A Methodology for Action-Selection using Value-Based Argumentation

Fahad Nawab[1], Trevor Bench-Capon, Paul Dunne
Department of Computer Science, University of Liverpool, Liverpool, UK

Abstract This paper describes a method for decision making using argumentation involving practical reasoning. The method is intended to produce the decision considered most likely to promote the agent's aims and aspirations. First, the problem situation is formulated in terms of an Action Based Alternating Transition System, representing the actions available to the agents relevant to the situation, and their consequences, taking into account the possible effects of the choices of the other relevant agents. Next arguments are constructed by instantiating an argumentation scheme designed to justify actions in terms of the values they promote and subjecting these instantiations to a series of critical questions to identify possible counter arguments. The resulting arguments are then organized into a Value-Based Argumentation Framework (VAF), so that a set of arguments acceptable to the agent can be identified. Finally the agent must select one of the acceptable actions to execute. The methodology is illustrated through the use of a detailed case study.

Keywords. Decision making based on argumentation, Reasoning about action using argument, Applications.

1. Introduction

When choosing what to do in a given situation an agent needs not only to identify its options and their likely effects, but to take into account factors outside of its control, such as the choices of other agents which can change the effects of its own actions. Moreover the choice will be determined by the short term and long term aims and aspirations of the agent, and perhaps also by its emotion and temperament. These factors will differ from agent to agent, and so different agents may rationally decide to pursue different courses of action. This paper will present a methodology for decision making for use by an autonomous agent that is designed to meet these requirements.

In selecting an action an agent needs to have regard to a range of considerations. It must take a view of the relevant features of the current situation in so far as they affect

[1] Corresponding Author: E-Mail: fahad@csc.liv.ac.uk

the range of options available, and the effects of selecting them. The reasons why the effects are desirable, and the priority given to the various desiderata by the agent will be important. The impact of any uncertainties in the current situation on the available actions and their effects need to be considered. Performing an action will often mean that some other action cannot be performed, while other actions may be enabled by performing an action. The side effects of actions need to be considered: in some cases any gains may be outweighed by losses. Conversely there may be beneficial side effects, increasing the attractiveness of the choice. Any method for action selection needs to take these considerations and more into account.

Our analysis is based on the use of argumentation schemes as a presumptive justification of action: in particular we use the argument scheme for justifying actions described in [3]. This scheme is also designed to allow for considerations stated in the preceding paragraph to be reasoned about through posing critical questions challenging the justification. Where the critical questions identify objections to the justification these can be seen as counter arguments. Resolving the conflicts between these arguments and their counter arguments will ensure that the considerations are given their due weight. Resolution is achieved using a Value-Based Argumentation Framework [4], which relates acceptance of the justifying arguments to the particular priority order the agents gives the motivations for action and so allows for the elements of subjective choice.

The approach to the problem can be seen as involving five stages:

- *Formulating the Problem*: produce a formal description of the problem to give all possible actions, values and all related factors that may influence the decision. This will be accomplished through an Action Based Alternating Transition System (AATS) as in [2].
- *Determining the Arguments*; on the basis of the AATS, arguments providing justifications of the various available actions are provided by instantiating the argument scheme introduced in [3]. Counterarguments are identified using a subset of the critical questions of [3], as interpreted in terms of an AATS in [2].
- *Building the Argumentation Framework*: In this step the arguments and attacks between them identified in the previous step are organized into an Argumentation Framework. Because the argument scheme used associates arguments with the values they promote or demote, arguments can be annotated with these values, yielding a Value Based Argumentation Framework (VAF) [4].
- *Evaluating the Argumentation Framework*. As described in [4], the arguments of VAFs are determined as acceptable or not with respect to a specific audience, characterized by the ordering on values subscribed to by the agent making the choice.
- *Sequencing the Actions*. The action deemed to be acceptable to the agent in the previous stage, must now be sequenced into a suitable order in which they should be performed.

The first section of this paper has given a brief background summarizing the approach being used. The second section will define the various elements that make up the proposed methodology. The third section will introduce the case study and provide

a detailed working through of the each of the five steps. Finally, we present some concluding remarks, observations and possible enhancements to this work.

2. Decision Making Framework

In this section we will describe the techniques we use in each of the five steps.

2.1 Formulating the Problem

In step one, the problem is formulated as an Action-Based Alternating Transition System (AATS). AATS were introduced in [10] as a foundation to formally describe a system in which several agents are able to act to change its state. In [10] an AATS with n agents is an (n+7) tuple; This was then extended by [2] to include the notion of values where Av is a set of values for each agent which is a subset of V and every transition from the set *Q* may either promote, demote, or be neutral with respect to those values.

Definition 1: AATS:
As extended by [2], an AATS is a (2n+8) tuple $S = < Q, q_0, Ag, Ac_1 \ldots, Ac_n, Av_1 \ldots, Av_n, ?, t, ?, p, d>$ where:
Q is a finite, non-empty set of *states*
$q_0 = q_x ? Q$ is the *initial state*
$Ag = \{1,...,n\}$ is a finite, non-empty set of *agents*
Ac_i is a finite, non-empty set of actions, for each $i ? Ag$ where $Ac_i \cap Ac_j = \emptyset$ for all $i ? j ? Ag$;
Av_i is a finite, non-empty set of values $Av_i ? V$, for each $i ? Ag$.
$? : Ac_{Ag} ? 2^Q$ is an *action precondition function*, which for each action $a ? Ac_{Ag}$ defines the set of states ? (a) from which a may be executed;
$t : Q \times J_{Ag} ? Q$ is a partial *system transition function*, which defines the state t (q, j) that would result by the performance of j from state q - note that, as this function is partial, not all joint actions are possible in all states (cf. the precondition function above);
? is a finite, non-empty set of *atomic propositions*;
$p : Q ? 2^?$ is an interpretation function, which gives the set of primitive propositions satisfied in each state: if $p ?$ p (q), then this means that the propositional variable p is satisfied (equivalently, true) in state q.
$d : Q \times Q \times Av_{Ag} ? \{+, –, =\}$ is a *valuation function* which defines the status (promoted (+), demoted (–) or neutral (=)) of a value $v_u ? Av_{Ag}$ ascribed by the agent to the transition between two states: d (q_x, q_y, v_u) labels the transition between q_x and q_y with one of $\{+, –, =\}$ with respect to the value $v_u ? Av_{Ag}$.

To represent a particular problem we first identify a set of propositions which we consider relevant to the problem. Each model of this set of propositions will be a potential state of the system. Next, we identify the relevant agents and the different possible actions the agents can perform, and how these will move us between these states, each transition representing a joint action of the agents involved. Finally we provide the values and relate them to the transitions between states.

2.2 Determining the Arguments

Our method of justifying actions is in terms of presumptive justification through the instantiation of an argument scheme, followed by a process of critical questioning to see whether the presumption can be maintained, as described in [9]. Specifically we use the argument scheme presented in [3] which extends the argument from sufficient reason of [9] to enable discrimination between the effects of an action (the *consequences*), the desired effects (the *goal*) and the reason why these effects are desired (the *value*). Thus our argument scheme is: *in the current state the agent should perform action A to reach a new state in which goal G is true, promoting value V.*

In [2] a realization of this scheme in terms an AATS is given:

Definition 2: Argument Scheme for practical Reasoning

In the initial state $q_0 = q_x$? Q, Agent i ? Ag should participate in joint action j_n ? J_{ag} where $j_{ni} = a_i$, such that t(q_x, j_n) is q_y, p_a ? p (q_y) and p_a ?p (q_x), and for some v_u ? Av_i, d(q_x, q_y, v_u) is +.

Critical questions are now used to addresses the factors which may lead to the presumptive justification being overturned. [3] identified sixteen critical questions that might be posed, but in our situation not all of them are relevant. For our purposes we need consider only six. These six critical questions are, as in [2], defined in terms of an AATS in Definition 3.

Definition 3: Relevant Critical Questions.

CQ1: *Are the believed circumstances true?*

q_0 ? q_x and q_0 ? ?(a_i).

CQ11: *Does doing the action preclude some other action which would promote some other value?*

In the initial state q_x ? Q, if agent i ? Ag participates in joint action j_n ? J_{Ag}, then t (q_x, j_n) is q_y and d(q_x, q_y, v_u) is +. There is some other joint action j_m ? J_{Ag}, where j_n ? j_m, such that t (q_x, j_m) is q_z, such that d(q_x, q_z, v_w) is +, where v_u ? v_w.

CQ2: *Assuming the circumstances, does the action have the stated consequences?*

t (q_x, j_n) is not q_y.

CQ7: *Are there alternative ways of promoting the same value?*

Agent i ? Ag can participate in joint action j_m ? J_{Ag}, where j_n ? j_m, such that t (q_x, j_m) is q_z, such that d(q_x, q_z, v_u) is +.

CQ8*: Does doing the action have a side effect which demotes the value?*

In the initial state q_x ? Q, if agent i ? Ag participates in joint action j_n ? J_{Ag}, then t (q_x,j_n) is q_y, such that p_b ? p(q_y), where p_a ? p_b, such that d(q_x, q_y, v_u) is –.

CQ9: *Does doing the action have a side effect which demotes some other value?*

In the initial state q_x ? Q, if agent i ? Ag participates in joint action j_n ? J_{Ag}, then t(q_x, j_n) is q_y, such that d(q_x, q_y, v_w) is –, where v_u ? v_w.

2.3 Building the Relationship Model

The previous step resulted in a number of arguments, associated with values, and a set of attack relations between them. These can be organized into a Value Based Argumentation Framework (VAF). A VAF can be defined as:

Definition 4: Value based Argumentation Framework.
A triple $<H(X,A), V, ?>$, *where* $H(X,A)$ *is an argument system,* $V = \{v_1, v_2, \ldots, v_k\}$ *a set of k* values, *and* $? : X \to V$ *a mapping that associates a value* $?(x) \in V$ *with each argument* $x \in X$. *An* audience, *a, for a VAF* $<H, V, ?>$, *is a total ordering of the values V. We say that* v_i is preferred to v_j in the audience *a, denoted* $v_i \succ_a v_j$, *if* v_i *is ranked higher than* v_j *in the total ordering defined by a.*

2.4 Evaluating the Model

Now we have built the VAF, the next step is to evaluate the attacks and determine which arguments will be acceptable to our agent. The strength of each argument is determined by the values associated with it. Given the ordering on values desired by the agent, we can determine what arguments will be acceptable to the agent, and determine the preferred extension with respect to the audience of which the agent is a member. This preferred extension, which will be unique and non-empty, represents the maximal set of acceptable arguments.

Definition 5: VAF related concepts.
Let $<H(X,A), V, ?>$ *be a VAF and a an audience.*

a. *For arguments x, y in* X, *x is a* successful attack *on y (or x* defeats *y)* with respect to the audience *a if:* $<x, y> \in A$ and *it is* not *the case that* $?(y) \succ_a ?(x)$.
b. *An argument x is* acceptable to the subset *S with respect to an audience a if: for every* $y \in X$ *that* successfully attacks *x with respect to a, there is some* $z \in S$ *that successfully attacks y with respect to a.*
c. *A subset R of* X *is* conflict-free with respect to the audience *a if: for each* $<x, y> \in R \times R$, *either* $<x, y> \notin A$ *or* $?(y) \succ_a ?(x)$.
d. *A subset R of* X *is* admissible *with respect to the audience a if: R is conflict free with respect to a and every* $x \in R$ *is acceptable to R with respect to a.*
e. *A subset R is a* preferred extension *for the audience a if it is a maximal admissible set with respect to a.*

2.5 Sequencing the Actions

The four previous steps have given us a set of actions acceptable to the agent given its priorities with respect to values. These actions are acceptable in the sense that they have survived the critique provided by the posing of critical questions, and have no attackers preferred to them. Often this set will contain multiple arguments, any of which could beneficially be performed in the current state. We suggest that these should be sequenced in terms of *safety*, by which we mean that unexpected consequences will not prevent the other actions in the set being performed, *opportunity*, by which we mean that the performance of an action may enable some desirable action which is not available in the current state, and *threat*, where a potentially bad side effect is brought into play.

3. Example Application

3.1 Formulating the Problem

Our case study concerns a problem which faces university Heads of Deptratment, and reflects the need to balance costs, Departmental and individual interests. Our agent is a Head of Department (HoD) in a university, and has requests relating to two specific conferences. He has three potential candidates and needs to decide which of them to send. Students 1 (S1) and 2 (S2) are new students. S1 is asking to go to a nearby conference, which will be cheaper financially, S2 is asking for a different conference which will cost more, but S2 has prepared a good paper that might help the department's reputation. Student 3 (S3) is an older, more established, student asking to be sent to the local conference and, although she has not prepared a paper, she is an excellent networker who is likely to impress other delegates and so present the department in a good light. The conferences are on different topics, so S2's paper would not be suitable for the local conference, but are both of equal standing. The budget will only allow two students to be sent.

Now, in the rest of this section we will set the different properties to allow the representation of the problem of our example as an AATS.

3.1.1 Propositions and Actions

We will first consider the different propositions that the agent will take into account in his decision making: whether there are currently funds available in the budget (Budget); whether a student can be sent to attend (Attendance S(1-3)), whether the student has written a paper (Paper S(1-3)) and, finally, whether the student has attended a conference before (Previous S(1-3)).

Now, we define all possible actions that the HoD can take in all circumstances. Those are either to ask any one of the three students to write a paper Write(S1), Write(S2) , Write(S3), or to agree to send a student to the requested conference Send(S1), Send(S2), Send(S3), These actions may change the state of Paper(Si) or Attendance(Si) respectively. If requested to write, the students may or may not succeed in writing a paper.

3.1.2 Values

Now, we list the list of related values to the HoD. We, then, link those values to the various transitions. Table 1 the different values that will be promoted or demoted in accordance with the changes in propositions.

Value	**Short**	**Promoted/Demoted if:**
Happiness	H(Si)	Promoted if Si attends
Happiness	H(Si)	Demoted if Si has written a paper and does not attend
Publication	P	Promoted if Si attends having written a paper
Experience	E(Si)	Promoted if Si has not attended before, and attends
Esteem	Est	Promoted if Si has attended a previous conference, has a paper and attends

Table 1: Values relevant to the audience (HoD)

By esteem, we mean the general enhancement of the reputation of the Department that comes from an impressive individual making an impact at a conference and raising the profile of the Department's research, the research links established, and such like. Note that happiness and experience are relative to individual students, whereas the other values are relative to the Department, although realized through properties of the individual students.

3.1.3 State Format

The states will be presented as follows: Budget, Attendance, Paper and Previous (B-XXX-XXX-XXX) where each X will be either 1 or 0 depending on whether or not the corresponding proposition is true in that state. Before we move into modeling the state transitions, let us look at the initial state q0. Budget is set to 3: the cheaper conference costs 1 and the expensive conference 2, so that we can send at most two students.. S1 and S3 will consume 1 point from the budget whenever chosen whereas S2 will consume 2 points. S3 has already attended a previous conference and S2 has a paper ready. Thus, q0 = (3-000-010-001). Figure 1 shows the initial state and some example transitions from that state. J0 is Send(S1), J1 is Send(S2), J2 is send(S3). Where a paper is requested and written, we have J3 for S1 and J5 for S3, while J4 represents a request which does not result in a paper. The transitions are also labeled with the values they promote or demote. Budget = 0 represents a terminal state, since no further actions are possible

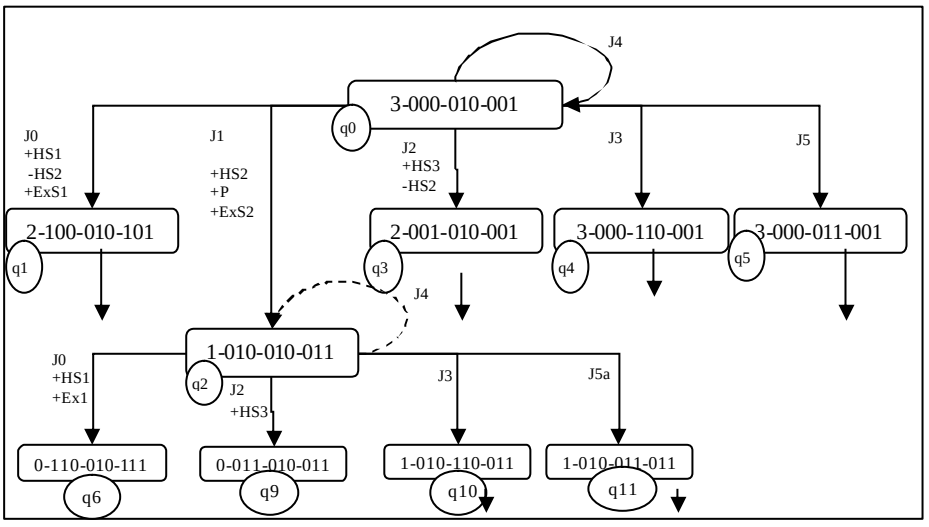

Figure 1: Transitions from initial state

3.1.4 Uncertainties

Consequences of our actions in reality are not always entirely predictable and we execute our actions hoping for a certain result which often will not come about because it has dependencies on other actions performed by other agents. When an agent performs an action where the results solely depend on itself, as with sending a student,

it is very easy to assume the resulting state. When, however, the HoD asks a student to write a paper, the student may or may not succeed. Whether their joint action will by J3 or J4 is thus out of his control, and so he cannot be certain which state will be reached. There may also uncertainty about the initial state: in Figure 1 we assume that S2 is the only one who has written a paper. But it may be that that the HoD is not sure of this, and any of the three students might have actually written a paper. This gives us seven different possible states, any of which could be the initial state. Obviously if an action is performed in a state other than the one assumed, the state reached may be different. These uncertainties will be considered through the mechanism of critical questions.

We now move to the second step where we start building arguments for and against performing the various actions.

3.2 Determining the Arguments

The AATS will allow us to evaluate each action at every state and relate the actions to propositions and the values they promote values. Table 2 shows the arguments that can be made for performing an available action in the initial state.

Arg	In State	Action	To get to State	Realize Goal				Promoting			
				Budget	Attend	Paper	Prev	H	P	E	Est
Arg1	Q0	J0	Q1		S1		S1	S1			
Arg2	Q0	J0	Q1		S1		S1			S1	
Arg3	Q0	J1	Q2		S2		S2	S2			
Arg4	Q0	J1	Q2		S2		S2		S1		
Arg5	Q0	J1	Q2		S2		S2			S2	
Arg6	Q0	J2	Q3		S3			S3			
Arg7	Q0	J3	Q4			S1					
Arg8	Q0	J5	Q5			S3					

Table 2: Arguments from state q0

We clearly can see from the table how these arguments differ with respect to the values promoted. The next step now is to use critical questions to identify which arguments are open to counter attack. We will use the identifying labels of critical questions as in [2].

3.3 Building the Relationship Model

3.3.1 CQ1: Are the stated circumstances true?

This question arises in this example from the fact that although the HoD believes that the initial state is 3-000-010-001 (q0) where S2 has written a paper, he cannot actually be sure that the other students have not also written papers and cannot be absolutely certain that S2 has in fact written a paper. This results in there being seven different possible initial states. So, in this case all the arguments in Table 2 are open to this attack.

The agent could assume that all states are possible and build up the argumentation model with all the possible states in mind. This will result in multiple Preferred Extensions (PEs), one for each possible initial state. The common elements in all PEs will then represent justifications of actions which are unaffected by the uncertainties with respect to what is true in the initial state. Should there be no arguments common

to all the PEs, it would be necessary to make choices as to what is to be believed in the original situation: for example it may be considered very unlikely that S2 would have misinformed the HoD about the status of his paper. It is also possible that the agent is able to confirm his beliefs before proceeding to make decisions. In our example this is the case: the HoD can ask S2 to show him the paper, and ask the others if they have a paper ready. In what follows, therefore, we can assume that the HoD is able to confirm his beliefs, and so objections arising from CQ1 can be discounted. In general, however, where complete information is not attainable, the question is an important one.

3.3.2 CQ11: Does the action preclude some other action which would promote some other value?

If we send S1 or S3 without asking them to write a paper, we lose our chance to promote publication, since we will no longer expect them to write a paper. Moreover, this will also lose the chance to promote esteem, which requires us to send S3 with a paper written. Thus Arg1, Arg2, and Arg6 are all attacked by an argument, A1a, that they prevent the promotion of publication. Arg6 is also attacked by an argument, Arg6a, that it precludes the promotion of esteem.

3.3.3 CQ2: Does the Action have the stated consequences?

This question occurs when we need to consider joint actions: cases where the agent is not in sole control of the state reached. In our example, this is represented by the possibility of the request to write a paper not being met. Thus Arg7 and Arg8 are attacked by Arg7a, that the joint action might turn out to be J4.

3.3.4 CQ8: Does the action have side effects which demote the value?

Sending any of the students other than S2, will demote the happiness of S2, since he has already written a paper. Supposing the HoD is impartial and so indifferent as to which student happiness is promoted in respect of, this will give an argument, Arg1b, against Arg1 and Arg6: while these arguments promote happiness, the actions they justify also demote it.

3.3.5 CQ9: Does the action have side effects which demote some other value?

If we do not send a student who has written a paper the happiness of that student will be demoted. This provides the basis for an argument against doing any action which involves not sending S2 for any reason other than the promotion of happiness. Thus Arg2 is subject to an attack from Arg2a since it would demote S2's happiness.

3.3.6 CQ7: Are there other ways to promote the same value?

Both Ar2 and Arg5 are based on the promotion of experience. This question indicates that they attack one another. Similarly happiness can be promoted by any of Ar1, Arg3 and Arg6. These also mutually attack, therefore.

Now, we have identified all possible attacks from the different arguments we are able to form the value based argumentation framework. Figure 2 is a graphical representation of the framework in the example.

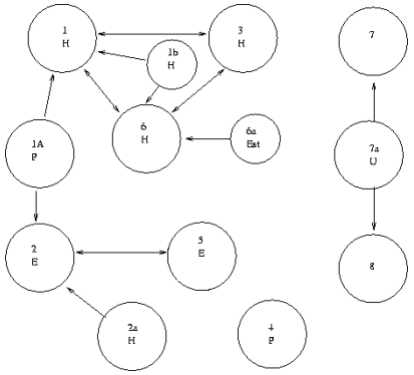

Figure 2: Value Based Argument Framework

3.4 Evaluating the Model

From the VAF in Figure 2 we can see that Arg4 has no arguments attacking it. Thus it will appear in every preferred extension, irrespective of the way we order our values. The status of Arg7 and Arg8 depend on our degree of confidence we have that the papers will be written if requested. Suppose we have this confidence, and so Arg7 and Arg8 are acceptable. In order to determine which of the remaining arguments are acceptable, we need to fix on an ordering of the values that the HoD wishes to promote at the particular time. Suppose that the value ordering is as follows: Esteem > Publication > Experience > Happiness. This gives us the ability to resolve the conflicts that we have in the model by eliminating unsuccessful attacks. Arg1b will defeat Arg1 and Arg6, leaving Arg3 for the preferred extension. Although Arg2 is not defeated by Arg2a, it is defeated by Arg1a, and so Arg5 survives. Thus our preferred extension is {Arg1a, Arg1b, Arg2a, Arg3, Arg4, Arg5, Arg6a, Arg7, Arg8}. In terms of actions we can justify sending S2, and requesting a paper from S1 and S3. The arguments in the preferred extension which do not justify actions are there to justify the rejection of other arguments.

3.5 Sequencing the Actions

From the last step we have identified three actions, which would take us to q2, q4, or q5, although if our confidence in the student's ability to produce a paper was misplaced, we could remain in q0. We must now choose some sensible sequence for these actions. This choice will need to consider both uncertainty about outcomes (Section 3.1.5) and eventually reaching the best state. This will need us to look ahead and consider what is possible in the states that would result from our action.

There are three issues we should consider here. First we need to consider whether the action is *safe*, in the sense that if it fails we do not move to a state where our desirable other actions are no longer possible. In our example, all our actions are safe, since those that can fail simply return us to the initial state. Next we must consider *opportunities*: what additional values can be promoted in the next state? If we ask S1 to write, we have the possibility of promoting publication (the chance to promote S1's experience already exists as Arg2), and, although we can already promote this by

sending S2, S1's publication will be an additional benefit. But if we ask S3 to write we can create the chance to promote esteem, as well as the additional publication. But there are also *threats*: if S1 and S3 write papers and are not sent, they will be unhappy. Since we have said we prefer esteem to experience, we will prefer the opportunities created by requesting a paper from S3, and so will prioritise this action over asking S1. Should we do this before sending S2? In the particular example this seems to not matter. We might, however, decide that if we sent S2, we would demotivate S2 and so reduce the likelihood of her producing a paper. Such calculation of the probability of the success of actions is outside the scope of this paper, but would be a suitable topic for further investigation. Suppose, however, we make this judgement: then we should request a paper from S3 before sending S2. If S3 does write the paper, we will move to another state in which the recommended actions will to send S2 and S3. If, however, S3 does not produce a paper: now we have no possibility of promoting S3, and no threat of making S3 unhappy, and so we should request a paper from S1, in the hope of making an opportunity to promote publication. If S1 does write, we should then send S1 and S2: and even if S1 does not write he should as this will still promote experience in respect of S1.

Note that had S3 been interested in the expensive conference, the situation represented in Figure 2 would differ, since Arg6a would now, via CQ11, attack, successfully given our value order, Arg3, Arg4 and Arg6, since we can send at most one student to the expensive conference. This would make requesting a paper from S3 the only choice. Note also that had the HoD preferred experience to esteem, he would ask S1 rather than S3, and send S1 whether or not a paper is written: once S2 has been sent, there is a straight choice between S1 and S3 so that sending S3 will preclude the promotion of experience, and so Arg6 will have an attacker based on CQ11.

4 Conclusion

In this paper we have given a detailed work through of a methodology for decision making in a situation where a number of competing interests need to be balanced. The example provides an illustration of how a set of relevant arguments can be generated and evaluated in accordance with the particular preferences of the decision maker. The example has drawn attention to two matters in particular: the need to decide upon the most appropriate sequence for actions which are justifiable in a given situation, and the need to estimate uncertainty with respect to actions the effect of which depend on what other agents will choose to do. With regard to sequencing actions we can see that even if the execution of a sequence of actions would result in the same final state there may still be reasons for choosing one ordering rather than another. This is because the paths taken to that state will differ, and so the intermediate states may give rise to different opportunities and threats, and because should the effects of actions be not what was expected, there will be differences in the ability to recover from these setbacks. We also need to take account of uncertainty: our confidence in what the other agent will do may differ from state to state. Thus different sequences may increase or diminish our confidence that the joint action will be as desired, and we should sequence our actions so as to achieve as much confidence ass possible when we choose to perform the action.

For future work, one issue worth exploring would be a technique for gauging uncertainty of the effects of actions, or working with the uncertainty. There are other interertsing directions also. We have assumed that the agent is able to provide a total order on its values, but, as noted in [8], such an ordering often emerges as part of the reasoning process. Possibilities for this would be either to stipulate that some arguments must be made acceptable, as in [5], or to allow reasoning about what the value order should be, as in [7]. Another direction is to envisage the agent as being an automated decision maker acting on behalf of a human decision maker. In such cases, it may be that the agent will make decisions which the human does not approve: such feedback would suggest that the value order used by the agent should be modified. But the nature of these modifications, and the timing of these modifications is not a straightforward matter: for example it could be that having preferred H(S1) to H(S2) on one occasion, a HoD would feel obliged to use the opposite preference on the next occasion. Seeing the problem in the context of an ongoing series of decisions would require additional considerations to be taken into account. Finally we need to accommodate the fact that these kind of decisions are not always entirely based on a rational assessment of their pros and cons: emotions can also play a role:

"emotions and feelings can cause havoc in the process of reasoning under certain circumstances. Traditional wisdom has told us that they can, and recent investigation of the normal reasoning process also reveal the potentially harmful influence of emotional biases. It is thus even more surprising and novel that the absence of emotion and feeling is no less damaging, no less capable of compromising the rationality that makes us distinctively human and allows us to decide in consonance with a sense of personal future, social convention, and moral principle." [6]

These considerations open further avenues for exploration which would be immensely interesting to explore.

References

[1] K. Atkinson. *What should we do? Computational representation of persuasive argument in practical reasoning*. PhD thesis, University of Liverpool, 2005.

[2] K. Atkinson and T. Bench-Capon. Action based alternating transition systems for arguments about action. In *Proceedings of the Twenty Second Conference on Artificial Intelligence* (AAAI 2007), Vancouver, Canada, pp. 24-29. AAAI Press, Menlo Park, CA, USA.

[3] K. Atkinson, P. McBurney, and T. Bench-Capon. Computational representation of practical arguments. *Synthese*, 152(2):157--206,pages 191–240, 2006.

[4] T. Bench-Capon. Persuasion in practical argument using value-based argumentation frameworks. *Journal of Logic and Computation*, 13(3):429–448, 2003.

[5] T.J.M. Bench-Capon, S. Doutre and P.E. Dunne (2007). Audiences in Argumentation Frameworks. *Artificial Intelligence*, Vol 171 no 1, pp42-71.

[6] A. Damasio. *Descartes' Error*. G P Putnam and Sons, 1994.

[7] S. Modgil. *Value Based Argumentation in Hierarchical Argumentation Frameworks*. In: Proc. 1st International Conference on Computational Models of Argument. IOS Press, Amsterdam, pp 297-310.

[8] J. R. Searle. *Rationality in Action*. The MIT Press, Cambridge, Massachusetts, 2003.

[9] D. N. Walton. *Argumentation Schemes for Presumptive Reasoning*. Lawrence Erlbaum Associates, Mahwah, New Jersey, 1996.

[10] M. Wooldridge and W. van der Hoek. On obligations and normative ability: Towards a logical analysis of the social contract. *J. Applied Logic*, 3(3-4):396–420, 2005.