Towards a Computational Account of Persuasion in Law

Katie Greenwood, Trevor Bench-Capon and Peter McBurney
Department of Computer Science
University of Liverpool
Liverpool L69 7ZF UK
{k.m.greenwood,tbc,p.j.mcburney}@csc.liv.ac.uk

ABSTRACT

In this paper we attempt to give an account of reasoning with legal cases contextualised within a general theory of persuasion in practical reasoning. We begin by presenting our general theory, concentrating on the variety of ways in which a particular position can be attacked. We then apply our theory to the legal domain, illustrating our approach by a case study based on the well known CATO system. From this we conclude that it is possible to see reasoning with legal cases as a particular instantiation of our general theory. We identify some points of interest for discussion, and conclude by stating our intended directions for future work.

1. INTRODUCTION

One way of looking at reasoning about legal cases is to see it as an exercise in persuasion. Two parties take opposing views on a case and advance arguments designed to persuade some third, supposedly neutral, person to adopt their point of view. In previous work [6, 7] we have examined persuasion dialogues in general and have produced a categorisation of ways in which positions can be attacked in such dialogues. Our approach takes as its starting point an influential typology of human dialogues, due to Walton and Krabbe [12]. In this paper we attempt to apply this work to the legal domain. However, we consider only the opening stages of an argument, leaving counter-attacks, shifts in the burden of proof, and similar issues to later work.

Section 2 will present the general theory. Section 3 will adapt the general theory to the specific example of law. Section 4 will illustrate the theory with examples drawn from a particular legal domain. Section 5 gives a detailed account of how the representation of section 4 can be used to generate arguments. Section 6 will offer a discussion of some points of interest and section 7 some concluding remarks.

ICAIL'03, June 24–28, 2003, Edinburgh, Scotland.

2. THE GENERAL THEORY

2.1 Elements of the Formalisation

The kind of persuasion that we consider is where one person is attempting to persuade another to adopt a course of action, and that other person is arguing against this. Persuasion is intended to be rational, and so reasons are advanced, and attacked, by each side. Moreover, the persuasion is intended to lead to action, so the debates are examples of practical reasoning.

We give the following as the general schema for a position motivating an action:

> (AS1) In the Current Circumstances R
> we should perform Action A
> to achieve New Circumstances S
> which will realize some goal G
> which will promote some value V.

Persuasion need not apply exclusively to future actions. In looking at past actions we may seek to justify them by persuading our audience that the action actually performed was the correct action to be performed at that time. For past actions, the schema is:

> (AS2) In the Circumstances R
> Action A was rightly performed
> to achieve Circumstances S
> which realized some goal G
> which promoted some value V.

As far as formalization of these schemas goes, we can treat both as equivalent. We need recognize no difference between resolving on a future action and justifying a past action. Moreover, an action may achieve multiple goals, and each goal may promote multiple values. For simplicity, we assume that the proponent of an action articulates an argument in the form of schema 1 for each goal realized and value promoted. We may then formalize the schemas as follows. We assume the existence of:

- A finite set of distinct actions, denoted *Acts*, with elements, A, B, C, etc.
- A finite set of propositions, denoted *Props*, with elements, p, q, r, etc.
- A finite set of states, denoted *States*, with elements, R, S, T, etc. Each element of *States* is an assignment of truth values $\{T, F\}$ to every element of *Props*.
- A finite set of propositional formulae, *Goals*, called goals, with elements G, H, etc.

- A finite set of values *Values*, with elements v, w, etc.
- A function *value* mapping each element of *Goals* to a pair $< v, sign >$, where $v \in$ *Values* and $sign \in \{+, -\}$.
- A ternary relation *apply* on *Acts* × *States* × *States*, with *apply(A, R, S)* to be read as: *"Performing action A in state R results in state S."*[1]

The argument schemas AS1 and AS2 contain a number of problematic notions which are not readily formalized in classical logic. We can, however, see that there are four classical statements which must hold if the argument represented by schema AS1 is to be valid:

Statement 1: R is the case.

Statement 2: *apply(A, R, S)* ∈ *apply*.

Statement 3: S $\models$ G (G is true in state S).

Statement 4: *value(G)* $=< v, + >$.

With appropriate change of tense, the same four statements must hold for AS2 to be valid.

2.2 Attacks on a Position

We may identify a number of different ways in which a position AS1 or AS2 can be attacked. In the next four sub-sections, we present a comprehensive list of these.

2.2.1 Denial of Premises

A proposal for a particular action A can first be attacked by denying one of the four statements which must obtain for the proposal to be valid.

Attack 1: R is not the case.

Attack 2: It is not the case that *apply(A, R, S)* ∈ *apply*.

Attack 3: It is not the case that S $\models$ G.

Attack 4: It is not the case that *value(G)* $=< v, + >$.

Except for the first attack, each of these attacks may be executed with differing degrees of force, and so we are able to distinguish variants of the main attack. For instance, we can identify seven variant attacks for **Attack 2**.

Attack 2a: It is not the case that *apply(A, R, S)* ∈ *apply*.

Attack 2b: It is not the case that *apply(A, R, S)* ∈ *apply*, and it is the case that *apply(A, R, T)* ∈ *apply*, where T ≠ S.

Attack 2c: It is not the case that *apply(A, R, S)* ∈ *apply*, and it is the case that *apply(A, R, T)* ∈ *apply*, where T ≠ S, but it is not the case that T $\models$ G.

Attack 2d: It is not the case that *apply(A, R, S)* ∈ *apply*, and it is the case that *apply(A, R, T)* ∈ *apply*, where T ≠ S, and it is the case that T $\models$ G, but it is not the case that *value(G)* $=< v, + >$.

Attack 2e: It is not the case that *apply(A, R, S)* ∈ *apply*, and it is the case that *apply(A, R, T)* ∈ *apply*, where T ≠ S, and it is the case that T $\models$ G, but *value(G)* $=< v, - >$.

[1] We remark that formalisms of actions and their effects have received a great deal of attention in AI, e.g., the situation calculus [9]. In future work, we intend to explore the connections between these formalisms and our approach.

Attack 2f: It is not the case that *apply(A, R, S)* ∈ *apply*, and it is the case that *apply(A, R, T)* ∈ *apply*, where T ≠ S, and it is the case that T $\models$ G, but *value(G)* $=< w, + >$, where w ≠ v.

Attack 2g: It is not the case that *apply(A, R, S)* ∈ *apply*, and it is the case that *apply(A, R, T)* ∈ *apply*, where T ≠ S, and it is the case that T $\models$ G, but *value(G)* $=< w, - >$, where w ≠ v.

Similarly, we may distinguish six variants of **Attack 3**:

Attack 3a: It is not the case that S $\models$ G.

Attack 3b: It is not the case that S $\models$ G and there is a goal H ∈ *Goals*, H ≠ G, such that S $\models$ H.

Attack 3c: It is not the case that S $\models$ G and there is a goal H ∈ *Goals*, H ≠ G, such that S $\models$ H and with *value(H)* $\neq <$ $v, + >$.

Attack 3d: It is not the case that S $\models$ G and there is a goal H ∈ *Goals*, H ≠ G, such that S $\models$ H and with *value(H)* $=<$ $v, - >$.

Attack 3e: It is not the case that S $\models$ G and there is a goal H ∈ *Goals*, H ≠ G, and a value w ∈ *Values*, w ≠ v, such that S $\models$ H and with *value(H)* $=< w, + >$.

Attack 3f: It is not the case that S $\models$ G and there is a goal H ∈ *Goals*, H ≠ G, and a value w ∈ *Values*, w ≠ v, such that S $\models$ H and with *value(H)* $=< w, - >$.

Likewise, we may distinguish four variants of **Attack 4**:

Attack 4a: It is not the case that *value(G)* $=< v, + >$.

Attack 4b: It is not the case that *value(G)* $=< v, + >$ and *value(G)* $=< v, - >$.

Attack 4c: It is not the case that *value(G)* $=< v, + >$ and there is a value w ∈ *Values*, w ≠ v, such that *value(G)* $=< w, + >$.

Attack 4d: It is not the case that *value(G)* $=< v, + >$ and there is a value w ∈ *Values*, w ≠ v, such that *value(G)* $=< w, - >$.

One may also attack the argument schemas AS1 and AS2 by denying that the proposed action A is an element of the set of actions, *Acts*, or that goal G is an element of the set of goals, *Goals*. We discuss these two attacks below in subsection 2.2.5.

As mentioned, the variant attacks for **Attacks 2, 3** and **4** are not all of the same strength, since they make different assertions about the existence of alternative states, goals or values. Accordingly, these variants may entail dialogical commitments for the attacker. For example, depending on the rules governing the interaction, **Attack 2b** may require an attacker to defend the existence of a state T alternative to S, where **Attack 2a** may not require this. Thus, the burden-of-proof may transfer from one participant to the other depending on the specific attack used. This issue we do not discuss further here, but will take up in future work.

2.2.2 Alternative Action for Same Effect

If **Statements 1** to **4** do hold, we can say that the position represents a reason for doing action A, and that there is a *prima facie* case for performing A. That A should be performed, however, is a stronger statement and may not hold.

First there may be alternative actions which also realize the same value v. We can offer alternatives to each of **Statements 2, 3** and **4**

Attack 5: There exists an action B $\in$ *Acts*, with B $\neq$ A, and *apply(B,R,S)* $\in$ *apply*.

Attack 6: There exists an action B $\in$ *Acts*, with B $\neq$ A, and *apply(B,R,T)* $\in$ *apply*, with T $\models$ G.

Attack 7: There exists an action B $\in$ *Acts*, with B $\neq$ A, and *apply(B,R,T)* $\in$ *apply*, with T $\models$ H, and *value(H)* $=< v, + >$.

These three attacks allow us to promote the same value by performing another action; some further reason will need to be advanced for choosing A over B. How that choice may be made will be discussed in subsection 2.2.6.

2.2.3 Side Effects of the Action

If **Attacks 1** through **7** are not possible, then there is a *prima facie* reason to choose action A, and there is no alternative action to promote value v. It is still possible, however, that we may wish not to perform A because we have reasons not to perform it. For example, this action A, in addition to promoting v, may have undesirable side effects. In each of the next three attacks, **Statements 1, 2, 3** and **4** are assumed to be true.

One undesirable side-effect may be that action A, in addition to realizing goal G which promotes value v, also realizes another goal H which demotes v, thus removing the motivation for realizing G.

Attack 8: There is a goal H $\in$ *Goals*, with H $\neq$ G, such that *apply(A,R,S)* $\in$ *apply* with S $\models$ H, and with *value(H)* $=< v, - >$.

Another undesirable side effect of action A would be for some other desirable value, not v, to be demoted:

Attack 9: There is a goal H $\in$ *Goals*, with H $\neq$ G, and there is a value w $\in$ *values*, with w $\neq$ v, such that *apply(A,R,S)* $\in$ *apply* with S $\models$ H, and with *value(H)* $=< w, - >$.

Similar are two objections which are most usefully directed against the justification of a past action. Here we do not say that it was not right to perform A, but argue that the position was not the only reason to choose A. This may because S may entail goals other than G, which may also promote the same or some other value.

Attack 10a: There is a goal H $\in$ *Goals*, with H $\neq$ G, such that *apply(A,R,S)* $\in$ *apply* with S $\models$ H, and with *value(H)* $=< v, + >$.

Attack 10b: There is a goal H $\in$ *Goals*, with H $\neq$ G, and there is a value w $\in$ *values*, with w $\neq$ v, such that *apply(A,R,S)* $\in$ *apply* with S $\models$ H, and with *value(H)* $=< w, + >$.

If it is important that the right action is chosen for the right reason, we may prefer explanation of a past action given by **Attack 10a** or by **Attack 10b** to the one originally presented in schema AS2.

2.2.4 Interference with Other Actions

Next we may attack the position for the proposed action A by recognizing that actions cannot be considered in isolation, but that performing one action may prevent the performance of others. This incompatibility may be expressed conceptually at the level of the actions themselves, or at the level of the goals that can be realized in the resulting states, or at the level of values promoted or demoted.

Attack 11a: It is the case that *apply(A,R,S)* $\in$ *apply*. There is a value w $\in$ *values* with w $\neq$ v. There is an action B $\in$ *Acts* with B $\neq$ A, such that *apply(B,R,T)* $\in$ *apply*, with T $\models$ H, and *value(H)* $=< w, + >$. However, there is no state X $\in$ *States* such that *apply(A&B,R,X)* $\in$ *apply*.

Attack 11b: It is the case that *apply(A,R,S)* $\in$ *apply*. There is a value w $\in$ *values* with w $\neq$ v. There is a goal H $\in$ *Goals*, such that *value(H)* $=< w, + >$. However, S $\models \neg$H.

Attack 11c: It is the case that *apply(A,R,S)* $\in$ *apply*. There is a value w $\in$ *values* with w $\neq$ v. However, if there is a goal J $\in$ *Goals*, with *value(J)* $=< w, + >$, then S $\models \neg$J.

Each of these attacks is based on the desirability of promoting a value w which is different from the value v promoted by action A. The first, **Attack 11a**, says that there is a particular action B which leads to a state T in which a goal H is realized, and this goal promotes value w, but that actions A and B are incompatible. In other words, actions A and B may not both be performed. The second attack, **Attack 11b**, says that the outcome state S of action A is inconsistent with a particular goal H which promotes value w. Thus any alternative action C which led to a state which realized goal H must be inconsistent with action A. The third attack, **Attack 11c**, says that the outcome state S of action A is inconsistent with any goal which promotes value w.

These three attacks are related. If action A realizes a state which is incompatible with the state realized by action B, then A and B cannot both be performed. In the case of **Attack 11a**, however, there may be a third action, C, which can realize state H and which is compatible with A. In the case of **Attack 11b** this cannot be so. None the less there may be a third action, C, compatible with A, which promotes w through realizing some other goal. **Attack 11c** excludes this possibility also. Thus **Attack 11c** entails **Attack 11b** entails **Attack 11a**.

2.2.5 Sources of Rational Disagreement

Now let us suppose we have two participants, and that any disagreement between them has some rational basis. Let us consider what this basis might be. First, they may disagree on the various elements of the original sets. Thus the set *Acts* contains actions. Some of these may not be possible, or may be believed to be impossible by one of the participants. Actions considered impossible by both participants need be of no concern, since we can assume that neither party will propose them.[2]

We thus have another potential attack on a position: namely that action A is not possible. Note that this can apply only to projected actions, not to past actions. Similarly, a proposition believed meaningful by one participant, may be meaningless to the other, and so an attack may be made on the propositions implicit in any proposal. Because states are defined as conjunctions of propositions, such an attack may be manifested in an attack on the states R or S. Likewise, the goal G or value v may be attacked as not being a valid goal or value, respectively.

Attack 12: It is not the case that A $\in$ *Acts*.

Attack 13a: It is not the case that R $\in$ *States*.

Attack 13b: It is not the case that S $\in$ *States*.

Attack 14: It is not the case that G $\in$ *Goals*.

Attack 15: It is not the case that v $\in$ *Values*.

An opponent of the proposer of argument schema AS1 may also disagree on the way in which goals are valued by the function *value*.

[2]We ignore, in this account, consideration of counter-factual and *straw-man* proposals, as in the debates of the British War Cabinet in May 1940 [8].

This does not give rise to any new attacks because such disagreement will be catered for by one of the variants of **Attack 4**. Similarly, protagonists may may also differ on the function *apply*, for example, if they have different theories of causality. Such differences will motivate the use of one or more of the variants of **Attack 2** given earlier.

2.2.6 Selection between Options

Our final discussion in this section considers how a decision-maker presented with a proposal for action by a protagonist and attacks against this proposal from an opponent, may decide between alternative action-options. In order to give grounds for such a choice we need to represent the decision-maker's preferences between actions, and his or her preferences between values. For a given decision-maker (or audience), we therefore define two relations:

actionPref is a transitive, irreflexive and antisymmetric binary relation on *Acts* $\times$ *Acts*. For actions A and B, we read *actionPref(A,B)* as: *"The decision maker prefers action A to action B."*

valuePref is a transitive, irreflexive and antisymmetric binary relation on *Values* $\times$ *Values*. For values v and w, we read *valuePref(v,w)* as: *"The decision maker prefers value v to value w."*

A rational attacker issuing **Attacks 5, 6** and **7** would only do so if the attacker believes the audience prefers action B to action A. Similarly, a rational attacker would only issue **Attacks 9, 10b, 11b** and **11c** would only do so if the attacker believes the audience prefers value w to value v. In the case of **Attack 11a**, the attacker may believe the audience prefers action B to action A, or value w to value v, or both. Neither an action nor a value preference would motivate the rational use of **Attack 10a**; this attack would be used when there is some strategic gain expected from explaining a past decision by reference to goal H rather than to goal G.

3. APPLICATION TO LAW

Now, how does this general theory translate to law? In order to give some illustration of our general thrust we shall assume the availability of a CATO-like analysis [1]. Here we have a set of *base level factors*, which are used to describe cases. Each base level factor (hereafter, simply, factor) relates, positively or negatively to one or more *abstract factors*. There is also a set of cases, each of which "contains" a subset of the base level factors, and possibly a decision for the plaintiff/defendant. Typically the envisaged situation in which the system will be used is that there is exactly one undecided case, and the purpose of the legal dispute is to decide that case.

Mapping this information onto the generic formalisation, we get the following:

There are two actions, deciding the undecided case for the plaintiff and deciding the undecided case for the defendant. Thus *Acts* comprises {*decide(case1,p), decide(case1,d)*}, where *case1* is the undecided case.

The set of propositions for a given case comprises for each factor a proposition saying that the factor obtains in that case. These propositions will be either true or false depending on the facts of the cases. Additionally each case has two associated propositions, one that it was decided for the plaintiff and one that it was decided for the defendant. After the case is decided, (i.e.in S) one of these propositions will be true and one false: before the case is decided (i.e. in R) they will both be false.

Our function apply will affect the status of these last two propositions: decide(case,p) will make *decided-for-p* true and *decide(case,d)* will make *decided-for-d* true for that case.

Thus far, our requirements correspond well to what can be found in [1]. Table 2 lists 26 base level factors which together with *decided-for-p* and *decided-for-d* make up *Props*. We can also find many case descriptions in [1] to enable us to construct C and S for particular cases. What, however, of goals and values? Our suggestion is that each factor relates to some behaviour on the part of the plaintiff or the defendant that the law wishes to encourage or discourage. It achieves this by deciding for the party who acts in the ways it wishes to encourage and against the party whose behaviour it wishes to discourage. For example, *F1 Disclosure-in-Negotiations* is present if the plaintiff discloses his secret to the defendant in negotiations, and favours the defendant. Deciding for the defendant in the presence of this factor will tend to discourage holders of secrets from disclosing them in negotiations. Cases, however, typically contain combinations of facts and the purpose of our goals is to enable combinations to be encouraged or discouraged. The goal is thus a conjunction of a subset of factors present in a case and the propositions recording the decision representing the view of the combination we wish to take. It is intended to represent the factors seen as relevant to the decision in the particular case, and should represent a sufficient reason so to decide the case. Thus, on their own, *F4 Agreed-not-to-disclose* and *F23 Waiver-of-confidentiality* favour the plaintiff and defendant respectively. But having *F4 & F23 & decided-for-d* as a goal will encourage potential defendants to seek waivers of confidentiality if there is such an agreement, and not to use this information without such a waiver, and to make potential plaintiffs wary of offering such waivers.

Finally we need values. The motive for encouraging or discouraging behaviour is to promote some socially desirable end. For example, marking F1 as an important consideration would promote the social end that people act with reasonable care for their own interests: if one has a secret one has a certain responsibility to keep it to oneself. The function $value$ will therefore map each factor to elements of some set of values, representing the social ends which motivate the encouraging or discouraging of certain behaviours. In our case study we relate values to the abstract factors found in [1]: this follows a suggestion of [10]. As discussed in Section 4, one value is selected to motivate the decision; if promotion of multiple values is desired, this can be achieved through instantiation of multiple instances of the appropriate argumentation schema. Thus we may restate the argument scheme AS1 as:

(LAS1) In the presence of these factors we should decide for this party to establish that this subset of factors lead to this resolution which will encourage these behaviours which will promote these social ends.

For decided cases we offer a rationale of the decision:

(LAS2) In the presence of these factors the case was decided for this party to establish that this subset of factors should lead to this resolution so as to encourage these behaviours which promote these social ends.

Note that there is an element of interpretation here in moving from the full set of factors available in the case to a selected subset claimed as the goal of the decision. Naturally this interpretation can be disputed. We can now see legal argumentation with cases as the proposal and defence of an interpretation of some past case or cases using LAS2 to establish the desirability of a goal, followed by the contention that deciding the new case should be decided for our side to achieve that goal using LAS1. In the next section we will formalise the analysis found in [1], so that we can give focus to the later discussion.

Table 1: Abstract Factors in CATO mapped to Values

Abstract Factor	Social Value	Short name
F102 Efforts-to-Maintain-Secrecy	People are not negligent concerning their own interests	NN
F104 Info-Valuable	People litigate only when there is a substantial claim	LL
F105 Info-Known-or Available	Enterprise is rewarded	RE
F111 Questionable-means	Dishonesty is punished	DP
F112 Info-Used	People litigate only when there is a substantial claim	LL
F115 Notice-of Confidentiality	Agreements are respected and upheld	CA
F120 Info-Legitimately-Obtained-or-Obtainable	Enterprise is rewarded	RE
F121 Express-Confidentiality-Agreement	Agreements are respected and upheld	CA
F122/3 Efforts-to-Maintain-Secrecy	People are not negligent concerning their own interests	NN

4. FORMALISATION OF CATO

To adapt the analysis used in CATO, we need to identify the elements required by our formalisation. We proceed as follows.

The decisions are as standard for any legal problem. Thus:

Acts = {decide(case1,p), decide(case1,d)}.

The base level factors of CATO, together with *decision-for-p* and *decision-for-d* form the set of propositions which can be true or false for any given case. The factors can be found in Table 2.

Goals will be a subset of the propositions, interpreted as a conjunction. This subset must include at least one of *decision-for-p* and *decision-for-d*.

We identify our values by reference to the abstract factors in [1]. We first identify the value promoted by the satisfaction of the abstract factors found in [1]. Table 1 records this association. Where factors are associated with two or more abstract factors in [1], we use a set of values. In previous work such as [5] factors were always associated with a single value. For our present needs this is too restrictive: we want the association of a factor with a value to be determined in the context of a particular case. Thus part of the interpretation of a case will involve deciding which value is promoted by the factor in this particular case. This answer may vary from case to case. The values associated with factors are given in Table 2.

In Table 2 we show the values associated with each factor. Each factor favours a particular party, indicated by (p) or (d), and deciding for that party promotes or demotes one or more values as indicated in the Values column. Deciding for the other party would reverse the promotion or demotion of that value.

One point of note here. While we have followed the Factor Hierarchy of [1] closely, we have added a relation between F17 and F25 to RE, even though CATO does not relate these factors to F105 Info-Known-or-available. We believe that it is correct to do so: if that a product *could* be reverse engineered (F16) contributes to F105, it seems to us that it *actually was* reverse engineered must *a fortiori* contribute to the information being known or available.

Table 2: CATO Factors and Associated Values

Factor	Label	Values
F1	Disclosure-in-negotiations (d)	NN+ DP- RE+
F2	Bribe-employee (p)	DP+
F3	Employee-sole-developer (d)	RE+
F4	Agreed-not-to-disclose (p)	CA+ NN+
F5	Agreement-not-specific (d)	CA+
F6	Security-Measures (p)	NN+
F7	Brought-Tools (p)	LL+
F8	Competitive-Advantage(p)	LL+
F10	Secrets-Disclosed-Outsiders (d)	NN+ RE+
F11	Vertical-Knowledge (d)	RE+
F12	Outsider-Disclosures-Restricted (p)	NN+
F13	Noncompetition-Agreement (p)	CA+
F14	Restricted-Material-Used (p)	CA+ DP+
F15	Unique-Product (p)	LL+ RE-
F16	Info-Reverse-Engineerable (d)	RE+
F17	Info-Independently-Generated (d)	RE+ DP+ LL-
F18	Identical-Products(p)	LL+
F19	No-security-measures (d)	NN+
F20	Info-Known-To-Competitors (d)	RE+
F21	Knew-Information-Confidential (p)	CA+
F22	Invasive-Techniques (p)	DP+
F23	Waiver-of-Confidentiality (d)	CA+
F24	Info-Obtainable-Elsewhere (d)	RE+ DP+
F25	Info-Reverse-Engineered (d)	RE+ DP+
F26	Deception (p)	DP+
F27	Disclosure-in-Public-Forum (d)	NN+ RE+

5. APPLICATION TO DOMAIN

For each of the attacks identified in the generic theory, we will now consider how they appear in the domain. We begin by considering the attacks generally, and then apply them to a specific example, the case of *Mason v Jack Daniel Distillery*. In which a bar owner, Mason, invented a Jack Daniels based cocktail called Lychburg Lemonade. The distillery wanted to use this secret recipe in a promotion, and did so, although negotiations with Mason broke down. Mason sued for the unauthorised use of his recipe.[3]

Attack 1. R is not the case. This attack denies that the case should be described as the opponent describes it. Thus there is a disagreement as to the truth value of one or more members of *Props*. An example from the CATO domain would be a claim that *F19 No-Security-Measures*, applies. Suppose in the case under consideration, some security measures had been taken but these were minimal. Should we see this an example of F19, favouring the defendant, or as *F6 Security measures*, favouring the plaintiff, or as neither F19 or F6? This issue relates to the move from HYPO style dimensions to factors in CATO (see [11] for a discussion of this

[3]The episode of the *Simpson's* television cartoon in which Homer invents the Flaming Mo is closely based on this case.

transition). F6 and F19 are in fact the extremes of what was a dimension in HYPO. In practice arguing about which side is favoured when a case represents an intermediate point on a dimension is quite common in case based argument. See, for example, [4] for more discussion of the need to allow for this style of dispute.

Attack 2. It is not the case that apply(A, R, S) ∈ apply. This method of attack does not seem possible since S is merely the conjunction of R with a decision, and our actions are guaranteed to bring this decision about. We therefore need consider none of the variants of 2.

Attack 3. It is not the case that S ⊨ G.. Again this attack is impossible, unless a simple error of logic has been made. We will ignore this possibility here, and give no further consideration to variants of 3.

Attack 4. It is not the case that value(G) $=< v, + >$. This attack focuses on the link between factors and the values they promote. This link may not be unambiguous, since where a factor relates to several values, we need to determine which is promoted in a particular case. For example, in a case with factor F2 and F25 present, F25 might be held to relate to DP, whereas if F6 were present instead of F2 it might be related to RE, exemplifying attack 4c, or 4a if the alternative valuation is not explicitly stated. In fact any of the four variants of Attack 4 are possible here.

Attack 5. There exists an action B ∈ Acts, with B ≠ A, and apply(B,R,S) ∈ apply. This attack cannot arise in our context, since S includes the decision for the particular party. We have only two actions, and they will self-fulfillingly realise different states.

Attack 6. There exists an action B ∈ Acts, with B ≠ A, and apply(B,R,T) ∈ apply, with B ⊨ G. Again, since the party decided for forms part of G, it cannot be that both S and T realise the same goal.

Attack 7. There exists an action B ∈ Acts, with B ≠ A, and apply(B,R,T) ∈ apply, with B ⊨ H, and value(H) $=< v, + >$. This attack is quite possible, since we may have factors favouring different parties but promoting the same value. Thus, in a given case situation, we may promote a particular value whichever decision we make.

Attack 8. There is a goal H ∈ Goals, with H ≠ G, such that apply(A,R,S) ∈ apply with S ⊨ H, and with value(H) $=< v, - >$. Again examples of this form of attack can be found. Suppose we have a pro-plaintiff factor F4 and a pro-defendant factor F5, both of which relate to the same value CA. Now deciding for the plaintiff promotes CA via F4 & decided-for-p, but also demotes CA via F5 & not decided-for-d.

Attack 9. There is a goal H ∈ Goals, with H ≠ G, such that apply(A,R,S) ∈ apply with S ⊨ H, and with value(H) $=< w, - >$, with $w \neq v$. Here some pro-defendant factor, say F25, promotes a different value from the pro-plaintiff factors (say F4 and F6) identified in G. S realises F25 & not decided-for-d, which demotes the value RE.

Attack 10a. There is a goal H ∈ Goals, with H ≠ G, such that apply(A, R, S) ∈ apply with S ⊨ H, and with value(H) $=< v, + >$. G is some subset of S, which is claimed to be the goal that is realised. Therefore this attack questions the way in which the decision is interpreted. Thus my claim may be that deciding a case with factors F6, F21, F1 and F16 for the plaintiff realised the goal F6 & F21 & decided-for-p. It is, however, possible that it be interpreted rather as all of F1 & F6 & F21 & F16 & decided-for-p, with corresponding consequences for future decisions. This is at the heart of the interpretation of past decisions: the significance of the case depends on how G is chosen.

Attack 10b. There is a goal H ∈ Goals, with H ≠ G, such that apply(A,R,S) ∈ apply with S ⊨ H, and with value(H) $=< w, + >$, with $w \neq v$. This attack accepts that A is the correct action, but disagrees as to its justification in terms of values. If we have a choice of pro-winner factors relating to different values, then we can make this attack.

Attack 11a. It is the case that apply(A,R,S) ∈ apply. The value w ∈ values, w ≠ v, is desirable and should be promoted. There is an action B ∈ Acts with B ≠ A, such that apply(B,R,T) ∈ apply, with T ⊨ H, and value(H) $=< w, + >$. However, there is no state X ∈ States such that apply(A&B,R,X) ∈ apply. Our two actions always exclude one another. Moreover, in the situation described for attack 11 deciding for the plaintiff, so promoting NN, excludes deciding for the defendant, which would promote RE. Thus this attack is always possible when we have factors favouring both parties which are related to different values.

Attack 11b. It is the case that apply(A,R,S) ∈ apply. The value w ∈ values, w ≠ v, is desirable and should be promoted. There is a goal H ∈ Goals, such that value(H) $=< w, + >$. However, S ⊨ ¬H. This attack is also possible, since S will exclude any goal with a decision for the other party, and so if a factor favouring this party promoting a different value is present, the attack can be made.

Attack 11c. It is the case that apply(A,R,S) ∈ apply. The value w ∈ values, w ≠ v, is desirable and should be promoted. However, if there is a goal J ∈ Goals, with value(J) $=< w, + >$, then S ⊨ ¬J.. This attack may or may not be possible: it is possible only if there is a both a factor favouring the other party related to a different value present, and there is no factor favouring the original party related to that value.

Attack 12. It is not the case that A ∈ Acts. This attack is not possible since the two available actions are considered possible in every case.

Attack 13a. It is not the case that R ∈ States. This attack is possible if there is disagreement as to whether particular features of a case are factors or not. Thus, for example, it might be disputed whether *F18 Identical-Products* was in fact relevant to any case.

Attack 13b. It is not the case that S ∈ States. This attack is also possible. Indeed if attack13a is possible, so too will attack 13b be possible.

Attack 14. It is not the case that G ∈ Goals. If attack 13 is possible, and the disputed proposition is included in G, then this attack is also possible.

Attacks 13 and 14 are possible, but require the formalisation to be disputed. In what follows we will assume that the disputants are agreed on the elements of *Props*, in which case we can ignore both of these attacks.

Attack 15. It is not the case that v ∈ Values. This denies that the

value concerned is worth promoting at all. Such disagreement is possible, although we generally see the disagreement as located in the relative ordering of values rather than whether a particular end is valued or not.

So, of our original fifteen attacks, five (attacks 2, 3, 5, 6 and 12) have been found to be inapplicable, and two (attacks 13 and 14) have been discounted as we assume agreement on the choice of factors. This leaves eight forms of attack which we can look for. So let us turn to examine some possible arguments relating to a particular example.

Suppose we consider a particular case, such as that of *Mason v Jack Daniel Distillery*, discussed at [1, p. 25]. Mason is represented by CATO as having five factors:

F1, F6, F15, F16, F21.

A possible precedent for Mason is *M. Bryce and Associates Inc v Gladstone*, which has the five factors:

F1, F4, F6, F18, F21.

An argument for the plaintiff in Mason citing Bryce would be:

(A1): In Bryce, given F1, F4, F6, F18, F21, it was correct to decide for the plaintiff, since this realises F1 & F6 & F21 & decided-for-p, promoting the values of NN, CA and DP. In Mason, given F1, F6, F15, F16 and F21, we should decide for the plaintiff so as to realise F1 & F6 & F21 & decided-for-p, promoting the values of NN, CA and DP.[4]

This argument is very similar to the first ply put forward in [1]. How can the defendant attack the position? We now provide an example of each of the attacks.

Attack 1: CATO represents Mason as containing F6, security measures taken. Reading the facts of the case, however, suggest that these were not as extensive as they might have been - certainly less extensive than those taken by Bryce, and so it might be argued that F6 does not apply to Mason.

Attack 4: The connection between the goal and the values promoted could be challenged. For F1, *Disclosure in negotiations*, we have a choice of three values that might be concerned. A1 selects DP. If we deny this, without proposing an alternative we have attack 4a. 4b cannot be made here, because no factor can be used to give a goal demoting DP. Suppose, however, we believe that in the context of Mason and/or Bryce, we should choose NN rather than DP to be the value promoted by F1. Now Mason's disclosure can be held to cancel out his security measures, leaving only CA, as promoted by F21 as the value promoted by G. This is an instance of attack 4d. No example of attack 4c appears here, as the alternative values related to F1 are both pro-defendant.

Attack 7: This attack can only be directed towards the current case, since we cannot alter the decision in Bryce. In order to make this attack we need to find some factor(s) which would promote the same values as were promoted in the original interpretation when conjoined with decided-for-d. This means that for each pro-plaintiff factor in the original G, we must find a pro-defendant factor related to the same value. This is not possible in A1, but suppose we had used the weaker argument A2, based on *Televation Telecommunications Inc v Saidon:*

(A2): In Televation, given F6,F10,F12,F16,F21, it was correct to decide for the plaintiff, since this realises F6 & decided-for-p, promoting the value of NN. In Mason, given F1, F6, F15, F16 and F21, we should decide for the plaintiff so as to realise F6 & decided-for-p, promoting the value of NN.

Now, by taking F1 to promote NN, we could argue that in Mason the goal F1 & decided-for-d would equally promote NN, and so the goal derived from Televation cannot be a reason to decide for the plaintiff in that case. Note that F1 is a distinction between Mason and Televation. It is, however, a different kind of distinction from that we will find in attack 10, since it makes the current case weaker for the plaintiff, rather than the precedent case stronger for the plaintiff.

Attack 8: In this attack we find a goal promoted by a decision for the plaintiff which demotes the value promoted in the precedent case. This can again be illustrated by using A2. Now, however, we say that F1 & decided-for-p would demote NN. Is there any real difference between these this attack and attack 7? The same distinction is relied on in both cases. There is a difference of emphasis since attack 7 is directed to providing a positive argument for deciding the other way, whereas attack 8 provides a negative argument against deciding in a particular way. These are very close in our current situation where we have only two actions to choose between so that not doing A and doing B are equivalent. Given a richer action set, the distinction may be important: moreover, we may find that these attacks have different effects when considered in the context of a complete dialogue.

Attack 9: To make this attack we must find a factor that would demote a particular value. Mason, unlike Bryce contains F16, and so we could claim that deciding for Mason would realise *F16 & decided-for-p*, demoting RE. Additionally we must claim that RE is more important than the values promoted by the goal claimed.

Attack 10: This attack could be directed at the interpretation of Bryce. Although the goal claimed is a possible interpretation of Bryce, it might equally be argued that goal realized by that decision was F1 & F4 & F6 & F21. This would promote the same values as the original claim. This goal cannot be realised in Mason, which does not contain F4. This is an example of attack 10a. For an example of attack 10b, we could argue that *Televation* should be explained by reference to the goal *F12 & decided-for-p*, promoting the value CA, and so explaining *Televation* in a way not helpful to Mason.

Attack 11: Just attacks 7 and 8 can be paired, this is extremely similar to attack 9 when we have only two actions. The claim here is that deciding for Mason would prevent the realisation of *F16 & decided-for-d*, which would promote RE. In fact this is the strongest form, attack 11c, since RE cannot be promoted by any goal containing decided-for-p on these facts.

Attack 15: This is difficult to argue here, since we have three values promoted, and all of them are valid according to our formalisation. Indeed such an attack represents a criticism of our formalisation more fundamental than attack 1, since we are questioning the *existence* of a value, rather than the *application* of a factor.

To summarize the above, we have found examples of each of our original eight attacks which we saw as applicable, and seen that two pairs, $\langle 7, 8\rangle$ and $\langle 9, 11\rangle$, reduce to essentially the same thing, leaving us with seven distinct attacks (counting 10a and 10b as distinct). Two of these can only be made by making a challenge to the representation; attack 1 challenges the factors used to represent a case, and attack 15 the values associated with the factors. The other five are possible attacks while not questioning the formalisation.

Thus we have five attacks that we might expect to be reflected in CATO: $4, \langle 7, 8\rangle, \langle 9, 11\rangle, 10a, 10b$. Four of these correspond to the

[4] Note that promotion of multiple values can be undertaken through multiple instantiations of the appropriate argumentation schema, as mentioned in Section 2.1.

distinguish case move in CATO, but each in a different way. A case may be distinguished if the precedent is stronger or the current case is weaker. Also distinctions may be capable of being downplayed or emphasized, giving four possible flavours of distinction. Each of these is represented by a different attack:

Attack $10a$: precedent stronger for p; can be downplayed.
Attack $\langle 7, 8\rangle$: current case weaker for p, can be downplayed.
Attack $10b$: precedent case stronger for p, can be emphasized.
Attack $\langle 9, 11\rangle$: current case weaker for p, can be emphasized.

This leaves attack 4 as new: what is happening here is that the value said to be promoted by the factor in question is different, a possibility raised by the ability of a factor to relate to several values, and the need to say which value is being promoted in the context of the particular decision.

In CATO eight argument moves are identified:

i. Analogising a case to a past case with a favourable outcome.
ii. Distinguishing a case with an unfavourable outcome.
iii. Downplaying the significance of a decision.
iv. Emphasising the significance of a distinction.
v. Citing a favourable case to emphasise strengths.
vi. Citing a favourable case to argue that weaknesses are not fatal.
vii. Citing a more on point counterexample to a case cited by an opponent.
viii. Citing an as on point counter example to a case cited by an opponent.

Of these our attacks correspond to only the first four. How then do we see the other four argument moves relating to our scheme? Both (v) and (vi) would not be expected to appear as attacks, since they are responses to attacks. We will suggest how such moves would appear in our framework. Before discussing this, however, we should consider the counter examples.

Counter examples should not be expected to appear as attacks on positions such as A1 and A2, because they involve citing a case with a different outcome, rather than questioning a particular position. Thus the counterexample response itself involves putting forward a position, just like the original claim. This will therefore comprise two parts: an interpretation of a past case and an application of that interpretation to the new case.

Two kinds of counterexample may be used.

Counter example to the interpretation. In the interpretation the goal comprises some set of factors conjoined with decided-for-p. The first kind of counterexample (which we call a *negative counterexample*) is a case which contains those same factors but which was decided for the defendant. Thus given the position in A2, we can say that the claimed goal is not vindicated by the body of case law and cite any case in which the plaintiff took security measures and yet there was a finding for the defendant, such as *Motorola Inc v Fairchild Camera and Instrument Corp*, which contains F2, F4, F5, F6 and was decided for the defendant.

The second kind of counterexample (*positive counterexample*) is where the defendant puts forward a whole position, containing both an interpretation of a past case and an application of that interpretation to the current case. For example in Mason, the defendant could use *Sandlin v Johnston* as a counterexample, by giving the position expressed in A3:

A3: In Sandlin, given F1, F10, F16, F19, F27, it was correct to decide for the defendant, since this realises F1 & F16 & decided-for-d, promoting the values of NN and RE. In Mason, given F1, F6, F15, F16 and F21, we should decide for the defendant so as to realise F1 & F16 & decided-for-d, promoting the values of NN and RE.

The most effective counterexample is capable of acting both as a negative and a positive counterexample, since this both defeats the opponent's interpretation and provides a case which can be interpreted as a ground for deciding the new case. Otherwise, counterexamples are most effective when combined with an another form of distinguishing attack. Thus for example a negative counterexample could be coupled with a proposed different interpretation of the goal achieved by the precedent. Thus a response to A2 might combine the Motorola counter example with Attack 10, giving:

The goal in Televation cannot be F6 & decided-for-p, since in Mortorola F6 & decided-for-d. In Televation, given F6, F10, F12, F16 and F21, it was correct to decide for the plaintiff, since this realises F12 & decided-for-p, promoting the values of CA.

This can then be elaborated with a *positive counterexample*, proposing a reason to decide Mason for the defendant, perhaps Sandlin as in A3 above.

The *positive counterexample* is important, since the other attacks only undermine the reasons proposed to decide for the plaintiff: motivating a decision for the defendant really requires a positive position such as A3.

There is no real reply to a negative counterexample. A positive counterexample, however, such as A3, can be subjected to all the kinds of attacks that the original position was subject to, through distinctions or a negative counterexample.

Attacks on any positive counterexamples form the first component of the plaintiffs rebuttal of the defendant's response. The rebuttal can then be extended by using moves corresponding to CATO's moves (v) and (vi), which involve citing cases to emphasize strengths or show that any weaknesses are not fatal.

Citing a favourable case simply involves putting forward additional cases which realise the goal. Thus any other cases with F1 & F6 & F21 & decided-for-p could be cited to rebut a response to A1.

Showing weaknesses not fatal is essentially providing a negative counterexample to the new goal proposed in one of the distinguishing attacks. One point should be noted about attack 10a, however: since the goal there points to a factor present in the precedent but absent from the current case, so, in the example, the relevant counter example must supply the goal F1 & F6 & F21 & notF4 & decided-for-p. Interesting this raises the need to include the absence of factors in goals. This in turn suggests the possibility of doing this when choosing the precedent to cite originally, so that we need not match only on factors present, but on factors absent as well. So for example the goal in A1 could have been F1 & F6 & F21 & notF20 & decided-for-p, if, as is possible, we considered the fact that the information was unknown to competitors to be important. This would avoid the need to include two factors which are the negations of one another, as for example F6 and F19, but would need to be handled with care, since, for example, notF5 makes little sense if there is no agreement at all.

These responses do not really address attack 4, which proposes a different interpretation of the value promoted by a factor. Currently we see no obvious way to respond to this kind of attack, other than by offering cases which exhibit the goal of the original position but not the goal introduced in the attack. Of course, since attack 4 is not present in CATO, we should expect no help from the moves identified in that system.

The above thus covers the argument moves identified in CATO. The use of values, however, suggest some additional moves in that we can cite a case which matches at the value level rather than the

factor level.

For example in Mason, suppose we had a case (call it Hypo) with the factors F1, F12, F16 and F18. Here we could advance A4:

A4: In Hypo, given F1, F12, F16 and F18 it was correct to decide for the plaintiff, since this realises F1 & F12 & F18 & decided-for-p, promoting the values of NN,LL and DP. In Mason, given F1, F6, F15, F16 and F21, we should decide for the plaintiff so as to realise F1 & F6 & F21 & decided-for-p, also promoting the values of NN,LL and DP.

Here we use the commonality of values to support the goal used in the current case: that none of the pro-plaintiff factors match no longer excludes Hypo as a valid precedent.

This kind of cite is possible as an original precedent, but is probably better used to find cases emphasising strengths: the idea here is that the original goal should be preferred to the counter proposal in the attack because the values that it supports can be shown to be generally held in esteem.

6. DISCUSSION

6.1 On Pointness

The notion of on-pointness is central to CATO (and its predecessor HYPO [2]), On pointness is a measure of the similarity and differences between cases: the greater the number of factors in common, and the fewer the number of different factors, the more on point is the case. The precedent chosen to cite is always the most on point to the current case, and counterexamples used in a response are meant to be at least as on point.

On pointness is defined [1, p. 214] as follows:

A case A is deemed more on-point to a problem (i.e. an undecided case), if the similarities between B and the problem are a proper subset of the similarities between A and the problem.

This means that our first focus is the intersection between the factors in the precedent and the current case. But the size of the intersection is not crucial: given a case situation with F1,F2,F3, and precedents with F1,F4 and F2,F3,F5, neither precedent is more on point than the other. Thus CATO, like HYPO, before it, develops a *Claims Lattice* in which each generation is as on point as its siblings and more on point than its children. What HYPO and CATO do is to choose one of the most on-point cases as the citation, and to use all the factors in their intersection when citing the case.

In our terms this becomes a *strategy* for choosing a case to cite, and selecting the goal once the case is chosen. This strategy has several things to recommend it, but is not forced in our account. Thus, in A2 above, our goal was chosen to contain only a subset of the intersection. What considerations might lead us to include or exclude available factors when determining the goal? The argument for inclusion is that it reduces the number of negative counterexamples that can be found. The more tightly specified the goal the fewer the available precedents that can be brought to bear on it. The corollary of this is that there will also be fewer cases available to cite in support when we come to the rebuttal. Similar considerations apply to differences between the current and cited case, in that the more distinctions that there are, the more attacks will be available.

The motive for allowing the goal to be less tightly specified than may be possible is that it allows us to focus consideration of relevant factors, without introducing distractions and potential red herrings. Thus if a similarity is not germane to our explanatory theory of the domain, we may decide not to use it. Our view therefore is that the notion of on-pointness supplies a plausible heuristic for selecting cases to cite and fixing on the goal, but need not be the only such heuristic. Our formalism gives the flexibility to explore this question further.

6.2 Evaluation of Attacks

Thus far we have identified how a position can be attacked. Attacks identify alternative ways of interpreting past cases and alternative decisions for current ones. But the existence of an alternative does not mean that this alternative should be adopted. Note that we do not want to say that an attack always succeeds, but rather have the audience relative framework of [3] in mind.

The role of the judge as we conceive it is to provide the audience which the arguments must persuade. There are a number of possibilities for his or her assessment. One is based on the ordering of values. Some of the attacks contain within them the notion that some value is more important than some other value. If this contention is rejected, the attack will fail. Note, however, that *actionPref* is not used in legal reasoning: the judge is neutral between the parties. The second possibility derives from the fact that some attacks require the attribution of a different value to a goal. Here the disagreement is as to which value is promoted by a given factor in the context of the particular cases. Again the judge must arbitrate on this matter. Thirdly, the move from case to goal may be in question, even though this makes no difference either to the order of values or the values promoted. Here the appeal is to which theory of the domain is considered more plausible, possibly determined by some notion of theory coherence. Fourthly, we may be questioning the representation of a case in terms of factors. Does a particularly lax set of security measures count as F6 or F19? Again the judge may be called upon to decide this issue.

Whether an attack succeeds or not thus depends on a judgement of the target audience. It would also be desirable if the party putting forward the attack also had a similar mechanism for adjudicating the above points, so that it would put forward the attacks it considered most likely to succeed. Of course, the two parties might well differ on their judgments, (as to the order of values, the value promoted by a factor, the goal governing a case, or where a line is drawn on a dimension) which is how we get the disagreement that leads to litigation and why we need to the decisive audience to be neutral between them.

6.3 Law as a Special Case of Persuasion

We took as our starting point a general theory of persuasion. In applying this to law we found that some of the scope that the framework provides was not needed. In particular this derives from two characteristics of the way in which we have represented the legal situation:

- There is always a choice between two actions, one of which must be adopted. The effect of this is to make two pairs of attack in the general framework indistinguishable from one another.
- The effect of an action is always known, and actions are guaranteed to succeed. This rules out two of the attacks identified in the general framework, which represent differing views on the effect of an action.

Law can also be special in the way arguments are constructed. The general framework does not prescribe where positions are supposed to come from, nor how they are to be presented, nor how goals are justified. In reasoning with precedents in law, however, the initial argument (and positive counterexamples) always involves putting forward some interpretation of a precedent case in order to provide some grounds for the goal, and then applying this interpretation to the new case.

We find it instructive to see this form of legal reasoning not as *sui generis*, but as a particular example of persuasive reasoning in general, adapted to meet the particular circumstances in which it is conducted. That these adaptations have perhaps become enshrined in conventions apparently peculiar to legal reasoning does not mean that anything radically different is happening from happens in ordinary persuasion.

6.4 Association between values and factors

Of particular interest is the way in which we have related factors and values. We have associated each factor with some subset of the available values, largely following the factor/abstract factor relationship in CATO. We have, however, required the choice of one of these values to be promoted by a factor in a particular goal. This allows us to see the valuation of a goal as a point of contention. There are, however, other ways in which this can be handled: moving one way we could allow a free choice of values to be associated with factors in goals, and moving the other we could enforce the association between a factor and a single value in our representation. Again the implications of different choices here are something that could be explored in future work.

7. CONCLUSIONS

In this final section we shall attempt to summarise the main points of this paper, and to suggest some directions for future investigation. Our intention has been to show how reasoning with legal cases can be seen in the context of a general theory of persuasive argument. Note that we have chosen to see law as an example of persuasion in *practical* reasoning, so that the decision for one of the parties to the case is an *action*, not the derivation of some *fact* about the case. We feel that this is important in order to capture the normative nature of the reasoning, that a decision for, say, the plaintiff is something that *should* be done and is capable of justification, rather than a property of the case capable of discovery.

We believe that we have successfully shown that legal reasoning with cases fits comfortably within our framework. Legal reasoning does not exemplify all the features that can be found in the general framework, but this is explicable in terms of some special features of the legal domain. In contrast we have not found features in the legal domain for which we cannot account in our general framework. This point has been illustrated by a detailed application of the framework to the domain of CATO, the system which currently employs the widest selection of argument moves.

The case study also suggests something about the style of legal reasoning, and how precedent cases are deployed in argument, and how they are related to a case under current consideration.

By taking the general framework as a starting point we are able to see what might potentially be done in a legal reasoning system. That is, we can see where restrictions to the general framework are part of the domain itself, and where they result from particular choices in the implementation of a given system. This would provide a basis for the comparison of systems, and the choices they embody. Thus for example we can see that attacks 5 and 6 are inapplicable to any legal system; attack 1 is available in HYPO but not in CATO, due to the move from dimensions to factors for representing cases (see [4, 11] for a discussion of this); and that while HYPO cannot discriminate between the various forms of distinguishing represented by attacks 7,8,10a, 10b,11 and 12, the use of abstract factor in CATO divides them into two groups, 7, 9 and 10a and 10b,11 and 12. Potentially it would be further possible to discriminate into the four groups 7/8, 10a, 10b and 9/11, but this is the limit: features of the domain render 7/8 and 9/11 effectively the same.

For future work we intend to continue to explore persuasion in general and to use law as one of our specific domains to test and evaluate what we do. Our intention is to specify a dialogue game which will include as moves the attacks discussed in this paper, and responses to these attacks. These will need to be supplemented by moves putting forward positions, counterexamples and supporting cases; one question is whether these position presenting moves are domain specific, or can be couched in general terms. This specification will be used as the basis for an implemented dialogue system: initially this will presuppose human participants, but our intention is to discover strategies and heuristics for autonomous computational participants to interact automatically. This is an ambitious programme, but the intermediate steps are in themselves useful as offering insight both into the topic of legal reasoning itself, and to approaches which have addressed it in the past.[5]

8. REFERENCES

[1] V. Aleven. *Teaching Case Based Argumentation Through an Example and Models*. Phd thesis, University of Pittsburgh, Pittsburgh, PA, USA, 1997.

[2] K. D. Ashley. *Modeling Legal Argument*. MIT Press, Cambridge, MA, USA, 1990.

[3] T. J. M. Bench-Capon. Agreeing to differ: modeling persuasive dialogue between parties without consensus about values. *Informal Logic*, 2003. *In press*.

[4] T. J. M. Bench-Capon and E. L. Rissland. Back to the future: dimensions revisited. In B. Verheij, A. Lodder, R. Loui, and A. Muntjewerff, editors, *Proceedings of JURIX 2001*, pages 41–52, Amsterdman, The Netherlands, 2001. IOS Press.

[5] T. J. M. Bench-Capon and G. Sartor. Theory based explanation of case law domains. In *Proc. Eighth ICAIL (ICAIL-2001)*, pages 12–21, New York, 2001. ACM Press.

[6] K. Greenwood, T. Bench-Capon, and P. McBurney. Argument over Proposals for Action. Technical Report ULCS-03-003, Department of Computer Science, University of Liverpool, Liverpool, UK, 2003.

[7] K. Greenwood, T. Bench-Capon, and P. McBurney. Strucuring dialogue between the people and their representatives. In *From e-Government to e-Governance (EGOV03)*, Prague, Czech Republic, 2003.

[8] J. Lukacs. *Five Days in London: May 1940*. Yale University Press, New Haven, CT, USA, 1999.

[9] J. McCarthy and P. J. Hayes. Some philosophical problems from the standpoint of artificial intelligence. In B. Melzer and D. Michie, editors, *Machine Intelligence 4*, pages 463–502. Edinburgh University Press, 1969.

[10] H. Prakken. An exercise in formalising teleological reasoning. In J. Breuker, R. Leenes, and R. Winkels, editors, *Proceedings of JURIX 2000*, pages 49–58, Amsterdam, The Netherlands, 2000. IOS Press.

[11] E. L. Rissland and K. D. Ashley. A note on dimensions and factors. *AI and Law*, 10(1–3):65–77, 2002.

[12] D. N. Walton and E. C. W. Krabbe. *Commitment in Dialogue: Basic Concepts of Interpersonal Reasoning*. SUNY Press, Albany, NY, USA, 1995.

[5] Miss Greenwood is grateful for support from the EPSRC.