A lightweight formal model of two-phase democratic deliberation

Trevor BENCH-CAPON [a], Henry PRAKKEN [b]

[a] *Department of Computer Science, University of Liverpool, UK*

[b] *Department of Information and Computing Sciences, Utrecht University and Faculty of Law, University of Groningen, The Netherlands*

Abstract. A formal two-phase model of democratic policy deliberation is presented, in which in the first phase sufficient and necessary criteria for proposals to be accepted are determined (the 'admissible' criteria') and in the second phase proposals are made and evaluated in light of the admissible criteria resulting from the first phase. Argument schemes for both phases are defined and formalised in a logical framework for structured argumentation. The *process* of deliberation is abstracted from and it is assumed that both deliberation phases result in a set of arguments and attack and defeat relations between them. Then preferred semantics is used to evaluate the acceptability status of criteria and proposals.

Keywords. E-democracy, deliberation, proposals, criteria

1. Introduction

Discussions on policy proposals often fall into two separate phases: first criteria that proposals should satisfy are determined and then specific proposals are put forward and evaluated against the criteria previously established. Possible benefits of such a separation are that thus the discussion has a clear structure and that no time and resources are wasted on evaluating proposals with unacceptable criteria. For contexts where these benefits are desirable, we present a formal two-phase model of democratic policy deliberation. In less organised contexts, where criteria and proposals are advanced in a less systematic fashion, our model can provide a useful tool to analyse and evaluate the discussion. In the first phase sufficient and necessary criteria for proposals to be accepted are determined (the 'admissible' criteria) and in the second phase proposals are made and evaluated in light of the admissible criteria resulting from the first phase. We abstract from the *process* of deliberation and so do not consider how the arguments are put forward, but simply assume that both deliberation phases result in a set of arguments and attack and defeat relations between them. We then apply preferred semantics [5] to evaluate the acceptability status of criteria and proposals. Preferred semantics is used since this arguably suits the inherently credulous nature of argumentation over action, where if action proposals are conflicting and the conflict cannot be resolved through logic alone, in the end a choice has to be made. The second phase is constrained by the first phase in that evaluating proposals in light of criteria can only make use of criteria that were accepted as admissible in the first phase.

The core of our model is two sets of argument schemes. The schemes for the first phase allow citizens to argue for and against necessary and sufficient criteria and to express preferences between criteria. The schemes for the second phase allow arguments for and against proposals to be made and evaluated in terms of how well they satisfy the admissible criteria resulting from the first phase. The language of the argument schemes is a light-weight formal one, to fit with the intended e-democracy applications.

The dialogue setting we assume is subjective but social. Central to our argument schemes is that criteria can be justified by saying that they satisfy something desirable. It is necessary, however, to distinguish what is desirable for particular individuals from what is desirable for the group as a whole. The distinction goes back to Rousseau [11]:

> There is often a great deal of difference between the will of all and the general will. The latter looks only to the common interest; the former considers private interest and is only a sum of private wills.

Ideally in a democracy people will indeed vote for what they consider to be the common good rather than from selfish motives. Of course, in practice this may not always be the case, as Rousseau [11] recognised:

> Finally, when the State, on the eve of ruin, maintains only a vain, illusory and formal existence, when in every heart the social bond is broken, and the meanest interest brazenly lays hold of the sacred name of "public good," the general will becomes mute: all men, guided by secret motives, no more give their views as citizens than if the State had never been; and iniquitous decrees directed solely to private interest get passed under the name of laws.

Often the truth lies between the two, with people torn between endorsing what is good for society as a whole and their own self interest. Often too there are differences between what people say in public and what they vote for in a secret ballot: income tax rises to fund services for the poorest receive far more support in opinion polls that at the ballot box. None the less we will in this paper assume that the users of our model are well intentioned and will deliberate about what is good for society as a whole and not about what is good for themselves, leaving individual goods for future work. We are aware that this is an idealised position, but we believe it is good to start with an idealised model and then see how it can be enriched. In our argument schemes therefore *it is desirable* is to be understood as *it is desirable for society as a whole*, and cannot be justified simply by some person or group actually desiring it.

This paper is organised as follows. In section 2 we briefly describe the formal setting: argumentation frameworks, an extension to allow the expression of preferences, and a means of providing structure for arguments. Section 3 describes the phase in which criteria are proposed and established. Section 4 describes the second phase in which proposals are evaluated against these criteria. Section 5 applies the model to an example and section 6 offers some discussion and conclusions.

2. The formal setting

We first briefly summarise the three formal frameworks we will use in this paper.

Dung's *abstract argument frameworks* [5] are a pair $AF = \langle \mathcal{A}, \mathit{defeat} \rangle$, where $\mathcal{A}$ is a set of arguments and *defeat* a binary relation on $\mathcal{A}$. A subset $\mathcal{B}$ of $\mathcal{A}$ is called *conflict-free* if no argument in $\mathcal{B}$ defeats an argument in $\mathcal{B}$ and it is called *admissible* if it is conflict-free and defends itself against any attack, i.e., if argument A_1 is in $\mathcal{B}$ and argument A_2 defeats A_1, then some argument in $\mathcal{B}$ defeats A_2. A *preferred extension* is then a maximally (wrt set inclusion) admissible set. Dung defines several other types of extensions but for our model they are not needed.

Modgil's *extended argumentation frameworks* refine those of Dung in two ways. First, instead of a defeat relation Modgil assumes a more basic *attack* relation, standing just for notions of syntactic conflict. Then Modgil allows *attacks on attacks* in addition to attacks on arguments. Intuitively, if argument C claims that argument B is preferred to argument A, and A attacks B, then C undermines the success of A's attack on B (i.e., A does not *defeat* B) by pref-attacking A's attack on B. Thus an *extended argumentation framework* is a triple $EAF = \langle \mathcal{A}, attack, \textit{pref-attack}, \rangle$, where $attack \subseteq \mathcal{A} \times \mathcal{A}$ and *pref-attack* $\subseteq \mathcal{A} \times attack$. Then defeat is made relative to a set of arguments: for any subset S of $\mathcal{A}$ and arguments A and B: B S-*defeats* A iff B attacks A and there is no argument C in S that pref-attacks this attack. Dung's theory of AFs is then reformulated with *defeat* replaced by S-*defeat*. Since arguments attacking attacks can themselves be attacked, as can these attacks, and so on, Modgil's extended argumentation frameworks can fully model argumentation about whether an argument defeats another.

Another refinement of Dung's abstract approach is the ASPIC framework for structured argumentation [10]. This framework assumes an unspecified logical language and defines arguments as inference trees formed by applying inference rules (which may be either strict or defeasible) to a knowledge base: the nature of the inference rules is also unspecified. The notion of an argument as an inference tree leads to three ways of attacking an argument: attacking an inference (*undercut*), attacking a conclusion (*rebuttal*) and attacking a premise (*undermining*), where rebutting and undercutting attacks can only be targeted at applications of *defeasible* inference rules. To resolve underminings and rebuttals, a preference relation $\prec$ on arguments (to be specified as input) is used, which leads to three corresponding kinds of defeat: undercutting, rebutting and undermining defeat. Basically, A successfully rebuts (undermines) B if A rebuts (undermines) B and $A \not\prec B$. Then A defeats B if A undercuts B or successfully rebuts or undermines B.

Finally, in [8] EAFS are instantiated with the ASPIC framework, where ASPIC's input ordering $\prec$ on arguments is replaced by pref-attacks on attacks. This results in a logical framework for structured argumentation with attacks on attacks, so that now arguments can be built that make explicit why an attack is attacked. Thus while an argument attacking an attack may simply be the expression of preference, it could also be a justification of that preference. Importantly admitting such arguments to the framework allows for them to be attacked by arguments expressing or justifying the contrary preference.

3. The criteria deliberation phase

In the criteria deliberation phase the participants propose and attack criteria for proposals. In this section we give a set of argument schemes (defeasible inference rules in the ASPIC framework) that can be used for these purposes. The distinction between sufficient and necessary criteria (relative to something that is desirable) is important, as these will be used differently in the second phase. The outcome of this phase is one or more sets of admissible criteria; these sets wil be used in the second phase to make and assess proposals.

3.1. Argument schemes for determining criteria

The first two argument schemes express variants of the positive and negative practical syllogism, to argue for sufficient and necessary criteria respectively.

> SCS: C should be a sufficient criterion since C satisfies D and D is desirable
> NCS: C should be a necessary criterion since D requires C and D is desirable

Note that by equivalence of 'D requires C' and 'Not-C results in Not-D' the NCS scheme can also be used to deal with negative side-effects of an action.

The remaining schemes generate rebutting attacks on uses of these two schemes (here the compatibility relation is assumed to be symmetric).

> $CN2N$: C should not be a necessary criterion since C' should be a necessary criterion and C' is not compatible with C.
> $CN2S$: C should not be a sufficient criterion since C' should be a necessary criterion and C' is not compatible with C.
> $CS2N$: C should not be a necessary criterion since C' should be a sufficient criterion and C' is not compatible with C.

Incompatible sufficient criteria do not conflict, as they are non-exclusive alternatives rather than rivals and so need not be true together. Note that the second premise of these conflict schemes allows a debate to be about whether two criteria are compatible. This allows, for instance, arguments utilising contextual reasons why criteria cannot be jointly satisfied, such as 'we cannot have a prime minister who is both a woman and conservative, since no current female conservative politician has sufficient status in her party'.

Applications of all schemes can be attacked on their premises. For example, a dispute as to what is desirable would result in arguments attacking and defending the premise that D is desirable. Our model allows for any set of argument schemes for attacking or supporting premises of the above schemes but in this paper we will not go into this. Furthermore, applications of SCS or NCS can be rebutted by an application of one of the three conflict schemes that contradicts its conclusion. For such attacks two questions arise: are these attacks symmetric and can preferences be used to determine whether an attack is successful, i.e., whether it results in defeat? Note that these questions are independent of each other, since in ASPIC, and other argumentation systems using preferences, an asymmetric attack of A on B is unsuccessful if B is preferred over A.

At an intuitive level it seems that all three kinds of attacks are symmetric and that they must be resolved with preferences based on the importance of the desires that generate the criteria involved in the conflict. So, for example, if we have (with the obvious abbreviations of the first premises)

$A_1 = NC(C_1)$ since D_1 requires C_1 and D_1 is desirable
$B_1 = NC(C_2)$ since D_2 requires C_2 and D_2 is desirable

Then if C_1 is not compatible with C_2, the following attacks can be constructed:

$A_2 =$ Not $NC(C_2)$ since $NC(C_1)$ and C_1 is not compatible with C_2
$B_2 =$ Not $NC(C_1)$ since $NC(C_2)$ and C_2 is not compatible with C_1

Note that the conclusion of A_1 provides a premises of A_2 and likewise for B_1 and B_2. Now we have that A_1 and B_2 are in conflict with each other and so are A_2 and B_1. Intuitively this is just one conflict, which should be resolved by comparing the importance of desires D_1 and D_2.

In the ASPIC logic [10] this can be achieved by formalising all schemes as defeasible inference rules. (SCS and NCS must be defeasible since in the present account the notion of a sufficient or necessary condition is a *ceteris-paribus* one, assuming that all other things are equal, namely, that there are no other desires that give rise to other conditions.) In our example, this yields the following attacks:

A_2 rebuts B_2 on B_1
B_2 rebuts A_2 on A_1
A_1 rebuts B_2 on B_2
B_1 rebuts A_2 on A_2

If ASPIC's argument ordering is defined in terms of a preference ordering on the desires used in applications of NCS or SCS then all these conflicts can be resolved with a single comparison between desires D_1 and D_2. For example, if $D_1 > D_2$ we have that $A_2 > B_1$ and $A_1 > B_2$ so (if there are no other relevant attacks) A_1 and A_2 are justified while B_1 and B_2 are overruled.

While this formalisation in ASPIC yields the intended outcomes, it formally multiplies an attack that intuitively is a single one. It would be interesting to investigate how the ASPIC logic could be extended with arguments about whether arguments rebut each other but this has to be left to a future occasion.

Finally, we assume that preferences on desires can be argued about as in [8]'s combination of [10] with [7]. In our example this yields, for example, that an argument concluding $D_2 > D_1$ pref-attacks A_2's attack on B_1.

3.2. The outcome of the criteria deliberation phase

With these schemes (plus possibly other schemes for supporting or attacking their premises) an extended argumentation framework is built. At the end of the first phase, the preferred extensions are determined and for each such extension the set of criteria conclusions it contains is an admissible criterion set. More precisely, let for any set E of arguments $\texttt{Conc}(E)$ be the set of all conclusions of any argument in E. Then S is an *admissible-criteria set* (ac-set for short) if there exists a preferred extension E such that

$$S = \{SC(C) \mid SC(C) \in \texttt{Conc}(E)\} \cup \{NC(C) \mid NC(C) \in \texttt{Conc}(E)\}.$$

We also keep track of the desires that motivate the criteria in the set, since this information will be needed in the second phase to compare proposals. Thus for each extension E and ac-set S induced by E we have for each sufficient criterion $s \in S$ that:

$$d(s) = \{d \mid s \textit{ satisfies } d \text{ and } d \textit{ is desirable} \text{ are in } \texttt{Conc}(E)\}$$

while for each necessary condition $n \in S$ this is changed to:

$$d(n) = \{d \mid d \textit{ requires } n \text{ and } d \textit{ is desirable} \text{ are in } \texttt{Conc}(E)\}$$

Note that if we also include schemes for accruing different arguments for the same conclusion (as in e.g. [9, 2]) the sets $d(s)$ and $d(n)$ can contain more than one element.

4. The proposal deliberation phase

In the proposal deliberation phase action proposals can be made and supported by sets of sufficient criteria that they satisfy, where such sets must be nonempty subsets of some ac-set resulting from the first phase. Proposal arguments can be attacked in two ways.

- They can be attacked on their premises by arguments claiming that the proposal does not satisfy some given sufficient criterion; whether this attack is symmetric depends on the nature of the attack.
- By alternative proposal arguments. These arguments can either take their premises from the same or from another admissible criterion set. This kind of attack is symmetric.

Note that in these attacks necessary criteria are not utilised. Necessary criteria, because they represent constraints on, rather than reasons for, action, are instead used in arguments that pref-attack attack relations. Such arguments summarise for each of the two conflicting proposal arguments which sets of sufficient conditions they satisfy and which sets of necessary conditions they violate, and they then expess a preference between the proposals based on preferences on these criteria (which may in turn be based on the desires motivating these criteria).

We abstract from the internal structure of proposals, for example, from whether a proposal concerns atomic or combined actions. We thus leave room for proposals that include other proposals (for example, to both raise taxes and cut social benefits). In much other work on argumentation over action (e.g.[1]) it is assumed that only one action can be performed in a situation but for democratic deliberation this assumption is not realistic, as our example shows. If a proposal that combines two actions overcommits since one of the actions satisfies the same sufficient criteria, then (if the debate is conducted properly) this will reflect itself in violation of a necessary criterion (such as 'don't put more financial burdens on citizens than necessary').

4.1. Argument schemes for proposal deliberation

The argument scheme for generating proposals has the following form:

PS: proposal P should be adopted since
P satisfies sufficient criterion s_1 and
... and
P satisfies sufficient criterion s_n and
$\{s_1, \ldots, s_n\}$ is a subset of ac-set $\mathcal{S}$.

The general scheme for preference arguments is as follows (where the set notation is a shorthand for a conjunction of statements on satisfaction of criteria, and $S(P)$ stands for the ac-set on which proposal P is based):

PrS: proposal P_2 is preferred over proposal P_1 since
P_1 satisfies sufficient criteria $S_1 = \{s_1, \ldots, s_m\}$
P_2 satisfies sufficient criteria $S_2 = \{s_n, \ldots, s_p\}$
P_1 violates necessary criteria $N_1 = \{n_1, \ldots, n_k\}$ from $S(P_1) \cup S(P_2)$
P_2 violates necessary criteria $N_2 = \{n_l, \ldots, n_q\}$ from $S(P_1) \cup S(P_2)$
$(S_1, N_2) < (S_2, N_1)$

The conclusion of this scheme is assumed to pref-attack P_1's attack of P_2. Like all other schemes, this scheme can be attacked on any of its premises, for example, on whether a proposal really violates some necessary criterion, or on the ordering on sets. In fact, this is the only way in which preference arguments can be attacked, since the idea is that if all premises hold, then the preference conclusion holds by definition. Thus all disagreement and defeasibility is located in the premises of this scheme.

A key issue is how the last premise of this set is verified. It makes sense to let this depend on the participants' preferences concerning individual criteria in this set, which will often in turn depend on the sets of desires motivating these criteria (the sets $d(s)$ and $d(n)$ defined in Section 3.2). In applications one approach is to let the participants build and attack arguments on preference relations between individual desires and to define strict inference rules for combining these preferences; these rules could adapt a suitable method from the literature (see e.g. [3]). The definition should at least ensure that a proposal that fully satisfies an ac-set (i.e., which satisfies at least one of its sufficient criteria and violates none of its necessary criteria) is preferred over a proposal that does not fully satisfy an ac-set. In any case, some of the preference arguments needed in this phase will already have been put forward in the first phase, to decide conflicts involving applications of the NCS scheme. To avoid duplication of these arguments, it makes sense in applications to import them and the arguments that pertain to them into the extended argumentation framework of the second phase and to disallow further supports or attacks with respect to these arguments.

Note that with the above schemes a proposal argument cannot simply be attacked by pointing at a sufficient criterion that it does not satisfy or a necessary criterion that it violates. Such attacks are always part of alternative proposal arguments. This is since otherwise there are too many asymmetric attacks and too few admissible proposals, since in the real world there may be no proposal that satisfies all sufficient criteria.

4.2. The outcome

The result is that during this discussion phase a second extended argumentation framework is built. At the end of the phase, all preferred extensions are automatically identified. Note that each extension will contain at most one proposal argument. If there is a unique extension then there is full agreement, otherwise there must be some external way to make a choice, for example, by a vote.

5. An example

For our example we will use an issue in UK Road Traffic policy, previously used as an e-participation example in [4]. The number of fatal road accidents is an obvious cause for concern, and in the UK there are speed restrictions on various types of road, in the belief that excessive speed causes accidents. The policy issue which we will consider is how to reduce road deaths.

A number of desirable things need to be considered when deciding on such a policy. Obviously saving lives is important, but also there are cost issues. There are also some general considerations: civil liberties pressure groups oppose measures that intrude on privacy and freedom to a disproportionate extent. On the other hand, respect for law and order should lead us to want to ensure compliance with any law, including speed limits.

So what criteria should we use to assess this aspect of Road Traffic policy? The law-and-order consideration would motivate an argument:

SCS_1: *Compliance with Speed Limits* should be a sufficient criterion since it satisfies *Obedience to the Law*, which is desirable.

Similarly saving lives would give rise to the argument:

SCS_2: *Fewer Accidents* should be a sufficient criterion since it satisfies *Saving Lives*, which is desirable.

We identified cost as a constraint, and hence it represents a necessary criterion:

NCS_1: *Minimum Expenditure* should be a necessary criterion since *Controlling the Budget* requires *Minimum Expenditure* and is desirable.

Additionally there are two constraints from the Civil Liberties perspective;

NCS_2: *Lack of Intrusion* should be a necessary criterion since *Privacy* requires *Lack of Intrusion* and is desirable.
NCS_3: *Lack of Regulation* should be a necessary criterion since *Freedom* requires *Lack of Regulation* and is desirable.

There is an incompatibility between *Compliance with Speed Limits* and *Lack of Regulation* and so we have instances of $CN2S$ and $CS2N$.

$CN2S_1$: *Compliance with Speed Limits* should not be a sufficient criterion since *Lack of Regulation* should be a necessary criterion and is incompatible with *Compliance with Speed Limits*.
$CS2N_1$: *Lack of Regulation* should not be a necessary criterion since *Compliance with Speed Limits* should be a sufficient criterion and is incompatible with *Lack of Regulation*.

We consider that all the desiderata are accepted as desirable, and that there is no consensus on the preference ordering, in the absence of further context.

We thus have two preferred extensions, and two sets of admissible criteria:

AC_1: $\{SC$(*Fewer Accidents*), NC(*Minimum Expenditure*), NC(*Lack of Intrusion*), NC(*Lack of Regulation*) $\}$
AC_2: $\{$ SC(*Compliance with Speed Limits*), SC(*Fewer Accidents*), NC(*Minimum Expenditure*), NC(*Lack of Intrusion*) $\}$

Next consider a set of proposals:

P_1: *Introduce More Speed Cameras*. This has a capital cost and a revenue stream (from fines), increases compliance and reduces accidents but is intrusive and does not lead to fewer regulations.
P_2: *Increase Traffic Police*. This has a revenue cost, increases compliance and reduces accidents but does not lead to fewer regulations.
P_3: *Educate the Public*. This has a revenue cost, increases compliance and reduces accidents.

We may assume that the most expensive is P_2, then P_3 and finally P_1. Statistics suggest that P_1 is the most effective in reducing accidents, then P_2 and finally P_3, while P_2 ensures greatest compliance, then P_1 and finally P_3. We assume all this can be established from statistics on previous use of the various schemes and so will not consider factual arguments further.

We now use PrS to compare the options $P_1 - P_3$ pairwise. Although many ways of comparing sets of criteria could be used, here we use a form of cancellation. Where both arguments satisfy a criterion, we count that criterion only for the proposal which satisfies it to a greater degree. Where the degree is equal, we exclude it from both.

P_1 v P_2: P_1 satisfies *Fewer Accidents* and violates *Lack of Intrusion* and *Lack of regulations*: P_2 satisfies *Compliance with Speed Limits* and violates *Minimum Expenditure*.

P_1 v P_3: P_1 satisfies *Fewer Accidents* and *Compliance with Speed Limits* but violates *Lack of Intrusion* and *Lack of Regulation*: P_3 violates *Minimum Expenditure*.

P_2 v P_3: P_2 satisfies *Fewer Accidents* and *Compliance with Speed Limits* but violates *Minimum Expenditure* and *Lack of Regulation*: P_3 violates nothing but is less effective in satisfying the sufficient conditions.

P_3 can now be discarded since it fails to satisfy either sufficient condition. Now using AC_1, P_1 is preferred since it satisfies the only sufficient condition. AC_2 is more problematic since P_2 now also satisfies a sufficient criterion and P_1 violates one fewer necessary criterion. If the most prized aims (as was the case in the UK at the time) are to save lives at reasonable cost, then P_1 is preferred. If on the other hand the priorities are to increase compliance while respecting civil liberties, P_2 will be chosen.

With the recent Government change in the UK, the issue has become how to reduce costs to cut what is considered an insupportable budget deficit. This changes the emphasis of the arguments. Now, the key sufficient criterion is SCS_3:

SCS_3: *Saving Money* should be a sufficient condition since it satisfies *Reducing the Deficit*, which is desirable.

Now cost becomes a reason for adopting a proposal, rather than only a constraint on what proposals can be adopted. But it is also *essential* that money be saved and so we still need an argument for the necessary criterion:

NCS_{1*}: *Saving Money* should be a necessary criterion since *Reducing the Deficit* requires *Saving Money* and is desirable

Unfortunately any way of saving lives requires expenditure, and so we have a $CN2S$ conflict between this and SCS_2. Because cost is now our most preferred objective, SCS_2 will be defeated. But since we do not wish to lose sight of saving lives altogether we may impose a new necessary condition:

NCS_5: *Not Increasing Accidents* should be a necessary condition since *Saving Lives* requires *Not Increasing Accidents* and is desirable.

In so far as saving money is possible without increasing accidents, NCS_5 allows this feature to be considered as a constraint on proposals: for example, NCS_5 will enable us to argue against any proposal to remove speed cameras (as opposed to not installing new ones), since this will increase accidents, without saving costs, although preference will still be given to proposals which save costs.

6. Discussion and Conclusion

In this paper we have proposed a formal two-phase model of democratic policy deliberation, with a clear separation between deliberation about criteria for proposals and about how proposals satisfy them. Related work especially concerns formal argumentation-based models for practical reasoning, such as [1, 4]. Unlike this and similar work, our separation in two phases allows that the premises of arguments in the second phase refer to the *outcome* of the first phase. For example, some premises of the PS and PrS schemes require that given sets resulting from the first phase are admissible. In this respect our approach is related to [6]'s modular assumption-based frameworks, in which premises of one module can refer to a consequence notion applied to another module. In future work it would be interesting to compare the pros and cons of both methods.

Another topic for future research is to extend the present model to allow for degrees of satisfaction of desires or constraints. Then preference arguments could consider degrees of satisfaction in a more principled and realistic way than was used in section 5.

References

[1] K. Atkinson, T.J.M. Bench-Capon, and P. McBurney. Computational representation of persuasive argument. *Synthese*, 152:157–206, 2006.

[2] T.J.M. Bench-Capon and H. Prakken. Justifying actions by accruing arguments. In P.E. Dunne and T.B.C. Bench-Capon, editors, *Computational Models of Argument. Proceedings of COMMA 2006*, pages 247–258. IOS Press, Amsterdam etc, 2006.

[3] J.-F. Bonnefon and H. Fargier. Comparing sets of positive and negative arguments: Empirical assessment of seven qualitative rules. In *Proceedings of the 17th European Conference on Artificial Intelligence (ECAI'06)*, pages 16–20, 2006.

[4] D. Cartwright and K. Atkinson. Using computational argumentation to support e-participation. *IEEE Intelligent Systems*, 24:42–52, 2009. Special Issue on Transforming E-government and E-participation through IT.

[5] 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.

[6] P.M. Dung and P.M. Thang. Modular argumentation for modelling legal doctrines in common law of contract. *Artificial Intelligence and Law*, 17:167–182, 2009.

[7] S. Modgil. Reasoning about preferences in argumentation frameworks. *Artificial Intelligence*, 173:901–934, 2009.

[8] S. Modgil and H. Prakken. Reasoning about preferences in structured extended argumentation frameworks. In P. Baroni, F. Cerutti, M. Giacomin, and G.R. Simari, editors, *Computational Models of Argument. Proceedings of COMMA 2010*, pages 347–358. IOS Press, Amsterdam etc, 2010.

[9] H. Prakken. A study of accrual of arguments, with applications to evidential reasoning. In *Proceedings of the Tenth International Conference on Artificial Intelligence and Law*, pages 85–94, New York, 2005. ACM Press.

[10] H. Prakken. An abstract framework for argumentation with structured arguments. *Argument and Computation*, 1:93–124, 2010.

[11] J.-J. Rousseau. The social contract. Or principles of political right, 1762. Translated by G.D.H. Cole, public domain. Available at `www.constitution.org/jjr/socon.htm`.