

# Persuasive Political Argument

**Katie Atkinson, Trevor Bench-Capon and Peter McBurney**

Department of Computer Science,

University of Liverpool,

Liverpool L69 3BX, UK.

{k.m.atkinson,tbc,p.j.mcburney}@csc.liv.ac.uk

## Abstract

In this paper we discuss how a computational version of argumentation involving practical reasoning can be applied to the domain of e-democracy. We discuss our previous work which proposed an argument scheme and associated critical questions to make use of presumptive reasoning in order to justify a proposal for action. We explain how this proposal can be made computational for use by BDI agents and we illustrate the approach with an example application of a recent political debate involving the Government's justification of a proposed action. We then go on to discuss some interesting issues arising from the example regarding the strength of the arguments and how they can be accrued to put forward a more convincing case for justifying the action. We also examine how this application compares with a previously implemented system which makes use of the same underlying theory.

## 1 Introduction

Debates about proposed actions are an inherent feature of a democratic society. Not only do government ministers debate the issues of the moment amongst themselves, but members of the public are also encouraged to engage in debates on current affairs. The passing of laws and motions is in principle subject to open-ended discussion and questioning before issues are put to a vote, though the current increasing use of "guillotine" measures comprises this, time circumscribing the debate. In discussing the issues the audience must be allowed to explore the reasons behind justifications presented to them, as well as being able to pose critical questions to the proponents of arguments in order to consider all aspects of the case. Reasons are what are crucial for gaining acceptance from particular audiences. The significance of the concept of an audience has been discussed by Perelman in [10] and has also recently been addressed in AI by Hunter in [8] and [9]. Hunter attempts to formalise Perelman's observation that we need to account for the fact that different audiences can have different perspectives on the same issue. Hunter proposes an approach to modelling audience preferences to enable arguments to be ranked to have a more empathetic effect upon

particular audiences. Such subjectivity in arguments with respect to a particular audience is obviously inherent in human reasoning and it must also be accounted for in practical reasoning in intelligent agents. A different take on audiences – in terms of the values they prize – is given in [4] and we follow this account.

One of the most widely debated, and perhaps most controversial debates to feature on the recent international political agenda has been the invasion of Iraq. Not only did this issue spark debate at national levels, it also received a great amount of time and attention at an international level. Debates of such importance require clarity about the issues and any arguments advanced by parties need to withstand critical arguments in order to be justifiable. In this paper we seek to take a previous theory developed by the authors of persuasion over action [2] and apply it to this particular scenario to show how it can be used by autonomous agents to effectively aid such reasoning. Section 2 of the paper recapitulates our general approach to practical reasoning through the use of an argument scheme and critical questions. This section then goes on to provide a description of how the argument scheme can be instantiated and how such instantiations can be challenged by autonomous agents based on the Belief-Desire-Intention (BDI) model [12]. We then briefly explain how the presumptive arguments and attacks generated are resolved through organisation into a Value-Based Argumentation Framework (VAF)[4], which enables the calculation of the dialectical status of the various arguments, with respect to the participants in the debate. Section 3 shows how this general theory is applied to the particular example. Our approach models the various participants in the debate as different agents. These agents subscribe to individual beliefs, goals and values, and therefore can represent the different viewpoints that can be brought to bear on the problem. Here we will show the relations between these views and how the arguments can be evaluated through the use of VAFs. In Section 4 we will discuss the notion of accrual and how unacceptable arguments can become acceptable for a particular agent through the support of cumulative arguments, which may individually be unacceptable. In Section 5 we will briefly compare the model of the debate presented here with a system previously implemented by the authors [3] which deals with the same political issue. Finally, in Section 6 we will offer some concluding remarks.

## 2 General Approach to Practical Reasoning with Agents

One way of addressing practical reasoning is as a species of presumptive argument: given an argument, we have a presumptive reason for performing the action. This presumption can, however, be challenged and withdrawn. Subjecting our argument to appropriate challenges is how we hope to identify and consider the alternatives that require consideration, and determine the best choice for us, in the particular context. One account of presumptive reasoning is in terms of argument schemes and critical questions, as given in [11]. The idea here is that an instantiated argument scheme gives a presumption in favour of its conclusion. Whether this presumption stands or falls depends on satisfactory answers being given to the critical questions associated with the scheme, posed in the particular debate.

In [2] we presented a scheme for practical reasoning. Our scheme extends the *sufficient condition scheme* of Walton [11] by unpacking his notion of a goal into three elements: the state of affairs brought about by the action; the goal (the desired features in that state of affairs); and, the value (the reason why those features are desirable). Our scheme is expressed as follows:

AS1: In the current circumstances R  
we should perform action A  
to bring about new circumstances S  
which will realise goal G  
which will promote value V

This presumptive justification may be attacked in a large variety of ways, as described in [7], as the various elements and the connections between them are open to question, and additionally there may be alternative possible actions, and side effects of the proposed action. We will define the subset of attacks used in the example of this paper later on in this section.

The computational setting for our approach is a multi-agent system, in which the agents form intentions based on their beliefs and desires. This is essentially the standard BDI agent model [12], except that we make a small extension by associating each desire with a value, the reason why it is desirable.

First we describe how an agent can construct a presumptive justification for action, instantiating argument scheme AS1.

A BDI agent has a set of beliefs, and we can therefore expect it to be able to reply “true”, “false” or “unknown” when queried about the status of a proposition. For a well-formed formula of standard propositional logic, if all the propositions it contains are given a truth value by the agent, the formula will evaluate either to true or false for that agent. In such cases we say the agent believes, respectively disbelieves, the formula. If some propositions are unknown by the agent, there are three possibilities. First it may be that all assignments to the unknown propositions will give models for the formula, in which case the agent believes the formula. Second it could be that no assignment which makes the formula true is possible, in which case the agent disbelieves it. Third it may be that some assignments are models and some are

not, in which case we say the agent can *assume* the formula to be true. It is important to be able to allow the agents to make assumptions since the knowledge of a particular agent is typically incomplete.

In many typical BDI architectures, e.g. [13] the agent also has a library of plans. These plans are designed to achieve some goal which the agent may wish to bring about. These goals are the *desires* of the agent. The process of practical reasoning is designed to select which desires to commit to achieving, namely to form *intentions*. Each plan will have a set of pre-conditions which when satisfied allow it to be performed, and a set of post-conditions which will become true when the plan is carried out. These post-conditions will include the goal for which the plan is undertaken, but there will typically be additional side effects. A plan is a sequence of atomic actions: here we do not consider the details of plans and take the execution of the entire plan to be what we have previously termed “performing an action”, since deciding which *plan* to perform is the aspect of practical reasoning which we are addressing here. Values are the reasons why the desires are held worth attempting to achieve: certain states of affairs will be held to promote or demote values to differing degrees. The desires of an agent will be states of affairs which promote some value of the agent. It may be necessary to augment standard BDI agents with functions mapping from states of affairs to values.

Now let us return to AS1. The current circumstances R are a conjunction of propositions which the agent believes, or can assume. The action A is some plan in the plan library of the agent which has pre-conditions which are, or can be assumed to be, satisfied in R. The circumstances S result from the application of the post-conditions to R. The goal G is the desire of the agent associated with the plan, and the value V is the value promoted by the realisation of G. These connections can allow us to discover in which ways the agent can, given its beliefs, plans and values, instantiate AS1.

In BDI terms AS1 becomes:

Given the current situation R, there is a plan A which if performed will bring about S, realising G which promotes V.

As well as instantiating AS1 to make a presumptive argument for executing A, agents can also attack such instantiations, using critical questions as described in [2]. We now describe the conditions under which agents can pose critical questions, for each of the attacks used later in this paper. A full list of attacks, and a more formal set of definitions can be found in [1]. The numbers given to the attacks refer to those used in the previous papers [1] and [7]. For each attack we here give the source critical question, a description of when it can be asked and a rendering of the argument it represents.

- **Attack 1a:**

*Source CQ:* Are the believed circumstances true? (CQ1).

*Description:* The agent can assume, but does not believe R.

*Argument:* R may not be true.

- **Attack 2a:**

*Source CQ:* Assuming R is true, will the plan result in S? (CQ2)

*Description:* The agent can assume, but does not believe that executing the plan in R will result in S.

*Argument:* The action may not have the desired consequences.

- **Attack 3a:**

*Source CQ:* Assuming R is true and the plan will result in S, will the post-conditions fulfil the desire? (CQ3)

*Description:* The agent can assume, but does not believe the desire is satisfied in S.

*Argument:* The plan may not fulfil the desire.

- **Attack 4a:**

*Source CQ:* Does the goal realise the value intended? (CQ4)

*Description:* The agent does not believe the value is promoted by the fulfilment of the desire.

*Argument:* The desire does not promote the value.

- **Attack 6:**

*Source CQ:* Are there alternative ways of fulfilling the same desire? (CQ6)

*Description:* There is another plan, the pre-conditions of which can be assumed to be satisfied, and the post-conditions of which will fulfil the desire.

*Argument:* There is an alternative way to satisfy the desire.

- **Attack 8:**

*Source CQ:* Does the plan have a side effect which demotes the value? (CQ8)

*Description:* The post-conditions of the plan can be assumed to realise a state of affairs which demotes the value.

*Argument:* There are side effects of the plan which may demote the value.

- **Attack 9:**

*Source CQ:* Does the plan have a side effect which demotes another value? (CQ9)

*Description:* The post-conditions of the plan can be assumed to realise a state of affairs which demotes some other value.

*Argument:* There are side effects of the plan which may demote some other value.

- **Attack 10:**

*Source CQ:* Does the plan promote some other value? (CQ10)

*Description:* The post-conditions of the plan can be assumed to realise a state of affairs which promotes some other value.

*Argument:* There are side effects of the plan which may promote some other value.

Note that attack 10 does not dispute that the plan should be executed, but questions the motive offered.

Now given a set of agents and a situation in which an action must be chosen we can first instantiate a number of presumptive justifications for each agent, and then see which attacks the agents can make on these justifications. The result will be a set of arguments and a set of attack relations between them, providing the key elements for an argumentation framework [6]. Moreover, since the arguments produced by instantiating AS1 are associated with values, we can see this as a VAF [4], and use the procedures in that paper to calculate the dialectical status of the arguments, with respect to the different audiences represented by the different agents. This process will be illustrated by a detailed working through of our example debate in the section below.

### 3 Political Example

The example we will use is based on the debate which took place in 2003 as to whether the UK should go to war with Iraq (set before Iraq was invaded). This was perhaps one of the most widely debated issues of recent years, and disagreement as to the motives and justification of the action taken remains to this day. Our example will model the viewpoint of the Government in putting forward its position on the issue and some of the attacks that this justification elicited from members of Parliament and the public.

#### 3.1 Context

In our reconstruction of the arguments we will use seven different agents to represent the different views put forward by

the parties involved. Firstly we have four agents advocating the action of invading Iraq but for different, though sometimes overlapping reasons. We will refer to these agents as: G, representing an agent named George; T, representing an agent named Tony; D, representing an agent named Donald; and C representing an agent named Colin. We will also make use of three agents who oppose the action of invading Iraq, again for different reasons. We will refer to these agents as: M, representing an agent named Michael; R, representing an agent named Robin; and J, representing an agent named Jacques.

We begin by instantiating our agents with the appropriate beliefs, desires and values, in accordance with the descriptions given in Section 2.

We use six possible propositions about the world to describe the given situation and these are as follows:

- P1: Saddam has weapons of mass destruction (WMD).
- P2: Saddam is a dictator.
- P3: Saddam will not disarm voluntarily.
- P4: Saddam is a threat to his neighbours.
- P5: Saddam is defying the UN.
- P6: Saddam is running an oppressive regime.

Our agents differ quite widely as to which propositions are believed true. Each agent subscribes to the propositions as shown in Table 1 with 1 representing belief in the proposition, -1 representing disbelief in the proposition and 0 representing unknown to show that the agent has subscribed to neither belief nor disbelief in the proposition.

Table 1: Propositions about the World

Agent	P1	P2	P3	P4	P5	P6
G	1	1	1	1	1	1
T	1	1	1	0	0	1
D	1	0	1	1	0	0
C	0	1	0	0	0	1
M	0	1	-1	-1	-1	1
R	-1	1	0	0	0	1
J	0	1	0	-1	0	1

We now identify the desires and values. We need to identify a set of desires for the agents, and give conditions under which the agents will accept that these desires are realised. We also need to associate these desires with a value, and a degree to which the satisfaction of the desire promotes the value. For now we list the set of desires, conditions and values in Table 2 below and we will go on to discuss degrees of promotion in the next section.

Based upon the beliefs and desires given in the above tables, each agent can provide one or more instantiation of AS1. The figures presented below give two argumentation frameworks to show the views of the agents. Initially we can see that there are two values involved in the debate: ‘world security’ and ‘human rights’. The Government’s argument provides two major justifications (which are instantiations of

Table 2: Possible Desires and Values in the Debate

No.	Desire	Value	Condition to be satisfied
1	No WMD	World security	Iraq has no WMD
2	No dictator	World security	Saddam deposed
3	Democracy in Iraq	Human rights	Saddam deposed
4	International agreement	Good world relations	All allies agree with the action
5	No human casualties	Respect for life	No war

AS1) which endorse the same action of invading Iraq. However, each justification provides different reasons and promotes different values, even though both justifications endorse the same action. We can therefore construct two argumentation frameworks to show the instantiations of AS1 representing the Government’s justifications and the attacks on these instantiations that can be made by the opposing agents. These attacks use the remaining three values. We present these argument schemes, frameworks and attacks below. In the figures, nodes represent arguments. They are labelled with a description of the argument, and on the right hand side, with letters representing the agents introducing the argument. Arcs are labelled with the number of the attack they represent. We then summarise what can be deduced from each framework.

### 3.2 Argument Based on Threat

Firstly we present argument scheme Arg1 showing the Government’s first justification of the action and the attacks made on it by opposing agents:

Arg1

R1: Saddam has WMD, Saddam is a dictator, Saddam will not disarm voluntarily, Saddam is a threat to his neighbours, Saddam is defying the UN

A1: we should invade Iraq

S1: which would get rid of the WMD and depose the dictator

G1: so this will remove the threat that Saddam poses to his neighbours and assert the authority of the UN

V1: which will promote world security.

This argument and the attacks that can be made on it by opposing agents, given their beliefs and desires, are represented in Figure 1 on the next page.

Looking at this argumentation framework we can see that the agents subscribe to the following arguments:

Agents G, T and D all put forward Arg1 to justify the action of invading Iraq. The first challenge to be made on this is executed by agent R who uses attack 1a to deny proposition 1 presented in R1 of Arg1. This attack states that agent R does not believe that Saddam has weapons of mass destruction (WMD). Consequently, agent R also does not believe that invading Iraq will get rid of the WMD and so he

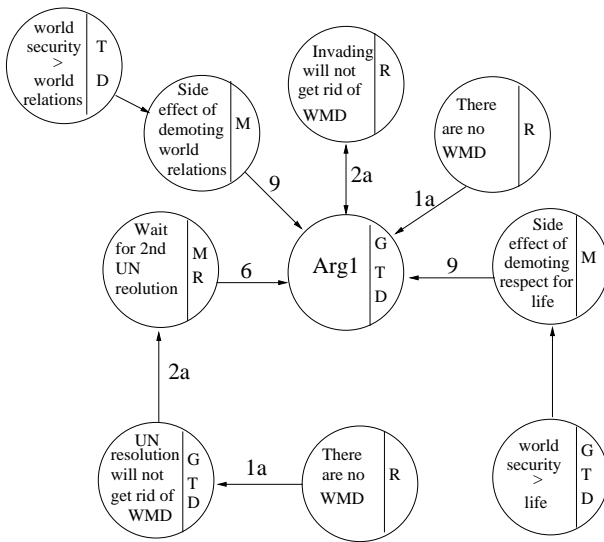


Figure 1. Arg1 and the attacks on it.

makes attack 2a as well. Agent M points out that there is a bad side effect of the action in that the unconsidered value of ‘good world relations’ will be demoted due to there being international disagreement about the proposed invasion. This is stated in attack 9. This is then attacked by agents T and D who state that they rank the value of world security higher than they rank the value of good world relations. Agents M and R then make a new attack to propose an alternative action to realise the goal. Using attack 6, the alternative action they propose here is waiting for a second UN resolution on the matter. However, this is counter-attacked by all agents supporting Arg1, through attack 2a stating that this newly proposed action will not achieve the goal, as waiting for a second UN resolution will not get rid of the WMD. This argument is then itself counter-attacked by agent R who again uses attack 1a to state that he does not believe that there are any WMD in Iraq in the first place. The final attack on Arg1 is made by agent M who believes that the action will have the detrimental side effect of demoting the value ‘respect for life’ and he uses attack 9 to state this. However, this is attacked by all proponents of Arg1 through their statement of value preference in which they rank world security as a more important value than respect for life, in this situation.

### 3.3 Argument Based on Regime Change

Now that all the agents’ arguments have been articulated regarding the justification in Arg1, we turn to Arg2 where agents C and T provide a second justification for the same action:

Arg2

- R2: Saddam is running an oppressive regime
- A2: we should invade Iraq
- S2: to depose Saddam
- G2: which will bring democracy to Iraq
- V2: which will promote human rights.

This argument and the attacks on it are represented in Figure 2 below:

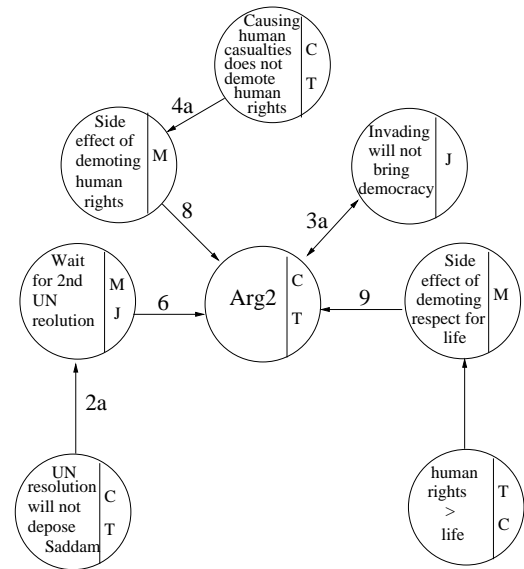


Figure 2. Arg2 and the attacks on it.

Looking at the above argumentation framework we can see that the agents subscribe to the following arguments:

Firstly, we can see from Arg2 that this argument is based on the fact that Saddam is running an oppressive regime and unlike in the last justification no agent disagrees outright with this fact, as we can see from Table 1. So, the first attack made on Arg2 is by agent J who, using attack 3a, states that the action of invading Iraq will not result in democracy being achieved. Agent M then makes attack 8 stating that there is a side effect of the action of demoting the value human rights. This is itself attacked by agents C and T who use attack 4a to state that causing human casualties does not demote human rights as they believe that human casualties may be a necessary evil involved in bringing about democracy. Next, as in the previous framework, attack 6 is used to propose the alternative action of waiting for a second UN resolution and this is put forward by agents M and J. As before, this argument is counter-attacked using attack 2a, this time to state that the alternative action will not achieve the goal of deposing Saddam and this is put forward by both agents supporting Arg2. However, this time no agent can attack this argument as agent M did in the previous framework, because they all believe the facts upon which the argument is based to be true. The final attack made on Arg2 is attack 9 in which agent M states that the action of invading Iraq again has the side effect of demoting the value of respect for life. Agents C and T both attack this by stating their belief that the value of human rights (in relation to the achievement of democracy in Iraq) is more important than the value of respect for life. This concludes the analysis of all the individual attacks used in each framework.

### 3.4 Discussion

We can see from the above analysis that all agents involved in the discussion have different, but sometimes overlapping reasons for their opinions in the debate. In framework 1 we are able to see that agents G, T and D all accept Arg1 on the basis that they believe Saddam has WMD which he is willing to use to detrimental effect. However, only agents T and D express a value preference of world security over world relations, which they use to defeat attack 9. From this we can see that agent G agrees that there may be the possible side effect of the action of demoting world relations, which is pointed out in attack 9. However, he does go on to defend an attack against the second use of attack 9 by stating that he ranks world security over respect for life. From this we can deduce that agent G only needs to have one goal (as opposed to all goals) of Arg1 satisfied in order to justify the action: asserting the authority of the UN is not important to G. As both agents T and D defend all attacks made by the opposing agents, they require all consequences and goals to be satisfied in order for them to be able to justify the action.

Of the opposing agents in framework 1, agent R rejects Arg1 on the basis that he believes the facts upon which it is based are false i.e. there are no WMD. Agent M rejects the argument on a different basis through reasons that reveal he does not support war rather than refuting the claim that there are WMD. Agent J does not feature in this framework as his views solely relate to arguments about the conveyance of democracy.

In framework 2 we can see that agent T supports this argument in addition to Arg1 and so he is the only agent who sees the need to justify both arguments in order to be able to justify the action. However, agent C also supports Arg2 and as he did not support Arg1 we conclude that he believes Arg2 to be sufficient on its own for the justification of the action.

Of the opposing agents M again reveals his anti-war attitude through the attacks he makes. Agent J disagrees with the result of the action showing that his attitude reflects the belief that democracy will not be achieved through invasion, which is the main thesis of his position. Agent R does not feature in this scenario as he is only interested in arguments resting on the basis of the evidence of WMD.

## 4 Accrual and Strength of Argument

In addition to the individual attacks in the frameworks there is also an attack that can be made between the two argument scheme instantiations Arg1 and Arg2, and this is attack 10. This is shown below in Figure 3.




Figure 3. Arg1 and Arg2 mutually attack each other.

An agent making attack 10 does not dispute that the action should be performed, but disputes the motive for performing it. In our example, G and D accept Arg1 but not Arg2 and

C accepts Arg2 but not Arg1. Thus G may attack Arg2 by saying that regime change is not a justification for invasion, but removing WMD is, and C may attack Arg1 by arguing the contrary. The case of T is different, since he accepts both arguments. If T accepts that both Arg1 and Arg2 are sufficient to justify invasion, he could be challenged to choose between world security, the value promoted by Arg1, and human rights, the value promoted by Arg2, so as to clarify his “real” reason for advocating invasion. In practice some politicians seemed to be in the position of T, and generally made the removal of WMD their lead justification, although subsequent to the failure to discover WMD, they cite regime change as sufficient in itself. If, however, desires can promote values more or less strongly, it may be that only one of the arguments is sufficient to justify the action. This will then be the “real” reason, and the other argument is superfluous. A different case is where neither argument is sufficient by itself to promote the action. Here attack 10 is inappropriate, since the two arguments are now intended to be mutually supporting and the action is justified only if both arguments stand. An additional pre-condition for attack 10 is thus that the attacking argument be sufficient to justify the action. The need in some cases to have mutually supporting arguments introduces the notion of accrual. To explore this notion we need to have some mechanism for distinguishing degrees of promotion, and determining when an argument is sufficient to justify the action. The descriptions in Section 2 can be extended to allow for both these ideas, and we illustrate them with an example in the next sub section.

### 4.1 Degrees of Promotion

The descriptions of how a BDI agent can instantiate and attack a justification for action, as given in Section 2, allow for differing degrees of promotion in a value to be expressed. This is intended to be flexible so as to accommodate a variety of concrete treatments. Assigning and combining things such as degrees of promotion presents conceptual and practical problems. What we propose here is admittedly rudimentary: we make no claims for its cognitive validity, but use it only to illustrate some points about accrual. We assume that agents can make a subjective assessment of the degree of promotion by responding to a question such as “from the standpoint of human rights, how important is it that a country be ruled democratically” with a qualitative assessment such as “utterly, very, somewhat, a little, or not at all”. We further assume that satisfying several desires promoting a value will promote that value to a degree greater than is achieved by satisfying only one of them. Our treatment is consistent with these assumptions, but remains rather *ad hoc*. We translate the qualitative assessments into numbers as follows:

- if the agent replies “not at all” the value is *not* promoted and we assign the number 0;
- if the agent replies “a little” the value is *weakly* promoted and we assign 0.3;
- if the agent replies “somewhat” the value is *moderately* promoted and we assign 0.5;
- if the agent replies “very” the value is *strongly* promoted and we assign 0.7;

- if the agent replies “utterly” the value is *fully* promoted and we assign 1.

For combinations of two desires promoting the same value, or two arguments proposing the same action we add the relevant numbers and subtract their product: i.e.  $combine(a,b)=(a+b)-(a*b)$ . This is the formula used for rule combination in MYCIN [5]. Again we make no claims for this other than that it satisfies the desired property of increasing the degree of satisfaction while remaining in range. We also, again quite arbitrarily, take 0.7 as the threshold which must be attained if the action is to be justified.

We now return to the example. Agent T needed to accept both Arg1 and Arg2 to convince him to act. Suppose he sees both arguments as moderately supporting the action, both with value 0.5. Separately neither is a sufficient reason to act, but together they support the action at 0.75, and so exceed the threshold. Note that if one argument had been no more than weakly supportive, the combined value would be only 0.65: some third argument offering at least weak support would be required to reach the threshold. In contrast, agents G and D see Arg1 as sufficient to justify the action on its own, and agent C sees Arg2 as sufficient in itself. For these agents there is a single reason to act and attack 10 will lead them to reject the other argument.

This illustrates one form of accrual, where two distinct arguments are involved. Arg1 illustrates a different form of accrual in that it is based on the satisfaction of two desires promoting the same value. If we separate these into Arg1a, based on the removal of Saddam’s threat, and Arg1b, based on asserting the authority of the UN, we can see that there are various possibilities. It may be that for some agents one of the desires promotes the value sufficiently to support the action on its own. Suppose that, as for agent G in Figure 1, the removal of the threat strongly promotes world security. Then Arg1a provides sufficient justification and G need not defend attacks directed at the part of the goal representing Arg1b. Agents D and T, in contrast, may see both parts of the goal as necessary (because both only moderately promote the value) and so must defend it against all attacks. Thus we can use two distinct types of accrual depending on whether one or more values are promoted. Different values require two arguments and we can ask which provides the *real* motive to act. One value can be represented as a single instance of the argument scheme. In this case, because only one value is involved, a person who believes that an argument based on one of the desires is sufficient need not reject the argument based on the other desire: the second argument is at worst superfluous and may strengthen the justification for the action.

It is often the case that in everyday practical reasoning we may be unsure as to whether to execute an action or not. However, if we find a separate reason (even if it is weak itself) to execute the action then this can compel us enough to carry out the action. This notion of accrual is also mentioned by Walton in [11] in his argument scheme for the ‘Argument From Sign’. Here he states that the more signs that are brought into the scheme, the more inferences can be drawn from them and the more the case builds up. This is particularly relevant for the political forum as the more evidence that is presented to

justify an action, the more likely ministers and the public are to accept the argument, enabling the Government to win their support.

The mechanisms we have used to implement degrees of promotion and strengths of justification have been simple and preliminary. They are intended only to illustrate how our model can capture the useful notion of accrual, and can distinguish different types of accrual. We hope to give a more principled account of this in future work, though the treatment of it may be domain specific and we will look to existing accounts to guide our future work on this area.

When we introduced the application of our theory to BDI agents in Section 2, we allowed them only to express belief in, disbelief in, or ignorance of a proposition. Some implementations might allow agents to express degrees of certainty in propositions. In such cases we could allow agents to assume a proposition to be true only if it attained a certain degree of likelihood. This could be important in domains where risks need to be taken into account, such as medicine. Actions with severe penalties if performed in error would require a greater degree of conviction before their pre-conditions could be assumed than those which have less serious potential consequences. Such considerations remain for future exploration.

## 5 Comparison with the PARMENIDES System

In earlier work [3] we presented an implementation of a system named Parmenides, which is based upon the same general theory of persuasion over action that we described earlier in this paper. The idea of Parmenides is to provide a system which makes use of our argument scheme and critical questions to solicit the public’s views on a particular issue. The topic used in Parmenides is the same as that discussed here: the invasion of Iraq. Parmenides provides a simple web-based interface to guide users through the justification of an action as presented by the Government. Users are given the opportunity to disagree with the individual elements of the position and propose alternatives. These points of disagreement each represent an individual attack from our theory of persuasion. In addition to this, after having critiqued the position presented to them, users can then go on to construct their own position regarding the matter. The responses of the users are written to a database to enable the proponent of the action (in this case the Government) to gather and analyse the information in order to identify what elements of an argument are more strongly supported than others. There are some interesting points to note when comparing Parmenides with the example application presented here.

Parmenides makes use of a subset of the attacks from our theory, as does the example presented here. These two subsets of attacks are extremely similar with the exception of three additional attacks being used in Parmenides (attacks 7, 11 and 15). The purpose of Parmenides was to build a system in which a user could critique a justification of an action in a particular domain and express their own views in the most complete way possible, using our theory. However, it would be perfectly acceptable to use the extra attacks found in Parmenides in the example presented here: they do not in fact

arise because we have limited ourselves to only seven agents, whose beliefs and desires do not happen to satisfy the pre-conditions for these attacks. Thus, it would be possible to reconstruct all the arguments made in Parmenides in the format we have presented in this paper, though this was not our original aim of the exercise.

One of the main motivations of Parmenides was to provide a system which facilitated debate between the Government and members of the public whilst being grounded in a firm model of argument that was transparent to the user. Conversely, the model presented in this paper is intended for use solely by autonomous computer agents. However, we believe that there may be a useful link between the two models. As mentioned earlier in this section, all the information entered into the Parmenides system is stored in a back-end database. Therefore, it would be possible to reconstruct new positions on the issue from the users' responses by introducing agents to represent their views. These new positions could then be used as input to generate presumptive arguments to be used by BDI agents, as demonstrated in this paper. This would allow us to gather a wide range of differing views on the topic and evaluate the warrant of each view. As part of the practical reasoning process this would ensure that all possible scenarios have been considered and thus aid us in choosing the best action and justification for the issue in question. This is would be an interesting avenue to pursue between our two systems and is something that we hope to explore in future work.

## 6 Concluding Remarks

In this paper we have used an approach to modelling practical reasoning and persuasive argument to represent the reasoning in a recent and important international political debate. Our account demonstrates how superficial agreement may conceal subtle but important differences in beliefs and aims. We recapitulated a previous theory of persuasion over action and showed how this can be made computational for use in BDI agents enhanced to deal with the notion of values. We applied this general theory to the political domain to show how real life issues can be debated with computational agents using our model. We went on to show how such argumentation can make effective use of the notion of accrual of arguments. This concept is an important feature of persuasive debate and we believe that our model is able to give some insight into this phenomenon. Finally, we drew some comparisons between this method of argumentation and a previous system built by the authors based on the same theory of persuasion over action. In future work we hope to apply the methods presented here to different domains, such as the medical one, where decisions are often based on the degree of certainty assumed to be attached to expected effects of actions. This will enable us to explore in more detail the notions of accrual and strength of arguments and how these differ across domains.

## Acknowledgements

Katie Atkinson is grateful for support from the EPSRC. Trevor Bench-Capon and Peter McBurney acknowledge partial support received from the European Commission, through Project ASPIC (IST-FP6-002307). We also thank the anonymous reviewers for their comments and suggestions.

## References

- [1] K. Atkinson, T. Bench-Capon, and P. McBurney. Attacks on a presumptive argument scheme in multi-agent systems: pre-conditions in terms of beliefs and desires. Technical Report ULCS-04-015, Department of Computer Science, University of Liverpool, UK, 2004.
- [2] K. Atkinson, T. Bench-Capon, and P. McBurney. Justifying practical reasoning. In F. Grasso, C. Reed, and G. Carenini, editors, *Proc. of CMNA workshop 2004*, pages 87–90, Valencia, Spain, 2004.
- [3] K. Atkinson, T. Bench-Capon, and P. McBurney. Parmenides: Facilitating democratic debate. In R. Traunmüller, editor, *Electronic Government 2004*, LNCS 3183, pages 313–316. Springer, Berlin, 2004.
- [4] T. Bench-Capon. Persuasion in practical argument using value based argumentation frameworks. *Journal of Logic and Computation*, 13 3:429–48, 2003.
- [5] B. G. Buchanan and E. H. Shortliffe. *Rule-Based Expert Systems: The MYCIN Experiments of the Stanford Heuristic Programming Project*. Addison-Wesley, Reading, MA, USA, 1984.
- [6] P. M. Dung. On the acceptability of arguments and its fundamental role in nonmonotonic reasoning, logic programming and n-person games. *Artificial Intelligence*, 77:321–357, 1995.
- [7] K. Greenwood, T. Bench-Capon, and P. McBurney. Towards a computational account of persuasion in law. In *Proc. of the Ninth Int. Conf. on AI and Law*, pages 22–31, New York, NY, USA, 2003. ACM Press.<sup>1</sup>
- [8] A. Hunter. Making argumentation more believable. In *Proceedings of the Nineteenth American National Conference on Artificial Intelligence (AAAI' 2004)*, pages 269–274, Cambridge, MA, USA, 2004. MIT Press.
- [9] A. Hunter. Towards high impact argumentation. In *Proceedings of the Nineteenth American National Conference on Artificial Intelligence (AAAI' 2004)*, pages 275–280, Cambridge, MA, USA, 2004. MIT Press.
- [10] C. Perelman. *Justice, Law, and Argument*. D. Reidel Publishing Company, Dordrecht, Holland, 1980.
- [11] D. N. Walton. *Argument Schemes for Presumptive Reasoning*. Lawrence Erlbaum Associates, Mahwah, NJ, USA, 1996.
- [12] M. J. Wooldridge. *Reasoning about Rational Agents*. MIT Press, Cambridge, MA, USA, 2000.
- [13] M. J. Wooldridge. *An Introduction to MultiAgent Systems*. John Wiley and Sons, New York, NY, USA, 2001.

---

<sup>1</sup>Greenwood is the maiden name of the first author of this paper.