Uncontested Semantics for Value-Based Argumentation

Paul E. DUNNE
Department of Computer Science, The University of Liverpool, U.K.

Abstract. We introduce an extension-based semantics for value-based argumentation frameworks (VAFs) that provides a counterpart to the recently proposed *ideal semantics* in standard – i.e. value–free – argumentation frameworks. A significant motivation for this so-called "*uncontested semantics*" is as a mechanism with which to refine the nature of objective acceptance: thus the set of uncontested arguments are not only considered justified irrespective of the value ordering endorsed by any audience but, in addition, collectively constitute a self-defending and internally consistent collection of beliefs within the framework. In this way the rationale underpinning objectively accepted arguments which fall outside this uncontested set must involve audience related features. In this paper we formalise the concept of uncontested arguments in VAFs, present a number of features distinguishing ideal and uncontested semantics, and analyse some basic complexity-theoretic issues.

Keywords. Computational properties of argumentation; argumentation frameworks; value-based argumentation; computational complexity

Introduction

The value-based argumentation model introduced by Bench-Capon [3] has proven to be an important approach with which to examine processes of practical reasoning and the rationale supporting why parties favour certain beliefs over alternatives, e.g. [16,4,1,2]. This model builds on the seminal approach to abstract argumentation pioneered by Dung [8] wherein a number of formalisations of the concept of "collection of justified beliefs" are proposed in terms of criteria defined on subsets of arguments in an abstract argumentation framework (AF). Two important classifications have been considered in terms of Dung's basic acceptability semantics: *credulous acceptance* – an argument is justified if it belongs to *at least one* admissible (i.e. self-defending and internally consistent) set of arguments; and *sceptical acceptance* – an argument is justified if it belongs to *every* maximal such set (or *preferred extension* in Dung's terminology). In recent work Dung, Mancarella and Toni [9,10] advance a new classification, *ideal acceptance*, under which an argument has not only to be sceptically accepted but also contained within an admissible set of sceptically accepted arguments.

In this paper we propose and examine an analogue of "ideal acceptability" tailored to value-based argumentation: the term *uncontested acceptability* being used as a general descriptor. Thus, as discussed in [4, pp. 50–51], the concept of credulous acceptance in AFs is similar to Bench-Capon's notion of *subjective acceptability* in VAFs; and that of sceptical acceptance akin to *objective acceptability* in VAFs.[1] In informal terms, an argument is considered to be uncontested if it is both objectively acceptable *and* contained in an admissible set in the standard sense of Dung [8], of objectively acceptable arguments. A significant motivation for the concept of uncontested acceptance derives from the rationale underpinning arguments, p, which are objectively accepted but *are not* uncontested, a status which can be interpreted in the following way: although the argument p is considered justified irrespective of the value priorities endorsed by an audience, the *reasons* and supporting cases differ between distinct audiences. As a very simple example consider the VAF of Fig. 1(a).

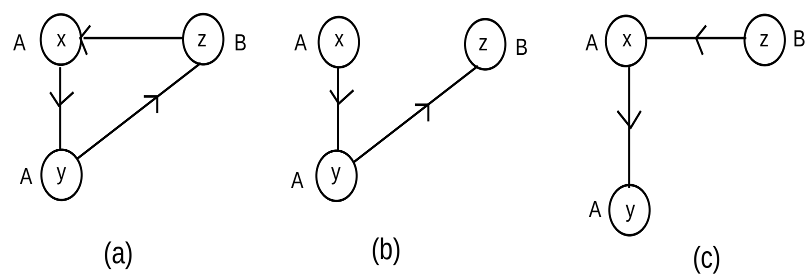

Figure 1. Different audiences accept z for different reasons

The argument z is objectively accepted since it is in the preferred extension ($\{x, z\}$) resulting for the audience that regards value A as more important than value B – the VAF of Fig. 1(b); and in the preferred extension $\{y, z\}$ deriving from the remaining audience which holds the opposite view, for which the relevant VAF is that shown in Fig. 1(c). The argument z, however, is not uncontested: $\{z\}$ does not form an admissible set. Overall the two different audiences have distinct rationalisations for accepting z: the audience preferring A to B accepts z by reason of "although z is attacked by the argument y, because we also accept x we can accept z since x attacks y"; the audience preferring B to A, however, regards z as indisputable and requires no further support to justify its acceptance.

After reprising the basic notions of Dung's argumentation frameworks and Bench-Capon's VAF development of these in Section 1 we continue in Section 2 by formally defining the concept of uncontested acceptability in VAFs. The relationship between this and the ideal semantics of Dung *et al.* [9,10] is examined in Section 2.1. One of our principal aims concerns complexity-theoretic analysis of uncontested semantics so that Section 3 presents a number of results regarding

[1]The parallel between credulous/subjective and sceptical/objective derives from the former being phrased in terms of "there exists $\mathcal{S}$ such that" (for suitable structures $\mathcal{S}$) and the latter in terms of "for all $\mathcal{S}$ such that". Aside from this syntactic similarity, however, the related semantics give rise to very different behaviours, cf. [4, Thm. 12]

both decision problems and the construction of uncontested sets in VAFs. Further work and conclusions are discussed in Section 4.

1. Preliminaries: AFs and VAFs

The following concepts were introduced in Dung [8].

Definition 1 *An* argumentation framework *(AF) is a pair $\mathcal{H} = \langle \mathcal{X}, \mathcal{A} \rangle$, in which $\mathcal{X}$ is a finite set of* arguments *and $\mathcal{A} \subseteq \mathcal{X} \times \mathcal{X}$ is the* attack relationship *for $\mathcal{H}$. A pair $\langle x, y \rangle \in \mathcal{A}$ is referred to as 'y* is attacked by *x' or 'x* attacks *y'. The convention of excluding "self-attacking" arguments is assumed, i.e. for all $x \in \mathcal{X}$, $\langle x, x \rangle \notin \mathcal{A}$. For R, S subsets of arguments in the AF $\mathcal{H}(\mathcal{X}, \mathcal{A})$, we say that $s \in S$ is* attacked *by R – written attacks(R, s) – if there is some $r \in R$ such that $\langle r, s \rangle \in \mathcal{A}$. For subsets R and S of $\mathcal{X}$ we write attacks(R, S) if there is some $s \in S$ for which attacks(R, s) holds; $x \in \mathcal{X}$ is* acceptable with respect to *S if for every $y \in \mathcal{X}$ that attacks x there is some $z \in S$ that attacks y; S is* conflict-free *if no argument in S is attacked by any other argument in S.*

A conflict-free set S is admissible *if every $y \in S$ is acceptable w.r.t S; S is a* preferred extension *if it is a maximal (with respect to $\subseteq$) admissible set; S is a* stable extension *if S is conflict free and every $y \notin S$ is attacked by S; S is an* ideal extension *([9,10]) of $\mathcal{H}$ if S is admissible* and *a subset of every preferred extension of $\mathcal{H}$.*

An AF, $\mathcal{H}$ is coherent *if every preferred extension in $\mathcal{H}$ is also a stable extension. $\mathcal{H}$ is* cohesive[2] *if its maximal ideal extension coincides with the intersection of all preferred extensions of $\mathcal{H}$.*

For $S \subseteq \mathcal{X}$,

$$\begin{array}{lll} S^- & =_{\text{def}} & \{\, p \;:\; \exists\, q \in S \;\; \textit{such that}\; \langle p, q \rangle \in \mathcal{A} \} \\ S^+ & =_{\text{def}} & \{\, p \;:\; \exists\, q \in S \;\; \textit{such that}\; \langle q, p \rangle \in \mathcal{A} \} \end{array}$$

An argument x is credulously accepted *if there is* some *preferred extension containing it; x is* sceptically accepted *if it is a member of* every *preferred extension.*

Bench-Capon [3] develops the concept of "attack" from Dung's model to take account of *values*.

Definition 2 *A* value-based argumentation framework *(VAF), is defined by a triple $\mathcal{H}^{(\mathcal{V})} = \langle \mathcal{H}(\mathcal{X}, \mathcal{A}), \mathcal{V}, \eta \rangle$, where $\mathcal{H}(\mathcal{X}, \mathcal{A})$ is an AF, $\mathcal{V} = \{v_1, v_2, \ldots, v_k\}$ a set of k* values*, and $\eta : \mathcal{X} \to \mathcal{V}$ a mapping that associates a value $\eta(x) \in \mathcal{V}$ with each argument $x \in \mathcal{X}$.*

An audience *for a VAF $\langle \mathcal{X}, \mathcal{A}, \mathcal{V}, \eta \rangle$, is a binary relation $\mathcal{R} \subset \mathcal{V} \times \mathcal{V}$ whose (irreflexive) transitive closure, $\mathcal{R}^*$, is asymmetric, i.e. at most one of $\langle v, v' \rangle$, $\langle v', v \rangle$ are members of $\mathcal{R}^*$ for any distinct v, $v' \in \mathcal{V}$. We say that v_i* is preferred

[2] The term *cohesive* is introduced in Dunne [11] although the property described had been raised in the original work of Dung *et al.* [9,10].

to v_j in the audience $\mathcal{R}$, *denoted* $v_i \succ_{\mathcal{R}} v_j$, *if* $\langle v_i, v_j\rangle \in \mathcal{R}^*$. *We say that* α *is a* specific *audience if* α *yields a* total *ordering of* $\mathcal{V}$.

For an audience, $\mathcal{R}$*, there will, generally, be a number of specific audiences consistent with* $\mathcal{R}^*$*. The notation,* $\chi(\mathcal{R})$ *is used to describe the set of all such specific audiences, i.e.*

$$\chi(\mathcal{R}) \quad =_{\text{def}} \quad \{\ \alpha\ :\ \forall\ v,\ v' \in \mathcal{V}\ \langle v, v'\rangle \in \mathcal{R}^*\ \Rightarrow\ v \succ_{\alpha} v'\ \}$$

Following, [4, p. 40], the audience $\mathcal{R} = \emptyset$ *is called the* universal audience *by reason of* $\chi(\emptyset)$ *containing every specific audience.*

A standard assumption from [3] which we retain in our subsequent development is the following:
Multivalued Cycles Assumption (MCA)
For any *simple cycle* of arguments in a VAF, $\langle \mathcal{X}, \mathcal{A}, \mathcal{V}, \eta\rangle$, – i.e. a finite sequence of arguments $y_1 y_2 \ldots y_i y_{i+1} \ldots y_r$ with $y_1 = y_r$, $|\{y_1, \ldots, y_{r-1}\}| = r - 1$, and $\langle y_j, y_{j+1}\rangle \in \mathcal{A}$ for each $1 \leq j < r$ – there are arguments y_i and y_j for which $\eta(y_i) \neq \eta(y_j)$.

In less formal terms, this assumption states every simple cycle in $\mathcal{H}^{(\mathcal{V})}$ uses at least two distinct values.

Using VAFs, ideas analogous to those introduced in Defn. 1 are given by relativising the concept of "attack" using that of *successful* attack with respect to an audience. Thus,

Definition 3 *Let* $\langle \mathcal{X}, \mathcal{A}, \mathcal{V}, \eta\rangle$ *be a* VAF *and* $\mathcal{R}$ *an audience. For arguments* x*,* y *in* $\mathcal{X}$*,* x *is a* successful attack *on* y *(or* x defeats y*)* with respect to the audience $\mathcal{R}$ *if:* $\langle x, y\rangle \in \mathcal{A}$ and *it is* not *the case that* $\eta(y) \succ_{\mathcal{R}} \eta(x)$.

Replacing "attack" by "successful attack w.r.t. the audience $\mathcal{R}$", in Defn. 1 yields definitions of "conflict-free", "admissible set" etc. relating to value-based systems, e.g. S is conflict–free w.r.t. to the audience $\mathcal{R}$ if for each x, y in S it is not the case that x successfully attacks y w.r.t. $\mathcal{R}$. It may be noted that a conflict-free set in this sense is not necessarily a conflict-free set in the sense of Defn. 1: for x and y in S we may have $\langle x, y\rangle \in \mathcal{A}$, provided that $\eta(y) \succ_{\mathcal{R}} \eta(x)$, i.e. the value promoted by y is preferred to that promoted by x for the audience $\mathcal{R}$.

The concept of successful attack w.r.t. an audience $\mathcal{R}$, leads to the following notation as a parallel to the sets S^- and S^+. Given $\langle \mathcal{X}, \mathcal{A}, \mathcal{V}, \eta\rangle$, $S \subseteq \mathcal{X}$ and an audience $\mathcal{R} \subset \mathcal{V} \times \mathcal{V}$,

$$\begin{array}{lll} S^-_{\mathcal{R}} & =_{\text{def}} & \{\ p\ :\ \exists\ q \in S\ \text{ such that } \langle p, q\rangle \in \mathcal{A} \text{ and } \langle \eta(q), \eta(p)\rangle \notin \mathcal{R}^*\} \\ S^+_{\mathcal{R}} & =_{\text{def}} & \{\ p\ :\ \exists\ q \in S\ \text{ such that } \langle q, p\rangle \in \mathcal{A} \text{ and } \langle \eta(p), \eta(q)\rangle \notin \mathcal{R}^*\} \end{array}$$

Bench-Capon [3] proves that every specific audience, α, induces a unique preferred extension within its underlying VAF: for a given VAF, $\mathcal{H}^{(\mathcal{V})}$, we use $P(\mathcal{H}^{(\mathcal{V})}, \alpha)$ to denote this extension: that $P(\mathcal{H}^{(\mathcal{V})}, \alpha)$ *is* unique and can be constructed efficiently, is an easy consequence of the following fact, implicit in [3].

Fact 1 *For any* VAF*,* $\mathcal{H}^{(\mathcal{V})}(\langle \mathcal{X}, \mathcal{A}, \mathcal{V}, \eta\rangle)$ *(satisfying* MCA*) and specific audience* α*, the framework induced by including only attacks in the set* $\mathcal{A}_{\alpha}$ *given by* $\mathcal{A} \setminus \{\langle x, y\rangle\ :\ \eta(y) \succ_{\alpha} \eta(x)\}$ *is* acyclic.

Proof: Suppose the contrary and let $y_1 y_2 \ldots y_r$ (with $y_r = y_1$) be any simple cycle in the VAF $\langle\langle \mathcal{X}, \mathcal{A}_\alpha \rangle, \mathcal{V}, \eta\rangle$ defined from $\mathcal{H}^{(\mathcal{V})}$ via the specific audience α. Since each of the attacks $\langle y_i, y_{i+1}\rangle$ for $1 \le i \le r-1$ occurs in $\mathcal{A} \cap \mathcal{A}_\alpha$ from the definition of $\mathcal{A}_\alpha$ we must have $\forall\ 1 \le i \le r-1\ \ \neg(\eta(y_{i+1})\ \succ_\alpha\ \eta(y_i))$. That is,

$$\forall\ 1 \le i \le r-1\ \ (\eta(y_i)\ \succ_\alpha\ \eta(y_{i+1})) \bigvee\ (\eta(y_i) = \eta(y_{i+1}))$$

With some minor abuse of notation, we write $v \succeq_\alpha w$ if $(v = w)\ \vee\ (v \succ_\alpha w)$, so that the expression above implies $\eta(y_1) \succeq_\alpha \eta(y_2) \succeq_\alpha \ldots \succeq_\alpha \eta(y_{r-1}) \succeq_\alpha \eta(y_1)$. Since α is a specific audience so that $\succeq_\alpha$ is a total ordering, the only possible choice of values which this behaviour could arise is $\eta(y_1) = \eta(y_2) = \ldots = \eta(y_i) = \ldots = \eta(y_{r-1})$ which contradicts the assumption that $\mathcal{H}^{(\mathcal{V})}$ satisfies MCA. □

Analogous to the concepts of credulous and sceptical acceptance, in VAFs the ideas of *subjective* and *objective* acceptance (w.r.t. an audience $\mathcal{R}$) arise, [4, p. 48]. The computational complexity of a number of decision problems in both standard AFs and VAFs has been considered in work of Dimopoulos and Torres [7], Dunne and Bench-Capon [13], and Dunne [12]. The results of these papers are summarised in Table 1.

Table 1. Computational complexity in AFs and VAF

Problem	Instance	Question	Complexity	
CA	$\langle \mathcal{H}(\mathcal{X},\mathcal{A}), x\rangle$	Is x credulously accepted?	NP–complete	[7]
SA	$\langle \mathcal{H}(\mathcal{X},\mathcal{A}), x\rangle$	Is x sceptically accepted?	Π_2^p–complete	[13]
COH	$\mathcal{H}(\mathcal{X},\mathcal{A})$	Is $\mathcal{H}$ coherent?	Π_2^p–complete.	[13]
SBA	$\langle \mathcal{H}^{(\mathcal{V})}, x, \mathcal{R}\rangle$	$\exists\ \alpha \in \chi(\mathcal{R})\ :\ x \in P(\mathcal{H}^{(\mathcal{V})}, \alpha)$?	NP–complete	[14,4] (with $\mathcal{R} = \emptyset$)
OBA	$\langle \mathcal{H}^{(\mathcal{V})}, x, \mathcal{R}\rangle$	$\forall\ \alpha \in \chi(\mathcal{R})\ :\ x \in P(\mathcal{H}^{(\mathcal{V})}, \alpha)$?	co-NP–complete	[14,4] (with $\mathcal{R} = \emptyset$)

In recent work of Dunne [12], the complexity of SBA and OBA w.r.t. the universal audience is shown to be unchanged under quite extreme restrictions on the form of instances.

Fact 2 *(Dunne [12])*

1. *Let* SBA$^{(T)}$ *be the decision problem* SBA *with instances restricted to those for which the graph structure* $\langle \mathcal{X}, \mathcal{A}\rangle$ *is a* binary *tree:* SBA$^{(T)}$ *is* NP*–complete.*
2. *Let* SBA$^{(T,\epsilon)}$ *be the decision problem* SBA$^{(T)}$ *in which instances are restricted to those in which* $|\mathcal{V}| \le |\mathcal{X}|^\epsilon$*:* $\forall\ \epsilon > 0$ SBA$^{(T,\epsilon)}$ *is* NP*–complete.*
3. *Suppose* SBA$^{(\mathcal{V},\le k)}$ *is the decision problem* SBA *restricted to instances for which* $\forall\ v \in \mathcal{V}\ |\eta^{-1}(v)| \le k$*, i.e. at most* k *arguments share a common value,* $v \in \mathcal{V}$*. Similarly,* SBA$^{(T),(\mathcal{V},\le k)}$ *is this problem with instances additionally restricted to trees:* SBA$^{(T),(\mathcal{V},\le 3)}$ *is* NP*–complete.*

Analogous co-NP–completeness results for OBA also hold for the restricted frameworks of Fact 2.

2. Uncontested Semantics in VAFs

In this paper we are concerned with a VAF based extension semantics that captures elements of the ideal semantics: the *uncontested semantics*.

Definition 4 *Let $\mathcal{H}^{(\mathcal{V})}$ be a* VAF *and $\mathcal{R}$ an audience. A set of arguments, S in $\mathcal{H}^{(\mathcal{V})}$ is an* uncontested extension w.r.t. $\mathcal{R}$ *if it is an admissible set in $\mathcal{H}(\langle \mathcal{X}, \mathcal{A}\rangle)$* and *every argument in S is objectively acceptable in $\mathcal{H}^{(\mathcal{V})}$ w.r.t. the audience $\mathcal{R}$.*

Our main interest in this section is to review various properties of this approach: characteristics in common with ideal extensions such as that described in Thm. 1; as well as points under which these forms differ.

Theorem 1 *Let $\mathcal{H}^{(\mathcal{V})}(\langle \mathcal{X}, \mathcal{A}, \mathcal{V}, \eta\rangle)$ be a* VAF*, $\mathcal{H}(\mathcal{X}, \mathcal{A})$ its supporting (value free)* AF*, and $\mathcal{R} \subset \mathcal{V} \times \mathcal{V}$ be any audience.*

If $\mathcal{U}_1$ and $\mathcal{U}_2$ are both uncontested extensions of $\mathcal{H}^{(\mathcal{V})}$ w.r.t. $\mathcal{R}$ then $\mathcal{U}_1 \cup \mathcal{U}_2$ is also an uncontested extension of $\mathcal{H}^{(\mathcal{V})}$ w.r.t. $\mathcal{R}$.

Proof: Suppose that $\mathcal{U}_1$ and $\mathcal{U}_2$ are uncontested extensions of $\mathcal{H}^{(\mathcal{V})}$ w.r.t. $\mathcal{R}$. Since, $\mathcal{U}_1 \subseteq \{x \in \mathcal{X} : \text{OBA}(\mathcal{H}^{(\mathcal{V})}, x, \mathcal{R})\}$ and $\mathcal{U}_2 \subseteq \{x \in \mathcal{X} : \text{OBA}(\mathcal{H}^{(\mathcal{V})}, x, \mathcal{R})\}$ it is certainly the case that $\mathcal{U}_1 \cup \mathcal{U}_2 \subseteq \{\ x\ :\ \text{OBA}(\mathcal{H}^{(\mathcal{V})}, x, \mathcal{R})\ \}$. It, therefore, suffices to show that $\text{ADM}(\mathcal{H}, \mathcal{U}_1 \cup \mathcal{U}_2)$, where $\text{ADM}(\mathcal{H}, S)$ is the predicate returning **true** if and only if the set S is admissible for $\mathcal{H}(\mathcal{X}, \mathcal{A})$. We first show that this set is conflict-free. Noting that both $\text{ADM}(\mathcal{H}, \mathcal{U}_1)$ and $\text{ADM}(\mathcal{H}, \mathcal{U}_2)$, the set $\mathcal{U}_1 \cup \mathcal{U}_2$ could only fail to be conflict-free if some attack in $\mathcal{A}$ involved an argument in $\mathcal{U}_1$ and an argument in $\mathcal{U}_2$. Without loss of generality suppose $\langle p_1, q_1\rangle \in \mathcal{A}$ with $p_1 \in \mathcal{U}_1$ and $q_1 \in \mathcal{U}_2$. Consider any specific audience $\alpha \in \chi(\mathcal{R})$, then from $\{p_1, q_1\} \subseteq P(\mathcal{H}^{(\mathcal{V})}, \alpha)$ (by objective acceptability w.r.t. $\mathcal{R}$ of both p_1 and q_1) for all such specific audiences we must have $\eta(q_1) \succ_\alpha \eta(p_1)$. The set $\mathcal{U}_2$, however, is admissible for $\langle \mathcal{X}, \mathcal{A}\rangle$ so there is some $q_2 \in \mathcal{U}_2$ for which $\langle q_2, p_1\rangle \in \mathcal{A}$. Now, in the same way as before, $\{p_1, q_1, q_2\} \subseteq P(\mathcal{H}^{(\mathcal{V})}, \alpha)$ so that $\eta(q_1) \succ_\alpha \eta(p_1) \succ_\alpha \eta(q_2)$. Continuing thus, using admissibility of $\mathcal{U}_1$ we find $p_2 \in \mathcal{U}_1$ with $\langle p_2, q_2\rangle \in \mathcal{A}$, so from the original premise $\langle p_1, q_1\rangle \in \mathcal{A}$ we identify subsets $\{p_1, \ldots, p_r\} \subseteq \mathcal{U}_1$ and $\{q_1, \ldots, q_r\} \subseteq \mathcal{U}_2$ for which

1. $\{p_1, p_2, \ldots, p_r, q_1, q_2, \ldots, q_r\} \subseteq P(\mathcal{H}^{(\mathcal{V})}, \alpha)$ for every $\alpha \in \chi(\mathcal{R})$.
2. $\langle p_i, q_i\rangle \in \mathcal{A}$ and $\langle q_{i+1}, p_i\rangle \in \mathcal{A}$ for each $1 \leq i \leq r$
3. $\eta(q_i) \succ_\alpha \eta(p_i) \succ_\alpha \eta(q_{i+1})$ $(1 \leq i \leq r)$ for all $\alpha \in \chi(\mathcal{R})$.

The sets $\mathcal{U}_1$ and $\mathcal{U}_2$ are finite[3] and thus the chain implied by (3) will eventually lead to a contradiction. We deduce, therefore, that $\mathcal{U}_1 \cup \mathcal{U}_2$ is a conflict-free set in $\mathcal{H}(\mathcal{X}, \mathcal{A})$. That it is also, admissible, follows easily: any attacker x of $\mathcal{U}_1 \cup \mathcal{U}_2$, either attacks some $y_1 \in \mathcal{U}_1$ (and thus counterattacked by $\mathcal{U}_1$) or some $y_2 \in \mathcal{U}_2$ (and so counterattacked by $\mathcal{U}_2$). □

Corollary 1 *For every* VAF*, $\mathcal{H}^{(\mathcal{V})}(\langle \mathcal{X}, \mathcal{A}, \mathcal{V}, \eta\rangle)$ and audience $\mathcal{R}$, there is a* unique, maximal *uncontested extension w.r.t. $\mathcal{R}$.*

[3]In fact, we could equally use the fact that $\mathcal{V}$ is finite to derive this contradiction, i.e. the analysis also applies to frameworks in which $\mathcal{X}$ is allowed to be an infinite set.

Proof: Trivial consequence of Thm. 1. □

We introduce the following notation for VAFs $\mathcal{H}^{(\mathcal{V})}(\langle \mathcal{X}, \mathcal{A}, \mathcal{V}, \eta \rangle)$ and audiences $\mathcal{R}$.

$$\begin{array}{lcl} \mathcal{U}_{\mathcal{R}} & = & \text{The maximal uncontested extension of } \mathcal{H}^{(\mathcal{V})} \text{ w.r.t. } \mathcal{R} \\ \mathcal{U} & = & \mathcal{U}_{\emptyset} \\ \mathcal{M} & = & \text{The maximal ideal extension of } \mathcal{H}(\mathcal{X}, \mathcal{A}) \end{array}$$

2.1. Ideal vs. Uncontested Semantics

One issue arising from our definitions is to what extent do the ideal semantics and uncontested semantics give rise to *distinct* subsets. The results comprising this section consider this issue.

Theorem 2 *Given a* VAF*, $\langle \mathcal{X}, \mathcal{A}, \mathcal{V}, \eta \rangle$ let $\mathcal{M}$ denote its maximal ideal extension w.r.t to the* AF *$\langle \mathcal{X}, \mathcal{A} \rangle$.*

a. *There are* VAF*s for which $x \in \mathcal{M}$ but $x \notin \mathcal{U}$.*
b. *There are* VAF*s for which $x \in \mathcal{U}$ but $x \notin \mathcal{M}$.*
c. *There are* VAF*s for which $\mathcal{U} \subset \{x \ : \ \text{OBA}(\langle \mathcal{X}, \mathcal{A}, \mathcal{V}, \eta \rangle, x, \emptyset)\}$.*

Proof: Consider the three systems of Fig. 2

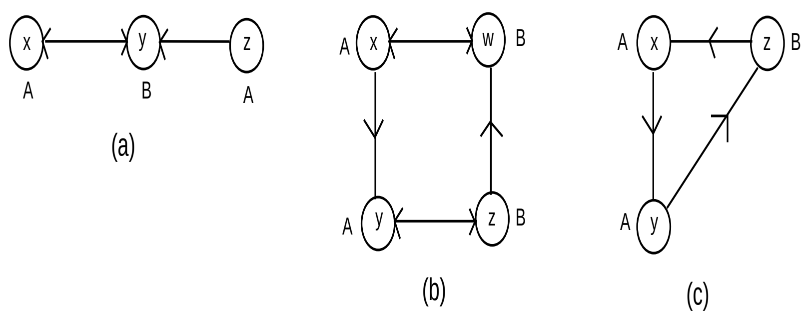

Figure 2. Cases in proof of Thm. 2

For (a), the maximal ideal extension for the AF of Fig. 2(a) is $\{x, z\}$, however $\mathcal{U} = \{z\}$ (since x is not in the preferred extension with respect to the specific audience $B \succ A$), so that $x \in \mathcal{M}$ but $x \notin \mathcal{U}$.

For (b), the VAF of Fig. 2(b) has $\mathcal{U} = \{x, z\}$: both specific audiences yielding the preferred extension $\{x, z\}$. In contrast $\mathcal{M} = \emptyset$ in the underlying AF: both $\{x, z\}$ and $\{y, w\}$ being preferred extensions of this.

Finally, for the VAF of Fig 2(c), we have $\mathcal{U} = \emptyset$ (since every non-empty subset of $\{x, y, z\}$ fails to be admissible) whereas the argument z is objectively accepted, so establishing (c). □

We have further indications that uncontested extensions describe radically different structures, in the failure of the following characterising lemmata, proven for ideal semantics in [11], to have an analogue in the uncontested semantics.

Fact 3 *Let $\mathcal{H}(\mathcal{X}, \mathcal{A})$ be an* AF *with maximal ideal extension $\mathcal{M} \subseteq \mathcal{X}$.*

a. *A subset S of $\mathcal{X}$ defines an ideal extension of $\mathcal{H}(\mathcal{X}, \mathcal{A})$ if and only if both (I1) and (I2) below hold:*

   I1. *S is an admissible set in $\mathcal{H}(\mathcal{X}, \mathcal{A})$.*
   I2. *No attacker of S is credulously accepted, i.e. $\forall\ y \in S^- \ \neg\text{CA}(\mathcal{H}, y)$.*

b. *For any $x \in \mathcal{X}$, $x \in \mathcal{M}$ if and only if both (M1) and (M2) below hold:*

   M1. *No attacker of x is credulously accepted, i.e. $\forall\ y \in \{x\}^- \ \neg\text{CA}(\mathcal{H}, y)$.*
   M2. *For each attacker, y of x, there is some attacker, z of y, for which $z \in \mathcal{M}$, i.e. $\forall\ y \in \{x\}^- \ : \ \{y\}^- \cap \mathcal{M} \neq \emptyset$.*

A natural reformulation of (M1) and (M2) in terms of VAFs is

U1. No attacker of x is subjectively accepted w.r.t. $\mathcal{R}$.
U2. For each attacker y of x, some attacker z of y is in $\mathcal{U}_{\mathcal{R}}$.

The following result demonstrates, however, that these fail to characterise maximal uncontested extensions.

Lemma 1 *There are* VAF*s, $\mathcal{H}^{(\mathcal{V})}(\langle \mathcal{X}, \mathcal{A}, \mathcal{V}, \eta \rangle)$ with maximal uncontested extensions, $\mathcal{U}$, that do not satisfy (U1). i.e. an argument $y \in \mathcal{U}^-$ is subjectively accepted, so that this is not a* necessary *condition for membership in the maximal uncontested extension.*

Proof: Consider the VAF, $\mathcal{H}^{(\mathcal{V})}$ of Fig. 3

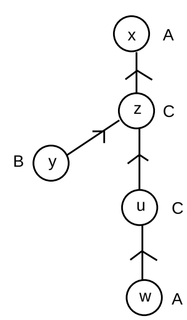

Figure 3. No attacker subjectively accepted is not a necessary condition

The maximal uncontested extension is formed by the set $\{x, y, w\}$: this is easily seen to be admissible since $\{y, w\}^- = \emptyset$ and the sole attacker, z, of x is counterattacked by y. Each argument in $\{x, y, w\}$ is also objectively accepted: that $\{y, w\} \subseteq P(\mathcal{H}^{(\mathcal{V})}, \alpha)$ for any specific audience α, is immediate from $\{y, w\}^- = \emptyset$. The argument x is in $P(\mathcal{H}^{(\mathcal{V})}, \alpha)$ for all specific audiences in which $A \succ_\alpha C$ (since the attack $\langle z, x \rangle$ does not succeed); the remaining specific audiences (in which $C \succ_\alpha A$) satisfy $u \in P(\mathcal{H}^{(\mathcal{V})}, \alpha)$, so that $\langle u, z \rangle$ is an attack in the acyclic AF induced by these, thereby providing u as a defence to the attack by z on x.

The argument $z \in \{x, y, w\}^-$ is, however, subjectively accepted using the specific audience $A \succ C \succ B$: $\langle y, z\rangle$ does not succeed with respect to this audience; the (successful) attack $\langle w, u\rangle$ provides w as a defence to the attack by u on z so that $P(\mathcal{H}^{(\mathcal{V})}, A \succ C \succ B) = \{x, y, w, z\}$. □

We observe that the constructions of Thm. 2 and Lemma 1 further emphasise the fact that there are a number of subtle differences between the divers acceptability semantics proposed for VAFs – i.e. Subjective, Objective and Uncontested – in comparison with superficially similar acceptability semantics in AFs, i.e. Credulous, Sceptical and Ideal. Thm. 2 and Lemma 1 thus develop the related comparison of Bench-Capon *et al.* [4, Thm. 12, pp. 50–51].

Despite Fact 3 failing to have an immediate counterpart when characterising uncontested extensions, cf. Lemma 1, it turns out that a very similar result can be obtained, albeit in a rather indirect manner. We first introduce the notion of a VAF being *k-terse*, where $k \geq 1$.

Definition 5 *For $k \in \mathbf{N}$, the* VAF*, $\mathcal{H}^{(\mathcal{V})}(\langle \mathcal{X}, \mathcal{A}, \mathcal{V}, \eta\rangle)$ is k-*terse *if every simple* directed *path of length k involves* at most *k different values from $\mathcal{V}$. Formally, $\forall\ x_1x_2x_3 \cdots x_{k+1} \in \mathcal{X}^{k+1}$ such that $\langle x_i, x_{i+1}\rangle \in \mathcal{A}$ for each $1 \leq i \leq k$ and $|\{x_1, x_2, \ldots, x_{k+1}\}| = k+1$, $|\{\ \eta(x_i)\ :\ 1 \leq i \leq k+1\}|\ \leq k$.*

Theorem 3 *Let $\mathcal{H}^{(\mathcal{V})}(\langle \mathcal{X}, \mathcal{A}, \mathcal{V}, \eta\rangle)$ be any 2-terse* VAF*, $\mathcal{R}$ an audience. The argument $x \in \mathcal{X}$ is in $\mathcal{U}_\mathcal{R}$ if and only if both of the following hold:*

U1. *No attacker, y, of x is subjectively accepted w.r.t. $\mathcal{R}$ in $\mathcal{H}^{(\mathcal{V})}$, i.e. $\forall\ y \in \{x\}^-$, $\neg$SBA$(\mathcal{H}^{(\mathcal{V})}, y, \mathcal{R})$.*
U2. *For every attacker, y, of x, at least one attacker, z of y, is in $\mathcal{U}_\mathcal{R}$, i.e. $\forall\ y \in \{x\}^-\ \{y\}^- \cap \mathcal{U}_\mathcal{R} \neq \emptyset$.*

Proof:
($\Rightarrow$) Suppose that $x \in \mathcal{U}_\mathcal{R}$ for the 2-terse VAF, $\mathcal{H}^{(\mathcal{V})}(\langle \mathcal{X}, \mathcal{A}, \mathcal{V}, \eta\rangle)$. We show that x satisfies both (U1) and (U2). To see that (U2) holds it suffices to observe that, since $\mathcal{U}_\mathcal{R}$ is an admissible set in $\langle \mathcal{X}, \mathcal{A}\rangle$ any attacker y of x must be counterattacked by some $z \in \mathcal{U}_\mathcal{R}$, thus for each $y \in \{x\}^-$ we have $\{y\}^- \cap \mathcal{U}_\mathcal{R} \neq \emptyset$. To see that x must satisfy (U1), suppose for the sake of contradiction that this were not the case, i.e. there is a 2-terse, VAF and audience $\mathcal{R}$ with maximal uncontested extension $\mathcal{U}_\mathcal{R}$ containing an argument x an attacker, y, of which is subjectively accepted w.r.t. $\mathcal{R}$. Since $\mathcal{U}_\mathcal{R}$ is an admissible set it must contain an argument z that attacks y, e.g. Fig 4 where $\eta(x) = V_x,\ \eta(y) = V_y$ and $\eta(z) = V_z$.

Consider any specific audience, $\alpha \in \chi(\mathcal{R})$, under which $y \in P(\mathcal{H}^{(\mathcal{V})}, \alpha)$: since $\{x, z\} \subseteq \mathcal{U}_\mathcal{R}$ it holds that OBA$(\mathcal{H}^{(\mathcal{V})}, x, \mathcal{R}) \wedge$ OBA$(\mathcal{H}^{(\mathcal{V})}, z, \mathcal{R})$ and thus, $\{x, y, z\} \subseteq P(\mathcal{H}^{(\mathcal{V})}, \alpha)$. It follows, therefore, that neither $\langle y, x\rangle$ nor $\langle z, y\rangle$ can be attacks in the AF arising from $\mathcal{H}^{(\mathcal{V})}$ with respect to the specific audience α,[4] i.e. $\eta(x) \succ_\alpha \eta(y)$ and $\eta(y) \succ_\alpha \eta(z)$. This, however, is only possible when $\eta(x)$, $\eta(y)$, and $\eta(z)$ are

[4]Notice that $\{x, y\} \subset P(\mathcal{H}^{(\mathcal{V})}, \alpha)$ also indicates that x cannot self-defend the attack by y, i.e. at most one of $\{\langle x, y\rangle, \langle y, x\rangle\}$ are in $\mathcal{A}$.

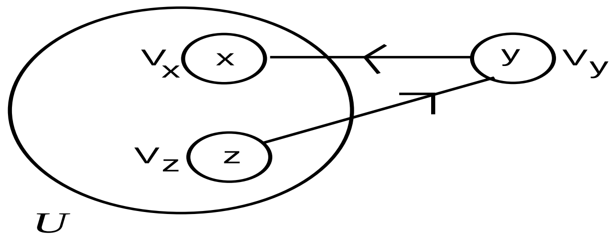

Figure 4. No attacker subjectively accepted is needed in 2-terse VAFs

all distinct values. This contradicts the assumption that $\mathcal{H}^{(\mathcal{V})}$ is 2-terse since the path $z \to y \to x$ involves three distinct values. As a result we deduce that $\neg\text{SBA}(\mathcal{H}^{(\mathcal{V})}, y, \mathcal{R})$ for every $y \in \{x\}^-$, i.e. that $x \in \mathcal{U}_\mathcal{R}$ implies that x satisfies (U1).

($\Leftarrow$) Suppose that x satisfies both (U1) and (U2). Consider the set $\mathcal{U}_\mathcal{R} \cup \{x\}$. Certainly this is conflict-free: if $\langle y, x\rangle \in \mathcal{A}$ for some $y \in \mathcal{U}_\mathcal{R}$ then x fails to satisfy (U1) since $\text{OBA}(\mathcal{H}^{(\mathcal{V})}, y, \mathcal{R})$; if $\langle x, y\rangle \in \mathcal{A}$ for some $y \in \mathcal{U}_\mathcal{R}$ then from admissibility we find $z \in \mathcal{U}_\mathcal{R}$ with $\langle z, x\rangle \in \mathcal{A}$ and again this contradicts x satisfying (U1). The set $\mathcal{U}_\mathcal{R} \cup \{x\}$ in addition to being conflict-free, is also admissible: any y that attacks $\mathcal{U}_\mathcal{R} \cup \{x\}$ either attacks $\mathcal{U}_\mathcal{R}$ (and so is counterattacked by some $z \in \mathcal{U}_\mathcal{R}$) or attacks x so that $y \in \{x\}^-$ so that since x satisfies (U2) we find $z \in \mathcal{U}_\mathcal{R} \cap \{y\}^-$ as a defence.

It must, however, also be the case that $\text{OBA}(\mathcal{H}^{(\mathcal{V})}, x, \mathcal{R})$: for suppose this were not so and for some specific audience, $\alpha \in \chi(\mathcal{R})$, $x \notin P(\mathcal{H}^{(\mathcal{V})}, \alpha)$. Since $P(\mathcal{H}^{(\mathcal{V})}, \alpha)$ is a *stable* extension, we must have $\{x\}^- \cap P(\mathcal{H}^{(\mathcal{V})}, \alpha) \neq \emptyset$: this, however, would contradict x satisfying (U1). In summary, from x satisfying (U1) and (U2), the set $\mathcal{U}_\mathcal{R} \cup \{x\}$ is both admissible in $\langle \mathcal{X}, \mathcal{A}\rangle$ *and* each of its arguments is objectively accepted w.r.t. $\mathcal{R}$ in $\mathcal{H}^{(\mathcal{V})}$, i.e. this set is an uncontested extension. From the fact that $\mathcal{U}_\mathcal{R}$ is *maximal* we deduce $\mathcal{U}_\mathcal{R} \cup \{x\} = \mathcal{U}_\mathcal{R}$, i.e $x \in \mathcal{U}_\mathcal{R}$ as required. $\square$

The property 2-terseness may seem rather too restrictive in order for the characterisation of Thm. 3 to be widely applicable. As the following result shows, this in fact is not necessarily the case.

Lemma 2 *(Path Dilation Lemma –* PDL*) Let $\langle \mathcal{X}, \mathcal{A}, \mathcal{V}, \eta\rangle$ be any* VAF*. There is a* VAF*, $\langle \mathcal{X} \cup \mathcal{Y}, \mathcal{B}, \mathcal{V}, \varepsilon\rangle$ such that*

PD1. $\forall x \in \mathcal{X}$, $\forall \alpha$, $x \in P(\langle \mathcal{X}, \mathcal{A}, \mathcal{V}, \eta\rangle, \alpha) \quad \Leftrightarrow \quad x \in P(\langle \mathcal{X} \cup \mathcal{Y}, \mathcal{B}, \mathcal{V}, \varepsilon\rangle, \alpha)$.
PD2. $\langle \mathcal{X} \cup \mathcal{Y}, \mathcal{B}, \mathcal{V}, \varepsilon\rangle$ *is 2-terse.*

Furthermore $\langle \mathcal{X} \cup \mathcal{Y}, \mathcal{B}, \mathcal{V}, \varepsilon\rangle$ is constructible in polynomial time from $\langle \mathcal{X}, \mathcal{A}, \mathcal{V}, \eta\rangle$.

Proof: Given $\langle \mathcal{X}, \mathcal{A}, \mathcal{V}, \eta\rangle$ as input, consider the result $\langle \mathcal{X} \cup \mathcal{Y}, \mathcal{B}, \mathcal{V}, \varepsilon\rangle$ of applying Algorithm 1

The typical transformation enacted by Algorithm 1 is illustrated in Fig. 5. Notice that we do *not* assume $\{x\}^- \cap \{x\}^+ = \emptyset$.

Observing that the effect of replacing a *single* $x_i \in \mathcal{X}$ by the structure described in Algorithm 1 reduces the total number of arguments contributing to

Algorithm 1 Construction of 2-terse VAF from $\langle \mathcal{X}, \mathcal{A}, \mathcal{V}, \eta \rangle$

$\mathcal{Y} := \emptyset$; $\mathcal{B} := \mathcal{A}$
for $x_i \in \mathcal{X}$ **do**
  $\varepsilon(x_i) := \eta(x_i)$
  if $\exists\ (y \in \{x_i\}^-) \wedge (z \in \{x_i\}^+)$ s.t. $|\{\eta(y), \eta(x_i), \eta(z)\}| = 3$ **then**
    $\mathcal{Y} := \mathcal{Y} \cup \{x_{i,1}^{in},\ x_{i,2}^{in},\ x_{i,1}^{out},\ x_{i,2}^{out}\}$
    $\varepsilon(x_{i,1}^{in}) := \varepsilon(x_i)$; $\varepsilon(x_{i,2}^{in}) := \varepsilon(x_i)$
    $\varepsilon(x_{i,1}^{out}) := \varepsilon(x_i)$; $\varepsilon(x_{i,2}^{out}) := \varepsilon(x_i)$
    $\mathcal{B} := \mathcal{B} \setminus \{\langle y, x_i \rangle : y \in \{x_i\}^-\}$
    $\mathcal{B} := \mathcal{B} \setminus \{\langle x_i, z \rangle : z \in \{x_i\}^+\}$
    $\mathcal{B} := \mathcal{B} \cup \{\langle y, x_{i,1}^{in} \rangle : \langle y, x_i \rangle \in \mathcal{A}\}$
    $\mathcal{B} := \mathcal{B} \cup \{\langle x_{i,2}^{out}, z \rangle : \langle x_i, z \rangle \in \mathcal{A}\}$
    $\mathcal{B} := \mathcal{B} \cup \{\langle x_{i,1}^{in}, x_{i,2}^{in} \rangle, \langle x_{i,2}^{in}, x_i \rangle\}$
    $\mathcal{B} := \mathcal{B} \cup \{\langle x_i, x_{i,1}^{out} \rangle, \langle x_{i,1}^{out}, x_{i,2}^{out} \rangle\}$
  end if
end for
return $\langle \mathcal{X} \cup \mathcal{Y}, \mathcal{B}, \mathcal{V}, \varepsilon \rangle$

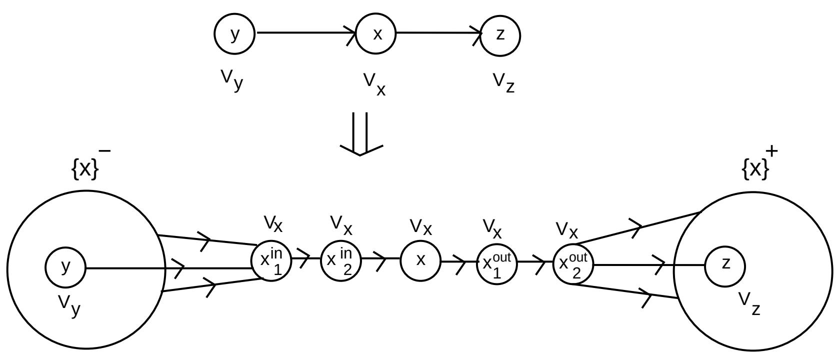

Figure 5. Path Dilation in VAFs resulting from Algorithm 1

paths which fail to be 2-terse, it is easily seen that the final VAF, $\langle \mathcal{X} \cup \mathcal{Y}, \mathcal{B}, \mathcal{V}, \varepsilon \rangle$ is 2-terse. To complete the proof we need to show that $x \in P(\langle \mathcal{X}, \mathcal{A}, \mathcal{V}, \eta \rangle, \alpha)$ if and only if $x \in P(\langle \mathcal{X} \cup \mathcal{Y}, \mathcal{B}, \mathcal{V}, \varepsilon \rangle, \alpha)$. Here it is only necessary to argue that the effect of replacing a single x preserves acceptability with respect to specific audiences. So suppose $x \in \mathcal{X}$ and let $\langle \mathcal{X} \cup \mathcal{Y}, \mathcal{B}, \mathcal{V}, \varepsilon \rangle$ be the VAF formed after a single iteration of the main loop in Algorithm 1. If no changes have occurred (that is, every length 2 path with middle argument x_1 is 2-terse), then it is certainly the case $\forall x\ x \in P(\langle \mathcal{X}, \mathcal{A}, \mathcal{V}, \eta \rangle, \alpha)$ if and only if $x \in P(\langle \mathcal{X} \cup \mathcal{Y}, \mathcal{B}, \mathcal{V}, \varepsilon \rangle, \alpha)$. So assume that x_1 has resulted in changes to the structure of $\langle \mathcal{X}, \mathcal{A}, \mathcal{V}, \eta \rangle$. Let $S = P(\langle \mathcal{X}, \mathcal{A}, \mathcal{V}, \eta \rangle, \alpha)$. We claim that

$$T \quad =_{\text{def}} \quad P(\langle \mathcal{X} \cup \mathcal{Y}, \mathcal{B}, \mathcal{V}, \varepsilon \rangle, \alpha) \quad = \quad \begin{cases} S \ \cup \ \{x_{1,1}^{in},\ x_{1,2}^{out}\} & \text{if } x_1 \in S \\ S \ \cup \ \{x_{1,2}^{in},\ x_{1,1}^{out}\} & \text{if } x_1 \notin S \end{cases}$$

To see this, let $\langle \mathcal{X}, \mathcal{A}_\alpha \rangle$ be the AF induced from $\langle \mathcal{X}, \mathcal{A}, \mathcal{V}, \eta \rangle$ by the specific audience α, and $\langle \mathcal{X} \cup \mathcal{Y}, \mathcal{B}_\alpha \rangle$ the AF induced from $\langle \mathcal{X} \cup \mathcal{Y}, \mathcal{B}, \mathcal{V}, \varepsilon \rangle$ by α so that

(similarly to the construction of Fact 1)

$$\begin{array}{rcl}\mathcal{A}_\alpha & = & \mathcal{A} \setminus \{\langle x,y\rangle \in \mathcal{A} \ : \ \eta(y) \succ_\alpha \eta(x)\} \\ \mathcal{B}_\alpha & = & \mathcal{B} \setminus \{\langle x,y\rangle \in \mathcal{B} \ : \ \varepsilon(y) \succ_\alpha \varepsilon(x)\}\end{array}$$

Observing that $\mathcal{Y} = \{x^{in}_{1,1}, x^{in}_{1,2}, x^{out}_{1,1}, x^{out}_{1,2}\}$, it is easily seen that $\mathcal{B}_\alpha$ contains every attack in $\mathcal{A}_\alpha$ with the exception of

$$\{\langle y, x_1\rangle \ : \ \neg(\eta(x_1) \succ_\alpha \eta(y)\} \bigcup \{\langle x_1, z\rangle \ : \ \neg(\eta(z) \succ_\alpha \eta(x_1)\}$$

and that the only attacks in $\mathcal{B}_\alpha$ that do *not* occur in $\mathcal{A}_\alpha$ are

$$\begin{array}{ll}\{\langle y, x^{in}_{1,1}\rangle \ : \ \neg(\varepsilon(x^{in}_{1,1}) \succ_\alpha \varepsilon(y)\} & \cup \ \{\langle x^{out}_{1,2}, z\rangle \ : \ \neg(\varepsilon(z) \succ_\alpha \varepsilon(x^{out}_{1,2})\} \\ & \cup \ \{\langle x^{in}_{1,1}, x^{in}_{1,2}\rangle, \langle x^{in}_{1,2}, x_1\rangle, \langle x_1, x^{out}_{1,1}\rangle, \langle x^{out}_{1,1}, x^{out}_{1,2}\rangle\}\end{array}$$

Recalling that S is the preferred extension of the AF, $\langle \mathcal{X}, \mathcal{A}_\alpha\rangle$, first suppose that $x_1 \in S$. In this case $S \cup \{x^{in}_{1,1}, x^{out}_{1,2}\}$ is certainly an admissible set in $\langle \mathcal{X} \cup \mathcal{Y}, \mathcal{B}_\alpha\rangle$: any attacker y of x_1 in $\langle \mathcal{X}, \mathcal{A}_\alpha\rangle$ attacks $x^{in}_{1,1}$ in $\langle \mathcal{X} \cup \mathcal{Y}, \mathcal{B}_\alpha\rangle$ and y is countered by some $z \in S$. If $z \in S \setminus \{x\}$ then $z \in S \cup \{x^{in}_{1,1}, x^{out}_{1,2}\}$ so that z defends $x^{in}_{1,1}$ against the attack by y. If $z = x$, then $\langle x^{out}_{1,2}, y\rangle \in \mathcal{B}_\alpha$ so that the attack $\langle y, x^{in}_{1,1}\rangle$ is countered by the attack $\langle x^{out}_{1,2}, y\rangle$. In consequence $S \cup \{x^{in}_{1,1}, x^{out}_{1,2}\}$ is an admissible set in $\langle \mathcal{X} \cup \mathcal{Y}, \mathcal{B}_\alpha\rangle$ if $x_1 \in S$, i.e. $S \cup \{x^{in}_{1,1}, x^{out}_{1,2}\} \subseteq T$. It must, however, also be a *maximal* such set in this AF. For suppose, $S \cup \{x^{in}_{1,1}, x^{out}_{1,2}\} \subset T$ so that T contains some argument, y say, not among $S \cup \{x^{in}_{1,1}, x^{out}_{1,2}\}$. It cannot be the case that $y \in \{x^{in}_{1,2}, x^{out}_{1,1}\}$ since $x_1 \in T$ and T is conflict-free. Thus $y \in \mathcal{X}$ but $y \notin S$. Since S is a stable extension of $\langle \mathcal{X}, \mathcal{A}_\alpha\rangle$, there must be some attacker, z, of y in S. If $z \in S \setminus x$ then $z \in S \cup \{x^{in}_{1,1}, x^{out}_{1,2}\}$ so that y could not belong to T. If $z = x$ then $\langle x^{out}_{1,2}, y\rangle \in \mathcal{B}_\alpha$ and, again, we cannot have $y \in T$. We deduce that if S is the unique preferred extension of $\langle \mathcal{X}, \mathcal{A}_\alpha\rangle$ and $x_1 \in S$ then $S \cup \{x^{in}_{1,1}, x^{out}_{1,2}\}$ is the unique preferred extension of $\langle \mathcal{X} \cup \mathcal{Y}, \mathcal{B}_\alpha\rangle$.

For the remaining possibility, suppose that $x_1 \notin S$ and consider $S \cup \{x^{in}_{1,2}, x^{out}_{1,1}\}$. Again this latter set is admissible: $x_1 \notin S$ requires some attacker, y, to be in S so that the attack by $x^{in}_{1,1}$ on $x^{in}_{1,2}$ is countered by the attack $\langle y, x^{in}_{1,1}\rangle \in \mathcal{B}_\alpha$ (it is immediate that $x^{in}_{1,2}$ defends $x^{out}_{1,1}$ from the attack $\langle x_1, x^{out}_{1,1}\rangle$). We thus have, $S \cup \{x^{in}_{1,2}, x^{out}_{1,2}\} \subseteq T$ with T the unique preferred extension of $\langle \mathcal{X} \cup \mathcal{Y}, \mathcal{B}_\alpha\rangle$. The set $S \cup \{x^{in}_{1,2}, x^{out}_{1,1}\}$ must, however, also be *maximal*. Were there some $y \in T$ not contained in it, it cannot be the case that $y \in \{x^{in}_{1,1}, x_1, x^{out}_{1,2}\}$ (recall that T is the *unique* preferred extension so that $S \cup \{x^{in}_{1,2}, x^{out}_{1,1}\}$ is a subset of T). Thus $y \notin \mathcal{Y}$, i.e. $y \in \mathcal{X} \setminus \{x\}$. As before, since $y \notin S$, it must be attacked by some $z \in S$ (since S is the stable extension of $\langle \mathcal{X}, \mathcal{A}_\alpha\rangle$) which suffices to guarantee that $y \notin T$, i.e. $T = S \cup \{x^{in}_{1,2}, x^{out}_{1,1}\}$ forms the unique preferred extension of $\langle \mathcal{X} \cup \mathcal{Y}, \mathcal{B}_\alpha\rangle$ whenever $x_1 \notin S$. □

3. Complexity of Uncontested Semantics

Results on the computational complexity of problems in the *ideal semantics* of Dung *et al.* [9,10] are presented by Dunne in [11] and summarised in Tables 2 and 3: The randomized reductions are in the sense described by Valiant and

Table 2. Decision questions for Ideal Semantics

Problem Name	Instance	Question
IE	$\langle\mathcal{H}(\mathcal{X},\mathcal{A}),S\rangle$	Is S an ideal extension?
IA	$\langle\mathcal{H}(\mathcal{X},\mathcal{A}),x\rangle$	Is x in the maximal ideal extension?
MIE$_\emptyset$	$\mathcal{H}(\mathcal{X},\mathcal{A})$	Is the maximal ideal extension empty?
MIE	$\langle\mathcal{H}(\mathcal{X},\mathcal{A}),S\rangle$	Is S the *maximal* ideal extension?
CS	$\mathcal{H}(\mathcal{X},\mathcal{A})$	Is $\mathcal{H}(\mathcal{X},\mathcal{A})$ *cohesive*?

Table 3. Complexity of ideal semantics in AFs

Decision Problem	Complexity (Randomized reduction)	Complexity ($\leq_m^p$)
IE	–	co-NP–complete
IA	$\mathrm{P}_{\|\|}^{\mathrm{NP}}$–complete	co-NP–hard
MIE	$\mathrm{P}_{\|\|}^{\mathrm{NP}}$–complete	D^p–hard
MIE$_\emptyset$	$\mathrm{P}_{\|\|}^{\mathrm{NP}}$–complete	NP–hard
CS	–	Σ_2^p–complete

Vazirani [18] and developed by Chang *et al.* [6]: all of the randomized hardness constructions succeed with probability $\geq 1-2^{-|\mathcal{X}|}$. It is further shown by [11] that the problem of *finding* the maximal ideal extension is complete (via standard many-one reducibility, $\leq_m^p$) for the function class $\mathrm{FP}_{||}^{\mathrm{NP}}$.

The problems arising for ideal argumentation described in Table 2 motivate those of Table 4

Table 4. Decision questions for Uncontested Semantics

Problem Name	Instance	Question
UE	$\langle\mathcal{H}^{(\mathcal{V})},S,\mathcal{R}\rangle$	Is S an uncontested extension?
UA	$\langle\mathcal{H}^{(\mathcal{V})},x,\mathcal{R}\rangle$	Is $x\in\mathcal{U}_\mathcal{R}(\mathcal{H}^{(\mathcal{V})})$?
MUE$_\emptyset$	$\langle\mathcal{H}^{(\mathcal{V})},\mathcal{R}\rangle$	Is $\mathcal{U}_\mathcal{R}(\mathcal{H}^{(\mathcal{V})})=\emptyset$?
MUE	$\langle\mathcal{H}^{(\mathcal{V})},S,\mathcal{R}\rangle$	Is $S=\mathcal{U}_\mathcal{R}(\mathcal{H}^{(\mathcal{V})})$?

Although almost all of the results presented are w.r.t. the *universal audience* we note the following: for *lower* bounds it suffices to demonstrate that these hold for a single audience, $\mathcal{R}$, thus the universal audience subsumes the cases where $\mathcal{R}$ is part of the instance; in addition where upper bound methods are given such will apply for arbitrary audiences and not only $\mathcal{R}=\emptyset$.

Theorem 4 UE *is co-*NP*–complete.*

Proof: To see that UE $\in$ co-NP it suffices to note that given $\langle\mathcal{H}^{(\mathcal{V})},S,\mathcal{R}\rangle$ as an instance of UE, S defines an uncontested extension of $\mathcal{H}^{(\mathcal{V})}$ if and only if S is admissible and every argument of S is objectively accepted w.r.t. $\mathcal{R}$. This can be testing by checking

$$\forall\ \alpha\in\chi(\mathcal{R})\quad \mathrm{ADM}(\mathcal{H},S)\ \wedge\ \bigwedge_{x\in S}\ (x\in P(\mathcal{H}^{(\mathcal{V})},\alpha))$$

which can be carried out in co-NP.[5]

To establish co-NP–hardness we use a reduction from UNSAT without loss of generality restricted to instances in 3-CNF. We will actually prove a stronger result: that UE is co-NP–hard even when instances are restricted to VAFs whose supporting AF is a *tree* and having every value in $\mathcal{V}$ associated with at most 3 arguments.

We start from the VAF, $\mathcal{T}_\Phi$, described in Dunne [12, Thm. 25], which is formed from a 3-CNF formula, $\Phi(z_1, \ldots, z_n)$, the argument, Φ, of which is subjectively accepted (w.r.t. the universal audience) if and only if $\Phi(z_1, \ldots, z_n)$ is satisfiable: $\mathcal{T}_\Phi$ is a tree and no value is associated with more than 3 arguments.[6] The instance, $\mathcal{F}_\Phi$ of UE is formed by adding two arguments – f_1 and f_2 – to $\mathcal{T}_\Phi$; attacks $\{\langle f_1, \Phi\rangle, \langle \Phi, f_2\rangle\}$; a new value, v_f to $\mathcal{V}_\Phi$ (the value set of $\mathcal{T}_\Phi$); and defining $\eta(f_1) = v_f$, $\eta(f_2) = v_f$. The instance of UE is completed by setting $S = \{f_1, f_2\}$ and the audience used is, again, $\mathcal{R} = \emptyset$, i.e. the universal audience. The construction is illustrated in Fig. 6.

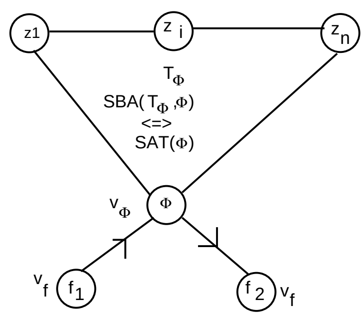

Figure 6. The VAF $\mathcal{F}_\Phi$: SBA$(\mathcal{T}_\Phi, \Phi) \Leftrightarrow$ SAT(Φ)

We now claim that $\langle \mathcal{F}_\Phi, \{f_1, f_2\}, \emptyset\rangle$ is accepted as an instance of UE if and only if $\Phi(z_1, \ldots, z_n)$ is unsatisfiable.

Suppose that Φ is unsatisfiable. The set $\{f_1, f_2\}$ is admissible, so it suffices to show that each of its constituent arguments is objectively accepted. Certainly $\{f_1, f_2\} \subseteq P(\mathcal{F}_\Phi, \alpha)$ for any specific audience in which $v_f \succ_\alpha v_\Phi$: $\{f_1\}^- = \emptyset$ and the attack $\langle \Phi, f_2\rangle$ is unsuccessful. For any specific audience in which $v_\Phi \succ_\alpha v_f$, f_1 is objectively acceptable (irrespective of whether Φ is satisfiable). When Φ is unsatisfiable and $v_\Phi \succ_\alpha v_f$ it is again the case that $f_2 \in P(\mathcal{F}_\Phi, \alpha)$: recalling that the value v_f does not occur amongst the values used in $\mathcal{T}_\Phi$, since Φ is unsatisfiable, there is no specific audience, β, for which $\Phi \in P(\mathcal{T}_\Phi, \beta)$ and thus the successful attack $\langle \Phi, f_2\rangle$ can always be countered using a suitable argument of $\mathcal{T}_\Phi$. We deduce that if Φ is unsatisfiable then $\{f_1, f_2\}$ is an uncontested extension.

[5] Note that given any audience $\mathcal{R} \subset \mathcal{V} \times \mathcal{V}$, testing $\alpha \in \chi(\mathcal{R})$ can be carried out in polynomial time, thus universal quantification could equally be applied to specific audiences in general, rather than simply those in $\chi(\mathcal{R})$.

[6] The precise specification of $\mathcal{T}_\Phi$ is not important for the purposes of the proof: only the property that SBA$(\mathcal{T}_\Phi, \Phi)$ if and only if $\Phi(z_1, \ldots, z_n)$ is satisfiable.

Conversely suppose that $\{f_1, f_2\}$ is an uncontested extension of $\mathcal{F}_\Phi$. In this case f_2 is objectively accepted, thus a member of $P(\mathcal{F}_\Phi, \alpha)$ for any specific audience. In particular, $f_2 \in P(\mathcal{F}_\Phi, \alpha)$ for all specific audiences in which $v_\Phi \succ_\alpha v_f$. We deduce that $\Phi \notin P(\mathcal{F}_\Phi, \alpha)$ and thus $\Phi \notin P(\mathcal{T}_\Phi, \alpha)$, i.e. $\neg$SBA$(\mathcal{T}_\Phi, \Phi, \emptyset)$ from which it follows that $\Phi(z_1, \ldots, z_n)$ is unsatisfiable. $\square$

We note that following the methods of [12, Corollary 7], this construction can be further developed to show UE is co-NP–hard for instances $\langle \mathcal{X}, \mathcal{A}, \mathcal{V}, \eta \rangle$ in which $\langle \mathcal{X}, \mathcal{A} \rangle$ is a *binary* tree with every $v \in \mathcal{V}$ associated with at most three arguments of $\mathcal{X}$.

Corollary 2

a. UA *is co*-NP*–hard.*
b. MUE$_\emptyset$ *is* NP*–hard.*
c. MUE *is* D^p*–hard.*

Proof:

a. For the VAF, $\mathcal{F}_\Phi$, described in the proof of Thm. 4, the argument f_2 belongs to its maximal uncontested extension if and only if Φ is unsatisfiable.
b. We employ a similar reduction to that given in the proof of Thm. 4, but using the VAF, $\mathcal{W}_\Phi$, described in [4, Thm. 8, pp. 65-66]: unlike the construction of $\mathcal{T}_\Phi$ from Dunne [12, Thm. 25], $\mathcal{W}_\Phi$ has an empty uncontested extension.[7] Modify $\mathcal{W}_\Phi$ as shown in Fig. 7 to give a new VAF, $\mathcal{D}_\Phi$. Then MUE$_\emptyset(\mathcal{D}_\Phi)$ if and only if $\Phi(z_1, \ldots, z_n)$ is satisfiable (otherwise the argument $\{f\}$ constitutes an uncontested extension.)

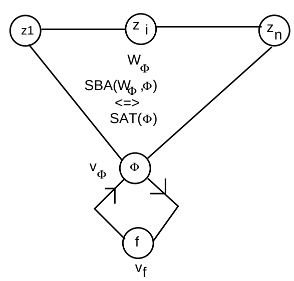

Figure 7. The VAF $\mathcal{D}_\Phi$

c. We show that the decision problem SAT-UNSAT whose instances are pairs of CNF-formulae, $\langle F, G \rangle$ accepted if and only if $F(z_1, \ldots, z_n)$ is satisfiable and $G(y_1, \ldots, y_n)$ is unsatisfiable, is polynomially reducible to MUE. Given an instance $\langle F, G \rangle$ of SAT-UNSAT, form the VAF shown in Fig. 8

[7]It should be noted that MUE$_\emptyset$ is trivial for any $\langle \mathcal{X}, \mathcal{A}, \mathcal{V}, \eta \rangle$ in which $\langle \mathcal{X}, \mathcal{A} \rangle$ is acyclic: such frameworks have at least one argument, x, for which $\{x\}^- = \emptyset$. Any such argument is objectively accepted and, on its own, defines an admissible set.

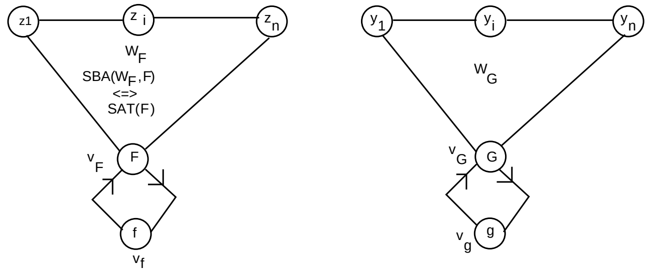

Figure 8. Reduction from SAT-UNSAT $\langle F, G\rangle$ to MUE

This consists of distinct copies of $\mathcal{D}_F$ and $\mathcal{D}_G$ as in the proof of part (b). Complete the instance of MUE by fixing $S = \{g\}$. The set S is a maximal uncontested extension if and only if F is satisfiable (so that $\mathcal{U}_F$ the maximal uncontested extension of $\mathcal{D}_F$ is empty) and G is unsatisfiable (so that $\mathcal{U}_G = \{g\}$).

□

Upper bounds for the problems UA, MUE$_\emptyset$ and MUE are immediate from the following result which proves that *constructing* the maximal uncontested extension can be performed in the complexity class $\text{FP}^{\text{NP}}_{||}$ of *function* computations realised by deterministic polynomial time algorithms that make a polynomially bounded number of *queries* to an NP oracle with all queries performed in parallel, i.e. non-adaptively so that the form of each query must be determined in advance of any invocation of the NP oracle.[8]

We require one preliminary result in proving this upper bound. Specifically the property of objectively accepted arguments w.r.t. $\mathcal{R}$ given in Lemma 3 and the consequence in Corollary 3. We note that although Lemma 3 is not difficult to derive, we are not aware of the property stated having been observed in earlier work.

Lemma 3 *Let $\langle \mathcal{X}, \mathcal{A}, \mathcal{V}, \eta\rangle$ be a* VAF*, $\mathcal{R} \subset \mathcal{V}\times\mathcal{V}$ an audience, and $\mathcal{X}^{\mathcal{R}}_{\text{OBA}}$ the set of objectively accepted arguments w.r.t. $\mathcal{R}$. The* AF *induced by $\mathcal{X}^{\mathcal{R}}_{\text{OBA}}$, i.e. $\langle \mathcal{X}^{\mathcal{R}}_{\text{OBA}}, \mathcal{E}\rangle$ in which $\mathcal{E} = \mathcal{A}\cap\{\langle x, y\rangle \;:\; x \in \mathcal{X}^{\mathcal{R}}_{\text{OBA}}$ and $y \in \mathcal{X}^{\mathcal{R}}_{\text{OBA}}\}$, is* acyclic.

Proof: Suppose the contrary holds and that for some VAF, $\langle \mathcal{X}, \mathcal{A}, \mathcal{V}, \eta\rangle$, and audience, $\mathcal{R}$, the AF $\langle \mathcal{X}^{\mathcal{R}}_{\text{OBA}}, \mathcal{E}\rangle$ contains a cycle, i.e. there is a sequence $x_1x_2\cdots x_r \in \mathcal{X}^r$ for which $x_1 = x_r$, and $\langle x_i, x_{i+1}\rangle \in \mathcal{E}$ for each $1 \le i \le r-1$. From the facts that $\langle x_i, x_{i+1}\rangle \in \mathcal{E}$ and $x_i \in \mathcal{X}^{\mathcal{R}}_{\text{OBA}}$ for each i, it must be the case that $\langle \eta(x_{i+1}), \eta(x_i)\rangle \in \mathcal{R}^*$, for otherwise $\eta(x_i) = \eta(x_{i+1})$ or there would be a specific audience $\alpha \in \chi(\mathcal{R})$ with $\eta(x_i) \succ_\alpha \eta(x_{i+1})$: in either case $\{x_i, x_{i+1}\} \not\subseteq P(\mathcal{H}^{(\mathcal{V})}, \alpha)$ contradicting the objective acceptability w.r.t $\mathcal{R}$ of both. From the assumption that a cycle exists and the constraints this imposes on $\mathcal{R}$, however, it follows that $\mathcal{R}^*$ forces, $\eta(x_r) \succ_{\mathcal{R}} \eta(x_{r-1}) \succ_{\mathcal{R}} \cdots \succ_{\mathcal{R}} \eta(x_2) \succ_{\mathcal{R}} \eta(x_1) = \eta(x_r)$. In this case, $\mathcal{R}$ could not be an audience. We deduce, therefore, that $\langle \mathcal{X}^{\mathcal{R}}_{\text{OBA}}, \mathcal{E}\rangle$ is acyclic. □

[8]For further background on $\text{FP}^{\text{NP}}_{||}$ and the related *decision* problem class, $\text{P}^{\text{NP}}_{||}$ we refer the reader to, e.g. [17, pp. 415–423], [15,19,20].

Corollary 3 *Let* $\mathcal{H}^{(\mathcal{V})}(\langle \mathcal{X}, \mathcal{A}, \mathcal{V}, \eta \rangle)$ *be a* VAF, $\mathcal{R}$ *an audience and* $\mathcal{P}(\mathcal{X}^{\mathcal{R}}_{\mathrm{OBA}})$ *the* unique *preferred extension of the* AF, $\langle \mathcal{X}^{\mathcal{R}}_{\mathrm{OBA}}, \mathcal{E} \rangle$ *defined in the statement of Lemma 3.*

$$\mathcal{U}_{\mathcal{R}}(\mathcal{H}^{(\mathcal{V})}) \subseteq \mathcal{P}(\mathcal{X}^{\mathcal{R}}_{\mathrm{OBA}})$$

Proof: Suppose to the contrary that there is an argument, $x_1 \in \mathcal{U}_{\mathcal{R}}$ for which $x_1 \notin \mathcal{P}(\mathcal{X}^{\mathcal{R}}_{\mathrm{OBA}})$. Since $\langle \mathcal{X}^{\mathcal{R}}_{\mathrm{OBA}}, \mathcal{E} \rangle$ is acyclic, the set $\mathcal{P}(\mathcal{X}^{\mathcal{R}}_{\mathrm{OBA}})$ constitutes a *stable* extension, and hence there is some $y_1 \in \mathcal{P}(\mathcal{X}^{\mathcal{R}}_{\mathrm{OBA}})$ such that $\langle y_1, x_1 \rangle \in \mathcal{E}$. Since $\mathcal{U}_{\mathcal{R}}$ defines an admissible subset of $\mathcal{X}^{\mathcal{R}}_{\mathrm{OBA}}$, there must be some $x_2 \in \mathcal{U}_{\mathcal{R}}$ for which $\langle x_2, y_1 \rangle \in \mathcal{A}$, furthermore it cannot be the case that $x_2 \in \mathcal{P}(\mathcal{X}^{\mathcal{R}}_{\mathrm{OBA}})$ as this set is conflict-free. The attack $\langle x_2, y_1 \rangle$ must also belong to $\mathcal{E}$ (from $\{x_2, y_1\} \subseteq \mathcal{X}^{\mathcal{R}}_{\mathrm{OBA}}$). Repeating this argument we identify subsets $\{x_1, x_2, \ldots, x_r\} \subseteq \mathcal{U}_{\mathcal{R}}$ and $\{y_1, y_2, \ldots, y_r\} \subseteq \mathcal{P}(\mathcal{X}^{\mathcal{R}}_{\mathrm{OBA}})$ for which $\langle y_i, x_i \rangle \in \mathcal{E}$ and $\langle x_{i+1}, y_i \rangle \in \mathcal{E}$. The sets $\mathcal{U}_{\mathcal{R}}$ and $\mathcal{P}(\mathcal{X}^{\mathcal{R}}_{\mathrm{OBA}})$ are both finite so the process described must eventually revisit some argument: but in this case we have identified a directed cycle in the AF, $\langle \mathcal{X}^{\mathcal{R}}_{\mathrm{OBA}}, \mathcal{E} \rangle$ contradicting Lemma 3. We deduce, therefore, that $\mathcal{U}_{\mathcal{R}} \subseteq \mathcal{P}(\mathcal{X}^{\mathcal{R}}_{\mathrm{OBA}})$ as claimed. □

Theorem 5 *Let* FMUE *be the (single-valued) function, defined as*

$$\mathrm{FMUE}(\mathcal{H}^{(\mathcal{V})}(\langle \mathcal{X}, \mathcal{A}, \mathcal{V}, \eta \rangle), \mathcal{R}) \quad =_{\mathrm{def}} \quad \mathcal{U}_{\mathcal{R}}(\mathcal{H}^{(\mathcal{V})})$$

i.e. given a VAF, $\mathcal{H}^{(\mathcal{V})}$*, and audience,* $\mathcal{R}$*, the function* FMUE *returns the maximal uncontested extension w.r.t* $\mathcal{R}$*:* FMUE *is* $\mathrm{FP}^{\mathrm{NP}}_{||}$*–complete.*

Proof: Hardness follows by considering the $\mathrm{FP}^{\mathrm{NP}}_{||}$–complete problem of computing, given a collection $\langle \Phi_1, \ldots, \Phi_m \rangle$ of CNF formulae, the bit sequence $\sigma_1\sigma_2 \ldots \sigma_m \in [0, 2^m - 1]$ corresponding to the characteristic function $\sigma_i = 1 \Leftrightarrow \mathrm{SAT}(\Phi_i)$: use m copies of the VAF described in Corollary 2(b) so that σ_i should be 1 if and only if f_i (the argument of $\mathcal{D}_{\Phi_i}$ associated with the corresponding CNF) is in the maximal uncontested extension.

Given a VAF, $\mathcal{H}^{(\mathcal{V})}(\langle \mathcal{X}, \mathcal{A}, \mathcal{V}, \eta \rangle)$, with argument set $\mathcal{X} = \{x_1, \ldots, x_n\}$ and audience $\mathcal{R}$ the sequence of binary values $\chi_1\chi_2\chi_3 \cdots \chi_n$ such that $\chi_j = 1$ if and only if $\neg\mathrm{OBA}(\mathcal{H}^{(\mathcal{V})}, x_j, \mathcal{R})$ can be determined with a single (length n) parallel query to an (NP) oracle for $\neg$OBA. We can thus compute the set, $\mathcal{X}^{\mathcal{R}}_{\mathrm{OBA}}$, of all objectively accepted arguments w.r.t. $\mathcal{R}$ in $\mathcal{H}^{(\mathcal{V})}$ via an $\mathrm{FP}^{\mathrm{NP}}_{||}$ algorithm: $\mathcal{X}^{\mathcal{R}}_{\mathrm{OBA}} = \{ x_i \ : \ \chi_i = 0\}$.

Now recalling from Corollary 3 that $\mathcal{U}_{\mathcal{R}} \subseteq \mathcal{P}(\mathcal{X}^{\mathcal{R}}_{\mathrm{OBA}})$, let $\mathcal{Y}^{\mathcal{R}}_{\mathrm{OBA}}$ denote the set $\mathcal{P}_(\mathcal{X}^{\mathcal{R}}_{\mathrm{OBA}})$ and consider the *bipartite* AF, $\mathcal{B}(\mathcal{Y}^{\mathcal{R}}_{\mathrm{OBA}}, \mathcal{X} \setminus \mathcal{X}^{\mathcal{R}}_{\mathrm{OBA}}, \mathcal{F})$ in which the set of attacks, $\mathcal{F}$, is

$$\mathcal{F} \quad = \quad \mathcal{A} \setminus \{ \ \langle x_i, x_j \rangle \ : \ x_i \notin \mathcal{X}^{\mathcal{R}}_{\mathrm{OBA}} \ \text{ and } x_j \notin \mathcal{X}^{\mathcal{R}}_{\mathrm{OBA}} \ \}$$

Note that the set of attacks $\mathcal{F}$ does indeed induce a bipartite graph on $(\mathcal{Y}^{\mathcal{R}}_{\mathrm{OBA}}, \mathcal{X} \setminus \mathcal{X}^{\mathcal{R}}_{\mathrm{OBA}})$ since $\mathcal{Y}^{\mathcal{R}}_{\mathrm{OBA}}$ must be a conflict-free set within the AF $\mathcal{H}(\mathcal{X}, \mathcal{A})$ underlying the VAF, $\mathcal{H}^{(\mathcal{V})}$.

The maximal uncontested extension, $\mathcal{U}_{\mathcal{R}}$, of $\mathcal{H}^{(\mathcal{V})}$, is the maximal admissible subset of $\mathcal{Y}^{\mathcal{R}}_{\mathrm{OBA}}$ within the AF, $\mathcal{H}(\mathcal{X}, \mathcal{A})$. In determining this subset, however, attacks involving only arguments outside $\mathcal{Y}^{\mathcal{R}}_{\mathrm{OBA}}$ are not relevant and thus the maximal admissible subset of $\mathcal{Y}^{\mathcal{R}}_{\mathrm{OBA}}$ considered with respect to the AF, $\mathcal{H}(\mathcal{X}, \mathcal{A})$ is identical to the maximal admissible subset of $\mathcal{Y}^{\mathcal{R}}_{\mathrm{OBA}}$ with respect to the bipartite framework $\mathcal{B}(\mathcal{Y}^{\mathcal{R}}_{\mathrm{OBA}}, \mathcal{X} \setminus \mathcal{X}^{\mathcal{R}}_{\mathrm{OBA}}, \mathcal{F})$. Applying the algorithm of Dunne [12, Thm. 6], given $\mathcal{B}(\mathcal{Y}^{\mathcal{R}}_{\mathrm{OBA}}, \mathcal{X} \setminus \mathcal{X}^{\mathcal{R}}_{\mathrm{OBA}}, \mathcal{F})$ this set can be identified by a (deterministic) polynomial time computation. It remains only to observe that $\mathcal{B}(\mathcal{Y}^{\mathcal{R}}_{\mathrm{OBA}}, \mathcal{X} \setminus \mathcal{X}^{\mathcal{R}}_{\mathrm{OBA}}, \mathcal{F})$ can be constructed in polynomial time from $\mathcal{H}^{(\mathcal{V})}$ given the set $\mathcal{X}^{\mathcal{R}}_{\mathrm{OBA}}$. □

Corollary 4 *The problems* UA*,* $\mathrm{MUE}_{\emptyset}$ *and* MUE *are all in* $\mathrm{P}^{\mathrm{NP}}_{||}$*.*

Proof: Given $\mathcal{H}^{(\mathcal{V})}$ its maximal uncontested extension, $\mathcal{U}^{\mathcal{H}}$ can be computed as described in Thm. 5. An instance $\langle \mathcal{H}^{(\mathcal{V})}, x\rangle$ of UA can be decided simply by checking $x \in \mathcal{U}^{\mathcal{H}}$; deciding an instance $\mathcal{H}^{(\mathcal{V})}$ of $\mathrm{MUE}_{\emptyset}$ involves testing $\mathcal{U}^{\mathcal{H}} = \emptyset$ and, similarly, checking the instance $\langle \mathcal{H}^{(\mathcal{V})}, S\rangle$ of MUE is carried out by testing $\mathcal{U}^{\mathcal{H}} = S$. □

The following provides a parallel for uncontested semantics to the properties of ideal semantics proven in [11, Thm. 6].

Theorem 6 *If* UA *is* NP*–hard then* UA *is* $\mathrm{P}^{\mathrm{NP}}_{||}$*-complete.*

Proof: We use the characterisation of $\mathrm{P}^{\mathrm{NP}}_{||}$–complete languages established by Chang and Kadin in [5, Thm. 9, p. 182], by which the result follows by showing UA to have the properties OR_{ω} and AND_{ω}, i.e. the languages $\mathrm{AND}_{\omega}(\mathrm{UA})$ and $\mathrm{OR}_{\omega}(\mathrm{UA})$ defined over m-tuples, $\mathbf{H}_m = \langle\langle \mathcal{H}^{(\mathcal{V})}_1, x_1, \mathcal{R}_1\rangle, \ldots, \langle \mathcal{H}^{(\mathcal{V})}_m, x_m, \mathcal{R}_m\rangle\rangle$ (for arbitrary $m \geq 1$) of distinct instances of UA by

$$\mathrm{AND}_{\omega}(\mathbf{H}_m) = \bigwedge_{i=1}^{m} \mathrm{UA}(\mathcal{H}^{(\mathcal{V})}_i, x_i) \quad ; \quad \mathrm{OR}_{\omega}(\mathbf{H}_m) = \bigvee_{i=1}^{m} \mathrm{UA}(\mathcal{H}^{(\mathcal{V})}_i, x_i)$$

are both polynomially reducible to UA.[9]

$\mathrm{AND}_{\omega}(\mathrm{UA}) \ \leq^p_m \ \mathrm{UA}$

Given $\langle\langle \mathcal{H}^{(\mathcal{V})}_1, x_1, \mathcal{R}_1\rangle, \ldots, \langle \mathcal{H}^{(\mathcal{V})}_m, x_m, \mathcal{R}_m\rangle\rangle$ an instance of $\mathrm{AND}_{\omega}(\mathrm{UA})$ construct the instance $\langle \mathcal{H}^{(\mathcal{V})}, z, \mathcal{R}\rangle$ of UA shown in Fig. 9 in which $\mathcal{R} = \cup_{i=1}^{m} \mathcal{R}_i$

In Fig. 9, $\{y_1, \ldots, y_m\}$ are new arguments, the value, $\eta(y_i)$ associated with y_i being that of its sole attacker x_i. The argument z has $\eta(z) = V_z$ with $V_z \notin \cup_{i=1}^{m} \mathcal{V}_{(i)}$ ($\mathcal{V}_{(i)}$ being the value set for the VAF $\mathcal{H}^{(\mathcal{V})}_i$). We claim that $\mathrm{UA}(\mathcal{H}^{(\mathcal{V})}, z, \mathcal{R})$ if and only if $\mathrm{UA}(\mathcal{H}^{(\mathcal{V})}_i, x_i, \mathcal{R}_i)$ for *all* $1 \leq i \leq m$.

[9]Note that, without loss of generality, it may be assumed that for $i \neq j$, $\mathcal{X}_i \cap \mathcal{X}_j = \mathcal{V}_i \cap \mathcal{V}_j = \emptyset$.

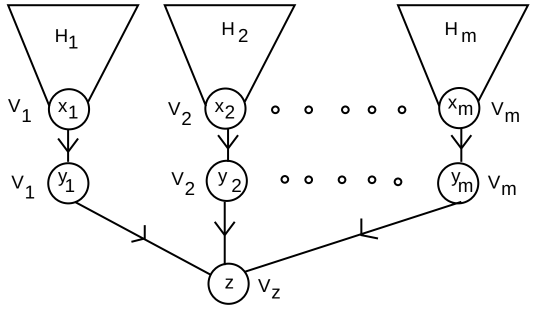

Figure 9. UA has property AND_ω

Suppose first that $\text{UA}(\mathcal{H}^{(\mathcal{V})}, z, \mathcal{R})$ and let $\mathcal{U}_\mathcal{R}$ be the maximal uncontested extension of $\mathcal{H}^{(\mathcal{V})}$ w.r.t. $\mathcal{R}$. In this case $z \in P(\mathcal{H}^{(\mathcal{V})}, \alpha, \mathcal{R})$ for every specific audience, $\alpha \in \chi(\mathcal{R})$, in particular for those specific audiences in which $V_i \succ_\alpha V_z$ for each V_i. It follows for each such audience, α, $\{x_1, \ldots, x_m\} \subseteq P(\mathcal{H}^{(\mathcal{V})}, \alpha)$ (otherwise for some i, no defence to the attack $\langle y_i, z\rangle$ would be present). We deduce that $\{z, x_1, \ldots, x_m\} \subseteq \mathcal{U}_\mathcal{R}$. In addition, however, defining $\mathcal{U}_{(i)} = \mathcal{U}_\mathcal{R} \cap \mathcal{X}_{(i)}$, we see that $\mathcal{U}_{(i)}$ is the maximal uncontested extension of $\mathcal{H}_i^{(\mathcal{V})}$ w.r.t. $\mathcal{R}_i$ and $x_i \in \mathcal{U}_{(i)}$, i.e. $\text{UA}(\mathcal{H}_i^{(\mathcal{V})}, x_i, \mathcal{R}_i)$ holds for each $1 \le i \le m$ as required.

Conversely, suppose that $\bigwedge_{i=1}^m \text{UA}(\mathcal{H}_i^{(\mathcal{V})}, x_i, \mathcal{R}_i)$ holds. Let $\mathcal{U}_{(i)}$ be the maximal uncontested extension of $\mathcal{H}_i^{(\mathcal{V})}$ w.r.t. $\mathcal{R}_i$ so that $x_i \in \mathcal{U}_{(i)}$. Consider the set of arguments $\mathcal{S} = \{z\} \cup_{i=1}^m \mathcal{U}_{(i)}$. Certainly $\mathcal{S}$ is admissible (with respect to the supporting AF of $\mathcal{H}^{(\mathcal{V})}$). Thus, to establish $\text{UA}(\mathcal{H}^{(\mathcal{V})}, z, \mathcal{R})$ it suffices to show $\text{OBA}(\mathcal{H}^{(\mathcal{V})}, z, \mathcal{R})$ given that $\text{OBA}(\mathcal{H}^{(\mathcal{V})}, x_i, \mathcal{R}_i)$ for each $1 \le i \le m$. It is certainly the case that $z \in P(\mathcal{H}^{(\mathcal{V})}, \alpha)$ for any audience satisfying $V_z \succ_\alpha V_i$ for each i. For the remaining audiences, should $V_i \succ_\alpha V_z$, then since $x_i \in P(\mathcal{H}_i^{(\mathcal{V})}, \alpha)$ (recall that $V_z \notin \mathcal{V}_{(i)}$), the attack $\langle y_i, z\rangle$ is countered by the attack $\langle x_i, y_i\rangle$. Thus $\text{OBA}(\mathcal{H}^{(\mathcal{V})}, z, \mathcal{R})$ holds and we deduce that $\mathcal{S}$ defines the maximal uncontested extension of $\mathcal{H}^{(\mathcal{V})}$, i.e. $\text{UA}(\mathcal{H}^{(\mathcal{V})}, z, \mathcal{R})$ holds as claimed.

$\text{OR}_\omega(\text{UA}) \ \le_m^p \ \text{UA}$

Given an m-tuple $\langle\langle\mathcal{H}_1^{(\mathcal{V})}, x_1, \mathcal{R}_1\rangle, \ldots, \langle\mathcal{H}_m^{(\mathcal{V})}, x_m, \mathcal{R}_m\rangle\rangle$ as an instance of $\text{OR}_\omega(\text{UA})$ construct the instance $\langle\mathcal{H}^{(\mathcal{V})}, z, \mathcal{R}\rangle$ of UA shown in Fig. 10 where, again $\mathcal{R} = \cup_{i=1}^m \mathcal{R}_i$

In this construction, $\{y_i, z_i \ : \ 1 \le i \le m\}$ together with $\{y, z\}$ are new arguments for which $\eta(y_i) = \eta(x_i) = V_i$ and the remaining new arguments are associated with a new value V_z. We claim that $\text{UA}(\mathcal{H}^{(\mathcal{V})}, z, \mathcal{R})$ if and only if at least one of $\text{UA}(\mathcal{H}_i^{(\mathcal{V})}, x_i, \mathcal{R}_i)$ holds.

Suppose first that $\text{UA}(\mathcal{H}^{(\mathcal{V})}, z, \mathcal{R})$ letting $\mathcal{U}_\mathcal{R}$ be the maximal uncontested extension so that $z \in \mathcal{U}_\mathcal{R}$. Since $\mathcal{U}_\mathcal{R}$ is an admissible set with $z \in \mathcal{U}_\mathcal{R}$ at least one of the arguments z_i must also belong to $\mathcal{U}_\mathcal{R}$ (in order to deal with the attack $\langle y, z\rangle$).

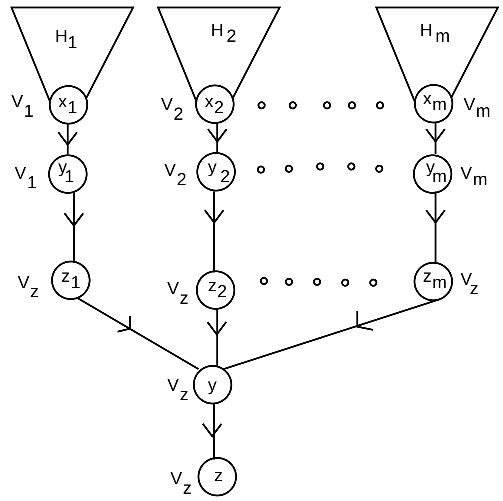

Figure 10. UA has property OR_ω

Similarly if $z_i \in \mathcal{U}_\mathcal{R}$ then $x_i \in \mathcal{U}_\mathcal{R}$ (otherwise the attack $\langle y_i, z_i\rangle$ is undefended). Hence from $\text{UA}(\mathcal{H}^{(\mathcal{V})}, z, \mathcal{R})$ we infer $\text{UA}(\mathcal{H}^{(\mathcal{V})}, x_i, \mathcal{R})$ (for some x_i) and now, noting that $V_z \notin \mathcal{V}_{(i)}$ and defining $\mathcal{U}_{(i)} = \mathcal{U}_\mathcal{R} \cap \cap \mathcal{X}_{(i)}$ we see that $\mathcal{U}_{(i)}$ defines the maximal uncontested extension of $\mathcal{H}_i^{(\mathcal{V})}$ w.r.t. $\mathcal{R}_i$ and contains x_i, i.e for some i, $\text{UA}(\mathcal{H}_i^{(\mathcal{V})}, x_i, \mathcal{R}_i)$ holds.

Conversely, without loss of generality, suppose that $\text{UA}(\mathcal{H}_1^{(\mathcal{V})}, x_1, \mathcal{R}_1)$ holds. The subset $\mathcal{U}_{(1)} \cup \{z_1, z\}$ is clearly admissible, so it suffices to show that $\text{OBA}(\mathcal{H}^{(\mathcal{V})}, z_1, \mathcal{R})$ since $\text{OBA}(\mathcal{H}^{(\mathcal{V})}, z, \mathcal{R})$ will follow from this. Certainly $z_1 \in P(\mathcal{H}^{(\mathcal{V})}, \alpha)$ for all audiences in which $V_z \succ_\alpha V_1$. For the remaining audiences (with $V_1 \succ_\alpha V_z$) from $\text{OBA}(\mathcal{H}_1^{(\mathcal{V})}, x_1, \mathcal{R}_1)$ we have $\text{OBA}(\mathcal{H}^{(\mathcal{V})}, x_1, \mathcal{R})$ and thus the attack $\langle y_1, z_1\rangle$ is countered by $\langle x_1, y_1\rangle$. Thus, $\text{OBA}(\mathcal{H}^{(\mathcal{V})}, z_1, \mathcal{R})$ so that $x_1 \in \mathcal{U}_{(1)}$ yields $\{z_1, z\} \subset \mathcal{U}_\mathcal{R}$, i.e. $\text{UA}(\mathcal{H}^{(\mathcal{V})}, z, \mathcal{R})$ as claimed. □

4. Conclusions and Further Development

This article presents an extension-based semantics for value-based argumentation frameworks which arises as a natural counterpart to the ideal semantics of Dung *et al.* [9,10] for the standard argumentation frameworks of [8]. It has been shown that although, in common with the ideal semantics, this form satisfies the property of defining a unique maximal extension w.r.t. any audience $\mathcal{R}$, nevertheless these give rise to significantly different behaviours: specifically, in general, these semantics are not coincident for a given VAF even in the case of the universal audience $\mathcal{R} = \emptyset$.

The motivation underlying our proposed formulation is in order to present one mechanism by which the nature of *objective acceptabililty* in VAFs may be further refined in the sense that those objectively accepted arguments falling outside the unique maximal uncontested extensions, although accepted by all relevant audiences, are so accepted on account of differing reasoning patterns germane to the audiences concerned. We have largely concentrated, in the present article, upon formal aspects of these semantics, in particular algorithmic and complexity-theoretic issues. These analyses represent only preliminary work and a number of questions remain unresolved forming the topic of work in progress. We conclude by briefly outlining some of these formal concerns.

One immediate issue is the gap between proven lower bounds (hardness) classifications and the $\mathrm{P}_{||}^{\mathrm{NP}}$ upper bound affecting several key decision questions. In the case of related questions in ideal semantics for AFs, Dunne [11, Cor. 7–10] deals with an apparently similar issue by establishing $\mathrm{P}_{||}^{\mathrm{NP}}$-hardness with respect to *randomized* reductions (as opposed to the more usual $\leq_m^p$ deterministic reducibility). It is unclear, however, whether the core construction used in the approach adopted in [11] can be translated into value-based settings: the proof that the so-called *Unique Satisfiability* problem is polynomially reducible to IA [11, Thm. 7]. Amongst other obstacles to such translations is the extensive use of Fact 3 in the correctness proof, a characterisation which, as we have seen in Lemma 1, fails to carry over to the uncontested semantics.

Further questions, which for reasons of space we have eschewed detailed treatment of in this paper, concern properties of *alternative* related semantics in VAFs. In particular, we have used the requirement that a subset, S of $\mathcal{X}_{\mathrm{OBA}}^{\mathcal{R}}$, qualifies as an uncontested extension of $\langle \mathcal{X}, \mathcal{A}, \mathcal{V}, \eta \rangle$ if S is admissible within the supporting AF, $\langle \mathcal{X}, \mathcal{A} \rangle$. There are, however, two other VAF based semantics defined in terms of subsets of $\mathcal{X}_{\mathrm{OBA}}^{\mathcal{R}}$ that may have properties of interest:

S1. $S \subseteq \mathcal{X}_{\mathrm{OBA}}^{\mathcal{R}}$ and S is admissible w.r.t. $\mathcal{R}$ in $\langle \mathcal{X}, \mathcal{A}, \mathcal{V}, \eta \rangle$.
S2. $S \subseteq \mathcal{X}_{\mathrm{OBA}}^{\mathcal{R}}$ and for all $\alpha \in \chi(\mathcal{R})$, S is admissible in the framework $\langle \mathcal{X}, \mathcal{A}_\alpha \rangle$ arising from $\langle \mathcal{X}, \mathcal{A}, \mathcal{V}, \eta \rangle$ by including only successful attacks w.r.t. α.

While it is straightforward to show that, for the universal audience, (S1) is equivalent to the uncontested semantics – i.e. S satisfies (S1) if and only if S is an uncontested extension – the behaviour described by (S2) is rather different.

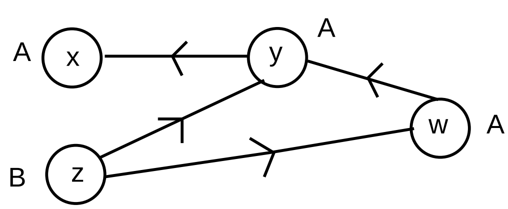

Figure 11. Uncontested semantics and (S2) semantics are distinct.

Thus, consider the example of Fig. 11, and the universal audience: $\mathcal{U} = \{x, z\}$, however, this set fails to satisfy the conditions imposed by (S2), since $\{x, z\}$ fails to be admissible w.r.t. the specific audience $A \succ B$. The maximal (S2)-extension is $\{z\}$ and, for any VAF and audience one may show that the maximal (S2)-extension w.r.t. $\mathcal{R}$ is unique and a subset of $\mathcal{U}_{\mathcal{R}}$. Further investigation of the properties of these semantics and their relationship to the uncontested semantics introduced in this paper forms the subject of current work in progress.

References

[1] K. Atkinson. Value-based argumentation for democratic decision support. In P. E. Dunne and T. J. M. Bench-Capon, editors, *Proc. 1st Int. Conf. on Computational Models of Argument*, volume 144 of *FAIA*, pages 47–58. IOS Press, 2006.

[2] K. Atkinson and T. J. M. Bench-Capon. Practical reasoning as presumptive argumentation using action based alternating transition systems. *Artificial Intelligence*, 2007. (in press), doi:10.1016/j.artint.2007.04.009.

[3] T. J. M. Bench-Capon. Persuasion in Practical Argument Using Value-based Argumentation Frameworks. *Journal of Logic and Computation*, 13(3):429–448, 2003.

[4] T. J. M. Bench-Capon, S. Doutre, and P. E. Dunne. Audiences in argumentation frameworks. *Artificial Intelligence*, 171:42–71, 2007.

[5] R. Chang and J. Kadin. On computing Boolean connectives of characteristic functions. *Math. Syst. Theory*, 28:173–198, 1995.

[6] R. Chang, J. Kadin, and P. Rohatgi. On unique satisfiability and the threshold behavior of randomised reductions. *Jnl. of Comp. and Syst. Sci.*, pages 359–373, 1995.

[7] Y. Dimopoulos and A. Torres. Graph theoretical structures in logic programs and default theories. *Theoretical Computer Science*, 170:209–244, 1996.

[8] 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.

[9] P. M. Dung, P. Mancarella, and F. Toni. A dialectical procedure for sceptical assumption-based argumentation. In P. E. Dunne and T. J. M. Bench-Capon, editors, *Proc. 1st Int. Conf. on Computational Models of Argument*, volume 144 of *FAIA*, pages 145–156. IOS Press, 2006.

[10] P. M. Dung, P. Mancarella, and F. Toni. Computing ideal sceptical argumentation. *Artificial Intelligence*, 2007. (in press), doi:10.1016/j.artint.2007.05.003.

[11] P. E. Dunne. The computational complexity of ideal semantics I: abstract argumentation frameworks. Technical Report ULCS-07-008, Dept. of Comp. Sci., Univ. of Liverpool, June 2007.

[12] P. E. Dunne. Computational properties of argument systems satisfying graph-theoretic constraints. *Artificial Intelligence*, 2007. (in press), doi:10.1016/j.artint.2007.03.006.

[13] P. E. Dunne and T. J. M. Bench-Capon. Coherence in finite argument systems. *Artificial Intelligence*, 141:187–203, 2002.

[14] P. E. Dunne and T. J. M. Bench-Capon. Complexity in value-based argument systems. In *Proc. 9th JELIA*, volume 3229 of *LNAI*, pages 360–371. Springer-Verlag, 2004.

[15] B. Jenner and J. Toran. Computing functions with parallel queries to NP. *Theoretical Computer Science*, 141:175–193, 1995.

[16] S. Modgil. Value based argumentation in hierarchical argumentation frameworks. In P. E. Dunne and T. J. M. Bench-Capon, editors, *Proc. 1st Int. Conf. on Computational Models of Argument*, volume 144 of *FAIA*, pages 297–308. IOS Press, 2006.

[17] C. H. Papadimitriou. *Computational Complexity*. Addison-Wesley, 1994.

[18] L. G. Valiant and V. V. Vazirani. NP is as easy as detecting unique solutions. *Theoretical Computer Science*, 47:85–93, 1986.

[19] K. Wagner. Bounded query computations. In *Proc. 3rd Conf. on Structure in Complexity Theory*, pages 260–277, 1988.

[20] K. Wagner. Bounded query classes. *SIAM Jnl. Comput.*, 19:833–846, 1990.