

Multi-Agent Argumentation for eDemocracy

Katie Atkinson ^a

Trevor Bench-Capon ^a

Peter McBurney ^a

^a *Department of Computer Science, Chadwick Building
University of Liverpool, Liverpool L69 3BX, UK.*
{katie,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 begin with a discussion of practical reasoning in general and we draw attention to differences between practical reasoning and reasoning about beliefs. We suggest that practical arguments should be treated as a species of presumptive reasoning, best handled using argument schemes and associated critical questions, in order to justify a proposal for action. We extend the argument scheme for practical reasoning and its critical questions proposed by Walton. We further 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 also examine how this application compares with a previously implemented system which makes use of the same underlying theory.

1 Introduction

Although many of the arguments that are deployed in everyday life are concerned with what it is sensible or practical to do, the topic of practical reasoning has been rather neglected by philosophers. Practical reasoning has, of course, been addressed (see, e.g., [14] for a collection of essays and [15] for a recent monograph), but it has received nothing like the attention that has been paid to reasoning about beliefs. When action has been considered, it has most often been in the context of ethics, considering what is morally right or wrong, rather than what is prudentially or practically useful. It has been similarly neglected in computer science, where practical reasoning has been treated as little different from deduction, standard backward chaining techniques being applied to rules with goals as consequents, and pre-conditions and actions as antecedents. The academic discipline which has perhaps focused most attention on the selection and justification of actions is economics. However, there the widespread adoption of an overly narrow definition of rationality¹ has hindered understanding of practical reasoning and even impeded progress in the discipline, as, for example, Nobel Memorial Laureate Amartya Sen has recently argued [16].

In this paper we first discuss some of the differences between reasoning about belief and reasoning about action which cause problems with approaches based on the practical syllogism. We then discuss the treatment of practical reasoning in [17] which makes use of an argument scheme and associated critical questions. We elaborate this scheme and extend the critical questions, and relate this to our previous work. We then describe an example application of this approach for use in BDI agents.

Section 2 of the paper examines the nature of the practical syllogism as a basis for practical reasoning in computer science and discusses the problems associated with this representation. Section 3 examines the use of argument schemes and critical questions as a way of embodying practical reasoning through presumptive argument. In particular, this is discussed in relation to the account of Walton [17] and here we propose our extension to one of his argument schemes and critical questions. Section 4 provides a short description of how agents based on the Belief-Desire-Intention (BDI) model [18] can instantiate the argument scheme and challenge particular instantiations of it. These presumptive arguments and attacks generated can then be resolved through organisation into a Value-Based Argumentation Framework (VAF) [6], which enables the calculation of the dialectical status of the various arguments, with respect to the participants in the debate. Section 5 shows how this general theory is applied to the particular example scenario. Our approach models

¹Oscar Lange, for example, defined *rationality* as the maximization of some quantity [12].

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 6 we will make some brief remarks on the notion of accrual of arguments. In Section 7 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 8 we will offer some concluding remarks.

2 Problems With The Practical Syllogism

Practical reasoning in computer science can predominately be seen as based on a form of the practical syllogism. An example from [11] is:

K1 I'm to be in London at 4.15
If I catch the 2.30, I'll be in London at 4.15
So, I'll catch the 2.30.

This, however, cannot be quite right. It may well be possible to accept both the premises and deny the conclusion. There are at least three bases for criticism:

C1 K1 represents a species of abduction, and so there may be alternative ways of achieving the goal.

C2 Performing an action typically excludes the performance of other actions, which might have other desirable results; these may be more desirable than the stated goal.

C3 Performing an action typically has a number of consequences. If some of these are undesirable, they may be sufficiently bad to lead us to abandon the goal.

In order to act on the basis of an argument such as K1 therefore, we need to consider alternative actions, alternative goals and any additional consequences, and then choose the best of these alternatives. Note the element of choice: we can choose our goals and actions in a way in which we cannot choose our beliefs, and different people may rationally make different choices. As Searle puts it:

“Assume universally valid and accepted standards of rationality, assume perfectly rational agents operating with perfect information, and you will find that rational disagreement will still occur; because, for example, the rational agents are likely to have different and inconsistent values and interests, each of which may be rationally acceptable.” [15, xv]

In a sense therefore any practical argument is directed to a specific person at a specific time, to encourage them towards a particular choice and the objectivity that we can find in factual matters cannot in general be attained in practical reasoning. An attempt to modify K1, similar to one put forward by Searle in [15] (although not regarded by him as satisfactory) is:

S1 I want, all things considered, to achieve E
The best way, all things considered, to achieve E is to do M
So, I will do M.

There are problems with this: we cannot in general consider all things, because we have limited reasoning resources and imperfect information. Nor is it easy to say what is meant by “best” here. In computer science there are often attempts to define best using some kind of utility function but, as Searle points out in [15], any preference ordering is more often the *product* of practical reasoning than an input to it. Coming to understand what we think is best is part of what we do in practical reasoning.

One way of addressing these problems is to regard practical reasoning as a species of presumptive argument. Given an argument like K1, 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 [17]. The idea here is that an argument scheme gives a presumption in favour of its conclusion. Whether this presumption stands or falls depends on satisfactory answers being given to the associated critical questions, posed in the particular situation.

3 Argument Schemes for Practical Reasoning

In [17] Walton gives a scheme for practical reasoning called the sufficient condition scheme²:

W1 G is a goal for agent *a*
 Doing action A is necessary for agent *a* to carry out goal G
 Therefore agent *a* ought to do action A.

He associates four critical questions with this scheme:

CQ1: Are there alternative ways of realising G?

CQ2: Is it possible to do A?

CQ3: Does agent *a* have goals other than G which should be taken into account?

CQ4: Are there other consequences of doing A which should be taken into account?

CQ1, CQ3 and CQ4 relate respectively to the criticisms C1, C2 and C3 identified above. We believe, however, that the argument scheme and the critical questions both need elaboration. Firstly, we believe that the notion of a goal as used in W1 is ambiguous.

Consider the following situation. I am in Liverpool. My friend X in London is about to go to Australia indefinitely, and I am eager to say farewell to him. To catch him before he leaves London, it is necessary that I arrive in London before 4.30. So I may say:

A1 I want to be in London before 4.30
 The 2.30 train arrives in London at 4.15
 So, I shall catch the 2.30 train.

Here I am justifying my action in terms of one of its consequences. Alternatively I may say:

A2 I want to see X before he leaves London
 The 2.30 train arrives in London at 4.15
 So, I shall catch the 2.30 train.

Here the action is not justified by its direct consequences, but by something else that follows from it. I do not really desire to be in London at all, except in so far as it is a means to the end of seeing X. Alternatively there is a third justification:

A3 Friendship requires that I see X before he leaves London
 The 2.30 train arrives in London at 4.15
 So, I shall catch the 2.30 train.

Now I justify my action not in terms of its direct consequences, nor in terms of a state of affairs which will result from the action, but in terms of the underlying social value³ which I hope to promote by the action.

In general we may write instead of:

W1a G is a goal for agent *a*

²In this scheme we label each of Walton's symbols for clarity.

³Values are distinguished from goals as subjective assessments of states of affairs and by being desirable in themselves, unlike goals, which are desired only in so far as they promote some value.

P1 Agent *a* wishes to achieve *S* so as to bring about *G* which promotes a value *V*

Note that the answers to CQ1 are different in the cases A1-3:

- In the case of A1, I must propose other ways of arriving in London on time, perhaps by driving;
- In the case of A2 I need not go to London at all; for example I could drive to Heathrow and say goodbye at the airport;
- In the case of A3 I need not meet with *X* at all; perhaps a telephone call and an apology will be enough to promote friendship.

Given this more refined notion of a goal we can extend CQ1 to:

CQ1a Are there alternative ways of realising the same consequences?

CQ1b Are there alternative ways of realising the same goal?

CQ1c Are there alternative ways of promoting the same values?

Similarly, the remaining critical questions can be elaborated, as discussed in [2].

Secondly, apart from the possibility of the action, Walton does not consider other problems with soundness of *W1*, presupposing that the second premise is to be understood in terms of what agent *a* knows or reasonably believes. In [8] we proposed an argument scheme which incorporates P1 and makes the factual context explicit:

AS1 In the current circumstances *R*
We should perform action *A*
Which will result in new circumstances *S*.
Which will realise goal *G*.
Which will promote some value *V*.

In this scheme we can identify three further critical questions and variants⁴. Again, see [2] for a complete discussion of these extended critical questions.

Our elaborated set of critical questions is shown in Table 1. Each critical question can be seen as an attack on the argument it is posed against. The fourth column shows the number of variants for each critical question.

We therefore believe that in an argument about a matter of practical action, we should expect to see one or more *prima facie* justifications advanced stating, explicitly or implicitly, the current situation, an action, the situation envisaged to result from the action, the features of that situation for which the action was performed and the value promoted by the action, together with negative answers to critical questions directed at those claims.

4 General Approach to Practical Reasoning with Agents

The computational setting for the application of 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 (see e.g., [18]), except that we make a small extension by associating each desire with a value, the reason why it is desirable. In [1] and [4] we formally described how our theory can be made computational for use by BDI agents. We have done this through provision of a formal set of definitions to show how an agent can put forward a position regarding the justification of an action. This is done by instantiating argument scheme AS1, with respect to a given agent's beliefs, desires, intentions, plans and values. As the BDI architecture has no element to represent values in our model we augment standard BDI agents with functions mapping from states of affairs to values. Thus, AS1 is instantiated by a BDI agent in the following manner:

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

⁴Variants differ according to the degree of commitment to alternatives on the part of the questioner, e.g., a feature of the circumstances may be simply denied, or an alternative proposed.

Table 1: Critical Questions associated with AS1

CQ	Characterisation	No. of variants
1	Disagree with the description of the current situation	2
2	Disagree with the consequences of the proposed action	7
3	Disagree that the desired features are part of the consequences	6
4	Disagree that these features promote the desired value	4
5	Believe that the consequences can be realised by some alternative action	1
6	Believe that the desired features can be realised through some alternative action	1
7	Believe that the desired value can be realised in an alternative way	2
8	Believe that the action has undesirable side effects which demote the desired value	1
9	Believe that the action has undesirable side effects which demote some other value	1
10	Agree that the action should be performed, but for different reasons	1
11	Believe that the action will preclude some more desirable action	3
12	Believe that the circumstances as described are not possible	1
13	Believe that the action is impossible	1
14	Believe that the consequences as described are not possible	1
15	Believe that the desired features cannot be realised	1
16	Disagree that the desired value is worth promoting	1

assumed to be, satisfied in R. The circumstances S result from the application of the post-conditions of A 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.

As well as instantiating AS1 to make a presumptive argument for executing action A, agents can also attack such instantiations, using critical questions as described in Section 3. In [1] and [4] we formally described the conditions under which agents can pose critical questions, for each of the attacks used later in this paper. Due to space limitations we will not reproduce all these definitions here, but instead we give two examples of informal descriptions of two attacks that occur in our theory and the critical question from which they are derived.

- **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.

So, 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 [7]. Moreover, since the arguments produced by instantiating AS1 are associated with values, we can see this as a VAF [6], 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 next section.

5 Political Example

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 compromises 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 [13] and has also recently been addressed in AI by Hunter in [9] and [10]. 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 [6] 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 dedicated to it 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 next section we take our account of persuasion over action and apply it to this particular scenario to show how it can be used by autonomous agents to effectively aid such reasoning.

5.1 Context

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.

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 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. 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 2 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 2: **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, which will be promoted through realisation of the desire. These are all shown in Table 3.

Table 3: **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

Based upon the beliefs and desires given in the above tables⁵, 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 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 agent(s) introducing the argument. Arcs are labelled with the number of the attack they represent. We then summarise what can be deduced from each framework.

5.2 Argument Based on Threat to World Security

Firstly we present argument scheme Arg1 showing the government’s first justification of the action:

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.

⁵As will be discussed in section 5.3, agent J does not accept that desire 3 from Table 3 is achieved by the condition stated. He believes this condition to be sufficient but not necessary, and he may have some other condition that he believes will satisfy the desire.

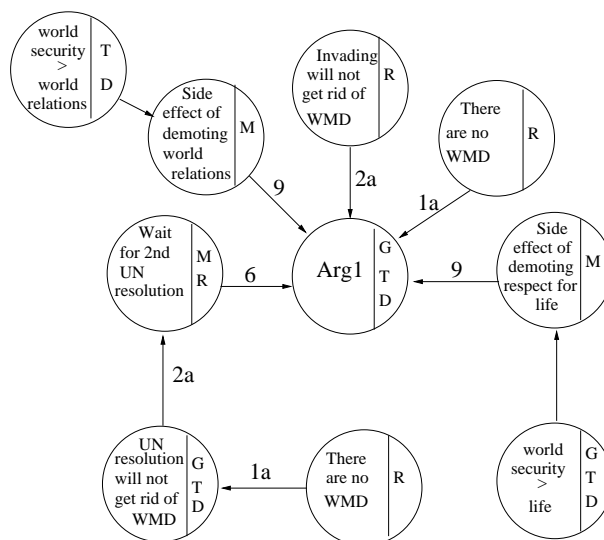


Figure 1. Arg1 and the attacks on it.

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), and this argument is given the value of 'truth', as it is a factual argument. This follows the use of VAFs in [6], where factual arguments are given the value 'truth' and this is ranked as the most important value by all audiences. Consequently, agent R also does not believe that invading Iraq will get rid of the WMD and so he makes attack 2a as well, which again takes the value of 'truth'. 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.

5.3 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:

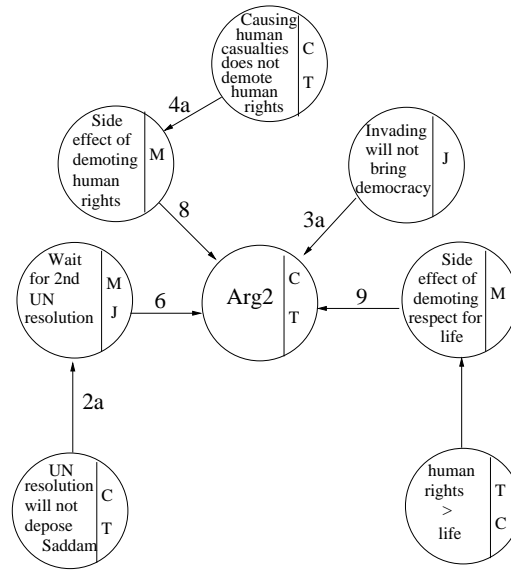


Figure 2. Arg2 and the attacks on it.

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 2. 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, since it requires more than just deposing Saddam to achieve democracy. 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.

5.4 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 the first instance of 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. Agents J and C do not feature in this framework as their views solely relate to arguments about the conveyance of democracy.

In framework 2 we can see that agent T supports Arg2 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.

6 Issues Arising From the Example

In this section we make brief remarks on an interesting issue, regarding the notion of accrual of arguments, that has arisen from the worked example in the previous section.

In addition to the individual attacks in the frameworks in Figures 1 and 2 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. The need in some cases to have mutually supporting arguments introduces the notion of accrual. In [5] we gave a brief outline of how our account could be extended to support this notion. We hope to further explore this concept, along with other similar issues regarding the strength of argument, degrees of promotion of values and uncertainty of beliefs, in future work.

7 Comparison with the PARMENIDES System

In earlier work [3] we presented an implementation of a system named Parmenides⁶, 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 was 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

⁶The system can be used at: <http://www.csc.liv.ac.uk/~katie/Parmenides.html>

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 16). 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 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 would be an interesting avenue to pursue between our two systems and is something that we hope to explore in future work.

8 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 provided a discussion of the problems inherent with practical reasoning in the form of the practical syllogism. In order to address these problems we suggested that practical arguments should be treated as a species of presumptive reasoning which is embodied in the notion of argument schemes and associated critical questions. We extended the argument scheme for practical reasoning, and its critical questions, proposed by Walton. We then described how this account 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 also drew attention to issues of accrual of arguments that can arise. 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 would also like to thank the anonymous referees for their comments and suggestions. Some of the work contained in this paper has previously been presented at the CMNA 2004 and CMNA 2005 workshops. We would like to thank the anonymous referees and participants of those workshops for their helpful comments.

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, *Proceedings of the Fourth International Workshop on Computational Models of Natural Argument (CMNA 2004)*, pages 87–90, Valencia, Spain, 2004.
- [3] K. Atkinson, T. Bench-Capon, and P. McBurney. PARMENIDES: Facilitating democratic debate. In R. Traummüller, editor, *Electronic Government 2004*, LNCS 3183, pages 313–316. Springer, 2004.
- [4] K. Atkinson, T. Bench-Capon, and P. McBurney. Arguing about cases as practical reasoning. In *Proceedings of the Tenth International Conference on Artificial Intelligence and Law (ICAIL 2005)*, pages 35–44, New York, NY, USA, 2005. ACM Press.
- [5] K. Atkinson, T. Bench-Capon, and P. McBurney. Persuasive political argument. In F. Grasso, C. Reed, and R. Kibble, editors, *Proceedings of the Fifth International Workshop on Computational Models of Natural Argument (CMNA 2005)*, pages 44–51, Edinburgh, Scotland, 2005.
- [6] T. Bench-Capon. Persuasion in practical argument using value based argumentation frameworks. *Journal of Logic and Computation*, 13 3:429–48, 2003.
- [7] 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.
- [8] K. Greenwood, T. Bench-Capon, and P. McBurney. Towards a computational account of persuasion in law. In *Proceedings of the Ninth International Conference on AI and Law (ICAIL 2003)*, pages 22–31, New York, NY, USA, 2003. ACM Press.
- [9] 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.
- [10] A. Hunter. Towards higher impact argumentation. In *Proceedings of the Nineteenth American National Conference on Artificial Intelligence (AAAI 2004)*, pages 275–280, Cambridge, MA, USA, 2004. MIT Press.
- [11] A. J. P. Kenny. *Practical Reasoning and Rational Appetite*. 1975. Reprinted in [14].
- [12] O. Lange. The scope and method of economics. *The Review of Economic Studies*, 13(1):19–32, 1945–1946.
- [13] C. Perelman. *Justice, Law, and Argument*. D. Reidel Publishing Company, Dordrecht, Holland, 1980.
- [14] J. Raz, editor. *Practical Reasoning*. Oxford University Press, Oxford, UK, 1978.
- [15] J. R. Searle. *Rationality in Action*. MIT Press, Cambridge, MA, USA, 2001.
- [16] A. Sen. *Rationality and Freedom*. Harvard University Press, Cambridge, MA, USA, 2002.
- [17] D. N. Walton. *Argument Schemes for Presumptive Reasoning*. Lawrence Erlbaum Associates, Mahwah, NJ, USA, 1996.
- [18] M. J. Wooldridge. *Reasoning about Rational Agents*. MIT Press, Cambridge, MA, USA, 2000.