0$ be a bilinear inequality in credibility and strategy variables in K(c) and S(u). The system $A(c,u,t) = \{E(S_i) - E(S_j) + t > 0\} \cup K(c) \cup S(u)$ can be reduced to a disjunction of linear systems in the credibility variables, $L(i,j,t) = \{L_1,\ldots,L_N\}$, with the following property: A(c,u,t) has a solution iff L(i,j,t) has a solution.

Hence, we can determine if a strategy is t-admissible$_2$ by reducing $\{E(S_i) - E(S_j) + t > 0\} \cup K(c) \cup S(u)$ $(i \neq j)$ to linear systems in the credibility variables for every i, and apply linear programming algorithms.

The definition of t-admissible$_2$ strategy seems reasonable, but we will make the concept of t-admissibility even stronger. As in the discussion about regular consistency for K(c) and S(u), we demand that in order to be t-admissible, there should not only be pointwise solutions.

Definition 13: Given an information frame I(c,u). The system A(c,u,t) is *regularly solvable given I(c,u)* iff $\Delta(\Omega(L_s)) \geq [r_i]^{K(c)}$ for some L_s in L(i,j,t). L(i,j,t) is as in the theorem above.

Definition 14: Given an information frame I(c,u). *S_i is a t-admissible$_3$ strategy given I(c,u) iff* for all j≠i, the systems A(c,u,t) are regularly solvable given I(c,u).

Depending on what expressions are included in the information frame, several other techniques are applicable [26].

3.6. The Concept of Proportion

Determining the set of admissible strategies is often not enough, since in non-trivial decision situations this set is too large, i.e. the admissible strategies are too numerous and the DMA cannot adequately discriminate between them. Moreover, when approaching an estimation problem, the autonomous agents as well as the DMA are encouraged to be deliberately imprecise, and thus values close to the boundaries of the different interval estimates seem to be the least reliable ones. Hence, a problem with the approach above is that the procedure for determining if a strategy is admissible is too insensitive to the different intervals involved. Therefore, we suggest the use of further discrimination principles.

One solution is to study in how large parts of the bases K(c) and S(u). the expression $E(S_i) - E(S_j) + t > 0$ is consistent. For example, a strategy could be admissible for 90 percent of the solution vectors to the strategy base, and for 80 percent of the solution vectors to the credibility base. This can be determined for example by a Monte Carlo method as suggested in [8]. However, this approach is inefficient and as a complement to Monte Carlo methods we suggest the use of the concept of *proportion*. The idea behind this concept is to investigate how much the different intervals can be decreased before the expression $E(S_i) - E(S_j) + t > 0$ ceases to be consistent with the bases. By this procedure we can study the stability of a result by gaining a better understanding of how important the boundary points are for the result. The best way to understand the concept of proportion is to regard it as a negative measure. In many situations all strategies become t-admissible$_3$, for, e.g., t = 0.1, and it is then interesting to investigate when a strategy ceases to be t-admissible$_3$. The strategy that should be chosen given the information frame could, for example, be the remaining one when the other strategies have ceased to be t-admissible$_3$. By integrating the concept of proportion with the procedures for handling admissibility above, a procedure that takes account of the amount of consistent instances of credibility and strategy variables where the strategies are admissible can be defined. The following presents an example of how this can be done. Needless to say, this is not the only possible candidate and alternative concepts are defined and investigated in [26].

The concept of proportion is defined in three steps. First we define a procedure for decreasing the size of all the intervals in a base maintaining regular consistency. The size of the steps could for instance be proportional to the sizes of the original intervals.

Definition 15: X is a system of linear equations, inequalities, and interval statements containing variables $\{x_1,\ldots,x_n\}$, and D = $[d_i]$ is a list of *n* numbers. A *D-reduction of X to X* is performed as follows:

- **a.** Let $X_0 = X$.
- **b.** For all $x_i \in \{x_1,\ldots,x_n\}$:
  Let $X'_i = X_{i-1}$, and replace the interval statement $(x_i - a_i \geq 0, -x_i + b_i \geq 0)$ in X'_i by $(x_i - a_i - d_i \geq 0, -x_i + b_i - d_i \geq 0)$, where $d_i \in$ D. Call this X_i and check if it is regularly consistent. If not, let $X_i = X_{i-1}$.
- **c.** Let $X^* = X_n$.

Hence, given a list of d_i's, a D-reduction decreases the size of all the intervals that can possibly be decreased while maintaining regular consistency. This procedure can now be iterated as far as possible. The number δ in the procedure below should be suitable for the given decision situation.

Definition 16: Given a system X of linear equations, inequalities, and interval statements, and a number δ in the interval [0, 0.5[. Let $[(x_i - a_i \geq 0, -x_i + b_i \geq 0)]$ be the list of interval statements in Ω(X), and let D = $[d_i]$ is a list where $d_i = \delta(b_i - a_i)$.

- **a.** A *δ_0-reduction of X to Y_0* is to let $Y_0 = \Omega(X)$.
- **b.** Given a δ_{i-1}-reduction of X to Y_{i-1}. A *δ_i-reduction of X to Y_i* is the D-reduction of Y_{i-1} to Y_i, where $Y_{i-1} \neq Y_i$.

Example (cont.): In the example, Ω(K(c)) = $\{c_1 \in [0.20, 0.70], c_2 \in [0.10, 0.30], c_3 \in [0.20, 0.70], c_4 \in [0, 0.50], \Sigma c_j = 1\}$. Assume that δ = 0.1. Then D becomes the list [0.05, 0.02, 0.05, 0.05]. Thus, a δ_1-reduction of K(c) is $\{c_1 \in [0.25, 0.65], c_2 \in [0.12, 0.28], c_3 \in [0.25, 0.65], c_4 \in [0.05, 0.45], \Sigma c_j = 1\}$. ■

The proportion expresses how far these reductions can be performed, and the following definitions are related to the concept of t-admissible$_3$ strategies.

Definition 17: Given an information frame I(c,u). The *δ_t-proportion of $E_i - E_j + t > 0$ given I(c,u)* is the number k = max({a : there is a δ_s-reduction of (K(c) ∪ L(i,j,t)) given $[r_i]^{K(c)}$}), L(i,j,t) as above.

Definition 18: Given an information frame I(c,u). The *δ_t-proportion of the strategy S_i given I(c,u)* is the number min({the δ_t-proportion of $E_i - E_j + t > 0$ given I(c,u) for every $j \neq i$}).

The comparison between the strategies shows in how large a δ_t-proportion the strategies are t-admissible$_3$. The smaller δ is, the better the estimate.

Example (cont.): Evaluating the example above results in the diagrams in Figs. 3–5. Fig. 3 shows the proportions of strategy S_1. The heights of the bars belonging to each strategy show the proportions of the information frame in which the strategy is admissible, i.e. how far the different intervals can be decreased before the strategies cease to be admissible. The variation along the horizontal axis indicates the different proportions with respect to t-values from –1 to 1. A reasonable strategy should have tall bars as far to the left as possible.

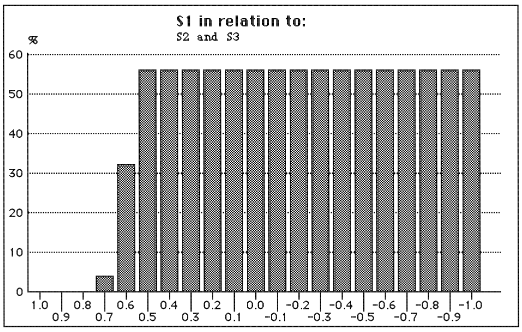

Fig. 3. The proportions of S_1.

Figs. 4 and 5 show the proportions of the strategies S_2 and S_3, respectively. Hence, strategy S_2 is definitely inferior to both S_1 and S_3, but we can see that strategy S_3 is only slightly better than S_1, and a further investigation is recommended in order to identify critical variables.

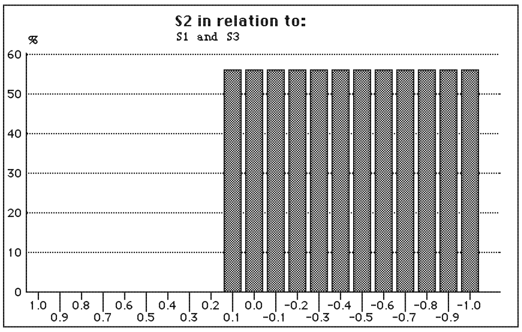

Fig. 4. The proportions of S_2.

.

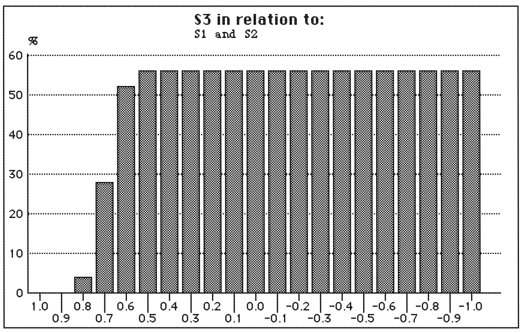

Fig. 5. The proportions of S_3.

It is natural to check how sensitive the different proportions are to changes in the information frame. The DMA can simultaneously vary any number of intervals to discover credibility or strategy variables that are especially critical.

Assume that the DMA wants to investigate whether it is meaningful to allocate resources to agent A_1 for collecting supplementary information about strategy S_3. Before doing that, the DMA wants to investigate how influential the report from the agent is. For instance, the DMA can restrict the maximum utility of u_{31} to 0.6 instead of 1 and evaluate the modified decision situation. Fig. 6 shows the result for strategy S_3. The strategy is now slightly worse than strategy S_1. The new information does not change the results shown in Fig. 3 or Fig. 4.

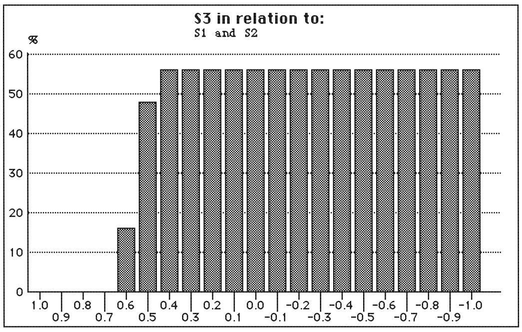

Fig. 6. The result of the modification.

The DMA may now interactively proceed in this way to investigate critical reports in order to gain a better understanding of the decision problem and finally reach a conclusion. ■

4. Concluding Remarks

In this paper, we present a method for the evaluation of how reports from sets of autonomous agents from a decision analytical viewpoint. We have shown how a decision making agent can make use of imprecise and possibly incomplete reports made by different autonomous agents when deciding which action strategy to adopt. The approach models a decision problem in a structured way by using informal and numerically imprecise statements. These statements are then translated into a suitable representation and the strategies are evaluated with respect to admissibility. The set of admissible strategies is usually too large after a first evaluation and the situation needs to be analysed with respect to further discriminating principles. Therefore, the concept of proportion is introduced as a complement to Monte Carlo approximations. The proportion indicates how much the different intervals can be reduced before different strategies cease to be admissible.

To allow for a flexible analysis of a decision situation, a tool such as the one described in this paper should contain the possibility of analysing the situation with respect to several parameters. Since our work includes efficient evaluation of non-trivial maximisation problems, our method and implementations thereof are well suited for use in the reasoning mechanisms of more sophisticated agent based information systems, and it is quite straightforward to include a multitude of decision rules in this framework.

The evaluation principle described is based on the principle on maximising the expected utility, since that principle is at the core of rational agent behaviour. However, this principle is not the only reasonable candidate for a decision rule. It has often been argued that various axiomatic theories proposed to support the principle [27] are too strong, cf.[28]. It has also been demonstrated that several axiom systems are too weak to imply the principle and that it seems very hard to design such a system [29]. A number of other criteria have been suggested in, e.g., [30], [31], and [32]. Therefore, it seems necessary to extend a framework based on the principle of maximising the expected utility by other criteria. One such extension is accomplished by using security levels for generalising rules such as maximin and minimax [33]. While a certain evaluation of a strategy results in an acceptable mean value, the consequences might be so dire that the strategy should be avoided. It might, for example, endanger the entire purpose of the system, and in that case even a report with a low credibility is too risky to neglect. In order to attain a high level of security when allowing the agents their autonomy and to be able to trust evaluations based on their reports, we argue that security levels should be imposed on the reports [34].

References

1. A.H. Bond and L. Gasser, *Readings in Distributed Artificial Intelligence*, (Morgan Kaufmann, 1988).
2. J.S. Rosenschein, Consenting Agents: Negotiation Mechanisms for Multi-Agent Systems, *Proc. 13th IJCAI* (1993) 792–799.
3. P. Vincke, *Multicriteria Decision Aid* (John Wiley and Sons, 1992), [First published in French 1989].
4. S.J. Russell and D. Subramanian, Provably Bounded-Optimal Agents, *Journal of Artificial Intelligence Research* **2** (1995) 575–609.
5. J. Doyle, Rationality and its Roles in Reasoning, *Computational Intelligence* **8:2** (1992) 376–409.
6. M. Boman, Rational Decisions and Multi-Agent Systems, *Proc. Working Notes for the AAAI Fall Symposium on Rational Agency*, ed. Fehling (MIT Technical Report, 1995).
7. P-E. Malmnäs, *From Qualitative to Quantitative Probability*, Stockholm Studies in Philosophy Vol. 7 (Almqvist & Wiksell Int, 1981).
8. P-E. Malmnäs, *Towards a Mechanization of Real Life Decisions*, in *Logic and Philosophy of Science in Uppsala*, eds. Prawitz and Westerståhl (Kluwer Academic Publishers, 1994).
9. L. Ekenberg, S. Oberoi and I. Orci, A Cost Model for Managing Information Security Hazards, *Computers & Security* **14** (1995) 707–717.
10. L. Ekenberg and M. Danielson, A Support System for Real-Life Decisions in Numerically Imprecise Domains, *Proc. International Conference on Operations Research '94* (Springer-Verlag, 1994) 500–505.
11. C.A.B. Smith, Consistency in Statistical Inference and Decision, *Journal of the Royal Statistic Society Series B* **xxiii** (1961) 1–25.
12. I.J. Good, *Subjective Probability as the Measure of a Non-measurable Set*, in *Logic, Methodology, and the Philosophy of Science,* eds. Suppes, Nagel and Tarski (Stanford University Press, 1962) 319–329.
13. G. Choquet, Theory of Capacities, *Ann. Inst. Fourier* **5** (1953/54) 131–295.
14. P.J. Huber and V. Strassen, Minimax Tests and the Neyman-Pearsons Lemma for Capacities, *Annals of Statistics* **1** (1973) 251–263.

15. P.J. Huber, The Case of Choquet Capacities in Statistics, *Bulletin of the International Statistical Institute* **45:4** (1973) 181–188.
16. A.P. Dempster, Upper and Lower Probabilities Induced by a Multivalued Mapping, *Annals of Mathematical Statistics* **xxxviii** (1967) 325–339.
17. I. Levi, On Indeterminate Probabilities, *The Journal of Philosophy* **71** (1974) 391–418.
18. P. Gärdenfors and N-E. Sahlin, Unreliable Probabilities, Risk Taking, and Decision Making, *Synthese* **53** (1982) 361–386.
19. S-J. Chen and C-L. Hwang, *Fuzzy Multiple Attribute Decision Making*, Lecture Notes in Economics and Mathematical Systems. Vol. 375 (Springer-Verlag, 1992).
20. Y-J. Lai and C-L. Hwang, *Fuzzy Multiple Objective Decision Making*, Lecture Notes in Economics and Mathematical Systems. Vol. 404 (Springer-Verlag, 1994).
21. G. Shafer, *A Mathematical Theory of Evidence* (Princeton University Press, 1976).
22. E.L. Lehmann, *Testing Statistical Hypothesis* (John Wiley and Sons, 1959).
23. K. Weichselberger and S. Pöhlman, *A Methodology for Uncertainty in Knowledge-Based Systems*, Lecture Notes in Artificial Intelligence (Springer-Verlag, 1990).
24. L. Ekenberg and M. Danielson, A Tool for Handling Uncertain Information in Multi-Agent Systems, *Proc. MAAMAW'94* (Springer-Verlag, 1995).
25. L. Ekenberg, M. Boman and M. Danielson, A Tool for Coordinating Autonomous Agents with Conflicting Goals, *Proc. ICMAS'95* (AAAI/MIT Press, 1995) 89–93.
26. M. Danielson and L. Ekenberg, A Framework for Analysing Decisions under Risk (1996) submitted.
27. P. Fishburn, Subjective Expected Utility: A Review of Normative Theories, *Theory and Decision* **13** (1981) 139–199.
28. M. Allais, *Fondements d'une Théorie Positive des Choix Comportant un Risque et Critique des Postulats et Axioms de L'Ecole Americane* (D. Reidel Publishing Company, 1953) [Translated into The Foundations of a Positive Theory of Choice involving Risk and a Criticism of the Postulates and Axioms of the American School in Expected Utility Hypothesis and the Allais Paradox 1979, D. Reidel Publishing Company].
29. P-E. Malmnäs, Axiomatic Justification of the Utility Principle, *Synthese* **99** (1994) 233–249.
30. A. Wald, *Statistical Decision Functions* (John Wiley and Sons, 1950).
31. L. Hurwicz, Optimality Criteria for Decision Making under Ignorance, *Cowles Commission Discussion Paper* **370** (1951).
32. L. Savage, The Theory of Statistical Decision, *Journal of the American Statistical Association* **46** (1951) 55–67.
33. L. Ekenberg, M. Danielson and M. Boman, Decision Support Using Security Levels for Autonomous Agents, to appear in *Decision Support Systems International Journal* (1996).
34. M. Boman and L. Ekenberg, Decision Making Agents with Relatively Unbounded Rationality, *Proc. DIMAS'95* (1995) I/28–I/35.