

An Multi Attack Argumentation Framework

Alexandros Vassiliades^{1,2}, Giorgos Flouris², Theodore Patkos², Antonis Bikakis³, Nick Bassiliades¹, and Dimitris Plexousakis²

¹ School of Informatics, Aristotle University of Thessaloniki

² Institute of Computer Science, Foundation for Research and Technology, Hellas

³ Department of Information Studies, University College London

{valexande, nbassili}@csd.auth.gr, {patkos, fgeo, dp}@ics.forth.gr,
a.bikakis@ucl.ac.uk

Abstract. This paper presents a novel abstract argumentation framework, called Multi-Attack Argumentation Framework (MAAF), which supports different types of attacks. The introduction of types gives rise to a new family of non-standard semantics which can support applications that classical approaches cannot, while also allowing classical semantics as a special case. The main novelty of the proposed semantics is the discrimination among two different roles that attacks play, namely an attack as a generator of conflicts, and an attack as a means to defend an argument. These two roles have traditionally been considered together in the argumentation literature. Allowing some attack types to serve one of those roles only, gives rise to the different semantics presented here.

1 Introduction

Many models for reasoning with arguments are grounded on Dung’s abstract argumentation framework (AAF) [8], where the only ingredients are a set of arguments and a binary attack relation on that set. The simplicity and intuitiveness of this model led to its wide acceptance, and, at the same time, helped reveal additional features that needed to be devised, in order to accommodate the requirements of diverse domains. For instance, AAFs are based on the strong assumption that all arguments and all attacks have the same strength; many variations have been proposed in order to overcome this limitation, which consider preferences on arguments [2] or attacks [11], add weights to arguments [1] or attacks [9], associate arguments with values [5], or impose hierarchies on arguments [12].

In this paper, we focus on accommodating a different need, namely to support a reasoning argumentation model where the attack relation among arguments can be of different types. This gives rise to a new class of semantics, which is based on a treatment of different attack types. Note that an attack, in the standard argumentation literature, is used both as a *conflict-generator* (i.e., creating conflicts which disallow conflicting arguments to be put in the same extension and which create the need for defense) and as a *defender* (i.e., defending other arguments against attacks). Separating these two roles, and allowing certain attacks to be treated in the non-classical way (e.g., allowing them to play only one

of these two roles), gives rise to various non-classical semantics which may find applications in different settings.

Consider, for example, the process of a trial. The presumption of innocence is so important that only specific types of attacks on the defendant’s claims should be taken into consideration by the jury, e.g., claims by eye witnesses, experts etc. Other attacks, e.g., on the defendant’s credibility, prior life choices etc., that generate doubts, should not be accounted for as evidence for conviction and need to be ignored. Interestingly, these latter types of attacks should still be considered relevant when placed against the claims of eye witnesses, experts etc, in order to ensure that the benefit of the doubt is given to the claimant. Figure 1 provides such an example. Intuitively, A_1 , A_3 and A_4 should all be acceptable; the latter two, because they do not receive any attacks so there is no reason not to accept them, and A_1 because questioning the credibility of the defendant (attack from A_3) should not be enough to lead to conviction. The same type of attack (i.e., credibility), however, should still be sufficient to defend the defendant’s claims from other attacks; in this case, for example, questioning the credibility of the witness (attack from A_4 to A_2) should be enough to defend the defendant from the witness’ testimony (attack from A_2 to A_1).

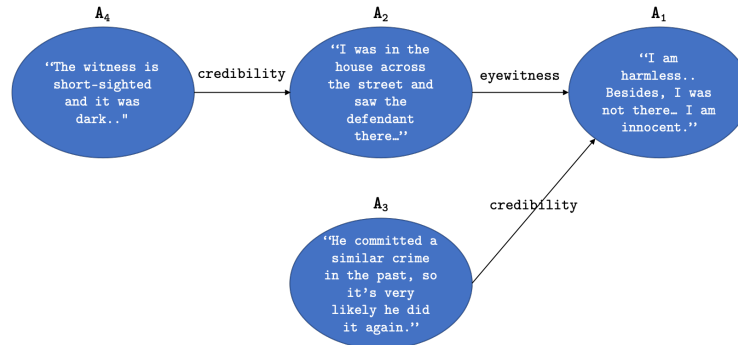


Fig. 1. Credibility attacks should be treated differently when they may lead to conviction than when they protect from it.

Notice that modeling different attack types is not always the same as modeling different argument types; while many of the existing AFs that characterize arguments and certain argument relations can indeed transfer this information to the attack relation, in the form of a strength value or a preference relation (see for example [7]), the inverse is not always possible. Attack types, such as rebuttals, undermines or undercuts, do not characterize the argument *per se*, but rather the relation between two arguments. Similarly, characterizing an attack as being of type irrelevant, i.e., arguing that a given argument is irrelevant in a given context, is not information inherent in the formulation of the argument, but, in a sense, on the placement of the argument in the argumentation tree.

Eliminating (filtering) attacks made by irrelevant arguments, before generating the sets of acceptable extensions of a dialogue, can therefore be considered a beneficial pre-processing step.

The objective of this paper is to provide a framework that supports the above scenarios, by introducing two important novelties. The first is the introduction of attack types, that allows treating different attacks in a different manner. The second is the separation among the two roles of attacks, namely as conflict-generators and as defenders. In particular, by seeing these two functions of attacks as separate, and allowing some of the attacks to be used for only one, or both, of these roles, we get three different semantics:

- *Loose semantics*, where certain attack types are used only as defenders, as e.g., in the case of credibility attacks in Figure 1.
- *Restricted semantics*, where certain attack types are ignored altogether (i.e., they have none of the two roles).
- *Firm semantics*, where certain attack types are used only as conflict-generators.

Following a related work analysis (Section 2), we formalise (Section 3) the three semantics explained above; the formalisation results to various types of extensions, in a manner similar to standard frameworks [8]. Then, we explore the properties of such semantics (Section 4), such as the existence of the different extensions, relationships among themselves and with the Dung semantics and others, and conclude in Section 5.

2 Related Work

The need to further refine the notion of attack in argumentation frameworks has led to several different extensions of Abstract Argumentation Frameworks. For example, Abstract Argumentation Frameworks with Recursive Attacks (AFRA) [3] and Extended Argumentation Frameworks (EAF) [12] extend the definition of attack, allowing attacks to be directed not only to arguments but also to other attacks. The difference between the two is that, while in EAFs only attacks whose target is an argument can be attacked, in AFRA any attack can be attacked. This idea is orthogonal to our approach that considers different types of attack, which are, however, all directed to arguments, and studying the combination of these two approaches, e.g. by allowing different types of attack that can be directed to arguments or attacks is an interesting research direction.

Commonsense Argumentation Frameworks [16], on the other hand, include two types of attacks, which differ in the type of arguments they are directed to, i.e. deductive arguments and commonsense arguments. They can therefore be considered as specializations of Multi-Attack Argumentation Frameworks, which we propose in this paper.

Some other studies have introduced weights or preferences on attacks following quantitative or qualitative approaches. For example, Weighted Argumentation Systems [9] assign weights to attacks as a way to describe their strength and use the idea of an inconsistency budget as a way to disregard attacks up

to a certain weight. The idea of weighted attacks is also used in [10], where the acceptability of arguments is not defined in terms of the standard Dung-style extensions, but in terms of numerical values derived from a set of equations describing the arguments and their attack relations. While social networks is indeed a domain where numerical weights can be derived from the reactions of the users, in many other domains (for example, the legal domain) such types of data may not be available.

A qualitative approach to represent preferences among attacks was proposed in [11]. Similarly to our approach, they define a framework with (an arbitrary number of) types of attack. These are partially ordered, and each attack is assigned one of these types. This allows for a finer grained definition of defence (compared to AAFs), which can roughly be described as follows: an argument is defended against an attack from a counter-argument, if the latter receives a stronger attack from another argument. It also allows for a finer definition of acceptability semantics, which take into account the relative difference of strength between defensive and offensive attacks.

All such preference-based approaches, which use either numerical values or priorities to represent the (relative) strength of attacks, have a common characteristic: any non-preferred attack is either ignored or invalidated. Our approach offers alternative ways to treat attacks, which take into account their roles in an argumentation system, i.e. whether they are used as offensive or defensive attacks. For example, according to the firm semantics, a defense is effective only if it is from an argument of a specific type, while according to the loose semantics, an offensive attack is effective if the attacker is of a specific type. Choosing the right semantics depends on the specific requirements and characteristics of the application domain.

Another approach that also considers different types of attack in abstract argumentation was proposed in [15]. The motivation is similar to ours, namely that each attack relation can represent a different criterion according to which the arguments can be evaluated one against another. The evaluation of arguments, however, is based on the aggregation of the different relations using methods from social choice theory, such as majority voting, and the use of the standard acceptability semantics in the aggregate argumentation framework. They do not, therefore, provide ways to treat certain criteria differently than others, which is one of the main characteristics of Multi-Attack Argumentation Frameworks.

Different types of attack are common in structured argumentation frameworks. For example, in *ASPIC*⁺ [13] arguments can be attacked in three different ways: on their uncertain premises (undermining), on their defeasible inferences (undercutting), or on the conclusions of their defeasible inferences (rebutting). Deductive argumentation [6] also supports different types of attack, which depend on the underlying logic. For example, choosing classical logic as the base logic provides seven different types of attack. The different types of attack in such frameworks are associated with the internal structure of arguments and cannot therefore be directly compared with Multi-Attack Argumentation Frameworks in which arguments are abstract. They can, however, easily be mapped

to the representation model of Multi-Attack Argumentation Frameworks, i.e. by mapping each of the different types of attack they support to a different attack type of MAAF. This mapping enables alternative ways to reason with structured arguments by treating differently the different types of attack, which may be meaningful in some domains.

3 Multi Attack Argumentation Frameworks (MAAFs)

We define a multi-attack argumentation framework as an argumentation framework where attacks are of multiple types. Formally:

Definition 1. A multi-attack argumentation framework (*MAAF for short*) is a tuple $\langle \mathcal{A}, \mathcal{T}, \mathcal{R} \rangle$, such that:

- \mathcal{A} is a set of arguments
- \mathcal{T} is a set of attack types
- $\mathcal{R} \subseteq \mathcal{A} \times \mathcal{A} \times \mathcal{T}$ is a set of *type-annotated attacks among arguments*

Note that \mathcal{A} and/or \mathcal{T} can be infinite, so \mathcal{R} can be infinite too. Intuitively an attack $(a, b, \tau) \in \mathcal{R}$ represents that a attacks b , and that the attack is of type τ . Note that the same two arguments may be related with attacks of different types, in which case each attack type is represented as a different triple in \mathcal{R} .

For any given set of types $\mathcal{T}_0 \subseteq \mathcal{T}$, we say that a attacks b w.r.t. \mathcal{T}_0 (denoted by $a \rightarrow_{\mathcal{T}_0} b$) if there exists $\tau \in \mathcal{T}_0$, such that $(a, b, \tau) \in \mathcal{R}$. For simplicity, we often write \rightarrow_{τ} to denote $\rightarrow_{\{\tau\}}$, and \rightarrow to denote $\rightarrow_{\mathcal{T}}$. We extend notation to sets of arguments, and, for $B, C \subseteq \mathcal{A}$, we write $B \rightarrow_{\mathcal{T}_0} C$ if and only if $\exists b \in B, c \in C$ such that $b \rightarrow_{\mathcal{T}_0} c$. For singleton sets, we often write $b \rightarrow_{\mathcal{T}_0} C$ and $B \rightarrow_{\mathcal{T}_0} c$ instead of $\{b\} \rightarrow_{\mathcal{T}_0} C$ and $B \rightarrow_{\mathcal{T}_0} \{c\}$, respectively.

The *restriction* of an MAAF to a specific set of types \mathcal{T}_0 is the AAF that is generated from the MAAF by considering only the attacks in \mathcal{T}_0 . Formally, given an MAAF $\langle \mathcal{A}, \mathcal{T}, \mathcal{R} \rangle$, the *restriction* of $\langle \mathcal{A}, \mathcal{T}, \mathcal{R} \rangle$ to \mathcal{T}_0 is an AAF $\langle \mathcal{A}', \mathcal{R}' \rangle$, where $\mathcal{A}' = \mathcal{A}$ and $\mathcal{R}' = \{(a, b) \mid (a, b, \tau) \in \mathcal{R} \text{ for some } \tau \in \mathcal{T}_0\}$.

The *flattening* of an MAAF is the AAF that is generated from the MAAF by ignoring types. Formally, for an MAAF $\mathcal{F} = \langle \mathcal{A}, \mathcal{T}, \mathcal{R} \rangle$, the *flattening* of \mathcal{F} is an AAF $\langle \mathcal{A}', \mathcal{R}' \rangle$, where $\mathcal{A}' = \mathcal{A}$ and $\mathcal{R}' = \{(a, b) \mid (a, b, \tau) \in \mathcal{R} \text{ for some } \tau \in \mathcal{T}\}$. Note that the flattening of \mathcal{F} is the same as the restriction of \mathcal{F} to \mathcal{T} .

3.1 Classes of extensions for MAAFs

To define MAAF extensions, we introduce three new classes of semantics: *firm*, *restricted* and *loose*. For each type of semantics defined in [8] (e.g., admissible, complete, etc), we define its counterpart for each class (e.g., firmly admissible, restrictedly stable, loosely complete, etc.). The three classes differ in how certain types of attack are considered. As already mentioned, the idea behind our semantics is the treatment of certain types of attacks as being conflict-generators only or attackers only. To do this, we consider a certain set of types, say \mathcal{T}_0 , which are treated in the “normal” manner. Different types of semantics can now result depending on the exact behaviour of the attacks in $\mathcal{T} \setminus \mathcal{T}_0$. In particular:

1. *Firm semantics* (e.g., admissible, complete etc) w.r.t. a certain set of attack types (say \mathcal{T}_0) requires a candidate extension to be defended against all types of attacks, and an attack can be defended only by attacks from \mathcal{T}_0 . In other words, attacks in \mathcal{T}_0 have the standard behaviour, but attacks in $\mathcal{T} \setminus \mathcal{T}_0$ act as conflict-generators only, not as defenders. We call them firm because, while they allow any type of argument to unleash offensive attacks, they only allow certain types of attack (those in \mathcal{T}_0) to defend an argument, making its defense more difficult.
2. *Restricted semantics* (e.g., admissible, complete etc) w.r.t. a certain set of attack types (say \mathcal{T}_0) require a candidate extension to be defended against attacks from \mathcal{T}_0 only, and an attack can be defended only by attacks from \mathcal{T}_0 . Thus, restricted semantics essentially consider only the attacks in \mathcal{T}_0 , both for the attacks and for defending against them, i.e., attacks in $\mathcal{T} \setminus \mathcal{T}_0$ are totally ignored. This brings them quite close to the notion of the restriction of an MAAF, a statement that will be made precise in Proposition 4.
3. *Loose semantics* (e.g., admissible, complete etc) w.r.t. a certain set of attack types (say \mathcal{T}_0) are the most “relaxed” ones, as they require a candidate extension to be defended only against attacks from \mathcal{T}_0 , while defense can happen by any type of attack. In other words, in loose semantics, attacks in $\mathcal{T} \setminus \mathcal{T}_0$ are treated as defenders only, and cannot generate attacks. Loose semantics allows attacks to be ignored, so they may result to extensions that are not defended against all attacks, specifically against attacks that are of types not in \mathcal{T}_0 .

In the following, we use shorthands to refer to the various types and classes of semantics. In particular, for the three classes of semantics, we use **fr** for firm, **re** for restricted, and **lo** for loose semantics. We also use θ as a catch-all variable that refers to any of these classes. Similarly, for types of extensions, we use **cf** for conflict-free, **ad** for admissible, **co** for complete, **pr** for preferred, **gr** for grounded, and **st** for stable. We also use σ as a catch-all variable to indicate any of these extension types. For example, we write **fr-co**-extension to refer to a firmly complete extension, and θ - σ -extension to refer to an extension of class θ and the type denoted by σ .

To formalise the above ideas, we first refine the notion of defense:

Definition 2. Consider an MAAF $\langle \mathcal{A}, \mathcal{T}, \mathcal{R} \rangle$, some $\mathcal{T}_0 \subseteq \mathcal{T}$, some $a \in \mathcal{A}$ and some set $\mathcal{E} \subseteq \mathcal{A}$. We define the notion of defense for the different classes of semantics as follows:

- \mathcal{E} firmly defends a (or **fr**-defends a) w.r.t. \mathcal{T}_0 if and only if $\mathcal{E} \rightarrow_{\mathcal{T}_0} b$ whenever $b \rightarrow a$
- \mathcal{E} restrictedly defends a (or **re**-defends a) w.r.t. \mathcal{T}_0 if and only if $\mathcal{E} \rightarrow_{\mathcal{T}_0} b$ whenever $b \rightarrow_{\mathcal{T}_0} a$
- \mathcal{E} loosely defends a (or **lo**-defends a) w.r.t. \mathcal{T}_0 if and only if $\mathcal{E} \rightarrow b$ whenever $b \rightarrow_{\mathcal{T}_0} a$

Figure 2 visualises the notion of defense for various cases.

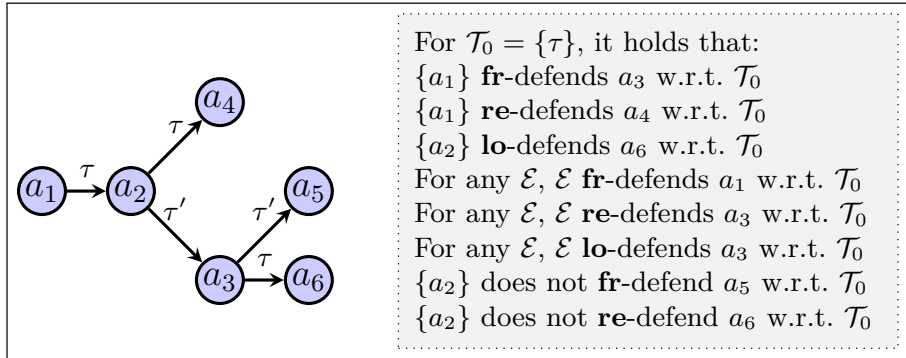


Fig. 2. The concept of **fr/re/lo**-defense visualised

3.2 Firm, restricted and loose extensions

Now we can recast the standard definitions for the different types of semantics given in [8], using the above ideas:

Definition 3. Consider an MAAF $\langle \mathcal{A}, \mathcal{T}, \mathcal{R} \rangle$ and some $\mathcal{T}_0 \subseteq \mathcal{T}$. A set $\mathcal{E} \subseteq \mathcal{A}$ is:

- Firmly conflict-free (**fr-cf**) w.r.t. \mathcal{T}_0 if and only if it is not the case that $\mathcal{E} \rightarrow \mathcal{E}$
- Restrictedly conflict-free (**re-cf**) w.r.t. \mathcal{T}_0 if and only if it is not the case that $\mathcal{E} \rightarrow_{\mathcal{T}_0} \mathcal{E}$
- Loose conflict-free (**lo-cf**) w.r.t. \mathcal{T}_0 if and only if it is not the case that $\mathcal{E} \rightarrow_{\mathcal{T}_0} \mathcal{E}$

Note how the intuition behind the different classes of semantics are applied in Definition 3: **lo-cf** and **re-cf** sets may include self-attacks, as long as they are not of types in \mathcal{T}_0 (because attacks in $\mathcal{T} \setminus \mathcal{T}_0$ are not conflict-generators in these semantics), whereas **fr-cf** sets cannot include any self-attack. As a result, the definition of **re-cf** and **lo-cf** coincides, since the notion of defense (where the two classes of semantics differ) is not relevant to that of conflict-freeness. Nevertheless, for purposes of uniformity and symmetry, we decided to include both definitions.

The same ideas are applied to admissible extensions, whose definition essentially mimics the ones typically used in AAFs, but considers the alternative notions of defense (Definition 2) for each case:

Definition 4. Consider an MAAF $\langle \mathcal{A}, \mathcal{T}, \mathcal{R} \rangle$ and some $\mathcal{T}_0 \subseteq \mathcal{T}$. For $\theta \in \{\mathbf{fr}, \mathbf{re}, \mathbf{lo}\}$, a set $\mathcal{E} \subseteq \mathcal{A}$ is a θ -**ad** extension w.r.t. \mathcal{T}_0 (in words: firmly/restrictedly/loosely admissible) if and only if:

- \mathcal{E} is θ -**cf**
- If $a \in \mathcal{E}$, then \mathcal{E} θ -defends a w.r.t. \mathcal{T}_0

Complete semantics' definition slightly deviates from the respective one in AAFs to accommodate the differences in the definition of conflict-freeness.

Definition 5. Consider an MAAF $\langle \mathcal{A}, \mathcal{T}, \mathcal{R} \rangle$ and some $\mathcal{T}_0 \subseteq \mathcal{T}$. For $\theta \in \{\mathbf{fr}, \mathbf{re}, \mathbf{lo}\}$, a set $\mathcal{E} \subseteq \mathcal{A}$ is a θ -**co** extension w.r.t. \mathcal{T}_0 (in words: firmly/restrictedly/loosely complete) if and only if:

- \mathcal{E} is θ -**ad**
- If \mathcal{E} θ -defends a w.r.t. \mathcal{T}_0 , and $\mathcal{E} \cup \{a\}$ is θ -**cf** w.r.t. \mathcal{T}_0 , then $a \in \mathcal{E}$

Note that, in the above definition, instead of only requiring that $a \in \mathcal{E}$ whenever \mathcal{E} θ -defends a , we have included the additional requirement that $\mathcal{E} \cup \{a\}$ is θ -**cf**, thereby deviating somewhat from the definition pattern used in AAFs for **co**-semantics [8]. This additional requirement is redundant in the AAF setting, because it results as a corollary of the weaker definition. The same is true in the MAAF setting, but only for the **fr** and **re** semantics (see Proposition 2 and the analysis that follows it). For this reason, and for purposes of uniformity and symmetry, we decided to include this extra requirement in Definition 5.

Grounded and preferred semantics are defined analogously:

Definition 6. Consider an MAAF $\langle \mathcal{A}, \mathcal{T}, \mathcal{R} \rangle$ and some $\mathcal{T}_0 \subseteq \mathcal{T}$. A set $\mathcal{E} \subseteq \mathcal{A}$ is a θ -**gr** extension w.r.t. \mathcal{T}_0 (in words: firmly/restrictedly/loosely grounded) if and only if \mathcal{E} is a minimal with respect to set inclusion θ -**co** extension w.r.t. \mathcal{T}_0 .

Definition 7. Consider an MAAF $\langle \mathcal{A}, \mathcal{T}, \mathcal{R} \rangle$ and some $\mathcal{T}_0 \subseteq \mathcal{T}$. A set $\mathcal{E} \subseteq \mathcal{A}$ is a θ -**pr** extension w.r.t. \mathcal{T}_0 (in words: firmly/restrictedly/loosely preferred) if and only if \mathcal{E} is a maximal with respect to set inclusion θ -**ad** extension w.r.t. \mathcal{T}_0 .

Stable semantics also follow a similar pattern:

Definition 8. Consider an MAAF $\langle \mathcal{A}, \mathcal{T}, \mathcal{R} \rangle$ and some $\mathcal{T}_0 \subseteq \mathcal{T}$. A set $\mathcal{E} \subseteq \mathcal{A}$ is:

- A firmly stable extension (**fr-st**) w.r.t. \mathcal{T}_0 if and only if:
 - \mathcal{E} is maximally **fr-cf** w.r.t. \mathcal{T}_0
 - $\mathcal{E} \rightarrow_{\mathcal{T}_0} a$ whenever $a \notin \mathcal{E}$
- A restrictedly stable extension (**re-st**) w.r.t. \mathcal{T}_0 if and only if:
 - \mathcal{E} is maximally **re-cf** w.r.t. \mathcal{T}_0
 - $\mathcal{E} \rightarrow_{\mathcal{T}_0} a$ whenever $a \notin \mathcal{E}$
- A loosely stable extension (**lo-st**) w.r.t. \mathcal{T}_0 if and only if:
 - \mathcal{E} is maximally **lo-cf** w.r.t. \mathcal{T}_0
 - $\mathcal{E} \rightarrow a$ whenever $a \notin \mathcal{E}$

Note that Definition 8 also deviates somewhat from the definition pattern of **st** semantics in standard AAFs. In particular, instead of requiring that \mathcal{E} is θ -**cf** (for the various θ), we have required that it is maximally θ -**cf**, i.e., a θ -**cf** set that is maximal among all other θ -**cf** sets. As with the **co** semantics, this stronger requirement is redundant in the AAF setting, and also in the MAAF setting for **fr** and **re** semantics (we have, however, included it in their definitions for uniformity), but is necessary for **lo** semantics (see Proposition 3 and the analysis that follows it).

4 Properties of MAAFs

We can show several properties with regards to the interplay among various types of θ - σ -extensions. To simplify presentation, all the following results assume an arbitrary MAAF $\mathcal{F} = \langle \mathcal{A}, \mathcal{T}, \mathcal{R} \rangle$ and some $\mathcal{T}_0 \subseteq \mathcal{T}$. Also, the reference to \mathcal{T}_0 is often omitted when obvious; e.g., we write that \mathcal{E} is a **lo-co** extension, to signify that \mathcal{E} is a **lo-co** extension w.r.t. \mathcal{T}_0 .

4.1 Initial results and special cases

We first show the analogous of Dung's fundamental lemma (Lemma 10 in [8]). Note the different formulation of this result for **lo** semantics¹:

Proposition 1. *For any given $\mathcal{E} \subseteq \mathcal{A}$, $a \in \mathcal{A}$, it holds that:*

1. *If \mathcal{E} is θ -ad, and \mathcal{E} θ -defends a , then $\mathcal{E} \cup \{a\}$ is θ -ad, for $\theta \in \{\mathbf{fr}, \mathbf{re}\}$*
2. *If \mathcal{E} is **lo-ad**, \mathcal{E} **lo-defends** a , and $\mathcal{E} \cup \{a\}$ is **lo-cf**, then $\mathcal{E} \cup \{a\}$ is **lo-ad***

Proposition 2 shows that the extra requirement of Definition 5 (compared to its counterpart in AAFs) is redundant for **fr** and **re** semantics:

Proposition 2. *For $\theta \in \{\mathbf{fr}, \mathbf{re}\}$, $\mathcal{E} \subseteq \mathcal{A}$, the following are equivalent:*

- *\mathcal{E} is a θ -co-extension w.r.t. \mathcal{T}_0*
- *The following hold for \mathcal{E} :*
 - *\mathcal{E} is θ -ad w.r.t. \mathcal{T}_0*
 - *If \mathcal{E} θ -defends a w.r.t. \mathcal{T}_0 , then $a \in \mathcal{E}$*

Note that the above equivalence does not hold for loose semantics. Indeed, the different formulation of Proposition 1 does not allow its use in the proof of Proposition 2. The MAAF visualised in Figure 3 provides a counter-example: $\{a, b\}$ is **lo-co**, despite the fact that $\{a, b\}$ **lo-defends** c and $c \notin \{a, b\}$. This is due to the extra requirement that we added in Definition 5; without it, neither $\{a, b\}$, nor $\{a, b, c\}$ would be **lo-co**, i.e., we would end up having a maximal **lo-ad** extension ($\{a, b\}$), that is not **lo-co**, which is against the intuition behind complete extensions.

Similarly, Proposition 3 shows that the extra requirement of Definition 8 (compared to its AAF counterpart) is redundant for **fr-st** and **re-st** semantics:

Proposition 3. *For $\theta \in \{\mathbf{fr}, \mathbf{re}\}$, $\mathcal{E} \subseteq \mathcal{A}$, the following are equivalent:*

- *\mathcal{E} is a θ -st-extension*
- *The following hold for \mathcal{E} :*
 - *\mathcal{E} is θ -cf*
 - *$\mathcal{E} \rightarrow_{\mathcal{T}_0} a$ whenever $a \notin \mathcal{E}$*




Fig. 3. Counter-example for the counterpart of Proposition 2 for **lo** semantics




Fig. 4. Counter-example for the counterpart of Proposition 3 for **lo** semantics

The example of Figure 4 shows a case where the counterpart of Proposition 3 would fail for **lo** semantics. The set $\{b\}$ attacks all other arguments, and is **lo-cf**, but not maximally so. Thus, it is not **lo-st**. On the contrary, $\{a, b\}$ is **lo-st**. This shows why the extra maximality condition that was added to Definition 8 (compared to its counterpart in AAFs) is necessary: without it, both $\{a, b\}$ and $\{b\}$ would be **lo-st**.

The next result shows that restricted semantics can be computed using the restriction of an MAAF:

Proposition 4. *Consider an MAAF $\mathcal{F} = \langle \mathcal{A}, \mathcal{T}, \mathcal{R} \rangle$ and some $\mathcal{T}_0 \subseteq \mathcal{T}$. Set $\mathcal{F}' = \langle \mathcal{A}', \mathcal{R}' \rangle$ the restriction of \mathcal{F} to \mathcal{T}_0 . For any given $\sigma \in \{\mathbf{cf}, \mathbf{ad}, \mathbf{co}, \mathbf{gr}, \mathbf{pr}, \mathbf{st}\}$ and $\mathcal{E} \subseteq \mathcal{A}$, \mathcal{E} is a **re- σ -extension** w.r.t. \mathcal{T}_0 if and only if \mathcal{E} is a σ -extension of \mathcal{F} .*

The following result describes a special case, showing essentially that our semantics is a generalisation of Dung's (i.e., that AAF semantics emerge as a special case of MAAFs):

Proposition 5. *Consider an MAAF $\mathcal{F} = \langle \mathcal{A}, \mathcal{T}, \mathcal{R} \rangle$ and set $\mathcal{T}_0 = \mathcal{T}$. Consider also the MAAF's flattening $\mathcal{F}' = \langle \mathcal{A}', \mathcal{R}' \rangle$, and its restriction to \mathcal{T}_0 , $\mathcal{F}'' = \langle \mathcal{A}'', \mathcal{R}'' \rangle$. Then, for any $\sigma \in \{\mathbf{cf}, \mathbf{ad}, \mathbf{co}, \mathbf{gr}, \mathbf{pr}, \mathbf{st}\}$, $\mathcal{E} \subseteq \mathcal{A}$ the following are equivalent:*

1. \mathcal{E} is a **lo- σ -extension** w.r.t. \mathcal{T}_0

¹ The proofs of all results appear in the Appendix.

2. \mathcal{E} is a **re**- σ -extension w.r.t. \mathcal{T}_0
3. \mathcal{E} is a **fr**- σ -extension w.r.t. \mathcal{T}_0
4. \mathcal{E} is a σ -extension of \mathcal{F}'
5. \mathcal{E} is a σ -extension of \mathcal{F}''

4.2 Relations among extension types, and existence results

The next proposition shows that the hierarchy of extensions that holds in the Dung setting, also holds for each class of extensions:

Proposition 6. *For any $\mathcal{E} \subseteq \mathcal{A}$:*

1. If \mathcal{E} is a θ -**ad**-extension w.r.t. \mathcal{T}_0 , then \mathcal{E} is a θ -**cf**-extension w.r.t. \mathcal{T}_0
2. If \mathcal{E} is a θ -**co**-extension w.r.t. \mathcal{T}_0 , then \mathcal{E} is a θ -**ad**-extension w.r.t. \mathcal{T}_0
3. If \mathcal{E} is a θ -**gr**-extension w.r.t. \mathcal{T}_0 , then \mathcal{E} is a θ -**co**-extension w.r.t. \mathcal{T}_0
4. If \mathcal{E} is a θ -**pr**-extension w.r.t. \mathcal{T}_0 , then \mathcal{E} is a θ -**co**-extension w.r.t. \mathcal{T}_0
5. If \mathcal{E} is a θ -**st**-extension w.r.t. \mathcal{T}_0 , then \mathcal{E} is a θ -**pr**-extension w.r.t. \mathcal{T}_0

Our next result shows that we can “incrementally” construct minimally-complete extensions starting from an **ad** one. The proof follows an iterative function, similar to the function F_{AF} used by Dung in [8]. However, for MAAFs, there are two subtleties.

First, F_{AF} (as defined in [8]) adds all acceptable arguments in each iteration; for the **lo** case, this could lead to a set that is not **lo-cf** (see, e.g., Figure 5: both b and c are acceptable by $\{a\}$, but $\{a, b, c\}$ is not **lo-cf**); thus, a more elaborate construction is needed.

Second, for infinite frameworks, the existence of a minimal fixpoint for F_{AF} (in [8]) is guaranteed by the implicit use of the Knaster-Tarski theorem ([14]), which requires an order preserving function. Although F_{AF} is order-preserving, our alternative is not.

To overcome these problems, the proof of Proposition 7 uses a more complex iterative function, employing ordinals. Importantly, this construction applies to all our semantics, as well as to standard AAFs, so it can be viewed also as an alternative proof for a well-known property of AAFs. Note also that the proof employs the Axiom of Choice.

Proposition 7. *Take any MAAF $\mathcal{F} = \langle \mathcal{A}, \mathcal{T}, \mathcal{R} \rangle$, some $\mathcal{T}_0 \subseteq \mathcal{T}$, and some $\mathcal{E}_* \subseteq \mathcal{A}$ such that \mathcal{E}_* is θ -**ad** (for $\theta \in \{\mathbf{fr}, \mathbf{re}, \mathbf{lo}\}$). Then, there exists some \mathcal{E} such that $\mathcal{E} \supseteq \mathcal{E}_*$, and the following hold:*

1. \mathcal{E} is θ -**co**.
2. For any \mathcal{E}' such that $\mathcal{E}_* \subseteq \mathcal{E}' \subset \mathcal{E}$, there exists $a \in \mathcal{E} \setminus \mathcal{E}'$ which is θ -defended by \mathcal{E}' and $\mathcal{E}' \cup \{a\}$ is θ -**cf**.
3. For any \mathcal{E}' such that $\mathcal{E}_* \subseteq \mathcal{E}' \subset \mathcal{E}$, \mathcal{E}' is not θ -**co**.

We next show that the existence of θ - σ extensions is guaranteed (except from θ -**st**), for all θ , analogously to the AAF case (see [8], [4]). Note that the proof for the infinite case in some of the semantics requires the Axiom of Choice:




Fig. 5. An MAAF with two **lo-gr** extensions

Proposition 8. For any MAAF $\mathcal{F} = \langle \mathcal{A}, \mathcal{T}, \mathcal{R} \rangle$, $\theta \in \{\mathbf{fr}, \mathbf{re}, \mathbf{lo}\}$, $\sigma \in \{\mathbf{cf}, \mathbf{ad}, \mathbf{co}, \mathbf{gr}, \mathbf{pr}\}$ and $\mathcal{T}_0 \subseteq \mathcal{T}$, there exists a θ - σ extension w.r.t. \mathcal{T}_0 in \mathcal{F} .

In AAFs, a **gr** extension is unique. The counter-example of Figure 5 shows that this is not the case for **lo-gr** extensions. However, for the other semantics (**fr**, **re**), the uniqueness of **gr** extensions is guaranteed:

Proposition 9. For any MAAF $\mathcal{F} = \langle \mathcal{A}, \mathcal{T}, \mathcal{R} \rangle$, $\theta \in \{\mathbf{fr}, \mathbf{re}\}$ and $\mathcal{T}_0 \subseteq \mathcal{T}$, there exists a unique θ -**gr** extension w.r.t. \mathcal{T}_0 in \mathcal{F} .




Fig. 6. Properties of MAAF extensions (apply to **fr**, **re**, **lo**, unless mentioned otherwise)

Propositions 6, 8 and 9 are summarised in Figure 6.

4.3 Relations among extension classes

The following propositions show the relation among **fr**, **re** and **lo** extensions, as well as the relation between these extensions and the extensions of the flattened AAF. This, along with Proposition 4, completes the picture with regards to the relationship among the different extension classes. We provide one proposition for each extension type (**cf**, **ad**, etc), starting with the simple case of **cf**:

Proposition 10. *Take an MAAF $\mathcal{F} = \langle \mathcal{A}, \mathcal{T}, \mathcal{R} \rangle$, its flattened AAF $\mathcal{F}_F = \langle \mathcal{A}, \mathcal{R}' \rangle$ and some $\mathcal{E} \subseteq \mathcal{A}$. Then:*

1. \mathcal{E} is **fr-cf** if and only if \mathcal{E} is **cf** in \mathcal{F}_F .
2. If \mathcal{E} is **fr-cf** then \mathcal{E} is **re-cf**.
3. \mathcal{E} is **re-cf** if and only if \mathcal{E} is **lo-cf**.

Interestingly, the direction of inference for the case of defense reverses (compared to the **cf** case) for the flattened AAF and the **re** class:

Proposition 11. *Take an MAAF $\mathcal{F} = \langle \mathcal{A}, \mathcal{T}, \mathcal{R} \rangle$, its flattened AAF $\mathcal{F}_F = \langle \mathcal{A}, \mathcal{R}_F \rangle$, some $\mathcal{E} \subseteq \mathcal{A}$ and some $a \in \mathcal{A}$. Then:*

1. If \mathcal{E} **fr-defends** a , then \mathcal{E} **re-defends** a .
2. If \mathcal{E} **re-defends** a , then \mathcal{E} **defends** a in \mathcal{F}_F .
3. If \mathcal{E} **defends** a in \mathcal{F}_F , then \mathcal{E} **lo-defends** a .

This reversal of the direction of inference (in Propositions 10, 11) leads to the following proposition:

Proposition 12. *Take an MAAF $\mathcal{F} = \langle \mathcal{A}, \mathcal{T}, \mathcal{R} \rangle$, its flattened AAF $\mathcal{F}_F = \langle \mathcal{A}, \mathcal{R}_F \rangle$ and some $\mathcal{E} \subseteq \mathcal{A}$. Then:*

1. If \mathcal{E} is **fr-ad**, then \mathcal{E} is **re-ad**.
2. If \mathcal{E} is **fr-ad**, then \mathcal{E} is **ad** in \mathcal{F}_F .
3. If \mathcal{E} is **re-ad**, then \mathcal{E} is **lo-ad**.
4. If \mathcal{E} is **ad** in \mathcal{F}_F , then \mathcal{E} is **lo-ad**.
5. If \mathcal{E} is **re-ad** and **cf** in \mathcal{F}_F , then \mathcal{E} is **ad** in \mathcal{F}_F .

For complete, grounded and preferred semantics, the situation is more complex:

Proposition 13. *Take an MAAF $\mathcal{F} = \langle \mathcal{A}, \mathcal{T}, \mathcal{R} \rangle$, its flattened AAF $\mathcal{F}_F = \langle \mathcal{A}, \mathcal{R}_F \rangle$ and some $\mathcal{E} \subseteq \mathcal{A}$. Then:*

1. If \mathcal{E} is **re-co** and **fr-ad**, then \mathcal{E} is **fr-co**.
2. If \mathcal{E} is **co** in \mathcal{F}_F and **re-ad**, then \mathcal{E} is **re-co**.
3. If \mathcal{E} is **lo-co** and **ad** in \mathcal{F}_F , then \mathcal{E} is **co** in \mathcal{F}_F .

Proposition 14. *Take an MAAF $\mathcal{F} = \langle \mathcal{A}, \mathcal{T}, \mathcal{R} \rangle$, its flattened AAF $\mathcal{F}_F = \langle \mathcal{A}, \mathcal{R}_F \rangle$ and some $\mathcal{E} \subseteq \mathcal{A}$. Then:*

1. If \mathcal{E} is **fr-gr** and **re-co**, then \mathcal{E} is **re-gr**.

2. If \mathcal{E} is **re-gr** and **co** in \mathcal{F}_F , then \mathcal{E} is **gr** in \mathcal{F}_F .
3. If \mathcal{E} is **gr** in \mathcal{F}_F and **lo-co**, then \mathcal{E} is **lo-gr**.

Proposition 15. *Take an MAAF $\mathcal{F} = \langle \mathcal{A}, \mathcal{T}, \mathcal{R} \rangle$, its flattened AAF $\mathcal{F}_F = \langle \mathcal{A}, \mathcal{R}_F \rangle$ and some $\mathcal{E} \subseteq \mathcal{A}$. Then:*

1. If \mathcal{E} is **re-pr** and **fr-ad**, then \mathcal{E} is **fr-pr**.
2. If \mathcal{E} is **pr** in \mathcal{F}_F and **re-ad**, then \mathcal{E} is **re-pr**.
3. If \mathcal{E} is **lo-pr** and **ad** in \mathcal{F}_F , then \mathcal{E} is **pr** in \mathcal{F}_F .

Finally, for stable semantics, the situation is similar to the case of admissible semantics:

Proposition 16. *Take an MAAF $\mathcal{F} = \langle \mathcal{A}, \mathcal{T}, \mathcal{R} \rangle$, its flattened AAF $\mathcal{F}_F = \langle \mathcal{A}, \mathcal{R}_F \rangle$ and some $\mathcal{E} \subseteq \mathcal{A}$. Then:*

1. If \mathcal{E} is **fr-st**, then \mathcal{E} is **re-st**.
2. If \mathcal{E} is **fr-st**, then \mathcal{E} is **st** in \mathcal{F}_F .
3. \mathcal{E} is **re-st**, if and only if \mathcal{E} is **lo-st**.
4. If \mathcal{E} is **re-st** and **cf** in \mathcal{F}_F , then \mathcal{E} is **st** in \mathcal{F}_F .

Further corollaries can be derived by combining the above results (Propositions 10, 11, 12, 13, 14, 15, 16) with Propositions 4 and 6, to connect the various types and classes of semantics among themselves, and with the semantics of restricted/flattened AAF. These are direct and omitted.

5 Discussion and Conclusion

In this paper we presented the semantics of multi-attack argumentation frameworks, i.e., frameworks which support multiple attack types among arguments. The important novelty of our semantics is the discrimination between two roles of attacks that have traditionally been considered inseparably: the role of conflict-generator, and the role of defender. The combination of these two aspects allowed us to define new classes of semantics, which model interesting real-life situations, have nice formal properties, and engulf standard models as a special case. An AAF cannot capture the aforementioned aspects to the extent that an MAAF does.

Note that, although MAAFs admit several types of attacks, during the computation of semantics, all attack types are split into two classes: those that are in \mathcal{T}_0 , and those that are not. Thus, we could define the same semantics by just allowing two different types. However, such a solution, albeit simpler, would have two disadvantages. The first is that it is intuitively better for the modeller to have several attack types, and then decide which ones are “normal” (to be placed in \mathcal{T}_0), and which ones are “special” (to be placed in $\mathcal{T} \setminus \mathcal{T}_0$). This approach has the additional advantage that the modeller can choose a different \mathcal{T}_0 depending on the application at hand. Second, our modelling allows more sophisticated semantics to be developed, e.g., by defining sets \mathcal{T}_1 , \mathcal{T}_2 and treating the attacks in \mathcal{T}_1 as defenders only, and attacks in \mathcal{T}_2 as conflict-generators only. This extension is part of our future work.

Acknowledgment

This project has received funding from the Hellenic Foundation for Research and Innovation (HFRI) and the General Secretariat for Research and Technology (GSRT), under grant agreement No 188.

References

1. Amgoud, L., Ben-Naim, J., Doder, D., Vesic, S.: Acceptability semantics for weighted argumentation frameworks. In: IJCAI-17. p. 56–62 (2017)
2. Amgoud, L., Vesic, S.: Rich preference-based argumentation frameworks. *International Journal of Approximate Reasoning* **55**(2), 585 – 606 (2014)
3. Baroni, P., Cerutti, F., Giacomin, M., Guida, G.: AFRA: Argumentation framework with recursive attacks. *IJAR* **52**(1), 19–37 (2011)
4. Baumann, R., Spanring, C.: Infinite argumentation frameworks. In: Eiter, T., Strass, H., Truszczyński, M., Woltran, S. (eds.) *Advances in Knowledge Representation, Logic Programming, and Abstract Argumentation: Essays Dedicated to Gerhard Brewka on the Occasion of His 60th Birthday*. pp. 281–295. Springer International Publishing (2015)
5. Bench-Capon, T.J.M.: Persuasion in Practical Argument Using Value-based Argumentation Frameworks. *Journal of Logic and Computation* **13**(3), 429–448 (2003)
6. Besnard, P., Hunter, A.: A logic-based theory of deductive arguments. *Artificial Intelligence* **128**(1-2), 203–235 (2001)
7. Cayrol, C., Lagasquie-Schiex, M.: From preferences over arguments to preferences over attacks in abstract argumentation: A comparative study. In: *25th IEEE International Conference on Tools with Artificial Intelligence*. pp. 588–595 (2013)
8. Dung, P.M.: On the acceptability of arguments and its fundamental role in non-monotonic reasoning, logic programming and n-person games. *Artificial Intelligence* **77**(2), 321–357 (1995)
9. Dunne, P.E., Hunter, A., McBurney, P., Parsons, S., Wooldridge, M.: Weighted argument systems: Basic definitions, algorithms, and complexity results. *Artificial Intelligence* **175**(2), 457–486 (Feb 2011)
10. Egilmez, S., Martins, J., Leite, J.: Extending social abstract argumentation with votes on attacks. In: *Theory and Applications of Formal Argumentation - Second International Workshop, TAFA*. vol. 8306, pp. 16–31. Springer (2013)
11. Martinez, D.C., Garcia, A.J., Simari, G.R.: An abstract argumentation framework with varied-strength attacks. In: *KR-08*. pp. 135–144 (2008)
12. Modgil, S.: Reasoning about preferences in argumentation frameworks. *Artificial intelligence* **173**(9-10), 901–934 (2009)
13. Modgil, S., Prakken, H.: A general account of argumentation with preferences. *Artificial Intelligence* **195**, 361–397 (2013)
14. Tarski, A.: A lattice-theoretical fixpoint theorem and its applications. *Pacific Journal of Mathematics* **5**(2), 285–309 (1955)
15. Tohmé, F., Bodanza, G., Simari, G.R.: Aggregation of attack relations: a social-choice theoretical analysis of defeasibility criteria. In: *FoIKS*. pp. 8–23 (2008)
16. Vassiliades, A., Patkos, T., Bikakis, A., Flouris, G., Bassiliades, N., Plexousakis, D.: Preliminary notions of arguments from commonsense knowledge. In: *11th Hellenic Conference on Artificial Intelligence*. pp. 211–214 (2020)

A Appendix

Proof of Proposition 1.

For the first result, $\mathcal{E} \cup \{a\}$ θ -defends $\mathcal{E} \cup \{a\}$, since, by our assumptions, \mathcal{E} θ -defends \mathcal{E} , and \mathcal{E} θ -defends a . So it suffices to show that, $\mathcal{E} \cup \{a\}$ is θ -**cf**.

Let us consider the case of firm semantics first. Suppose that $\mathcal{E} \cup \{a\}$ is not **fr-cf**. Then, there exist $a_1, a_2 \in \mathcal{E} \cup \{a\}$ such that $a_1 \rightarrow a_2$. We consider four cases, all of which lead to a contradiction, thus proving the point:

1. If $a_1, a_2 \in \mathcal{E}$, then \mathcal{E} is not **fr-cf**, a contradiction.
2. If $a_1 \in \mathcal{E}, a_2 = a$, then, since \mathcal{E} **fr**-defends a , it follows that there exists some $a_3 \in \mathcal{E}$ such that $a_3 \rightarrow_{\mathcal{T}_0} a_1$, a contradiction by case #1.
3. If $a_1 = a, a_2 \in \mathcal{E}$, then, since \mathcal{E} is an **fr-ad**-extension, it follows that there exists $a_3 \in \mathcal{E}$, such that $a_3 \rightarrow_{\mathcal{T}_0} a$, i.e., $a_3 \rightarrow a$, a contradiction by case #2.
4. If $a_1 = a_2 = a$, then, since \mathcal{E} **fr**-defends a , it follows that there exists some $a_3 \in \mathcal{E}$ such that $a_3 \rightarrow_{\mathcal{T}_0} a$, i.e., $a_3 \rightarrow a$, a contradiction by case #2.

The case of restricted semantics is completely analogous and omitted.

For the second result, using the same reasoning we note that $\mathcal{E} \cup \{a\}$ **lo**-defends $\mathcal{E} \cup \{a\}$. Given that $\mathcal{E} \cup \{a\}$ is **lo-cf** by our assumptions, the result follows. \square

Proof of Proposition 2.

By Proposition 1 when \mathcal{E} is θ -**ad**, and \mathcal{E} θ -defends a , then $\mathcal{E} \cup \{a\}$ is θ -**cf**, for $\theta \in \{\mathbf{fr}, \mathbf{re}\}$. The result then follows trivially. \square

Proof of Proposition 3.

It suffices to show that when \mathcal{E} is θ -**cf**, and $\mathcal{E} \rightarrow_{\mathcal{T}_0} a$ whenever $a \notin \mathcal{E}$, then \mathcal{E} is maximally θ -**cf**. Indeed, suppose that \mathcal{E}' is θ -**cf** and $\mathcal{E}' \supset \mathcal{E}$. Then, take some $a \in \mathcal{E}' \setminus \mathcal{E}$. By our hypothesis, $\mathcal{E} \rightarrow_{\mathcal{T}_0} a$, i.e., $\mathcal{E}' \rightarrow_{\mathcal{T}_0} \mathcal{E}'$, a contradiction by our hypothesis that \mathcal{E}' is θ -**cf**. \square

Proof of Proposition 4.

Since $\mathcal{A} = \mathcal{A}'$, take any $a, b \in \mathcal{A}$, $\mathcal{E} \subseteq \mathcal{A}$. Then, apparently:

- a attacks b in \mathcal{F}' if and only if $a \rightarrow_{\mathcal{T}_0} b$ in \mathcal{F}
- \mathcal{E} defends a in \mathcal{F}' if and only if \mathcal{E} **re**-defends a in \mathcal{F}

Using the above two statements and Propositions 2, 3 (necessary for the case of **co**- and **st**-extensions respectively), it is easy to show the result. \square

Proof of Proposition 5.

Since $\mathcal{T}_0 = \mathcal{T}$, we note that $a \rightarrow_{\mathcal{T}_0} b$ if and only if $a \rightarrow b$. The equivalence among #1, #2, #3 is then obvious by the respective definitions on θ -extensions. Moreover, the equivalence among #2 and #4 is obvious from Proposition 4, whereas the equivalence among #4 and #5 follows from the fact that $\mathcal{F}' = \mathcal{F}''$. \square

Proof of Proposition 6.

For **re** semantics, all results follow from Proposition 4 and the corresponding results on the AAF (e.g., [8]), so let us consider the case of **fr** and **lo** semantics.

#1, #2 and #3 are obvious by the respective definitions.

For #4, let $\theta \in \{\mathbf{fr}, \mathbf{lo}\}$, and take \mathcal{E} to be a θ -**pr**-extension. Then it is θ -**ad**. Suppose that it is not θ -**co**. Then, there is some $a \notin \mathcal{E}$, such that \mathcal{E} θ -defends a and $\mathcal{E} \cup \{a\}$ is θ -**cf**. But then, it is easy to see that $\mathcal{E} \cup \{a\}$ is θ -**ad**, which is a contradiction by the definition of θ -**pr**-extensions and the fact that $\mathcal{E} \cup \{a\} \supset \mathcal{E}$. For #5, let us consider the case of firm semantics first, and take \mathcal{E} to be an **fr-st**-extension. Then, it is **fr-cf** (and maximally so). We will show that it is also **fr-ad**. Indeed, take some $a, b \in \mathcal{A}$, such that $a \in \mathcal{E}$ and $b \rightarrow a$. Then $b \notin \mathcal{E}$ (since \mathcal{E} is **fr-cf**), thus $\mathcal{E} \rightarrow_{\tau_0} b$ (since \mathcal{E} is **fr-st**), which implies that \mathcal{E} **fr**-defends a . Thus, \mathcal{E} is also **fr-ad**. It is also maximal, because \mathcal{E} is maximally **fr-cf**. Therefore, \mathcal{E} is an **fr-pr**-extension.

For the **lo** case, take \mathcal{E} to be a **lo-st**-extension. Then, it is **lo-cf** (and maximally so). We will show that it is also **lo-ad**. Indeed, take some $a, b \in \mathcal{A}$, such that $a \in \mathcal{E}$ and $b \rightarrow_{\tau_0} a$. Then $b \notin \mathcal{E}$ (since \mathcal{E} is **lo-cf**), thus $\mathcal{E} \rightarrow b$ (since \mathcal{E} is **lo-st**), which implies that \mathcal{E} **lo**-defends a . Thus, \mathcal{E} is also **lo-ad**. It is also maximal, because \mathcal{E} is maximally **lo-cf**. Therefore, \mathcal{E} is a **lo-pr**-extension. \square

Proof of Proposition 7.

We will prove the claim constructively. First, we will describe a construction over \mathcal{F} , and then we will show that this construction generates some \mathcal{E} with the above properties. The proof is broken down in steps, represented as claims proved individually below. The last claim (Claim 5) shows the result.

Construction. We assume a well-order $<$ over \mathcal{A} (its existence is guaranteed by the Axiom of Choice). For a given set $E \subseteq \mathcal{A}$, we denote by $\min_{<} E$ the minimal element of E according to $<$.

Moreover, for $E \subseteq \mathcal{A}$, set $E^{\mathbf{v}} = \{a \in \mathcal{A} \setminus E \mid E: \theta\text{-defends } a, E \cup \{a\}: \theta\text{-cf}\}$, i.e., the arguments that are defended by E , and do not conflict with E .

We define the function: $\phi : 2^{\mathcal{A}} \mapsto 2^{\mathcal{A}}$ as follows:

$$\phi(E) = \begin{cases} E & , \text{ when } E^{\mathbf{v}} = \emptyset \\ E \cup \{\min_{<}(E^{\mathbf{v}})\} & , \text{ when } E^{\mathbf{v}} \neq \emptyset \end{cases}$$

Finally, we define a function \mathcal{G} recursively on the ordinals as follows:

$$\begin{aligned} \mathcal{G}(\beta) &= \mathcal{E}_* & , \text{ when } \beta = 0 \\ \mathcal{G}(\beta + 1) &= \phi(\mathcal{G}(\beta)) & , \text{ when } \beta \text{ is a successor ordinal} \\ \mathcal{G}(\beta) &= \bigcup \{\mathcal{G}(\gamma) \mid \gamma < \beta\} & , \text{ when } \beta \text{ is a limit ordinal} \end{aligned}$$

Claim 1. For two ordinals β, γ , if $\beta < \gamma$, then $\mathcal{G}(\beta) \subseteq \mathcal{G}(\gamma)$.

Proof of Claim 1. We will use transfinite induction on γ .

If $\gamma = 0$, then the result holds trivially as there is no β for which $\beta < \gamma$. Suppose that the result holds for all $\gamma < \delta$; we will show that it holds for $\gamma = \delta$.

If δ is a successor ordinal, then there exists some δ^- such that $\delta = \delta^- + 1$. Clearly, by the definition of \mathcal{G} and ϕ , $\mathcal{G}(\delta) \supseteq \mathcal{G}(\delta^-)$. Furthermore, by the inductive hypothesis, $\mathcal{G}(\delta^-) \supseteq \mathcal{G}(\beta)$, which shows the result.

If δ is a limit ordinal, then the result follows directly by the definition of \mathcal{G} . \circ

Claim 2. For any ordinals β , $\mathcal{G}(\beta) \supseteq \mathcal{E}_*$.

Proof of Claim 2. If $\beta = 0$ the result follows by the definition of \mathcal{G} . If $\beta > 0$, the result follows by Claim 1. \circ

Claim 3. For any ordinal β , $\mathcal{G}(\beta)$ is θ -**ad**.

Proof of Claim 3. We will use transfinite induction over β . For $\beta = 0$, the result follows by our assumption on \mathcal{E}_* . Now suppose that it holds for all $\beta < \gamma$. We will show that it holds for $\beta = \gamma$.

If γ is a successor ordinal, then take γ^- such that $\gamma = \gamma^- + 1$. Then, by definition, $\mathcal{G}(\gamma) = \phi(\mathcal{G}(\gamma^-))$. By the inductive hypothesis $\mathcal{G}(\gamma^-)$ is θ -**ad**. Moreover, by the definition of ϕ , $\phi(E)$ is θ -**ad** whenever E is θ -**ad**, so $\mathcal{G}(\gamma)$ is θ -**ad**.

If γ is a limit ordinal, then suppose that $\mathcal{G}(\gamma)$ is not θ -**cf**. Then, there exist $a_1, a_2 \in \mathcal{G}(\gamma)$ such that $\{a_1, a_2\}$ is not θ -**cf**, and, thus, there exist ordinals δ_1, δ_2 such that $\delta_1 < \gamma$, $\delta_2 < \gamma$, $a_1 \in \mathcal{G}(\delta_1)$, $a_2 \in \mathcal{G}(\delta_2)$. If $\delta_1 = \delta_2$ then $\mathcal{G}(\delta_1)$ is not θ -**cf**, a contradiction by the inductive hypothesis. If $\delta_1 < \delta_2$ then $\mathcal{G}(\delta_2) \supseteq \mathcal{G}(\delta_1)$ (by Claim 1), so $a_1, a_2 \in \mathcal{G}(\delta_2)$, a contradiction by the inductive hypothesis. The case of $\delta_2 < \delta_1$ is analogous. Thus, $\mathcal{G}(\gamma)$ is θ -**cf**.

Now consider some $a \in \mathcal{G}(\gamma)$. Then, by the definition of \mathcal{G} , there exists some $\delta < \gamma$ such that $a \in \mathcal{G}(\delta)$. Since $\mathcal{G}(\delta)$ is θ -**ad** by the inductive hypothesis, it follows that $\mathcal{G}(\delta)$ θ -defends a , so, given that $\mathcal{G}(\gamma) \supseteq \mathcal{G}(\delta)$ (Claim 1), we conclude that $\mathcal{G}(\gamma)$ θ -defends a . Thus, $\mathcal{G}(\gamma)$ is θ -**ad**. \circ

Claim 4. There exists ordinal β such that $\mathcal{G}(\beta) = \mathcal{G}(\beta + 1)$.

Proof of Claim 4. By Claim 1, we conclude that \mathcal{G} is an increasing function from the ordinals into 2^A . It cannot be strictly increasing, as if it were we would have an injective function from the ordinals into a set, violating Hartogs' lemma. Therefore the function must be eventually constant, so for some β , $\mathcal{G}(\beta) = \mathcal{G}(\beta + 1)$. \circ

Claim 5. There exists some \mathcal{E} such that $\mathcal{E} \supseteq \mathcal{E}_*$, and the following hold:

1. \mathcal{E} is θ -**co**.
2. For any \mathcal{E}' such that $\mathcal{E}_* \subseteq \mathcal{E}' \subset \mathcal{E}$, there exists $a \in \mathcal{E} \setminus \mathcal{E}'$ which is θ -defended by \mathcal{E}' and $\mathcal{E}' \cup \{a\}$ is θ -**cf**.
3. For any \mathcal{E}' such that $\mathcal{E}_* \subseteq \mathcal{E}' \subset \mathcal{E}$, \mathcal{E}' is not θ -**co**.

Proof of Claim 5. By Claim 4, there exists ordinal β such that $\mathcal{G}(\beta) = \mathcal{G}(\beta + 1)$. Set $\mathcal{E} = \mathcal{G}(\beta)$. By Claim 2, $\mathcal{E} \supseteq \mathcal{E}_*$, so it is an adequate choice. We will show that \mathcal{E} satisfies the required properties.

For the first result, note that by Claim 3, \mathcal{E} is θ -**ad**. Moreover, $\mathcal{E} = \mathcal{G}(\beta) = \mathcal{G}(\beta + 1) = \phi(\mathcal{G}(\beta)) = \phi(\mathcal{E})$, which implies that $\mathcal{E}^\bullet = \emptyset$, which, in tandem with the fact that \mathcal{E} is θ -**ad** leads to the conclusion that \mathcal{E} is θ -**co**.

For the second result, take some \mathcal{E}' such that $\mathcal{E}_* \subseteq \mathcal{E}' \subset \mathcal{E}$.

Set $S = \{\gamma \mid \mathcal{G}(\gamma) \not\subseteq \mathcal{E}'\}$. We observe that $\beta \in S$, so $S \neq \emptyset$. Set $\delta = \min_{<} S$. Obviously, $\delta = \beta$ or $\delta < \beta$.

If $\delta = 0$, then $\mathcal{G}(\delta) = \mathcal{E}_* \subseteq \mathcal{E}'$, a contradiction.

If δ is a successor ordinal, then take δ^- such that $\delta = \delta^- + 1$. Thus, $\mathcal{G}(\delta) = \phi(\mathcal{G}(\delta^-))$. By construction, $\mathcal{G}(\delta^-) \subseteq \mathcal{E}'$ and $\mathcal{G}(\delta) \not\subseteq \mathcal{E}'$, therefore $\mathcal{G}(\delta) = \mathcal{G}(\delta^-) \cup \{a\}$, for some a for which $\mathcal{G}(\delta^-)$ θ -defends a and $\mathcal{G}(\delta^-) \cup \{a\}$ is θ -**cf**. If $a \in \mathcal{E}'$, then $\mathcal{G}(\delta) \subseteq \mathcal{E}'$, a contradiction by the choice of δ , so $a \notin \mathcal{E}'$. Moreover, $a \in \mathcal{G}(\delta)$.

If $\delta = \beta$ then $\mathcal{G}(\delta) = \mathcal{E}$, so $a \in \mathcal{E}$. If $\delta < \beta$ then $a \in \mathcal{G}(\delta) \subseteq \mathcal{G}(\beta)$ (by Claim 1), so $a \in \mathcal{E}$. We conclude that $a \in \mathcal{E} \setminus \mathcal{E}'$. Thus, we have found some a with the required properties.

If δ is a limit ordinal, then, by the definition of δ , $\mathcal{G}(\delta') \subseteq \mathcal{E}'$ for all $\delta' < \delta$. Therefore, $\mathcal{G}(\delta) = \bigcup_{\delta' < \delta} \mathcal{G}(\delta') \subseteq \mathcal{E}'$, a contradiction by the choice of δ .

The third result follows from the second: indeed, as there exists $a \in \mathcal{E} \setminus \mathcal{E}'$ which is θ -defended by \mathcal{E}' and $\mathcal{E}' \cup \{a\}$ is θ -cf, it cannot be the case that \mathcal{E}' is θ -co. \square

Proof of Proposition 8.

For the case where $\theta = \mathbf{re}$, the proof follows directly by Proposition 4 and the related results from the AAF literature. So suppose that $\theta \in \{\mathbf{fr}, \mathbf{lo}\}$.

We first note that \emptyset is θ -cf and θ -ad w.r.t. \mathcal{T}_0 , so the claim is true for $\sigma \in \{\mathbf{cf}, \mathbf{ad}\}$.

Let us now turn our attention to the case where $\sigma = \mathbf{pr}$. Our proof follows the lines of the respective proof in [4]. Set $\mathcal{AD} = \{\mathcal{E} \mid \mathcal{E} \text{ is } \theta\text{-ad}\}$ ($\mathcal{AD} \neq \emptyset$, as shown above). We will show that, any \subseteq -chain $(\mathcal{E}_i)_{i \in I}$ in \mathcal{AD} possesses an upper bound. Indeed, set $\mathcal{E} = \bigcup \mathcal{E}_i$. Obviously $\mathcal{E} \supseteq \mathcal{E}_i$, so it is an upper bound; it remains to show that $\mathcal{E} \in \mathcal{AD}$, i.e., that \mathcal{E} is θ -ad.

Now suppose that \mathcal{E} is not θ -cf. Then there exist $a_1, a_2 \in \mathcal{E}$ that attack each other ($a \rightarrow b$ for $\theta = \mathbf{fr}$, $a \rightarrow_{\mathcal{T}_0} b$ for $\theta = \mathbf{lo}$). By the definition of \mathcal{E} , there exist $\mathcal{E}_i, \mathcal{E}_j$ such that $a_1 \in \mathcal{E}_i, a_2 \in \mathcal{E}_j$ for some $i, j \in I$. It is the case that $\mathcal{E}_i \subseteq \mathcal{E}_j$ or $\mathcal{E}_i \subseteq \mathcal{E}_j$, so suppose, without loss of generality, that $\mathcal{E}_i \subseteq \mathcal{E}_j$. Then $a_1, a_2 \in \mathcal{E}_j$, a contradiction, since \mathcal{E}_j is θ -ad (thus θ -cf). Thus, \mathcal{E} is θ -cf. It remains to show that \mathcal{E} defends all $a \in \mathcal{E}$. Indeed, take some $a \in \mathcal{E}$. Then, $a \in \mathcal{E}_i$ for some $i \in I$, and, thus \mathcal{E}_i θ -defends a , which implies that \mathcal{E} θ -defends a , since $\mathcal{E} \supseteq \mathcal{E}_i$. Thus, any \subseteq -chain $(\mathcal{E}_i)_{i \in I}$ in \mathcal{AD} possesses an upper bound, which, by Zorn's Lemma, implies that \mathcal{AD} has a maximal element, i.e., that there exists a θ -pr extension. By proposition 6, this implies that there exists a θ -co extension as well.

For θ -gr extensions, note that \emptyset is θ -ad, so applying Proposition 7 for $\mathcal{E}_* = \emptyset$ we ensure the existence of some \mathcal{E} which is minimally θ -co, i.e., \mathcal{E} is θ -gr. \square

Proof of Proposition 9.

Given that \emptyset is θ -ad, we can apply Proposition 7 for $\mathcal{E}_* = \emptyset$ to get some \mathcal{E} which is minimally θ -co, i.e., \mathcal{E} is θ -gr. Now suppose that there is a second θ -gr extension, say \mathcal{E}' ($\mathcal{E}' \neq \mathcal{E}$). Obviously, $\mathcal{E} \not\subseteq \mathcal{E}'$ and $\mathcal{E}' \not\subseteq \mathcal{E}$. Set $\mathcal{E}_0 = \mathcal{E} \cap \mathcal{E}'$. It follows that $\emptyset \subseteq \mathcal{E}_0 \subset \mathcal{E}$, so by Proposition 7 again there exists some $a \in \mathcal{E} \setminus \mathcal{E}_0$ which is θ -defended by \mathcal{E}_0 and $\mathcal{E}_0 \cup \{a\}$ θ -cf. Moreover, $\mathcal{E}_0 \subset \mathcal{E}'$, so a is θ -defended by \mathcal{E}' . Thus, \mathcal{E}' is θ -gr, thus θ -co, and also \mathcal{E}' θ -defends a , so by Proposition 2, $a \in \mathcal{E}'$, a contradiction by the choice of a . \square

Proof of Proposition 10.

The first case is direct from Definition 3 and the definition of \mathcal{F}_F . The second case is direct using proof by contradiction and the fact that $\mathcal{E} \rightarrow_{\mathcal{T}_0} \mathcal{E}$ implies $\mathcal{E} \rightarrow \mathcal{E}$. The third is direct from Definition 3. \square

Proof of Proposition 11.

The first case follows from the fact that $b \rightarrow_{\mathcal{T}_0} c$ implies that $b \rightarrow c$ for any $b, c \in \mathcal{A}$. For the second and third cases, note that $a \rightarrow b$ if and only if $(a, b) \in \mathcal{R}_F$,

and that $a \rightarrow_{\tau_0} b$ implies that $a \rightarrow b$. From these, and the definition of defense in AAFs and MAAF, the results follow easily. \square

Proof of Proposition 12.

The first four cases are direct from Propositions 10, 11. For the fifth case, note that, since \mathcal{E} is **re-ad**, it follows that for all $a \in \mathcal{E}$, \mathcal{E} **re-defends** a for τ_0 , and, thus, by Proposition 11, \mathcal{E} **defends** a in \mathcal{F}_F . Combining this with the fact that \mathcal{E} is **cf** in \mathcal{F}_F , we get the result. \square

Proof of Proposition 13.

For the first case, it suffices to show that, if \mathcal{E} **fr-defends** a w.r.t. \mathcal{T}_0 , then $a \in \mathcal{E}$. Indeed, if \mathcal{E} **fr-defends** a , then, by Proposition 11, \mathcal{E} **re-defends** a , so, given that \mathcal{E} is **re-co**, it follows that $a \in \mathcal{E}$. The proofs for the other cases are analogous. \square

Proof of Proposition 14.

For the first case, we note that \emptyset is **fr-ad**, so applying Proposition 7 for $\mathcal{E}_* = \emptyset$, we will get a **fr-co** extension (say \mathcal{E}) that is minimal among **fr-co** extensions, thus it is the (only) **fr-gr** extension of \mathcal{F} . By Proposition 7 again, we observe that, for any $\mathcal{E}' \subset \mathcal{E}$, there exists some $a \in \mathcal{E} \setminus \mathcal{E}'$ such that \mathcal{E}' **fr-defends** a , i.e., \mathcal{E}' **re-defends** a , i.e., \mathcal{E}' is not **re-co**. Thus, \mathcal{E} is **re-gr**.

The second case is totally analogous.

The third case uses a similar proof (and the same reasoning, except that the existence of a is guaranteed by the results in [8] (instead of Proposition 7). \square

Proof of Proposition 15.

For the first case, suppose that \mathcal{E} is not **fr-pr**. Then, there exists some $\mathcal{E}' \supset \mathcal{E}$ such that \mathcal{E}' is **fr-pr**. But then, \mathcal{E}' is **fr-ad** so (by Proposition 12) \mathcal{E}' is **re-ad**, a contradiction by the fact that \mathcal{E} is **re-pr**. The other cases are analogous. \square

Proof of Proposition 16.

For the first: observe that, by Proposition 10, \mathcal{E} is maximally **fr-cf** if and only if \mathcal{E} is maximally **re-cf**. Then, the result is obvious by Definition 8.

For the second: we obtain by Proposition 10 that \mathcal{E} is **cf** in \mathcal{F}_F . Also, since \mathcal{E} is **fr-st**, $\mathcal{E} \rightarrow_{\tau_0} a$ for all $a \notin \mathcal{E}$, thus $\mathcal{E} \rightarrow a$ in \mathcal{F}_F . We conclude that \mathcal{E} is **st** in \mathcal{F}_F .

For the third: we observe that, by Proposition 10, \mathcal{E} is maximally **re-cf** if and only if it is maximally **lo-cf**. Now take some $a \notin \mathcal{E}$. If \mathcal{E} is **re-st**, then $ext \rightarrow_{\tau_0} a$, so $\mathcal{E} \rightarrow a$, so \mathcal{E} is **lo-st**. If \mathcal{E} is **lo-st**, then $\mathcal{E} \rightarrow a$, and suppose that it is not the case that $\mathcal{E} \rightarrow_{\tau_0} a$. Then, $\mathcal{E} \cup \{a\} \supset \mathcal{E}$ and **lo-cf**, a contradiction.

For the fourth: since \mathcal{E} is **re-st**, we get that $\mathcal{E} \rightarrow_{\tau_0} a$ whenever $a \notin \mathcal{E}$, thus $\mathcal{E} \rightarrow a$ in \mathcal{F}_F for all $a \notin \mathcal{E}$, and \mathcal{E} is **cf** in \mathcal{F}_F by the hypothesis, so \mathcal{E} is **st** in \mathcal{F}_F . \square