A

Multi-Agent Systems

This appendix deals with an application area of computational decision analysis, the area of multi-agent systems. The content of the appendix is joint work with Magnus Boman, DSV, and Love Ekenberg, IIASA. The text is partly derived from [EDB96b] and [EDB97].

Distributed AI (DAI) emerged as a research field in its own right around 1980 [BG88] and a partition is often made into distributed problem-solving systems (DPSs) and multi-agent systems (MASs). Both parts of DAI are important to software 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 immediate aid from a coordinator. 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 [R93].

Theories of intelligent agents offer means for dealing with the complexity inherent in developing distributed systems, and the advances in DAI over the last five years have affected the design methods of distributed software in several ways. One main issue in DAI is how a group of agents can cooperate in order to solve different tasks and how such a system of agents can be coordinated. Some aspects of decision theory have influenced the area of MASs [RS95], partly as a result of philosophical aspects of agent rationality [D92], and partly because of interest in extending the principle of maximising the expected value in efficient real-life applications [B95].

The idea in this appendix is to demonstrate that a method for evaluating reports from sets of autonomous agents from a decision analytical viewpoint can be built around the DELTA concepts. A decision-making agent (DMA) may make use of imprecise and possibly incomplete reports made by different autonomous agents when co-ordinating its activities and deciding which action strategy to adopt. In a manner similar to the standard DELTA method, these reports are translated into a suitable representation and the strategies are evaluated. The set of non-dominated strategies is usually too large after a first evaluation and the situation needs to be analysed with respect to further discriminating principles. To allow the DMA to make a flexible analysis of its decision situation, a method such as the one described here ought to contain the possibility of analysing the situation in several respects. Since DELTA includes efficient evaluation of non-trivial decision problems, the 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.

This particular application considers a decision problem with respect to the contents and the credibilities of the received reports. These two aspects are modelled in an agent decision frame consisting of two systems of translated interval statements, similar to an ordinary decision frame. Once it is decided that a set of agents should achieve some goal, and some semantic mapping has been provided for any syntactically heterogeneous subsets of information deemed to be of interest, then the possibility of a disagreement must be considered. This is the problem of coordinating incomplete and possibly conflicting reports made by autonomous agents, with the purpose of reaching a decision on which action to take.

Rational decision-making is weakly defined in [S76b] as the process of choosing among a finite number of acts by a series of steps that

(i) lists the acts,
(ii) determines all their consequences, and
(iii) makes a comparative evaluation.

Although the definition is of little use as such, its weaknesses make it suitable for use as a proviso for some points made in this appendix. Note that the term *act* is loosely used, and the concept of *strategy* is used here instead. A more detailed discussion can be found in [L92]. A classical problem concerning (iii) is that there exists no absolute notion of rational decision-making. Rather, rationality is usually interpreted as meaning that any agent behaviour being sub-optimal with respect to the goal is either accidental or unavoidable. To explicate this interpretation, one may turn to the first presidential address of AAAI [N81] which has been influential in spreading the agent metaphor. Drawing upon ideas put forward by McCarthy in the late 1950s, Newell suggested his principle of rationality: "*If an agent has knowledge that one of its actions will lead to one of its goals, then the agent will select that action […] The principle of rationality provides, in effect, a general functional equation for knowledge. The problem for agents is to find systems at the symbol level that are solutions to this functional equation, and hence can serve as representations of knowledge* […] *The principle of rationality corresponds at the symbol level to the processes (and associated data structures) that attempt to carry out problem-solving to attain the agent's goals.*" [op. cit. pp.8–14].

The concept of rationality was initially treated in the MAS area as merely another property that agents could have, along with, e.g., autonomy, mobility, and benevolence (cf. Chapter 2 of [G86]). This development undoubtedly came about as a reaction to the view proposed earlier by traditional AI that cognitive capabilities are more important than an agent's means to communicate, react, or adapt. In the most extreme MAS frameworks, rationality is treated as an emergent feature of an agent system [B86].

The prime evaluation principle suggested is based on the principle of maximising the expected value (PMEV) since that principle is at the core of rational agent behaviour.[1] In the last few years, several researchers within DAI have equated rationality with the use of PMEV as a decision rule (see, e.g., [GD93]). However, this principle is not the only reasonable candidate for a decision rule. There are many reasons not to identify rationality with the PMEV, some of them well-known to game theorists [BE95]. The unrealistic assumption that the perfectly rational (or even hyper-rational, see [R92], p.107) players of the game have full knowledge of the game structure, and of the rationality of their opponents, is necessary to attain the desired equilibria [BC92]. Even if one accepts that game-theoretical decision rules cannot always provide useful advice to agents in non-ideal games, a view now seemingly assumed in computer science [R93], there remain difficult problems to face [M92].

As mentioned in Chapter 5, a number of other rules have been suggested by various researchers. One conclusion from that chapter is that it seems plausible to *supplement* a method based on PMEV with other rules. The strategies might be evaluated relative to a set of security levels considering how risky the strategies are. Moreover, it can be investigated in which parts of the hull those conditions are met. This is accomplished by using contractions for security levels as well.

Agent Modelling

In the agent model that underlies this approach, the DMA[2] 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

[1] To be more precise, it should be called the principle of maximising the reported value. The credibilities represent importance weights given to individual reports. The aggregation should therefore be considered a weighted report value rather than an expectation.

[2] A DMA may be a human coordinator as well as another agent process.

opinions on the strategies to the DMA, see Figure A.1. Each of these agents may itself play the role of decision making agent, and the theory is independent of whether there is a specific coordinating agent or not. In other words, the focus in this appendix on a particular DMA is a matter of convenience. 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.

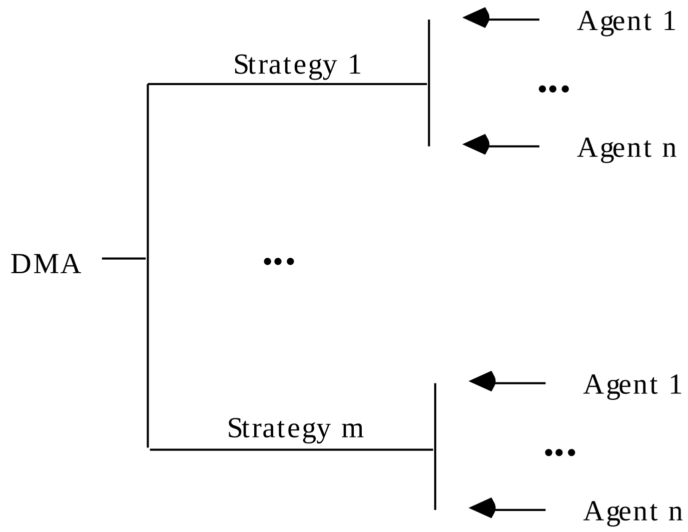

Figure A.1 A multi-agent decision model

In a situation modelled as in Figure A.1, some agents may be more reliable than others when evaluating the strategies involved, since different agents may have different capabilities to determine the values. The DMA may also have access to assessments expressing how trustworthy the different agents are. In the model, 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 evaluation method allows for vague and numerically imprecise information. Thus, the DMA may rank the credibilities of the different autonomous agents as well as quantify them in imprecise terms. The autonomous agents have a similar expressibility regarding their respective opinions about the strategies under consideration.

Example A.1: Assume a 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 opinions concerning the strategies for managing a system communications resource. The DMA has to decide whether to keep all time slots open for negotiation, to allocate some fixed bandwidth for high-volume users, or to lease out some of the bandwidth to neighbouring systems. Call these strategies S_1, S_2, and S_3, respectively. Further, assume that the agents A_1 through A_4 have reported to the DMA the following value statements.[3] The values involved could, for example, be monetary. In that case, they are linearly transformed to real values in the interval [0,1].

Statements according to agent A_1:

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

Statements according to agent A_2:

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

Statements according to agent A_3:

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

Statements according to agent A_4:

[3] The agents may have evaluated the prospective strategies using any number of well-established datacom traffic models. Here, only the evaluation of the total throughput situation is considered.

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

The rest of this appendix describes how the DMA may use the DELTA method in evaluating multi-agent problems such as the one above. A significant feature of the method is that it encourages the agents not to present report statements with an unrealistic degree of precision. Essentially, the model consists of a set of agents, a set of strategies, and two systems of statements concerning the credibilities and values involved. The sets of credibility statements and value reports are transformed into bases of linear constraints. The properties of those bases are discussed next.

Credibility Bases

A *credibility base* K with m agents is expressed in the credibility variables $\{c_1,\ldots,c_m\}$, stating the relative credibility of the different agents. The term c_k denotes the credibility assessment of agent A_k. A credibility base contains expressions about the credibility of each agent. To make the qualitative statements of credibility computable, they are translated in a manner similar to the standard DELTA method. Here, four types of possible credibility statements will be discussed. For a longer discussion of the parameters involved in the translations, refer to the corresponding treatment of probability statements in Chapter 4.

1. The credibility of A_k equals a number r, is at least r, is at most r.

   Example: The credibility of A_k is greater than r.

Translation: $c_k \in [r+\eta_1, r+\lambda_1]$

2. The credibility of A_k is between some real numbers.

   Example: The credibility of A_k is between r_1 and r_2.

   Translation: $c_k \in [r_1-\varepsilon_1, r_2+\varepsilon_1]$

3. The credibility of A_k 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.

   Example: The credibility of A_k is equal to the credibility of A_j.

   Translation: $c_k - c_j \in [-\varepsilon_2, \varepsilon_2]$

4. Agent A_k is credible, the opinion of agent A_k is worth considering, agent A_k is not credible, etc.

   Example: Agent A_k is credible.

   Translation: $c_k \in [r_3, r_4]$

In order for the credibility statements to be normalised, the constraint $\sum_k c_k = 1$ is added to the constraints above. The conjunction of constraints of the four types above, together with the normalisation, is the credibility base.

Example A.1 (cont'd): The DMA has estimated the credibility of A_1 through A_4 as numbers in the interval [0,1]. The translation of the statements into a credibility base results in the following expressions.

$c_1 \in [0.20, 0.90]$
$c_2 \in [0.10, 0.30]$
$c_3 \in [0.20, 0.70]$
$c_4 \in [0.00, 0.50]$

The credibilities are subject to the normalisation constraint $\sum_k c_k = 1$. 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). Since no other weight is affected, the hull of this base is $\{\langle 0.20, 0.70\rangle, \langle 0.10, 0.30\rangle, \langle 0.20, 0.70\rangle, \langle 0.00, 0.50\rangle\}$. ■

Report Bases

A report base R contains statements about individual agents' opinions of the values of different strategies, i.e., it consists of a number of interval constraints and core intervals that represent various strategy statements. It is expressed in value variables $\{v_{11},\ldots,v_{1n},\ldots,v_{m1},\ldots,v_{mn}\}$ stating the values of the strategies according to the different agents. The term v_{ik} denotes the value of strategy S_i according to the report of agent A_k. Five types of possible report statements are handled.
Given an autonomous agent A_k:

1. The value of the strategy S_i equals *r*, is at least *r*, etc.
   Example: The value of S_i is greater than r.
   Translation: $v_{ik} \in [r+\eta_1, r+\lambda_1]$
2. The value of strategy S_i is between some real numbers.
   Example: The value of S_i is between r_1 and r_2.
   Translation: $v_{ik} \in [r_1-\varepsilon_1, r_2+\varepsilon_1]$
3. The strategy S_i is as desirable (or undesirable) as strategy S_k, more desirable than S_k, the value of S_i is approximately equal to the value of S_k.
   Example: The strategy S_i is as desirable as S_j.
   Translation: $v_{ik} - v_{jk} \in [-\varepsilon_2, \varepsilon_2]$
4. The difference in value between S_i and S_j is not less than the difference in value between S_m and S_n.[4]
   Translation: $(v_{ik} - v_{jk}) - (v_{mk} - v_{nk}) \in [\eta_1, \lambda_1]$
5. The strategy S_i is desirable, S_i is fairly desirable, S_i is undesirable, etc.
   Example: The strategy S_i is desirable.
   Translation: $v_{ik} \in [r_3, r_4]$

Example A.1 (cont'd): The reports provided by the agents are translated into the following expressions.[5]

$v_{11} \in [0.50, 0.70]$ $v_{33} \in [0.50, 0.70]$

[4] For simplicity, it is assumed that the value of S_i is greater than the value of S_j, and that the value of S_m is greater than the value of S_n.

[5] The constants in the translations are chosen to keep the presentation simple.

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

This report base is then subject to evaluations using aggregate rules or security levels. ■

Agent Decision Frames

A credibility base K together with a report base R constitute an *agent decision frame* ⟨S,K,R⟩, where S is the set of strategies. This is in analogy to the ordinary decision frame ⟨C,P,V⟩ in the standard DELTA method. The mapping onto ordinary frames is straightforward. The strategies correspond to consequence sets, and the report elements are analogous to the consequences. Further, the credibilities have properties similar to probabilities, and report values are almost the same as values in the ordinary frame.

The mapping is not perfect, though. At first, it seems that credibilities map directly onto probabilities in that they have a similar role, distributing mass over the report values. But if credibilities are allowed to be assigned per strategy for each agent, then a more credible report about v_{ik} from the agent A_k might be forced to assume a lower credibility than a less credible report about v_{jk} from the same agent due to other agents also being more credible when giving reports about strategy S_i and the credibilities being normalised to sum to one.[6] Thus, only one credibility assessment per agent ought to be allowed. Still, since it is a normalised mass to be distributed, it might be more reasonable to interpret credibilities as weights instead. If there are no credible reports, the agents' credibilities must sum to one, and conversely, if all reports are very credible, they must still sum to one. This is not in accordance with the common interpretation of credibility. Finally, if an agent A_i has

[6] If there would be no normalisation, then the aggregated value would not make sense.

low credibility and another agent A_j has a much higher credibility, the statement $v_{ik} > v_{jk}$ has the same effect regardless. These discrepancies must be accounted for in an agent decision model. Such problems notwithstanding, the DELTA method is well-suited for multi-agent systems.

Comparing Strategies

Relative to a particular agent decision frame, which strategy should be chosen? The problem formulation is mathematically almost equivalent to the decision frame in Chapters 4–6, thus rendering the method and computational machinery of those chapters suitable for this task as well. As is the case for ordinary decision frames, for agent frames it is often not enough to determine the set of non-dominated (admissible) strategies, 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 a problem, the autonomous agents as well as the DMA are encouraged to be deliberately imprecise, and thus values close to the boundaries of the interval constraints seem to be the least reliable ones. This is a typical case for applying the contraction principle as described in Chapters 4–5, and in the example, the effects of contraction can be seen. Note that no core is specified, and the contraction goes from the hull inwards to the degree of 80%.

Example A.1 (cont'd): Entering the information into DELTA results in the agent decision frame in Table A.1.

```
Frame 'ExA1' in folder 'PhD' has 3 strategies
S1 (Strategy 1)
S2 (Strategy 2)
S3 (Strategy 3)

Each strategy is valued by 4 agents
A1 (Agent 1)
A2 (Agent 2)
A3 (Agent 3)
A4 (Agent 4)

The credibility base contains 4 constraints
 1: C1 ∈ [0.200,0.900]
```

```
 2: C2 ∈ [0.100,0.300]
 3: C3 ∈ [0.200,0.700]
 4: C4 ∈ [0.000,0.500]

Credibility hull    Symmetry hull
C1 ∈ [0.200,0.700]  [0.200,0.494]
C2 ∈ [0.100,0.300]  [0.100,0.218]
C3 ∈ [0.200,0.700]  [0.200,0.494]
C4 ∈ [0.000,0.500]  [0.000,0.294]

The report base contains 10 constraints
 1: V1.1 ∈ [0.500,0.700]
 2: V2.1 ∈ [0.100,0.700]
 3: V3.1 ∈ [0.300,1.000]
 4: V1.2 ∈ [0.100,0.500]
 5: V2.2 ∈ [0.400,0.700]
 6: V3.3 ∈ [0.500,0.700]
 7: V1.4 ∈ [0.500,0.700]
 8: V2.4 ∈ [0.000,0.700]
 9: V1.3 - V2.3 ∈ [0.000,1.000]
10: V2.4 - V3.4 ∈ [0.000,1.000]

Report hull
V1.1 ∈ [0.500,0.700]
V1.2 ∈ [0.100,0.500]
V1.3 ∈ [0.000,1.000]
V1.4 ∈ [0.500,0.700]
V2.1 ∈ [0.100,0.700]
V2.2 ∈ [0.400,0.700]
V2.3 ∈ [0.000,1.000]
V2.4 ∈ [0.000,0.700]
V3.1 ∈ [0.300,1.000]
V3.2 ∈ [0.000,1.000]
V3.3 ∈ [0.500,0.700]
V3.4 ∈ [0.000,0.700]

Focal point
Cred:  0.347  0.159  0.347  0.147
Agent     A1     A2     A3     A4
S1:    0.600  0.300  0.500  0.600
S2:    0.400  0.550  0.500  0.350
S3:    0.650  0.500  0.600  0.350
```

Table A.1 The agent decision frame

Evaluating the frame above results in Tables A.2–A.4.[7] Table A.2 shows the contraction of strategy S_1.

```
Contraction       0%     20%     40%     60%     80%
S1      min:  0.166   0.247   0.322   0.392   0.458
```

[7] The output from DELTALIB (see Chapter 3) is numeric. DMAs, especially software agents, often desire to receive the evaluation results in the form of matrices or tables instead of graphs in order to perform numerical computations on them.

	mid:	0.518	0.518	0.518	0.518	0.518
	max:	0.828	0.758	0.691	0.629	0.572

Table A.2 The contraction of S_1

Tables A.3–A.4 show the contractions of the strategies S_2 and S_3, respectively. Hence, strategy S_2 is inferior to both S_1 and S_3, but strategy S_3 is slightly better than S_1. A further investigation is recommended in order to identify critical variables.

Contraction		0%	20%	40%	60%	80%
S2	min:	0.060	0.143	0.223	0.300	0.375
	mid:	0.451	0.451	0.451	0.451	0.451
	max:	0.848	0.767	0.686	0.607	0.528

Table A.3 The contraction of S_2

Contraction		0%	20%	40%	60%	80%
S3	min:	0.186	0.269	0.349	0.424	0.496
	mid:	0.565	0.565	0.565	0.565	0.565
	max:	0.914	0.840	0.768	0.699	0.631

Table A.4 The contraction of S_3

It is natural to ask how sensitive the different contractions are to changes in the agent frame. The DMA can simultaneously vary any number of intervals to discover credibility or value 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 additional information about strategy S_3. Before doing that, the DMA can investigate how influential the report from the agent would be. For instance, the DMA can restrict the maximum value of v_{31} to 0.6 instead of 1 and evaluate the modified decision situation. Table A.5 shows the result for strategy S_3. The strategy is now slightly worse than S_1. The new information does not change the results in Tables A.2 or A.3.

Contraction		0%	20%	40%	60%	80%
S3	min:	0.186	0.257	0.324	0.386	0.443
	mid:	0.495	0.495	0.495	0.495	0.495
	max:	0.745	0.695	0.645	0.596	0.546

Table A.5 The result of modifying S_3

Thus, it is reasonable to allocate resources to collect more information about strategy S_3 from agent A_1. 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. ■

Security Levels

The intuition behind security levels is that they provide limits beyond which a strategy is undesirable. Thus, a DMA might regard a strategy as undesirable if it has access to a report in which a credible agent assigns a low value to the strategy.

Example A.2: Suppose that the DMA has stipulated that a strategy S_i is undesirable **iff**

- according to some agent A_j, the value of strategy S_i is less than 0.45
- the credibility of that agent A_j is greater than 0.65.

Assume that v_{12} is in the interval [0.40, 0.60] and that c_2 is in the interval [0.20, 0.70]. Then S_1 is below the threshold and is thus undesirable. It is advisable to investigate how much the different intervals can be decreased while the security levels are still violated. In this manner, the stability of the result can be studied. For example, it can be seen that strategy S_1 ceases to be undesirable when the left end-point of the interval of v_{12} is increased by 0.05. Consequently, the result is quite unstable. ■

The example contained a very simplistic approach to limiting undesirable outcomes. To be more sophisticated and utilise the DELTA method, the concept of security level as defined in Chapter 5 is applied. There, an observation regarding security levels was made, which is here turned into a definition of agent security levels and is put to use in testing which strategies might be undesirable.

Definition A.1: Given an agent decision frame $\langle S,K,R \rangle$ and two real numbers $r,s \in [0,1]$, a strategy S_j *violates agent security level s for value threshold r* **iff** for $K_j = \{k$ ™ $v_{jk} \geq r\}$ $\sum_{k \in K_j} c_k \leq 1 - s$.

This is best illustrated by an example which evaluates the security levels using weak first order dominance.[8]

Example A.1 (cont'd): Using the definitions above, it may now be investigated to what extent the different strategies are undesirable. Figure A.2 shows, for each strategy and a value threshold of 0.10, the worst possible credibility assignments consistent with the frame for different degrees of contraction, i.e. the security levels violated by weak dominance. In the figure, the K- and R-bases are contracted simultaneously, but this is not the only option. The K-base might be left uncontracted, studying only the R-base under contraction, and conversely, the R-base might be untouched while contracting K.

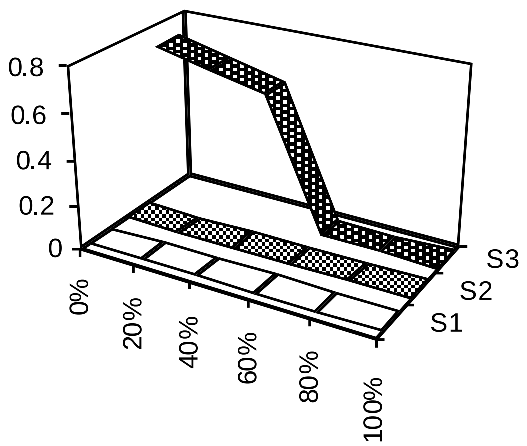

Figure A.2 Value threshold 0.10

From the figure, it can be seen that the strategies S_1 and S_2 are not undesirable in any part of the decision frame. Strategy S_3 is undesirable in the original frame and remains so until it is contracted by more than 60%. For instance, when the decision frame is con-

[8] See Chapter 5 for an explanation of weak dominance.

tracted by 40%, the greatest joint credibility for the bad reports of this strategy is 0.58.

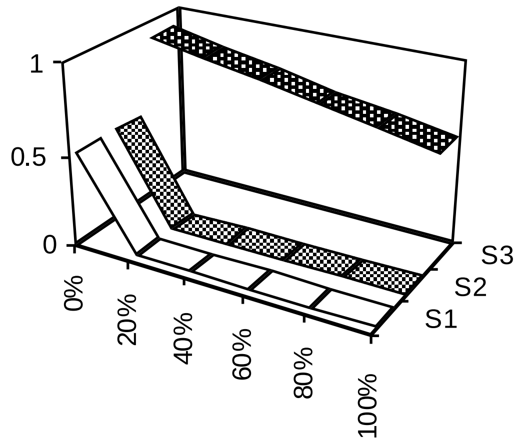

Figure A.3 Value threshold 0.20

Figures A.3 and A.4 show the evaluations for the value thresholds 0.20 and 0.50 respectively. As can be seen in Figure A.3, the strategies S_1 and S_2 are now undesirable in some parts of the decision frame. However, they cease to be undesirable at contractions of at least 20%.

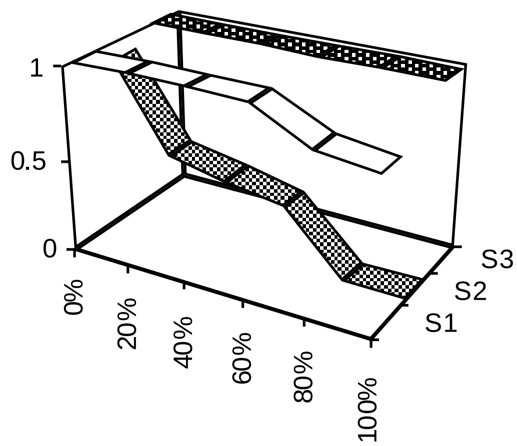

Figure A.4 Value threshold 0.50

Figure A.4 shows that for very high value thresholds, S_3 is undesirable regardless of the degree of contraction. Thus, it can be seen that the results of the evaluation are strongly dependent on boundary values, and consequently they should be further investigated in sensitivity analyses. While S_1 and S_3 were preferable to S_2 already in the primary evaluation above, S_3 seems to be too dangerous to adopt as the main strategy for the time slot allocation. Thus, the agent selects strategy S_1 – keeping all the time slots open for negotiation. ■

By using security levels, the decisions made by the DMA will be more reliable and predictable than if such levels were not imposed on the reports. The trust the DMA can put in the results will increase considerably as it is able to set the levels and thresholds according to its appreciation of the particular decision problem.

Long you live and high you fly
Smiles you give and tears you cry
All you touch and all you see
Is all your life will ever be

Run, rabbit run
Dig that hole, forget the sun
And when at last the work is done
Don't sit down, it's time to dig another one

For long you live and high you fly
But only if you ride the tide
And balance on the biggest wave
You race towards an early grave

– R. Waters