On the realisability of weak argumentation semantics

Theodoros Doukas, Matthias Thimm
Artificial Intelligence Group, University of Hagen, Germany

Abstract

We analyse representatives of the class of non-admissible semantics, particularly the undisputed, strongly undisputed, weakly admissible, and weakly preferred semantics, in terms of *realisability* and their *signatures*. More specifically, we determine properties of the extension sets they produce, as well as structural features of frameworks that realise extension sets under these semantics. We describe two classes of extension sets for which we show that they are not realisable under undisputed semantics. We also identify approaches that have proven to be useful in the construction of frameworks that aim to realise given extension sets under classical semantics and transfer them to the non-admissible case. While a full characterisation of the signatures of the non-admissible semantics remains elusive, we provide plausible upper bounds and discuss the challenges in establishing concrete lower bounds.

1 Introduction

Abstract argumentation frameworks [10] model scenarios in which arguments are exchanged in the course of a dispute. Only the interdependencies between the arguments are considered; the arguments themselves are not assessed, neither in terms of their truthfulness nor their applicability to the topic of the debate. The subject of the modelling is merely the arguments as entities without futher internal structure, which attempt to invalidate each other through directed attacks. Accordingly, arguments and their interactions can be represented as a directed graph that contains the arguments as nodes and in which an edge represents an attack from one argument to another with the aim of invalidating the attacked argument. In this scenario, the question now arises as to which sets of arguments can be considered valid in the sense that they represent a justifiable point of view. Firstly, it seems imperative that a valid set of arguments is *conflict-free*, i. e., that it does not contain any arguments that attack other arguments in the set. To safeguard against attacks from the outside, one can further demand that all attacking arguments should themselves be counter-attacked from the selected set, leading to the notion of *admissibility* (we will provide formal definitions in Section 2). Based on the notion of admissibility, formal semantics to abstract argumentation frameworks, i. e., functions that assign to each argumentation framework a set of sets of arguments (*extensions*) as plausible outcomes, can be defined [10, 1].

One of the consequences arising from requiring admissibility of an extension has already been questioned in Dung's seminal paper [10, p.351], using the simple example from Figure 1: since argument b does not defend itself against the attack by a, the set $\{b\}$ is not considered admissible; however, since a already invalidates itself, the question is justified whether b should actually need to defend itself at all under these circumstances.

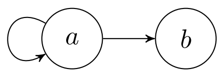

Figure 1: In this simple framework, b is attacked by a nonsensical attacker a.

This example, along with similar ones, inspired the development of so-called *weak* semantics. These aim to weaken the concept of admissibility to the point where attacks by arguments that are, in some sense, not considered serious, no longer need to be countered by extensions.

This new class of semantics includes those based on the concept of *weak admissibility* developed by Baumann, Brewka and Ulbricht [3, 4]; key ideas were initially introduced by Kakas and Mancarella [17]. It also includes semantics defined in terms of *undisputed sets* by Thimm [21]. Another approach is Dondio's also called "weakly admissible" semantics [9], which Baumann *et al.* compare along with other alternative approaches to their own [4]. While the various proposed semantics all address the same problem, albeit from different angles, they may well differ in the sets of extensions they produce in response to a particular argumentation scenario.

The discovery and development of an increasing number of semantics raises the need to systematically analyse and compare their properties. To this end, Baroni and Giacomin [2] enumerated various principles and studied the fulfilment of these principles by the individual semantics; this consideration was extended by Van der Torre and Vesic [22] to include further principles and semantics. Dvořák and Woltran [15] as well as Dvořák and Spanring [13] investigate the question whether there exist translations between frameworks so that for two semantics, the set of extensions of the original framework under the first semantics is identical or at least sufficiently similar to the set of extensions of the translated framework under the second semantics. Finally, Dunne *et al.* [12] explore *signatures* of semantics: they investigate properties that the extension sets of various semantics necessarily possess, and show that some of these properties are also sufficient to guarantee the existence of frameworks capable of realizing such sets of extensions. The study of signatures has so far mainly been carried out for semantics which are based on the notion of admissibility.

The present work aims to analyse realisability for weak argumentation semantics, by investigating properties of their extension sets and by assessing construction methods for frameworks that attempt to realise the extension sets of some of the representatives of this more novel class of semantics. We ask ourselves the same questions that Dunne *et al.* [12] have already answered for classical semantics: for which sets of argument sets are there frameworks that have these as their extensions? For this purpose, we describe properties that extension sets necessarily exhibit under weak semantics on the one hand, and frameworks that generate given sets of argument sets as extensions on the other. In doing so, we focus our attention on weak admissibility [3], weakly preferred semantics [3], undisputed semantics [21], as well as strongly undisputed semantics [21].

We first provide some necessary background on abstract argumentation, weak argumentation semantics and realisability in Section 2. We then consider *upper bounds* for signatures in Section 3, i. e., we identify necessary properties of the sets of extensions of the aforementioned semantics. Afterwards, in Section 4 we consider *lower bounds* for those signatures. We do this constructively by exploring different construction methods for argumentation frameworks that are able to produce a desired set of extensions. Despite our in-depth analysis, a full characterisation of the considered signatures remains open. We conclude this paper with a discussion on this and other aspects in Section 5.[1]

2 Preliminaries

This section recalls the necessary background on abstract argumentation (Section 2.1), weak argumentation semantics (Section 2.2), and realisability (Section 2.3).

2.1 Abstract argumentation

Following Dung [10], we start with a basic set $\mathfrak{A}$, whose elements $a \in \mathfrak{A}$ we call *arguments*.

Definition 1. An *argumentation framework* is a pair $F = (\mathcal{A}, \mathcal{R})$ consisting of a finite subset $\mathcal{A} \subseteq \mathfrak{A}$ and a relation $\mathcal{R} \subseteq \mathcal{A} \times \mathcal{A}$.

We call $\mathcal{R}$ the *attack relation* (on F); the set of all argumentation frameworks (over $\mathfrak{A}$) is denoted by $\mathfrak{F}_{\mathfrak{A}}$.

[1] Many of the examples given in this paper were computed automatically; the source code used in these computations is available at `https://github.com/tdoukas/aaf`.

For $F = (\mathcal{A}, \mathcal{R})$ and $(a, b) \in \mathcal{R}$, we write $a \rightarrow_{\mathcal{R}} b$, or shorter $a \rightarrow b$ if the reference to the relation $\mathcal{R}$ is unambiguous, and say that a *attacks* b (in F). A set $S \subseteq \mathcal{A}$ attacks an argument $a \in \mathcal{A}$ ($S \rightarrow a$) if any member $s \in S$ attacks a. Likewise, an argument $a \in \mathcal{A}$ is said to attack a set $S \subseteq \mathcal{A}$ if a attacks some $s \in S$; we then write $a \rightarrow S$. Extending the notation in the obvious way to two sets $S, T \subseteq \mathcal{A}$, we write $S \rightarrow T$ to denote that there are $s \in S, t \in T$ such that $s \rightarrow t$. Additionally, given a set $S \subseteq \mathcal{A}$, $S^+ = \{a \in \mathcal{A} \mid S \rightarrow a\}$ denotes the set of all arguments attacked by S, while $S^- = \{a \in \mathcal{A} \mid a \rightarrow S\}$ is the set of all attackers of S. Finally, we call $S^\oplus = S \cup S^+$ the *range of* S.

Definition 2. Let $(\mathcal{A}, \mathcal{R}) \in \mathfrak{F}_\mathfrak{A}$. A set $D \subseteq \mathcal{A}$ *defends* a set $S \subseteq \mathcal{A}$ if $S^- \subseteq D^+$; equivalently, for every attacker $a \in \mathcal{A}$ with $a \rightarrow S$, we have $D \rightarrow a$.

This concept of defence is fundamental to the following definition of admissibility.

Definition 3. Let $(\mathcal{A}, \mathcal{R}) \in \mathfrak{F}_\mathfrak{A}$ be an argumentation framework. A set $S \subseteq \mathcal{A}$ is *conflict-free* iff it does not contain two arguments $a, b \in S$ with $a \rightarrow b$; furthermore, S is *admissible* iff it is conflict-free and defends itself against any attacks from within $\mathcal{A}$.

Next, we introduce mappings from argumentation frameworks to sets of sets of arguments.

Definition 4. A *semantics* is a mapping $\sigma : \mathfrak{F}_\mathfrak{A} \rightarrow 2^{2^\mathfrak{A}}$ with $F = (\mathcal{A}, \mathcal{R}) \mapsto \sigma(F) \subseteq 2^\mathcal{A}$. For a given framework $F \in \mathfrak{F}_\mathfrak{A}$, we call $\sigma(F)$ the set of σ-*extensions* of F.

Certain semantics along with their extensions are given special names. The below defined semantics were described by Dung [10]; we refer to them as the *classical* semantics (we do no consider the grounded semantics here, though).

Definition 5. Let $F = (\mathcal{A}, \mathcal{R}) \in \mathfrak{F}_\mathfrak{A}$. An extension is said to be:

- *preferred*, iff it is a $\subseteq$-maximal admissible subset of $\mathcal{A}$;
- *complete*, iff it is admissible and includes all arguments that it defends;
- *stable*, iff it is conflict-free and attacks all arguments (of $\mathcal{A}$) it does not contain.

The designations also apply to the semantics themselves (e. g., the semantics that assigns to a framework the set of its stable extensions is called the *stable semantics*).

We write $\mathsf{cf}(F)$ resp. $\mathsf{adm}(F)$ to denote the set of conflict-free resp. admissible sets of F. Likewise we denote by $\mathsf{pr}(F)$, $\mathsf{com}(F)$, and $\mathsf{stb}(F)$ the sets of preferred, complete, and stable extensions of F, respectively.

Example 1. For the framework shown in Figure 2 (call it F), we find $\mathsf{stb}(F) = \{\{b, d\}\}$, $\mathsf{pr}(F) = \mathsf{stb}(F) \cup \{\{c, d\}\}$, $\mathsf{com}(F) = \mathsf{pr}(F) \cup \{\{d\}\}$, and $\mathsf{adm}(F) = \mathsf{com}(F) \cup \{\emptyset, \{b\}, \{c\}\}$.

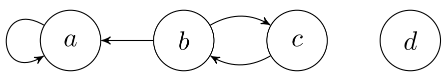

Figure 2: An example framework (F), for which $\emptyset \neq \mathsf{stb}(F) \subsetneq \mathsf{pr}(F) \subsetneq \mathsf{com}(F) \subsetneq \mathsf{adm}(F)$.

Finally, for a framework $F = (\mathcal{A}, \mathcal{R}) \in \mathfrak{F}_\mathfrak{A}$ and a subset of arguments $S \subseteq \mathcal{A}$, the *restriction of* F *to* S

$$F|_S = (S, \mathcal{R} \cap (S \times S))$$

consists only of the arguments of S and the attacks among them.

2.2 Weak argumentation semantics

In the following, we introduce representatives of a class of semantics that only require defence against attacks that are in a certain sense considered serious. The following definition is from [3].

Definition 6. Let $F = (\mathcal{A}, \mathcal{R}) \in \mathfrak{F}_{\mathfrak{A}}$ and let $S \subseteq \mathcal{A}$. The *reduct of F with respect to S* is the argumentation framework

$$F^S = F|_{\mathcal{A} \setminus S^{\oplus}} ,$$

that is, the "remainder" of F after removing the arguments of S, and all arguments attacked by S, along with any attacks they participate in.

Example 2. In the framework $F = (\mathcal{A}, \mathcal{R})$ from Figure 3, for $S = \{a\}$ we have $S^{\oplus} = \{a, c\}$, $\mathcal{A} \setminus S^{\oplus} = \{b\}$, and $F|_{A \setminus S^{\oplus}} = (\{b\}, \emptyset)$, while for $S' = \{b\}$, we have $S'^{\oplus} = \{a, b, c\}$, $\mathcal{A} \setminus S'^{\oplus} = \emptyset$, and $F|_{\mathcal{A} \setminus S'^{\oplus}} = (\emptyset, \emptyset)$.

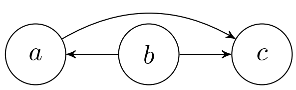

Figure 3: In this example, the reduct operation with respect to $\{b\}$ produces the empty framework.

Building on the concept of the reduct, *weak admissibility* [3] is defined as follows.

Definition 7. Let $F = (\mathcal{A}, \mathcal{R}) \in \mathfrak{F}_{\mathfrak{A}}$ and let $S \subseteq \mathcal{A}$. The set S is *weakly admissible* in F ($S \in \mathsf{wadm}(F)$) if it is both conflict-free and not attacked from any weakly admissible set of the reduct F^S, so that[2]

$$S \in \mathsf{cf}(F) \quad \text{and} \quad \forall a \in \mathcal{A} : \left(a \to S \implies a \notin \bigcup \mathsf{wadm}\left(F^S\right) \right).$$

We call S *weakly preferred* ($S \in \mathsf{wpr}(F)$) if it is a $\subseteq$-maximal weakly admissible extension.

Admissible sets defend all of their elements, leaving no attackers in the reduct, so that the body of Definition 7 holds trivially for admissible sets as well [4]:

Proposition 1. Admissibility implies weak admissibility. □

Example 3. Figure 4 shows a series of frameworks in which the argument b is attacked from an odd cycle of arguments. Consider for example $S = \{a_1, \ldots, a_5\}$ in the rightmost framework (call it F). Every non-empty, conflict-free subset of S (for example, $\{a_1, a_4\}$) has attackers (continuing the example, $\{a_3, a_5\}$) against which it does not completely defend (the attack by a_3 is undefended), so no non-empty subset of S is admissible; the same applies analogously to the other frameworks in Figure 4. We therefore argue that the attacks on b each time originate from a position that itself is unsustainable, and in this sense consider the attacks on b not serious.

Let us now check for weak admissibility of the attacked set $\{b\}$ in F. The reduct $F_1 = F^{\{b\}}$ consists of the odd cycle S and the attacks among S. Looking at possible attackers of $\{b\}$ such as the conflict-free subset $S' = \{a_1, a_4\} \subset S$, we want to assess weak admissibility of S' in F_1 and find that $F_1^{S'} = (\{a_3\}, \emptyset)$, in which $\{a_3\}$ is indeed admissible and attacks S' in F_1, so that S' cannot be weakly admissible in F_1 itself. The cases for other conflict-free subsets of S are analogous and consequently, $\{b\}$ is weakly admissible.

The recursive nature of Definition 7 suggests that reasoning under weakly admissible semantics is computationally challenging. Indeed, Dvořák, Ulbricht, and Woltran [14] show that all standard decision problems, such as the verification whether a given set is weakly admissible, are PSPACE-complete. An alternative approach that is designed to be more tractable computationally is that of undisputed sets [21], which we introduce in the following.

[2] For a set of sets $\mathbb{S}$, we intend $\bigcup \mathbb{S}$ to mean $\bigcup_{S \in \mathbb{S}} S$.

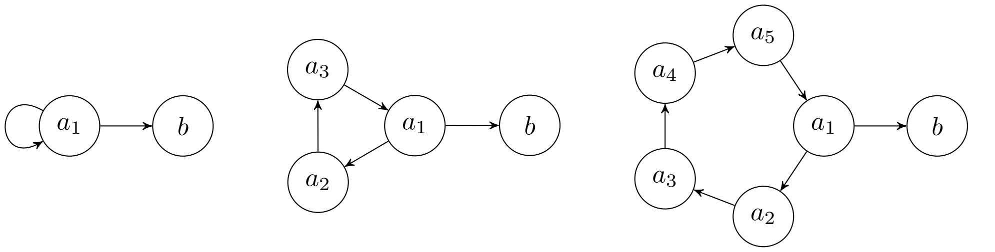

Figure 4: Odd cycles of length 1, 3, and 5 attack an argument b; none of the attacks on b are considered serious.

Definition 8 ([21])**.** An argumentation framework $F \in \mathfrak{F}_{\mathfrak{A}}$ is *vacuous with respect to a semantics* σ if the σ-extensions of F contain at most the empty set, i.e. $\sigma(F) \subseteq \{\emptyset\}$.

For a framework that is vacuous with respect to a given semantics, no subset of its arguments is acceptable under that semantics.

Definition 9 ([21])**.** Let $F = (\mathcal{A}, \mathcal{R}) \in \mathfrak{F}_{\mathfrak{A}}$ and let σ, τ be two arbitrary semantics. A set $S \subseteq \mathcal{A}$ is a σ^τ*-extension* if S is a σ-extension, and the reduct F^S is vacuous with respect to τ. A mapping $\mathfrak{F}_{\mathfrak{A}} \to 2^{2^{\mathfrak{A}}}$ is a σ^τ*-semantics* if it assigns to a framework the set of its σ^τ-extensions; we denote such a mapping simply by σ^τ.

We build on the previous definition and define the following concrete semantics.

Definition 10 ([21])**.** We call

- $\mathsf{ud} = \mathsf{cf}^{\mathsf{adm}}$ the *undisputed* semantics, and
- $\mathsf{sud} = \mathsf{cf}^{\mathsf{ud}}$ the *strongly undisputed* semantics.

It is worth noting that vacuous reduct semantics provide a template for describing other new as well as existing semantics (cf. also [5]).

Example 4. We return to Figure 4 and want to assess undisputedness of $\{b\}$ in the rightmost framework, which we call F. The set $\{b\}$ is conflict-free, and we have seen earlier (in Example 3) that the reduct $F^{\{b\}}$ contains no non-empty admissible sets, so $\{b\}$ is undisputed.

Finally, to provide some intuition into the distinction between undisputed and strongly undisputed semantics, consider the following example (from [21]).

Example 5. In the framework (F) in Figure 5, $\{c\}$ is undisputed since $\mathsf{adm}(F^{\{c\}}) = \{\emptyset\}$; we may argue however that b presents a valid and undefended attack on c since b is attacked only by the nonsensical a. In contrast, strongly undisputed semantics does not accept $\{c\}$, since in the reduct $F^{\{c\}} = (\{a, b\}, \{(a, a), (a, b)\})$, $\{b\}$ is undisputed.

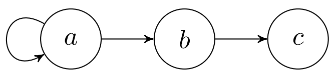

Figure 5: Here, $\{c\}$ is undisputed, but not strongly undisputed.

2.3 Realisability and signatures

The notion of *realisability* [12] is concerned with the expressivity of abstract argumentation semantics and is formally defined as follows.

Definition 11. Let $\mathbb{S} \subseteq 2^{\mathfrak{A}}$ be a collection of sets of arguments, and let $\sigma : \mathfrak{F}_{\mathfrak{A}} \to 2^{2^{\mathfrak{A}}}$ be a semantics. The set $\mathbb{S}$ is *realisable under* σ iff there is an argumentation framework F with $\sigma(F) = \mathbb{S}$.

The next definition captures the main topic of our consideration.

Definition 12. Let $\sigma : \mathfrak{F}_{\mathfrak{A}} \to 2^{2^{\mathfrak{A}}}$ be a semantics. The *signature* Σ_σ *of* σ is the set

$$\Sigma_\sigma = \{\sigma(F) \mid F \in \mathfrak{F}_{\mathfrak{A}}\}.$$

In other words, the signature Σ_σ of a semantics σ aggregates all extension sets that are realisable under σ. For some semantics, their signatures turn out to have exact characterisations; we reproduce some of the respective results of Dunne *et al.* [12] below. For that, we need some further notions.

Definition 13. Let $\mathbb{S} \subseteq 2^{\mathfrak{A}}$ be a collection of sets of arguments. The *downward-closure dcl*($\mathbb{S}$) *of* $\mathbb{S}$ is the set of all subsets of the individual elements of $\mathbb{S}$, i. e., $\text{dcl}(\mathbb{S}) = \{S \subseteq \bigcup \mathbb{S} \mid \exists S' \in \mathbb{S} : S \subseteq S'\}$.

Example 6. The downward-closure of $\mathbb{S} = \{\{a, b\}, \{a, c\}, \{b, c\}\}$ is $\text{dcl}(\mathbb{S}) = \{\emptyset, \{a\}, \{b\}, \{c\}\} \cup \mathbb{S}$.

Definition 14. Let $\mathbb{S} \subseteq 2^{\mathfrak{A}}$. We call $\mathbb{S}$:

- *downward-closed*, if $\text{dcl}(\mathbb{S}) = \mathbb{S}$;
- *incomparable*, if for all $S, S' \in \mathbb{S}$, $S \subseteq S'$ implies $S = S'$;
- *conflict-sensitive*, if for all $A, B \in \mathbb{S}$ with $A \cup B \notin \mathbb{S}$ there exist $a, b \in A \cup B$ so that $\{a, b\} \nsubseteq S$ for all $S \in \mathbb{S}$; i.e., there are two arguments in $A \cup B$ that never occur jointly in any set in $\mathbb{S}$;
- *tight*, if for all $A \in \mathbb{S}$ and $b \in \bigcup \mathbb{S}$ with $A \cup \{b\} \notin \mathbb{S}$ there exists $a \in A$ so that $\{a, b\} \nsubseteq S$ for all $S \in \mathbb{S}$; i.e., there is an argument in A that never occurs jointly with b in any extension of $\mathbb{S}$.

The following result from [12] follows straightforwardly.

Proposition 2. Tightness implies conflict-sensitivity. □

The reverse implication however does not hold.

Example 7. The set $\mathbb{S} = \{\emptyset, \{a\}, \{a, b\}\}$ is conflict-sensitive: for all $A, B \in \mathbb{S}$, we have $A \cup B \in \mathbb{S}$. It is not tight however since for $A = \emptyset$, we have $A \cup \{b\} = \{b\} \notin \mathbb{S}$, but since A is empty, there is no argument in A that never occurs jointly with b in any extension of $\mathbb{S}$.

Now some existing results on the characterisation of signatures are as follow [12].

Theorem 1. With respect to a set of arguments $\mathfrak{A}$, the following characterisations for semantics hold:

- $\Sigma_{\mathsf{cf}} = \{\mathbb{S} \subseteq 2^{\mathfrak{A}} \mid \mathbb{S}$ is non-empty, downward-closed, and tight$\}$;
- $\Sigma_{\mathsf{adm}} = \{\mathbb{S} \subseteq 2^{\mathfrak{A}} \mid \mathbb{S}$ contains $\emptyset$ and is conflict-sensitive$\}$;
- $\Sigma_{\mathsf{pr}} = \{\mathbb{S} \subseteq 2^{\mathfrak{A}} \mid \mathbb{S}$ is non-empty, incomparable, and conflict-sensitive$\}$;
- $\Sigma_{\mathsf{stb}} = \{\mathbb{S} \subseteq 2^{\mathfrak{A}} \mid \mathbb{S}$ is incomparable and tight$\}$. □

Some relationships between signatures of different semantics are summarised in the following result from [12].

Theorem 2. The following relations hold:

- $\Sigma_{\mathsf{cf}} \subsetneq \Sigma_{\mathsf{adm}}$;

- $\{\mathbb{S} \cup \{\emptyset\} \mid \mathbb{S} \in \Sigma_{\mathsf{pr}}\} \subsetneq \Sigma_{\mathsf{adm}}$. □

The question whether a certain extension set is realisable by a given semantics is generally answered constructively. In the following we briefly recapitulate framework construction methods from [12].

Definition 15. Let $\mathbb{S} \subseteq 2^{\mathfrak{A}}$ be a set of sets of arguments. The *Canonical Argumentation Framework* is the framework $F^{\mathsf{cf}}_{\mathbb{S}} = (\mathcal{A}, \mathcal{R})$ where

$$\mathcal{A} = \bigcup_{S \in \mathbb{S}} S$$

and

$$\mathcal{R} = \{(a, b) \mid a, b \in \mathcal{A} \text{ and } \forall S \in \mathbb{S} : \{a, b\} \nsubseteq S\}.$$

In $F^{\mathsf{cf}}_{\mathbb{S}}$, an argument a attacks an argument b if and only if a and b do not occur jointly in any of the sets of $\mathbb{S}$; this ensures that conflict-free semantics as well as certain derived semantics applied to $F^{\mathsf{cf}}_{\mathbb{S}}$ encompass at least all elements of $\mathbb{S}$. More concretely:

Proposition 3. Let $\sigma \in \{\mathsf{cf}, \mathsf{adm}, \mathsf{stb}\}$, and let $\mathbb{S} \in \Sigma_\sigma$. Then $\mathbb{S} \subseteq \sigma(F^{\mathsf{cf}}_{\mathbb{S}})$.

Proof. For $S \in \mathbb{S}$, no attacks are constructed between $a, b \in S$, so $S \in \mathsf{cf}(F^{\mathsf{cf}}_{\mathbb{S}})$. Since all attacks are reciprocal, S defends itself and thus $S \in \mathsf{adm}(F^{\mathsf{cf}}_{\mathbb{S}})$. Finally, assume $\mathbb{S}$ is tight, and let $S \in \mathbb{S}$, and $a \in \bigcup \mathbb{S}$ with $a \notin S$. Then one $s \in S$ must never occur jointly with a in any $S \in \mathbb{S}$; thus by construction of $F^{\mathsf{cf}}_{\mathbb{S}}$, $s \to a$, so $S \in \mathsf{stb}(F^{\mathsf{cf}}_{\mathbb{S}})$.[3] □

Preferred semantics is notably absent in Proposition 3; indeed, only a weaker statement holds in this case. We first give a counterexample.

Example 8. The extension set $\mathbb{S} = \{\{a, b, c\}, \{c, d, e\}, \{a, e\}\}$ is conflict-sensitive, since a and d, and b and e respectively, never occur jointly; it is incomparable, since no extension is a subset of another extension. We have

$$F^{\mathsf{cf}}_{\mathbb{S}} = (\{a, b, c, d, e\}, \{(a, d), (b, d), (b, e), (d, a), (d, b), (e, b)\}),$$

and $\mathsf{pr}(F^{\mathsf{cf}}_{\mathbb{S}}) = \{\{a, b, c\}, \{c, d, e\}, \{a, c, e\}\}$; this includes a superset of $\{a, e\}$, but not the extension itself.

Proposition 4. For $\mathbb{S} \in \Sigma_{\mathsf{pr}}$ holds:

$$\forall S \in \mathbb{S}\ \exists S' \in \mathsf{pr}(F^{\mathsf{cf}}_{\mathbb{S}}) : S \subseteq S'.$$

Proof. Proposition 3 has shown $S \in \mathsf{adm}(F^{\mathsf{cf}}_{\mathbb{S}})$, so $S \subseteq S'$ for some $S' \in \mathsf{pr}(F^{\mathsf{cf}}_{\mathbb{S}})$. □

$F^{\mathsf{cf}}_{\mathbb{S}}$ is then further modified to contain additional arguments and attacks so that admissible or preferred semantics produce exactly the elements of $\mathbb{S}$ as their extensions (as long as $\mathbb{S}$ has the adequate necessary properties), as defined below.

Definition 16. Let $\mathbb{S} \subseteq 2^{\mathfrak{A}}$ and let $F^{\mathsf{cf}}_{\mathbb{S}}$ be the canonical argumentation framework for $\mathbb{S}$. The *canonical defence framework* $F^{\mathsf{adm}}_{\mathbb{S}}$ results from applying the following modifications to $F^{\mathsf{cf}}_{\mathbb{S}} = (\mathcal{A}, \mathcal{R})$:

1. For every $a \in \mathcal{A}$, let $S^a_1, \ldots, S^a_n \in \mathbb{S}$ be the extensions that contain a, and define

$$\mathcal{D}^a = \{\{s_1, \ldots, s_n\} \mid s_i \in S^a_i \setminus \{a\}\}.$$

   Note that this implies $\mathcal{D}^a = \emptyset$ if, for any i, $S^a_i = \{a\}$.[4]

2. For each $D \in \mathcal{D}^a$, introduce a new argument z^a_D. This creates a total of $|\mathcal{D}^a| = (|S^a_1| - 1) \cdot \ldots \cdot (|S^a_n| - 1)$ new arguments.

[3] Dunne *et al.* [12] prove partial statements of this proposition in various places; we have summarised them here.

[4] No set $\{s_1, \ldots, s_n\}$ satisfies $s_i \in \emptyset$ for some $i \in 1, \ldots, n$.

3. Introduce attacks so that: (i) each z^a_D attacks itself; (ii) z^a_D is attacked by all $s \in D$; and (iii) z^a_D attacks a.

The purpose of the above construction is to defend the membership of each $a \in \mathcal{A}$ in each of the sets S^a_i. To this end, $\mathcal{D}^a$ collects all combinations of arguments that appear jointly with a in extensions of $\mathbb{S}$, and for each such combination $D \in \mathcal{D}^a$, an attacker z^a_D of a is introduced, that is self-defeating and against which the members of all extensions that contain a (except a itself) defend. For a conflict-sensitive $\mathbb{S}$ that contains the empty set, this is sufficient to single out exactly the members of $\mathbb{S}$ as the set of admissible extensions; or as the set of preferred extensions, if $\mathbb{S}$ is instead non-empty, conflict-sensitive, and incomparable [12]. To summarise, we recall the formal result from [12]:

Proposition 5. Let $\mathbb{S} \subseteq 2^{\mathfrak{A}}$, and let $F^{\mathsf{adm}}_{\mathbb{S}}$ be the canonical defence framework for $\mathbb{S}$. If $\mathbb{S}$ is conflict-sensitive and contains $\emptyset$, then $\mathsf{adm}(F^{\mathsf{adm}}_{\mathbb{S}}) = \mathbb{S}$. If $\mathbb{S}$ is non-empty, conflict-sensitive and incomparable, then $\mathsf{pr}(F^{\mathsf{adm}}_{\mathbb{S}}) = \mathbb{S}$. ☐

Example 9. Figure 6 shows the defence framework for $\mathbb{S} = \{\emptyset, \{a, b\}, \{b, c\}\}$. We have $\mathcal{D}^a = \{\{b\}\}$, $\mathcal{D}^b = \{\{a, c\}\}$, and $\mathcal{D}^c = \{\{b\}\}$, as well as $z^a_{\{b\}} = z_1$, $z^b_{\{a,c\}} = z_2$, and $z^c_{\{b\}} = z_3$.

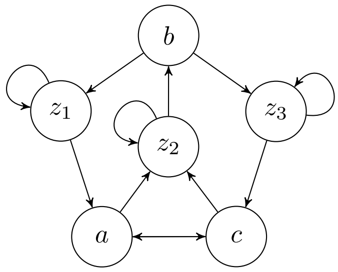

Figure 6: The defence framework $F^{\mathsf{adm}}_{\mathbb{S}}$ for $\mathbb{S} = \{\emptyset, \{a, b\}, \{b, c\}\}$.

In Definition 15 we introduced attacks between arguments in order to prevent them to occur jointly in any extension of the constructed framework. In view of later applications, we would like to point out that this is indeed the only mechanism that is able to separate extensions under admissibility.

Proposition 6. Let $F \in \mathfrak{F}_{\mathfrak{A}}$, and let $X, Y \in \mathsf{adm}(F)$ so that $X \cup Y \notin \mathsf{adm}(F)$. Then $X \cup Y \notin \mathsf{cf}(F)$.

Proof. $X \cup Y$ defends itself against all attacks directed towards either X or Y. If no attack originates from within $X \cup Y$, then $X \cup Y$ is already admissible. ☐

The next construction method aims at stable semantics. The method starts from the canonical argumentation framework and subsequently excludes unwanted extensions.

Definition 17. Let $\mathbb{S} \subseteq 2^{\mathfrak{A}}$ be a tight and incomparable extension set and let $F^{\mathsf{cf}}_{\mathbb{S}}$ be its canonical argumentation framework. The *stable canonical framework* $F^{\mathsf{stb}}_{\mathbb{S}}$ results from applying the following modification to $F^{\mathsf{cf}}_{\mathbb{S}}$:

1. Compute $\mathbb{X} = \mathsf{stb}(F^{\mathsf{cf}}_{\mathbb{S}}) \setminus \mathbb{S}$.
2. For each $X \in \mathbb{X}$, introduce a single self-attacking argument z_X, and for each $a \in (\bigcup \mathbb{S}) \setminus X$, add an attack $a \to z_X$.

This construction excludes all unwanted sets $X \in \mathbb{X}$ from the stable extensions, because they do not attack the argument z_X which they do not contain (compare Definition 5), in contrast to the desired sets in $\mathbb{S}$. Additionally, z_X is not eligible for inclusion in an extension since it defeats itself. This is summarised in the following proposition from [12].

Proposition 7. Let $\mathbb{S} \subseteq 2^{\mathfrak{A}}$ be incomparable and tight, and let $F^{\mathsf{stb}}_{\mathbb{S}}$ be the stable canonical framework for $\mathbb{S}$. Then, $\mathsf{stb}(F^{\mathsf{stb}}_{\mathbb{S}}) = \mathbb{S}$. ☐

Further construction ideas are presented by Dvořák and Woltran [15] as well as Dvořák and Spanring [13], this time in the context of translations between frameworks, which (roughly speaking) are intended to preserve extensions across different semantics.

Definition 18. A *framework translation* is a mapping $\mathfrak{F}_{\mathfrak{A}} \to \mathfrak{F}_{\mathfrak{A}}$. With respect to two semantics $\sigma, \sigma' : \mathfrak{F}_{\mathfrak{A}} \to 2^{2^{\mathfrak{A}}}$ and an arbitrary argumentation framework $F = (\mathcal{A}, \mathcal{R}) \in \mathfrak{F}_{\mathfrak{A}}$, a framework translation T is called *exact* if $\sigma(F) = \sigma'(T(F))$.

For example, any given framework $F = (\mathcal{A}, \mathcal{R}) \in \mathfrak{F}_{\mathfrak{A}}$ can be exactly translated to a new framework F' so that $\mathsf{adm}(F) = \mathsf{com}(F')$, as follows [15]: (i) for each argument $a \in \mathcal{A}$, add a new argument z_a; (ii) add reciprocal attacks between a and z_a; and (iii) make z_a attack itself. The following concrete example illustrates this.

Example 10. The admissible sets of the left framework in Figure 7 are identical to the complete extensions of framework to the right, namely $\{\emptyset, \{a\}, \{c\}, \{a, c\}\}$.

Figure 7: A translation between two frameworks that preserves the extension set from admissible to complete semantics.

The above construction exploits the fact that complete extensions are required to contain all arguments they defend: since every original argument has its own attacker which only the original argument defends against, the admissible sets of the original framework are exactly the complete extensions in the translation.

We are interested in framework translations primarily in order to see whether the constructions contain ideas that may also be useful for the non-admissible case. In doing so, we are particularly looking at exact translations, since they need to make sure that extensions are shaped exactly as required, often by employing mechanisms to filter out undesired elements. Such mechanisms are indeed at work in all of the constructions we have investigated so far, as we point out next.

Observation 1. The classical framework construction methods studied so far exhibit the following features.

- *A mechanism to separate extensions,* as employed in the construction of the canonical argumentation framework, whose conflict-free extensions exactly match its admissible extensions.
- *The use of a base framework and its subsequent customization:* in the case of the canonical frameworks, the canonical argumentation framework serves as base construction, while the translation algorithm we have examined augments an existing framework.
- *Filtering strategies that retain only desired extensions,* interlinking the extensions of the base framework with additional constructs that utilise the defining properties of the respective semantics in a deliberate manner.
- *Suppression of arguments* which are essential in the construction but should not appear in the extensions themselves, by making them self-attacking.

3 Upper bounds for the realisability of weak argumentation semantics

In this section, we take a closer look at the necessary properties, that is, properties that apply commonly to all extension sets of a certain weak semantics. In doing so, we focus on weakly admissible and weakly preferred semantics, as well as undisputed and strongly undisputed semantics. These semantics are related to each other as follows [3, 21].

Proposition 8. For every argumentation framework $F \in \mathfrak{F}_{\mathfrak{A}}$, the following subset relations hold:

- $\mathsf{stb}(F) \subseteq \mathsf{wpr}(F) \subseteq \mathsf{wadm}(F) \subseteq \mathsf{cf}(F)$;
- $\mathsf{stb}(F) \subseteq \mathsf{sud}(F) \subseteq \mathsf{ud}(F) \subseteq \mathsf{cf}(F)$, as well as $\mathsf{pr}(F) \subseteq \mathsf{ud}(F)$. □

Although these subset relationships do not necessarily imply the transition of properties of extension sets between any two semantics involved, this can certainly be the case in individual instances; for example, the property of preferred extension sets being universally defined does transfer to undisputed semantics. In any case, we investigate to what extent these relationships can provide indications regarding the common properties of the semantics being analysed here.

Another lead we follow is the symmetry of the definitions of preferred and weakly preferred semantics (compare Definitions 5 and 7); we investigate whether some of the classical properties from Theorem 1 carry over to the corresponding non-admissible cases.

In Section 3.1 we first analyse the compliance of weak semantics with the classical properties from Definition 14. In Section 3.2 we discuss limits of the expressiveness of undisputed semantics, and in Section 3.3 we summarise upper bounds for the signatures of all weak semantics.

3.1 Properties of signatures

We start with considering existing properties of signatures (Definition 14) and check whether weak semantics satisfy these properties.

As Theorem 1 has already shown, admissible extension sets are always conflict-sensitive. This is not necessarily true for weakly admissible extension sets, as the following example shows.

Example 11. Consider the framework (F) in Figure 8. Here we have

$$\mathsf{wadm}(F) = \{\emptyset, \{a\}, \{b\}, \{c\}, \{a,b\}, \{a,c\}, \{b,c\}\},$$

which does not contain $\{a,b,c\}$ despite the fact that $\{a,b\}, \{b,c\}, \{a,c\} \in \mathsf{wadm}(F)$.

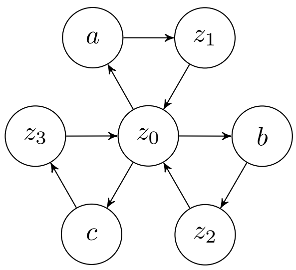

Figure 8: The weakly admissible extension set of this framework is not conflict-sensitive.

Dunne *et al.* [12] also introduce the following relaxation of the concept of conflict-sensitivity.

Definition 19. A set of argument-sets $\mathbb{S} \subseteq 2^{\mathfrak{A}}$ is called *com-closed* if, for every subset $\mathbb{S}' \subseteq \mathbb{S}$ and all $a, b \in \bigcup \mathbb{S}'$ holds that a, b occur jointly in some set of $\mathbb{S}$, then there is exactly one $\subseteq$-minimal set in $\mathbb{S}$ that encompasses all of $\bigcup \mathbb{S}'$ (this set is called the *unique completion set*).

Again, weakly admissible extension sets do not comply with the property of com-closure.

Example 12. The weakly admissible extensions of the framework in Figure 8 do not contain a completion set that encompasses $\{a,b,c\}$.

The property of com-closure generalises conflict-sensitivity by requiring that, in the absence of evidence for conflict,[5] not the exact union of two extensions need to be present in the extension set, but the union should at least be included in another (unique) extension. The framework of Figure 8 contradicts attempts to further relax this requirement when applied to weak admissibility, e.g. by dropping uniqueness of the completion-set. We conjecture at this point that no analogous property (involving the requirement that unions of extensions must appear as subset of another extension) holds for weak admissibility.

With regard to the classical properties of Definition 14 we further note that weakly admissible extension sets are neither necessarily tight (this would imply conflict-sensitivity) nor necessarily downward-closed. The latter statement is evidenced by the following example.

Example 13. For the framework (F) in Figure 9, we have $\mathsf{wadm}(F) = \{\emptyset, \{b\}, \{c\}, \{a, c\}\}$, which is not downward-closed since $\{a\} \notin \mathsf{wadm}(F)$.

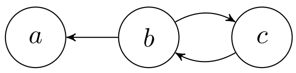

Figure 9: The weakly admissible extension set for this framework is not downward-closed.

Below we state the only property common to all weakly admissible extension sets remaining so far.

Proposition 9. Every weakly admissible extension set contains the empty set.

Proof. The empty set is conflict-free and has no attackers. □

We now turn to preferred extensions and recall that every preferred extension set is non-empty, incomparable and conflict-sensitive, cf. Theorem 1. Only the properties of non-emptiness and incomparability transfer to the weakly preferred semantics.

Proposition 10. Weakly preferred extension sets are non-empty and incomparable.

Proof. Non-emptiness follows from the definition, and incomparability is implied by $\subseteq$-maximality. □

The following example rules out the other classical properties from Definition 14 for weakly preferred semantics.

Example 14. Looking again at the framework in Figure 8, we find that its weakly preferred extension set is $\{\{a, b\}, \{a, c\}, \{b, c\}\}$, which is neither conflict-sensitive, tight, com-closed, nor downward-closed.

We proceed to undisputed and strongly undisputed sets and recall the following basic result from [21],

Proposition 11. Undisputed extension sets are non-empty, and strongly undisputed extension sets are incomparable (but not necessarily non-empty). □

Undisputed and strongly undisputed extension sets are neither necessarily com-closed nor conflict-sensitive (and therefore, also not tight).

Example 15. For the framework F in Figure 8, we have $\mathsf{sud}(F) = \{\{a, b\}, \{a, c\}, \{b, c\}\}$ and $\mathsf{ud}(F) = \{\emptyset, \{a\}, \{b\}, \{c\}\} \cup \mathsf{sud}(F)$.

Neither ud nor sud are com-closed.

Example 16. In Figure 9 (call the framework F), we have $\mathsf{ud}(F) = \mathsf{sud}(F) = \{\{b\}, \{a, c\}\}$.

We summarise our findings in Table 1.

[5] In the same sense as in the definition of conflict-sensitivity (Definition 14). The term (“evidently in conflict”) will be formalised in Definition 20.

	wadm	wpr	ud	sud	References
conflict-sensitivity	×	×	×	×	Ex. 11, 14, 15.
tightness	×	×	×	×	Prop. 2.
com-closedness	×	×	×	×	Ex. 12, 14, 15.
downward-closedness	×	×	×	×	Ex. 13, 14, 16.
incomparability	×	✓	×	✓	Ex. 11, 15; Prop. 10, 11.
$\emptyset$-inclusion	✓	×	×	×	Prop. 9; Ex. 14, 15, 16.
non-emptiness	✓	✓	✓	×	Prop. 9, 11.

Table 1: Classical properties of extension sets satisfied or violated by weak semantics.

3.2 Limits to the expressiveness of undisputed semantics

The analyses in the previous subsection focused on properties of signatures of classical semantics. In the following, we present two properties that have no classical counterpart and that limit the expressive power of undisputed semantics. We start with some basic observations.

Lemma 1. Every undisputed extension set contains at least one preferred extension. This preferred extension is the empty set if and only if the empty set is undisputed.

Proof. Let $F \in \mathfrak{F}_{\mathfrak{A}}$. If $\emptyset \notin \mathsf{ud}(F)$, then $\mathsf{adm}(F^{\emptyset}) \not\subseteq \{\emptyset\}$, so there is a non-empty $S \in \mathsf{pr}(F)$, for which $S \in \mathsf{ud}(F)$ by Proposition 8. If $\emptyset \in \mathsf{ud}(F)$, then $\mathsf{adm}(F^{\emptyset}) \subseteq \{\emptyset\}$ and $\mathsf{pr}(F) = \{\emptyset\}$. ☐

The following concept extracts a key element from the definition of conflict-sensitivity (Definition 14).

Definition 20. Let $\mathbb{S} \subseteq 2^{\mathfrak{A}}$, and let $X, Y \in \mathbb{S}$. We say that X and Y are *evidently in conflict with respect to* $\mathbb{S}$ when

$$\exists x \in X, y \in Y : \forall S \in \mathbb{S} : \{x, y\} \not\subseteq S.$$

For semantics that produce conflict-free extensions, lack of evidence of conflict guarantees that no attacks exist between two sets.

Lemma 2. For $F \in \mathfrak{F}_{\mathfrak{A}}$, let $\sigma : \mathfrak{F}_{\mathfrak{A}} \to 2^{2^{\mathfrak{A}}}$ so that every $S \in \sigma(F)$ is conflict-free, and let $X, Y \in \sigma(F)$. If X and Y are not evidently in conflict with respect to $\sigma(F)$, then $X \cap Y^+ = \emptyset = Y \cap X^+$.

Proof. The negation of Definition 20 states that for all pairs $(x, y) \in X \times Y$ we have some $S \in \sigma(F)$ so that $\{x, y\} \subseteq S$; so no member of X can attack any member of Y, and vice versa. ☐

We continue with two technical lemmas that we will use immediately after.

Lemma 3. Let $S, A_1, \ldots, A_n$ be sets so that $\bigcap_i A_i = \emptyset$ and $S \cap A_i \neq \emptyset$ for each A_i ($i \in \{1, \ldots, n\}$). Then, $S \setminus A_i \neq \emptyset$ for at least one A_i.

Proof. From $S \setminus A_i = \emptyset$ follows $S \subseteq A_i$; requiring this for all A_i implies $S \subseteq \bigcap_i A_i$, and since $S \neq \emptyset$, we have a contradiction. ☐

Lemma 4. For $F = (\mathcal{A}, \mathcal{R}) \in \mathfrak{F}_{\mathfrak{A}}$, let $S \in \mathsf{adm}(F)$, and let $S' \subseteq \mathcal{A}$, so that $S \cap S'^+ = \emptyset$. Then $S \setminus S' \in \mathsf{adm}(F^{S'})$.

Proof. Since S' does not attack any argument of S, the arguments of $S \setminus S'$ are still present in the reduct $F^{S'}$. Moreover, all attackers of S that S' defended against are removed in $F^{S'}$, so that $S \setminus S'$ defends itself in $F^{S'}$. ☐

Finally, the following definition captures a class of extension sets that cannot be realised by undisputed semantics.

Definition 21. A set of sets $\mathbb{S} = \{A_1, \ldots, A_n\} \subseteq 2^{\mathfrak{A}}$ ($n > 1$) is *disjointly supported* if, for $i \in \{1, \ldots, n\}$, each $A_i \neq \emptyset$, $\bigcap_i A_i = \emptyset$, and the A_i are pairwise not in evident conflict.

Example 17. The set $\{\{a,b\},\{b,c\},\{a,c\}\}$ is disjointly supported.

Disjointly supported extension sets are unrealisable by wadm, simply because they do not contain the empty set. For ud, the reason is more subtle.

Proposition 12. No disjointly supported extension set is realisable by ud.

Proof. Let $\mathbb{S} = \{A_1, \dots, A_n\} \subseteq 2^{\mathfrak{A}}$, so that each $A_i \neq \emptyset$, $\bigcap_i A_i = \emptyset$, and no A_i, A_j are evidently in conflict. Assume that there is a $F \in \mathfrak{F}_{\mathfrak{A}}$ with $\mathsf{ud}(F) = \mathbb{S}$. From Lemma 1 we obtain the existence of a non-empty $S \in \mathbb{S}$ with $S \in \mathsf{pr}(F)$. Furthermore, we have $S \cap A_i^{\oplus} \neq \emptyset$ for all A_i, otherwise $S \in \mathsf{adm}(F^{A_i})$, contradicting $A_i \in \mathsf{ud}(F)$. We also have $S \cap A_i^{+} = \emptyset$ by Lemma 2; this leaves $S \cap A_i \neq \emptyset$ for all A_i. Lemma 3 then states that for at least one $A_j \in \mathbb{S}$ we have $S \setminus A_j \neq \emptyset$, and by Lemma 4 we have $S \setminus A_j \in \mathsf{adm}(F^{A_j})$, again contradicting $A_j \in \mathsf{ud}(F)$. □

Below we describe a second class of extension sets that is unrealisable by undisputed semantics.

Definition 22. We call a set of sets $\mathbb{S} \subseteq 2^{\mathfrak{A}}$ a *distinct pairing* if $\emptyset \notin \mathbb{S}$, and

$$\forall S \in \mathbb{S}\ \exists S' \in \mathbb{S} : S \setminus S' \neq \emptyset \ \text{ and } \ S \cup S' \in \mathbb{S}.$$

Example 18. The set $\{\{a\},\{b\},\{a,b\},\{a,c\}\}$ is a distinct pairing, but not disjointly supported. Conversely, the set $\{\{a,b\},\{b,c\},\{a,c\}\}$ from Example 17 is not a distinct pairing. Furthermore, $\{\{a\},\{b\},\{a,b\}\}$ is both disjointly supported and a distinct pairing.

Proposition 13. Undisputed semantics cannot realise distinct pairings.

Proof. Let $\mathbb{S}$ be a distinct pairing, and assume we have a framework $F \in \mathfrak{F}_{\mathfrak{A}}$ so that $\mathsf{ud}(F) = \mathbb{S}$. Since $\emptyset \notin \mathbb{S}$, by Lemma 1 one extension $S \in \mathbb{S}$ must be preferred. The prerequisite gives us $S' \in \mathbb{S}$ that cannot attack S since $S \cup S' \in \mathbb{S}$; so Lemma 4 applies, and $\emptyset \neq S \setminus S' \in \mathsf{adm}(F^{S'})$, contradicting the undisputedness of S'. □

3.3 Upper bounds of signatures of weak semantics

The results from the previous subsections allow us to give first estimates for the signatures of weak semantics.

Proposition 14. Upper bounds for the signatures of weak semantics are given as follows.

- $\Sigma_{\mathsf{wadm}} \subseteq \{\mathbb{S} \subseteq 2^{\mathfrak{A}} \mid \emptyset \in \mathbb{S}\}$;
- $\Sigma_{\mathsf{wpr}} \subseteq \{\mathbb{S} \subseteq 2^{\mathfrak{A}} \mid \mathbb{S}$ is non-empty and incomparable$\}$;
- $\Sigma_{\mathsf{ud}} \subseteq \{\mathbb{S} \subseteq 2^{\mathfrak{A}} \mid \mathbb{S}$ is non-empty, not disjointly supported, and not a distinct pairing$\}$;
- $\Sigma_{\mathsf{sud}} \subseteq \{\mathbb{S} \subseteq 2^{\mathfrak{A}} \mid \mathbb{S}$ is incomparable$\}$.

Proof. This follows immediately from Propositions 9, 10, 11, 12, and 13. □

In order to establish lower bounds for the characterisations of weak semantics, we will explore construction methods for frameworks that produce extension sets meeting the necessary conditions outlined in this section. Along the way we will uncover some evidence suggesting that the upper bound for undisputed semantics given above may not, in fact, be the smallest (Section 4.8). The remainder of this work focuses on the investigation of these construction methods.

4 Lower bounds for the realisability of weak argumentation semantics

In this section, we consider construction methods for showing how certain extension sets can be realised. We start in Section 4.1 with a brief discussion on the suppression of auxiliary arguments (introduced for the sake of the construction) in the context of weak argumentation semantics. We continue in Section 4.2 that is concerned with the reasons why certain unions of extensions may not be extensions themselves. In Section 4.3 we provide some first general constructions that provide overestimations of the extension sets to be realised. Afterwards, in Section 4.4 we discuss methods how to avoid (or *filter*) these overestimations and Section 4.5 discusses the practical limits of these methods. Section 4.6 presents further methods for filters, leading to a new general construction approach in Section 4.7. Section 4.8 summarises the methods developed so far and Section 4.9 discusses a characterisation of a special case.

4.1 Suppression of arguments

All classical construction methods have in common that newly added arguments, which are essential for the effectiveness of the construction, but should not appear in the extensions themselves, can simply be filtered out by making them self-attacking. Unfortunately, this strategy does not readily translate to non-admissible semantics. In the case of weakly admissible semantics, the reason for this is given by the following result from [3]. For that, we define $\mathcal{A}^\circ = \{a \in \mathcal{A} \mid (a,a) \notin \mathcal{R}\}$ and

$$F^\circ = F|_{\mathcal{A}^\circ}.$$

In other words, F° is the same as F but with all self-attacking arguments removed.[6]

Proposition 15. Let $F = (\mathcal{A}, \mathcal{R}) \in \mathfrak{F}_\mathfrak{A}$. For $\sigma \in \{\mathsf{wadm}, \mathsf{wpr}\}$, we have $\sigma(F) = \sigma(F^\circ)$. □

This means that self-attacking arguments can be neglected for both wadm and wpr, so they cannot have the intended effect in the construction algorithms previously considered. The same is not true for undisputed and strongly undisputed semantics however, as the next example shows.

Example 19. In the framework from Figure 10 (F), we have $\mathsf{ud}(F) = \{\emptyset, \{c\}\}$ and $\mathsf{sud}(F) = \emptyset$, but $\mathsf{ud}(F^\circ) = \mathsf{sud}(F^\circ) = \{\emptyset\}$.

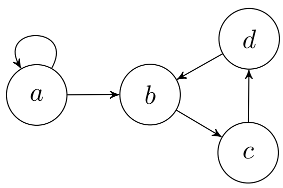

Figure 10: Undisputed and strongly undisputed semantics applied to this framework are not invariant with respect to removal of self-attacking arguments.

Although self-attacking arguments can never be part of an undisputed or strongly undisputed extension, they can still have an effect on the admissibility of the reduct (as was the case in the above example with the set $\{c\}$). Later on (in Proposition 21) we will use this feature in a systematic way.

4.2 Separating extensions

We turn to weakly admissible semantics and investigate mechanisms that can cause two extensions to be separated from each other in an extension set, in the sense that their union is not an extension itself.

[6] An essentially identical definition appears in [8, Definition 6].

Observation 2. In the classical case, the only reason for the union of two admissible extensions not to be admissible as well is that one extension attacks the other so that the union would not be conflict-free (Proposition 6). Non-admissible semantics introduce another mechanism that is able to separate extensions, namely odd cycles, as shown in the following example.

Example 20. Figure 11 shows the three interlinked odd cycles that we already have encountered in previous examples. Call this framework F. To the left, we highlight the reduct $F^{\{a,b\}}$: accepting the arguments $\{a, b\}$ leaves z in an odd cycle, so that it presents no serious threat. In the right picture we see that $\{a, b\} \cup \{b, c\} = \{a, b, c\}$ cannot be jointly accepted, since in the reduct $F^{\{a,b,c\}}$, $\{z\}$ would be admissible. This alone is reason enough for $\{a, b, c\}$ not to be accepted as $\{\mathsf{ud}, \mathsf{sud}\}$-extension; it disqualifies as $\{\mathsf{wadm}, \mathsf{wpr}\}$-extension as well because z attacks a, b, c.

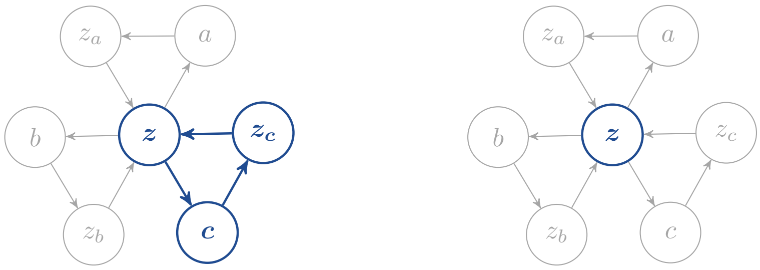

Figure 11: Mechanics of odd cycles in non-admissible semantics.

We would like to formalise the mechanics at work in Example 20. In order to do this, we need the following two results. The first is the *Modularisation Theorem* by Baumann *et al.* [4].

Theorem 3. Let $F \in \mathfrak{F}_{\mathfrak{A}}$, let $X \in \mathsf{wadm}(F)$, and let $Y \in \mathsf{wadm}(F^X)$. Then $X \cup Y \in \mathsf{wadm}(F)$. □

The second prerequisite concerns the composition of the reduct, which turns out not to be commutative, in the sense that $(F^X)^Y = (F^Y)^X$ does not generally hold. The below example conveys the intuition.

Example 21. Consider the framework from Figure 12 (call it F), and let $X = \{a\}$ and $Y = \{b\}$. We have $F^X = (\{d\}, \emptyset) = (F^X)^Y \neq (F^Y)^X = (\emptyset, \emptyset) = F^{X \cup Y}$.

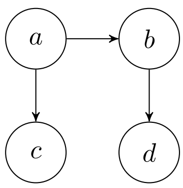

Figure 12: The reduct of this framework with respect to $\{a\}$ eliminates $\{a, b, c\}$, and the subsequent reduct with respect to $\{b\}$ cannot affect d any more.

The following lemma, along with its accompanying corollary, state that the order of reduct construction is not significant as long as the involved sets do not attack each other.[7]

Lemma 5. Let $F = (\mathcal{A}, \mathcal{R}) \in \mathfrak{F}_{\mathfrak{A}}$, and let $X, Y \subseteq \mathcal{A}$. If $X^+ \cap Y = \emptyset$, then $\left(F^X\right)^Y = F^{X \cup Y}$.

Proof. Let $F^X = (\mathcal{A}_X, \mathcal{R}_X)$ and $\mathcal{A}_X = \mathcal{A} \setminus (X \cup X^+_{\mathcal{R}})$, where we write $X^+_{\mathcal{R}}$ instead of X^+, as we need to be careful about which attack relation we refer to. We then have $(F^X)^Y = (\mathcal{A}_{X,Y}, \mathcal{R}_{X,Y})$, where

$$\begin{aligned} \mathcal{A}_{X,Y} &= \big(\mathcal{A} \setminus (X \cup X^+_{\mathcal{R}})\big) \setminus (Y \cup Y^+_{\mathcal{R}_X}) \\ &= \mathcal{A} \setminus (X \cup Y \cup X^+_{\mathcal{R}} \cup Y^+_{\mathcal{R}_X}), \end{aligned}$$

[7]Baumann *et al.* [4, Proposition 3.3] give a similar proposition, but require $X \cap Y = \emptyset$ as well as $X \cup Y \in \mathsf{cf}(F)$.

while, for $F^{X\cup Y} = (\mathcal{A}_{XY}, \mathcal{R}_{XY})$, we have[8]

$$\begin{aligned}\mathcal{A}_{XY} &= \mathcal{A} \setminus (X \cup Y \cup (X\cup Y)^+_{\mathcal{R}}) \\ &= \mathcal{A} \setminus (X \cup Y \cup X^+_{\mathcal{R}} \cup Y^+_{\mathcal{R}}).\end{aligned}$$

Finally, since $X^+_{\mathcal{R}} \cap Y = \emptyset$, we have $Y^+_{\mathcal{R}_X} = Y^+_{\mathcal{R}} \setminus X^+_{\mathcal{R}}$ (the difference between $Y^+_{\mathcal{R}}$ and $Y^+_{\mathcal{R}_X}$ is only due to arguments commonly attacked by X and Y, not due to attacks on Y), so $X^+_{\mathcal{R}} \cup Y^+_{\mathcal{R}} = X^+_{\mathcal{R}} \cup Y^+_{\mathcal{R}_X}$. Then $\mathcal{A}_{X,Y} = \mathcal{A}_{XY}$ and $\mathcal{R}_{X,Y} = \mathcal{R}_{XY}$ follows straightforwardly. □

Corollary 1. Let $F = (\mathcal{A}, \mathcal{R}) \in \mathfrak{F}_{\mathfrak{A}}$, and let $X, Y \subseteq \mathcal{A}$, so that $X \cup Y \in \mathsf{cf}(F)$. Then $(F^X)^Y = (F^Y)^X = F^{X\cup Y}$. □

We can now describe a structural feature that is necessarily present in situations as in Example 20.

Proposition 16. Let $F \in \mathfrak{F}_{\mathfrak{A}}$, and let $X, Y \in \mathsf{wadm}(F)$ so that $X \cup Y \in \mathsf{cf}(F)$ but $X \cup Y \notin \mathsf{wadm}(F)$. Then there are sets $Z_X, Z_Y \in \mathsf{wadm}(F^{X\cup Y})$ where $Z_X \to X$ and $Z_Y \to Y$.

Proof. Since X and Y are conflict-free, by definition of weak admissibility there is a set $Z \in \mathsf{wadm}(F^{X\cup Y})$ so that $Z \to X \cup Y$; in fact, there may be many such sets. Say that none of these sets attack Y. Because of conflict-freeness of $X \cup Y$, by Corollary 1 we have $F^{X\cup Y} = (F^X)^Y$; so Y is not attacked from $\bigcup \mathsf{wadm}((F^X)^Y)$, i.e., $Y \in \mathsf{wadm}(F^X)$. From Theorem 3 it then follows that $X \cup Y \in \mathsf{wadm}(F)$, a contradiction. □

Example 22. Let F be the argumentation framework from Figure 11, and consider $X = \{a, b\}$ and $Y = \{b, c\}$, both of which are weakly admissible extensions of F. The union $X \cup Y$ however is not weakly admissible. Since $X \cup Y$ is conflict-free, Proposition 16 requires the existence of attacks on both X and Y originating from a weakly admissible set of $F^{X\cup Y} = (\{z\}, \emptyset)$; indeed, $\{z\}$ fits this requirement.

We continue our investigation of structural features with the case of undisputed semantics. While weak admissibility requires that actual attacks originate from weakly admissible sets of the reduct in order to inhibit the existence of extensions, for undisputedness the mere presence of a non-empty admissible set in the reduct suffices; consequently, we may expect to find different structural features than those from Proposition 16. Let us look at an example.

Example 23. Consider the framework (F) of Figure 13. Its only weakly admissible extension is the empty set, but its undisputed extensions are $\emptyset$, $\{a\}$, and $\{b\}$, though not $\{a, b\}$. Note that the reduct $F^{\{a\}}$ (depicted to the left) contains no non-empty admissible set; the same is true for $F^{\{b\}}$. In $F^{\{a,b\}}$ however (shown in the picture to the right), $\{z_2, z_3\}$ is admissible. Neither $\{a\}$ nor $\{b\}$ are weakly admissible, because both are attacked from weakly admissible sets of their reduct ($\{z_3\}$ and $\{z_2\}$ respectively).

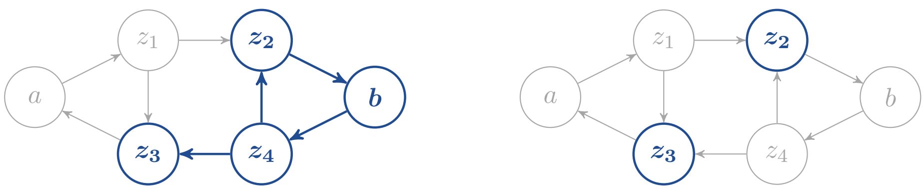

Figure 13: In this framework, $\{a, b\}$ is not among the undisputed extensions.

We now formalise two structural features of frameworks that separate extensions under undisputed semantics.

Proposition 17. Let $F \in \mathfrak{F}_{\mathfrak{A}}$, and let $X, Y \in \mathsf{ud}(F)$ so that $X \cup Y \in \mathsf{cf}(F)$, but at the same time, $X \cup Y \notin \mathsf{ud}(F)$. Then both X and Y have attackers in F against which they do not defend.

[8] Read X, Y as "X, then Y", and XY as "X and Y simultaneously".

Proof. The prerequisites demand $X \nsubseteq Y$ and $Y \nsubseteq X$, otherwise $X \cup Y \in \mathsf{ud}(F)$. As $X \cup Y \in \mathsf{cf}(F)$, the remainder $X \setminus Y = X \setminus Y^{\oplus}$ is non-empty. Assume that X is admissible in F; any threats that $X \cap Y$ defends against then disappear in the reduct F^Y, so that the remainder $X \setminus Y^{\oplus}$ remains admissible in F^Y, contradicting $Y \in \mathsf{ud}(F)$. So $X \notin \mathsf{adm}(F)$. The situation for Y is symmetrical. □

Example 24. Applying Proposition 17 to Figure 13, we have $X = \{a\}$, $Y = \{b\}$, and find $z_3 \to X$ and $z_2 \to Y$; neither X nor Y defend themselves.

Proposition 18. Let $F \in \mathfrak{F}_{\mathfrak{A}}$, and let $X, Y \in \mathsf{ud}(F)$ so that $X \cup Y \in \mathsf{cf}(F)$, but at the same time, $X \cup Y \notin \mathsf{ud}(F)$. Then there is a nonempty admissible extension $S \in \mathsf{adm}(F^{X \cup Y})$ that is attacked by both X^+ and Y^+.

Proof. The existence of a nonempty $S \in \mathsf{adm}(F^{X \cup Y})$ is guaranteed by $X \cup Y \in \mathsf{cf}(F)$ and $X \cup Y \notin \mathsf{ud}(F)$. All of the arguments of S are also present in F^X. Since $S \notin \mathsf{adm}(F^X)$, there must be an attack $a \to S$ in F^X that S does not defend against. After removal of $Y^{\oplus}$ and its related attacks from F^X we obtain $(F^X)^Y = F^{X \cup Y}$ from Corollary 1, where this same attack $a \to S$ is no longer threatening the admissibility of S. The formation of the reduct cannot have added any structures, particularly, it cannot have introduced a defence $S \to a$; we conclude that the attack must have originated from Y^+ ($a \to S$ was present in F^X, but disappeared in $F^{X \cup Y}$; and since $S \in \mathsf{adm}(F^{X \cup Y})$, $a \notin Y$). For reasons of symmetry we have both $X^+ \to S$ and $Y^+ \to S$. □

Example 25. In the framework (F) of Figure 13, let $X = \{a\}$, $Y = \{b\}$, and let $S = \{z_2, z_3\} \in \mathsf{adm}(F^{X \cup Y})$. We find $X^+ = \{z_1\} \to S$ as well as $Y^+ = \{z_4\} \to S$.

4.3 Base frameworks

We now address the question of how to construct argumentation frameworks that realise a given extension set. We first consider construction methods that *overestimate* the extension set, i. e., methods where the resulting argumentation frameworks possess *at least* the required extensions, possibly along with additional extensions.

Definition 23. With respect to a semantics $\sigma : \mathfrak{F}_{\mathfrak{A}} \to 2^{2^{\mathfrak{A}}}$ and a set of extension sets $\mathfrak{S} \subseteq 2^{2^{\mathfrak{A}}}$ we call the mapping $\mathbb{S} \mapsto F_{\mathbb{S}}$ a *base framework* if for all $\mathbb{S} \in \mathfrak{S}$ we have $\mathbb{S} \subseteq \sigma(F_{\mathbb{S}})$.

We have seen an example of a base framework in Proposition 3:

Example 26. For each $\sigma \in \{\mathsf{cf}, \mathsf{adm}, \mathsf{stb}\}$ and for all $\mathbb{S} \in \Sigma_\sigma$ (from Theorem 1), the construction $\mathbb{S} \mapsto F^{\mathsf{cf}}_{\mathbb{S}}$ of the canonical argumentation framework (Definition 15) is a base framework with respect to σ and Σ_σ.

As in the above example, we hope to find base frameworks for the non-admissible cases as well, so that these can be modified, or *filtered* respectively, in a second step in order to arrive at a realising argumentation framework. So let us consider the canonical argumentation framework from Definition 15, as well as the canonical defence framework from Definition 16, and see whether these are suitable as base frameworks under weak semantics.[9]

Proposition 19. For all $\mathbb{S} \subseteq 2^{\mathfrak{A}}$ we have $\mathbb{S} \subseteq \mathsf{wadm}(F^{\mathsf{cf}}_{\mathbb{S}}) = \mathsf{wadm}(F^{\mathsf{adm}}_{\mathbb{S}})$.

Proof. $F^{\mathsf{cf}}_{\mathbb{S}}$ constructs reciprocal attacks only between arguments that do not appear jointly in any $S \in \mathbb{S}$, so each $S \in \mathbb{S}$ defends itself and $S \in \mathsf{adm}(F^{\mathsf{cf}}_{\mathbb{S}})$. Furthermore, $\mathsf{adm}(F^{\mathsf{cf}}_{\mathbb{S}}) \subseteq \mathsf{wadm}(F^{\mathsf{cf}}_{\mathbb{S}})$ holds by Proposition 1. Finally, by construction we have $(F^{\mathsf{adm}}_{\mathbb{S}})^{\circ} = F^{\mathsf{cf}}_{\mathbb{S}}$ for any $\mathbb{S} \subseteq 2^{\mathfrak{A}}$, therefore $\mathsf{wadm}(F^{\mathsf{cf}}_{\mathbb{S}}) = \mathsf{wadm}(F^{\mathsf{adm}}_{\mathbb{S}})$ by Proposition 15. □

[9] We introduced the base framework as a mapping in order to emphasise the dependency on the parameter $\mathbb{S}$. In the following, we generally do not differentiate terminologically between the mapping and the constructed framework.

The range of values of $\mathbb{S}$ considered here corresponds to the upper bound for wadm from Proposition 9, minus the empty set; the empty set however is present in any weakly admissible extension set regardless of the framework. Thus, $\mathbb{S} \mapsto F_{\mathbb{S}}^{\mathsf{cf}}$ and $\mathbb{S} \mapsto F_{\mathbb{S}}^{\mathsf{adm}}$ are base frameworks with respect to weakly admissible semantics. For weakly preferred, undisputed, and strongly undisputed semantics the situation is different, as demonstrated by the following two examples.

Example 27. Consider $\mathbb{S} = \{\{a, b\}, \{b, c\}, \{a, c\}\}$; this extension set is realisable under weakly preferred and strongly undisputed semantics, as witnessed by the framework (F) from Figure 11, where we have $\mathsf{wpr}(F) = \mathsf{sud}(F) = \mathbb{S}$. However, we have $F_{\mathbb{S}}^{\mathsf{cf}} = (\{a, b, c\}, \emptyset)$, and $\mathbb{S} \not\subseteq \mathsf{wpr}(F_{\mathbb{S}}^{\mathsf{cf}}) = \mathsf{sud}(F_{\mathbb{S}}^{\mathsf{cf}}) = \{\{a, b, c\}\}$. As for the canonical defence framework constructed for $\mathbb{S}$, we depict this in Figure 14; for this framework, we have $\mathbb{S} \not\subseteq \mathsf{sud}(F_{\mathbb{S}}^{\mathsf{adm}}) = \{\{a, b, c\}\}$.

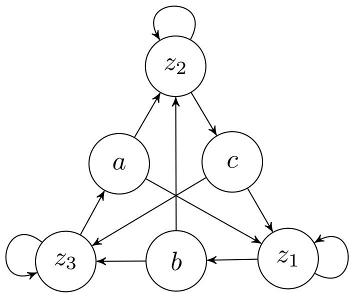

Figure 14: The canonical defence framework $F_{\mathbb{S}}^{\mathsf{adm}}$, constructed for $\mathbb{S} = \{\{a, b\}, \{b, c\}, \{a, c\}\}$.

Example 28. The extension set $\mathbb{S} = \{\emptyset, \{b\}\}$ is realised under undisputed semantics by the framework from Figure 1; yet we have $F_{\mathbb{S}}^{\mathsf{cf}} = F_{\mathbb{S}}^{\mathsf{adm}} = (\{b\}, \emptyset)$ and $\mathbb{S} \not\subseteq \mathsf{ud}(F_{\mathbb{S}}^{\mathsf{cf}}) = \{\{b\}\}$.

Weakly admissible semantics possess an even simpler base framework than the canonical argumentation framework.

Proposition 20. $\mathbb{S} \mapsto F_{\mathbb{S}}^{\emptyset} = (\bigcup \mathbb{S}, \emptyset)$ is a base framework with respect to weakly admissible semantics and $\mathbb{S} \in 2^{\mathfrak{A}}$.

Proof. In $F_{\mathbb{S}}^{\emptyset} = (\mathcal{A}, \emptyset)$, all subsets of $\mathcal{A}$ are unattacked and conflict-free. □

A small modification yields a base framework for undisputed semantics.

Proposition 21. Let $z \notin \bigcup \mathbb{S}$. $\mathbb{S} \mapsto F_{\mathbb{S}}^{\mathsf{ud}} = (\bigcup \mathbb{S} \cup \{z\}, \{(z, z)\} \cup \{(z, s) \mid s \in \bigcup \mathbb{S}\})$ is a base framework with respect to undisputed semantics and $\mathbb{S} \subseteq 2^{\mathfrak{A}}$.

Proof. We have $\mathsf{ud}(F_{\mathbb{S}}^{\mathsf{ud}}) = 2^{(\bigcup \mathbb{S})}$, because every $S \subseteq \bigcup \mathbb{S}$ is conflict-free, and its complement $\bigcup \mathbb{S} \setminus S$ is attacked by the self-attacking z and thus is not admissible. □

Example 29. Figure 15 shows the base framework $F_{\mathbb{S}}^{\mathsf{ud}}$ for any $\mathbb{S}$ with $\bigcup \mathbb{S} = \{a, b, c\}$.

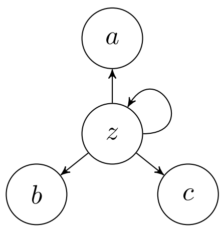

Figure 15: The undisputed extension set of this framework is the set of all subsets of the original arguments $\{a, b, c\}$.

The base frameworks for wadm and ud described in Propositions 20 and 21 are not suitable for wpr and sud. The reason for this is that both constructions only take into account the union of arguments and do not involve the structure of the extension set in any other way in the construction. While both wadm and ud are able to produce the maximal extension set 2^A for any $A \subseteq \mathfrak{A}$, thus subsuming any extensions of interest, wpr and sud only produce incomparable extension sets, and the constructions that worked well for wadm and ud lack features that would separate extensions from another.

4.4 Filtering

The question that we ask in this subsection is this: Given a framework F for which a semantics produces an extension set that is a superset of the set that we desire to obtain, how can this framework be modified to remove the unwanted extensions? More formally, we are looking for a mapping as in the following definition.

Definition 24. With respect to a semantics σ, a signature Σ_σ, and a set of argumentation frameworks $\mathfrak{F}' \subseteq \mathfrak{F}_\mathfrak{A}$, a *filter* is a mapping $f : \mathfrak{F}' \times 2^{2^\mathfrak{A}} \to \mathfrak{F}_\mathfrak{A}$, so that $\sigma(f(F, \mathbb{X})) = \sigma(F) \setminus \mathbb{X}$, as long as $\sigma(F) \setminus \mathbb{X} \in \Sigma_\sigma$. A filter is *total* with respect to its first argument if $\mathfrak{F}' = \mathfrak{F}_\mathfrak{A}$, otherwise it is *partial*.

In the above definition, $\mathbb{X}$ is the set of unwanted extensions. We have already seen such mappings at work.

Example 30. For $\sigma = \mathsf{stb}$ we used the following total filter (cf. Definition 17):

$$\begin{aligned} f_{\mathsf{stb}}((\mathcal{A}, \mathcal{R}), \mathbb{X}) &= (\mathcal{A}', \mathcal{R}') \\ \text{where } \mathcal{A}' &= \mathcal{A} \cup \{z_X \mid X \in \mathbb{X}\}, \\ \mathcal{R}' &= \mathcal{R} \cup \{(z_X, z_X) \mid X \in \mathbb{X}\} \cup \{(a, z_X) \mid X \in \mathbb{X} \ \wedge \ a \in \mathcal{A} \setminus X\}. \end{aligned}$$

The concept of a filter shares similarities with framework translations which we described in Subsection 2.3; but instead of attempting to preserve extension sets between different semantics, we seek to remove certain extensions under the same semantics. The following terminology is borrowed from the context of framework translations [13].

Definition 25. A filter f is called *covering* if for every $F = (\mathcal{A}, \mathcal{R})$ and every $\mathbb{X} \subseteq 2^\mathfrak{A}$ the resulting argumentation framework $f(F, \mathbb{X}) = (\mathcal{A}', \mathcal{R}')$ satisfies $\mathcal{A} \subseteq \mathcal{A}'$ and $\mathcal{R} \subseteq \mathcal{R}'$.

In other words, the mapping performed by a covering filter does only augment, but not alter existing structures. We observe that filters applied in the classical cases generally appear to be covering.

Observation 3. The filter from Example 30 is covering, and the construction of the canonical defence framework (Definition 16) employs a covering filter as well.[10]

The fact that these filters are covering seems to make a significant contribution to ensuring that their designs remain comprehensible and universally applicable. In a first attempt to develop a corresponding understanding of non-admissible filter mechanisms, let us look at the problem posed in the following example.

Example 31. Consider the following problem: for the framework (F) in Figure 16, we have $\mathsf{wadm}(F) = \{\emptyset, \{a\}, \{a, b\}\}$. Under weakly admissible semantics, how could a covering filter f proceed in order to eliminate the extension $\{a\}$, so that $\mathsf{wadm}(f(F, \{\{a\}\})) = \{\emptyset, \{a, b\}\}$?

Possible solutions are shown in Figure 17; all frameworks depicted in there have the same weakly admissible extension set of $\{\emptyset, \{a, b\}\}$. Every solution required at least one additional argument, although solutions exist that utilise more than one auxiliary argument. The solutions are also (locally) minimal in the sense that removing an argument or an attack renders the respective solution invalid.

[10] Additionally, the framework translations featured in the work of Dvořák and Spanring [13] are generally covering; some are even *embedding*, meaning that no additional attacks between original arguments are introduced by the translation.

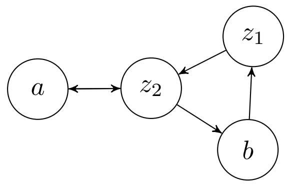

Figure 16: What modifications are needed to filter out one of this framework's weakly admissible extensions?

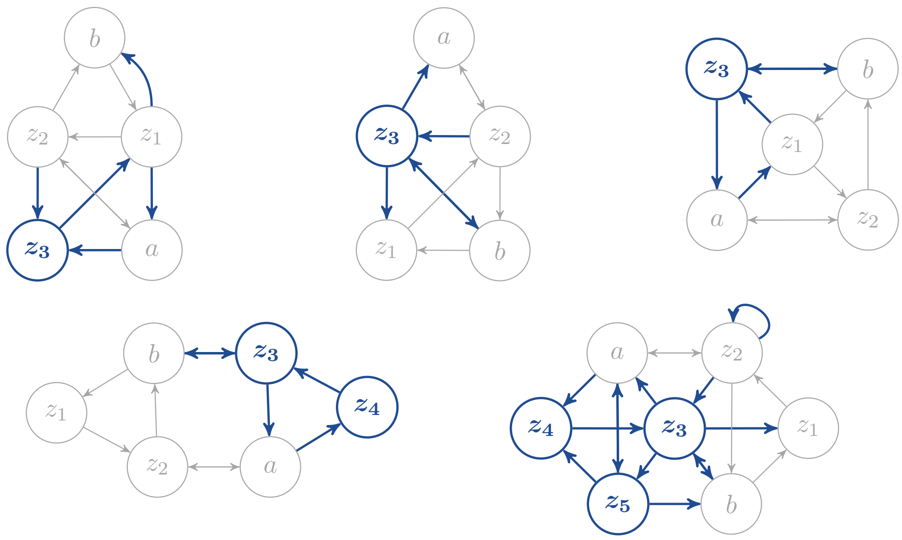

Figure 17: Solutions to the problem posed in Example 31. Added structures have been emphasized; all solutions are strict augmentations.

The solutions do not look particularly intuitive; rather, the fact that they actually do represent solutions seems to be more of a coincidence, and no obvious comprehensive and universally applicable pattern appears to emerge.

Let us see if the situation is any different for undisputed semantics.

Example 32. Figure 18 shows the counterpart of Figure 16, this time for the case of undisputed semantics. The undisputed extension set is identical to the weakly admissible one of the previous example, namely $\{\emptyset, \{a\}, \{a, b\}\}$, and we are looking again for a way to filter the extension $\{a\}$ while covering the original framework.

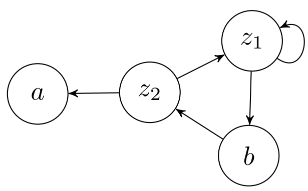

Figure 18: This framework has the undisputed extension set $\{\emptyset, \{a\}, \{a, b\}\}$.

Several minimal solutions are presented in Figure 19. All frameworks shown there have the same set of undisputed extensions, namely $\{\emptyset, \{a, b\}\}$.

Much like in the previous Example 31, here as well it appears difficult to recognise patterns that could provide a template for a universally applicable filter method.

The impression left by our investigations so far is ambivalent. On the one hand, the classical recipe of creating a base framework and then applying a filter seems to have ample expressive

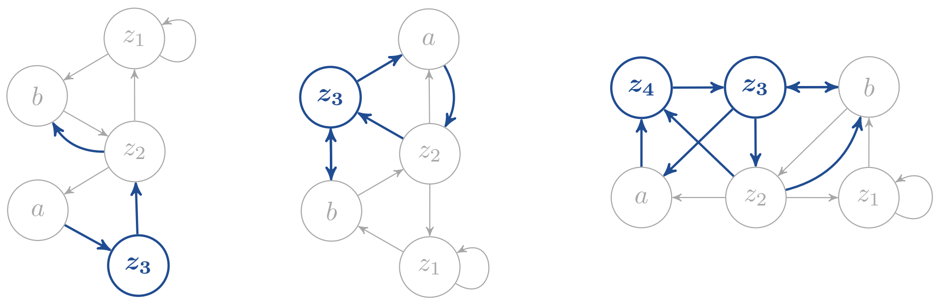

Figure 19: Some solutions to the problem from Example 32.

power even in the non-admissible case. On the other hand, the filter algorithms appear to be much more complex than in the classical case, provided they exist and can be found at all.

Next we explore strategies and construction methods that may be useful for the realisation of non-admissible extension sets. In doing so, we focus primarily on the undisputed semantics, as it is the conceptually simplest compared to the other non-admissible semantics that we studied so far.

4.5 Minimal scenarios

Examples 31 and 32 demonstrated that a filter problem can typically have many solutions which differ in the number of attacks and arguments being added to an argumentation framework. We would like to better understand what the minimum elements are that have to be present in a solution; to this end, we examine in the following the expressive power of a filter that only adds attacks, but no auxiliary arguments.

Example 33. For a set of arguments $A \subseteq \mathfrak{A}$, let $F = (A \cup \{z\}, \{(z,z)\} \cup \{(z,a) \mid a \in A\})$ be the base framework $F^{\mathsf{ud}}_{\mathbb{S}}$ from Proposition 21, where $\bigcup \mathbb{S} = A$, and $z \notin A$ is the auxiliary argument attacking all arguments in A. Let $a, b \in A$; adding the following attacks causes the stated filter operations on the framework's undisputed extension set:

- $a \to a$: removes $\{S \in \mathsf{ud}(F) \mid a \in S\}$;
- $a \to b$: removes $\{S \in \mathsf{ud}(F) \mid a \in S \wedge b \in S\}$;
- $a \to z$: removes $\{S \in \mathsf{ud}(F) \mid S \neq \bigcup \mathbb{S}\}$.

The operations from Example 33 apply only one attack each. There are also other, less obvious filter operations that use several attacks simultaneously.

Example 34. Figure 20 shows some filter operations on the base framework $F^{\mathsf{ud}}_{\mathbb{S}}$, instantiated for $\bigcup \mathbb{S} = \{a, b, c\}$; each filter operation uses multiple attacks but does not introduce any new arguments. To provide some intuition for item (c), we determine the extension set by observing that b and c mutually exclude each other and therefore cannot appear in the same extension. Meanwhile, a unconditionally defeats the already self-defeating z, and thus appears in every extension.

The obvious question now is whether filters of the type being analysed here offer sufficient expressive power to remove arbitrary elements from the extension set. The following case study shows that, in general, this question must be answered negatively.

Case Study 1. We conduct an exhaustive search over all attacks that can possibly be applied to $F^{\mathsf{ud}}_{\mathbb{S}}$ for a problem size of 3, i.e., $\bigcup \mathbb{S} = \{a, b, c\}$, and record their effect on $\mathsf{ud}(F^{\mathsf{ud}}_{\mathbb{S}})$. The number of scenarios to consider is $2^{4^2-4-3} - 1 = 511$, accounting for the four attacks that are invariably present in $F^{\mathsf{ud}}_{\mathbb{S}}$, as well as the exclusion of self-attacking arguments a, b, c, since self-attacking arguments are never included in any extension. From the 256 subsets of $2^{\{a,b,c\}}$, 38 do not contain all of the

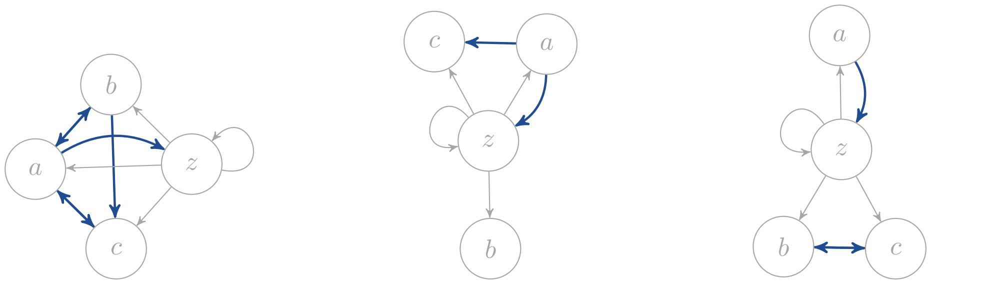

(a) $\mathsf{ud}(F) = \{\{a\}, \{b\}, \{c\}\}$ (b) $\mathsf{ud}(F) = \{\{a, b\}\}$ (c) $\mathsf{ud}(F) = \{\{a, b\}, \{a, c\}\}$

Figure 20: Filter operations that apply multiple attacks, and their outcomes.

arguments $\{a, b, c\}$ (as we are specifically interested in extension sets with three arguments). Of the remaining 218 sets, 65 are disjointly supported or distinct pairings. Subtracting these, we arrive at 153 interesting subsets of $2^{\{a,b,c\}}$; which can be assigned to 47 different isomorphism classes.[11] We would like to see how many of these isomorphism classes can be realised by our construction.

Result. Only 30 out of the 153 interesting extension sets, belonging to 10 different isomorphism classes, were realised (when using $F^{\mathsf{ud}}_{\mathbb{S}}$ as the base framework), although each of the 511 attack scenarios had an effect on the undisputed extension set. The results are summarised in Table 2, where isomorphic extension sets have been omitted.

Realised extension set	Attacks, augmenting $F^{\mathsf{ud}}_{\mathbb{S}}$
$\{\{a, b, c\}\}$	$a \to z$
$\{\{a, c\}, \{b, c\}\}$	$a \to b,\ a \to z,\ b \to a,\ b \to z$
$\{\{a\}, \{b, c\}\}$	$a \to b,\ a \to c,\ b \to a,\ b \to z$
$\{\{b\}, \{a, c\}, \{b, c\}\}$	$a \to b,\ a \to z,\ b \to a$
$\{\{a\}, \{c\}, \{b, c\}\}$	$a \to b,\ a \to c,\ a \to z,\ c \to a$
$\{\{a\}, \{b\}, \{c\}\}$	$a \to b,\ a \to c,\ a \to z,\ b \to a,\ b \to c,\ c \to a$
$\{\{a\}, \{b\}, \{c\}, \{b, c\}\}$	$a \to b,\ a \to c,\ a \to z,\ b \to a,\ c \to a$
$\{\emptyset, \{a\}, \{b\}, \{c\}\}$	$a \to b,\ a \to c,\ b \to c$
$\{\emptyset, \{a\}, \{b\}, \{c\}, \{b, c\}\}$	$a \to b,\ a \to c$
$\{\emptyset, \{a\}, \{b\}, \{c\}, \{a, c\}, \{b, c\}\}$	$a \to b$

Table 2: Results of Case Study 1. extension sets isomorphic to the ones presented here are not included.

As we will see, the low number of extension sets realised is at least partly due to the base framework offering insufficient potential for modification. Below we introduce a base framework that allows for a wider range of attack options.

Proposition 22. For a given $\mathbb{S} \subseteq 2^{\mathfrak{A}}$, let $A = \bigcup \mathbb{S}$, and for each $a \in A$, introduce a new argument z_a. The construction

$$\mathbb{S} \mapsto F^{\mathsf{ud}+}_{\mathbb{S}} = (A \cup \{z_a \mid a \in A\}, \{(z_a, a) \mid a \in A\} \cup \{(z_a, z_a) \mid a \in A\})$$

is a base framework with respect to ud, and we have $\mathsf{ud}(F^{\mathsf{ud}+}_{\mathbb{S}}) = 2^{(\bigcup \mathbb{S})}$ for any $\mathbb{S} \subseteq 2^{\mathfrak{A}}$.

Proof. Every $S \subseteq A$ is conflict-free; S is only attacked from self-attacking arguments $\{z_a \mid a \in A\}$, so that the reduct with respect to S is adm-vacuous. $\square$

Example 35. Figure 21 shows the base framework $F^{\mathsf{ud}+}_{\mathbb{S}}$ for $\bigcup \mathbb{S} = \{a, b, c\}$.

Compared to $F^{\mathsf{ud}}_{\mathbb{S}}$, the framework $F^{\mathsf{ud}+}_{\mathbb{S}}$ offers additional attack options.

[11] Two extension sets are *isomorphic* if the one extension set results from a one-to-one renaming of the arguments of the other extension set.

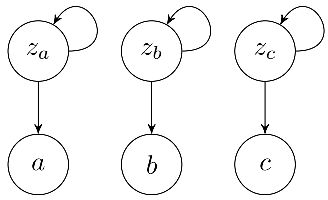

Figure 21: A base framework for ud, constructed from three arguments.

Example 36. Adding the following attacks has the stated effects on the undisputed extension set of $F_{\mathbb{S}}^{\mathsf{ud}+}$:

- $a \to z_a$: removes $\{S \in \mathsf{ud}(F) \mid a \notin S\}$;
- $a \to z_b$ $(a \neq b)$: removes $\{S \in \mathsf{ud}(F) \mid a \in S \wedge b \notin S\}$.

Repeating the previously conducted case study using this new base framework, we expect to find a wider range of extension sets being realised.

Case Study 2. We use the same setup as in Case Study 1 but exchange the base framework for $F_{\mathbb{S}}^{\mathsf{ud}+}$. The number of attack scenarios is now considerably higher, namely $2^{6^2-6-3} - 1 \approx 10^8$.

Result. 96 out of the 153 interesting extension sets (in 26 isomorphism classes) were realised; except for a tiny fraction of around 10^{-4}, all attack scenarios had an effect on the undisputed extension set of the original framework. The results include all realised sets from Case Study 1; we summarise them in Table 3.

The inability to realise all conceivable and interesting extension sets is still, at least in part, caused by the limited expressiveness of our construction, as the following example shows.

Example 37. The extension set $\{\emptyset, \{a, b, c\}\}$ was among the unrealised extension sets of Case Study 2. Figure 22 however shows that a realising framework exists, which is not an augmentation of the previously considered base framework.

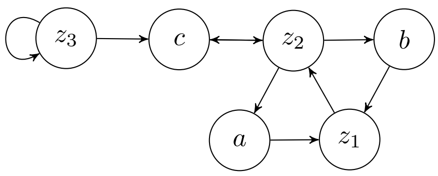

Figure 22: The undisputed extension set of this framework is $\{\emptyset, \{a, b, c\}\}$.

The minimal scenarios that we investigated in this subsection had the advantage that they allowed an exhaustive search over all possible modifications of the base framework. Between Case Studies 1 and 2 we saw some evidence that an increased complexity of the base framework may likely give way to an increased expressiveness of the resulting construction. Unfortunately, exhaustive analyses of even more complex scenarios than in Case Study 2 already become computationally unfeasible, so we will turn to a different strategy in the next subsection.

4.6 Transformations and building blocks

The consideration of minimal scenarios in the previous subsection soon led us to a complexity limit beyond which exhaustive searches were no longer feasible. In order to uncover further construction patterns, we pursue a more structured strategy in this subsection: instead of inserting arbitrary

Realised extension set	Attacks, augmenting $F_{\mathbb{S}}^{\mathsf{ud}+}$
$\{\{a,b,c\}\}$	$a \to z_a$, $a \to z_b$, $a \to z_c$
$\{\{a,c\},\{b,c\}\}$	$a \to b$, $a \to z_a$, $a \to z_c$, $b \to a$, $b \to z_c$
$\{\{a\},\{b,c\}\}$	$a \to b$, $a \to c$, $a \to z_a$, $b \to a$, $b \to z_c$
$\{\{b\},\{a,c\},\{b,c\}\}$	$a \to b$, $a \to z_a$, $a \to z_c$, $b \to a$
$\{\{a\},\{c\},\{b,c\}\}$	$a \to b$, $a \to c$, $a \to z_a$, $c \to a$
$\{\{a\},\{b\},\{c\}\}$	$a \to b$, $a \to c$, $a \to z_a$, $b \to a$, $b \to c$, $c \to a$
$\{\{a\},\{b\},\{c\},\{b,c\}\}$	$a \to b$, $a \to c$, $a \to z_a$, $b \to a$, $c \to a$
$\{\emptyset,\{a\},\{b\},\{c\}\}$	$a \to b$, $a \to c$, $b \to c$
$\{\emptyset,\{a\},\{b\},\{c\},\{b,c\}\}$	$a \to b$, $a \to c$, $z_c \to z_b$
$\{\emptyset,\{a\},\{b\},\{c\},\{a,c\},\{b,c\}\}$	$a \to b$, $z_c \to z_b$
$\{\{b,c\},\{a,b,c\}\}$	$b \to z_b$, $b \to z_c$
$\{\{c\},\{b,c\},\{a,b,c\}\}$	$a \to z_b$, $c \to z_c$
$\{\{c\},\{a,c\},\{b,c\}\}$	$a \to b$, $c \to z_c$
$\{\emptyset,\{a\},\{b,c\}\}$	$a \to b$, $a \to c$, $b \to z_a$, $b \to z_b$, $c \to a$
$\{\{c\},\{a,c\},\{b,c\},\{a,b,c\}\}$	$c \to z_c$
$\{\emptyset,\{c\},\{b,c\},\{a,b,c\}\}$	$a \to z_b$, $b \to z_c$
$\{\emptyset,\{c\},\{a,c\},\{b,c\}\}$	$a \to b$, $a \to z_c$, $b \to z_c$
$\{\{a\},\{b\},\{a,c\},\{b,c\}\}$	$a \to b$, $a \to z_a$, $b \to a$
$\{\emptyset,\{b\},\{a,c\},\{b,c\}\}$	$a \to b$, $a \to z_c$, $b \to a$, $c \to z_b$
$\{\emptyset,\{a\},\{c\},\{b,c\}\}$	$a \to c$, $b \to a$, $b \to z_c$
$\{\emptyset,\{c\},\{a,c\},\{b,c\},\{a,b,c\}\}$	$a \to z_c$, $b \to z_c$
$\{\emptyset,\{b\},\{c\},\{b,c\},\{a,b,c\}\}$	$a \to z_b$, $a \to z_c$
$\{\emptyset,\{b\},\{c\},\{a,c\},\{b,c\}\}$	$a \to b$, $a \to z_c$, $z_c \to z_b$
$\{\emptyset,\{a\},\{b\},\{a,c\},\{b,c\}\}$	$a \to b$, $b \to a$, $c \to z_a$
$\{\emptyset,\{b\},\{c\},\{a,c\},\{b,c\},\{a,b,c\}\}$	$a \to z_c$, $z_c \to z_b$
$\{\emptyset,\{a\},\{b\},\{c\},\{a,c\},\{b,c\},\{a,b,c\}\}$	$a \to z_c$, $b \to z_a$, $z_a \to c$, $z_b \to a$

Table 3: Results of Case Study 2, showing only representatives of their respective isomorphism classes. The first ten results (separated by the dashed line) realise the same extension sets as in Case Study 1.

attacks into a given base framework, we define a number of mappings, which we refer to as *transformations*, that each alter a given framework in a certain manner, and then proceed to combine these transformations.

The first transformation adds reciprocal attacks between two arguments.

Transformation 1 (Reciprocal attacks)**.** For $(\mathcal{A},\mathcal{R}) \in \mathfrak{F}_{\mathfrak{A}}$ and $a, b \in \mathcal{A}$, we define the mapping

$$\alpha(a,b) : (\mathcal{A},\mathcal{R}) \mapsto \begin{cases} (\mathcal{A}, \mathcal{R} \cup \{(a,b),(b,a)\}), & \text{if } a,b \in \mathcal{A}; \\ (\mathcal{A},\mathcal{R}), & \text{otherwise,} \end{cases}$$

with the special case

$$\hat{\alpha}(a) = \alpha(a,a).$$

The second transformation is a merge operation.

Transformation 2 (Merge)**.** Let $(\mathcal{A},\mathcal{R}) \in \mathfrak{F}_{\mathfrak{A}}$, let $S \subseteq \mathfrak{A}$, and let $z \in \mathfrak{A} \setminus \mathcal{A}$. We define a replace operator

$$r_S^z : \mathfrak{A} \to \mathfrak{A}, \quad a \mapsto r_S^z(a) = \begin{cases} z, & \text{if } a \in S; \\ a, & \text{otherwise,} \end{cases}$$

which we use in the mapping

$$\mu(S,z) : (\mathcal{A},\mathcal{R}) \mapsto (\{r_S^z(a) \mid a \in \mathcal{A}\}, \{(r_S^z(a), r_S^z(b)) \mid (a,b) \in \mathcal{R}\}).$$

The intuition is that $\mu(S,z)$ merges (possibly disjoint) parts of the framework, eliminating all arguments in S and replacing them with the single argument z. Attacks to and from S are then translated to attacks to and from z.

Example 38. Figure 23 illustrates an instance of the merge transformation μ.

Figure 23: Applying the transformation $\mu(\{b,e\},z)$ to the left framework yields the resulting framework to the right.

The final transformation removes arguments from the framework.

Transformation 3 (Removal)**.** For $F = (\mathcal{A}, \mathcal{R}) \in \mathfrak{F}_{\mathfrak{A}}$ and $S \subseteq \mathfrak{A}$, we define

$$\delta(S) : F \mapsto F|_{\mathcal{A} \setminus S} .$$

To ease presentation, we focus on the case in which we attempt to realise undisputed extension sets containing three arguments a, b, c. Our starting point this time is the framework from Figure 24, which is not a base framework (its undisputed extension set is $\{\emptyset\}$).

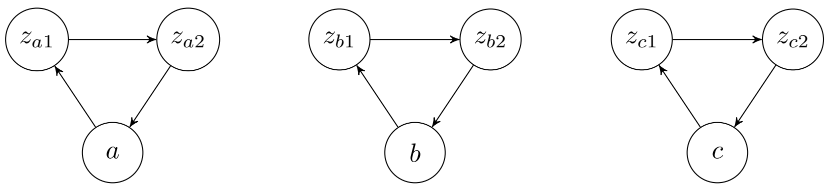

Figure 24: These three odd cycles serve as starting scenario to which subsequent transformations are applied.

Case Study 3. We apply the above defined transformations to the framework from Figure 24, which we call F_0. For the concrete case of the three arguments a, b, c, we identify the following *building blocks*, i.e. single and composed transformations which yield distinctive undisputed extension sets when applied to F_0:

- $\hat{\alpha}(z_{x2})$ for $x \in \{a, b, c\}$;
- $\alpha(x, z_{y2})$ for $x, y \in \{a, b, c\}$ and $x \neq y$;
- $\alpha(x, y) \circ \hat{\alpha}(z_{x1}) \circ \hat{\alpha}(z_{y1}) \circ \hat{\alpha}(z_{x2}) \circ \hat{\alpha}(z_{y2})$ for $(x, y) \in \{(a, b), (b, c), (c, a)\}$;
- $\mu(\{z_{x2}, z_{y2}\}, z)$ for $(x, y) \in \{(a, b), (b, c), (c, a)\}$ and a new argument z;
- $\delta(\{z_{x1}, z_{x2}, z\})$ for $x \in \{a, b, c\}$.

This way we can instantiate 18 concrete building blocks in total, which can then arbitrarily be combined, so that we arrive at $2^{18} = 262144$ different possibilities to consider. For each of these possibilities we compute the resulting framework and its undisputed extension set, and record any realisations that have not been encountered in Case Studies 1 and 2.

Result. The ten new construction recipes that were discovered this way are listed in Table 4. In summary, construction patterns for 35 different isomorphism classes representing extension sets that contain three arguments were found, ten of which are new constructions that were not realised in the previous case studies. Adding these to the 26 isomorphism classes that have been previously realised, we now arrive at realisations for 36 out of the 47 three-argument classes that are of interest.

Realised extension set	Transformations applied to F_0
$\{\emptyset, \{a, b, c\}\}$	$\alpha(b, z_{c2}) \circ \alpha(a, z_{b2}) \circ \hat{\alpha}(z_{a2})$
$\{\{b\}, \{a, b, c\}\}$	$\delta(\{z_{b1}, z_{b2}, z\}) \circ \alpha(a, z_{c2}) \circ \hat{\alpha}(z_{a2})$
$\{\emptyset, \{b\}, \{a, b, c\}\}$	$\alpha(a, z_{c2}) \circ \alpha(a, z_{b2}) \circ \hat{\alpha}(z_{b2}) \circ \hat{\alpha}(z_{a2})$
$\{\emptyset, \{a, b\}, \{a, c\}\}$	$\alpha(b, c) \circ \hat{\alpha}(z_{b1}) \circ \hat{\alpha}(z_{c1}) \circ \hat{\alpha}(z_{b2}) \circ \hat{\alpha}(z_{c2}) \circ$ $\alpha(c, z_{a2}) \circ \alpha(b, z_{a2})$
$\{\emptyset, \{a, c\}, \{a, b, c\}\}$	$\alpha(a, z_{c2}) \circ \alpha(b, z_{a2}) \circ \hat{\alpha}(z_{b2}) \circ \hat{\alpha}(z_{a2})$
$\{\emptyset, \{b\}, \{a, c\}, \{a, b, c\}\}$	$\alpha(a, z_{c2}) \circ \hat{\alpha}(z_{b2}) \circ \hat{\alpha}(z_{a2})$
$\{\emptyset, \{a, b\}, \{a, c\}, \{a, b, c\}\}$	$\alpha(c, z_{a2}) \circ \alpha(b, z_{a2}) \circ \hat{\alpha}(z_{c2}) \circ \hat{\alpha}(z_{b2})$
$\{\emptyset, \{b\}, \{c\}, \{a, b, c\}\}$	$\mu(\{z_{b2}, z_{c2}, z\}, z) \circ \alpha(a, z_{b2}) \circ \hat{\alpha}(z_{a2})$
$\{\emptyset, \{a\}, \{b\}, \{c\}, \{a, b\}, \{a, c\}, \{b, c\}\}$	$\mu(\{z_{a2}, z_{c2}, z\}, z) \circ \mu(\{z_{a2}, z_{b2}, z\}, z)$
$\{\emptyset, \{a\}, \{b\}, \{c\}, \{a, b\}, \{a, c\}, \{b, c\},$ $\{a, b, c\}\}$	$\hat{\alpha}(z_{c2}) \circ \hat{\alpha}(z_{b2}) \circ \hat{\alpha}(z_{a2})$

Table 4: Results of Case Study 3. With regard to the extension sets we again only list representatives of their respective isomorphism classes.

4.7 The cycle hub framework

We now examine another base framework, which is inspired from Case Study 3, particularly, from the merge operator (Transformation 2) applied to multiple odd cycles. First however we want to name a construction that we have already encountered several times so far.

Definition 26. For a given set of arguments A, the *cycle hub framework* $F_A^\triangle$ is constructed as follows. We start with $F = (\mathcal{A}, \mathcal{R})$ where $\mathcal{A} = A$, $\mathcal{R} = \emptyset$.

1. For each $a \in A$, add the arguments z_a, z'_a to F, where $z_a, z'_a \notin \mathcal{A}$ are new arguments, and add the attacks $a \to z_a$, $z_a \to z'_a$, $z'_a \to a$. This creates the odd cycles scenario from Figure 24 for an arbitrary set of arguments A.

2. Introduce a new argument $z_0 \notin \mathcal{A}$, and apply $\mu(\{z'_x \mid x \in A\}, z_0)$ to F.

Example 39. The left framework of Figure 25 shows the cycle hub framework $F_A^\triangle$ constructed for the three arguments $A = \{a, b, c\}$.

Proposition 23. For a finite $A \subset \mathfrak{A}$, $\mathsf{ud}(F_A^\triangle) = 2^A \setminus A$.

Proof. For any $S \subsetneq A$, $(F_A^\triangle)^S$ consists of a number of interlinked odd cycles around the central argument z_0, and there are obviously no non-empty admissible subsets. For $S = A$, $(F_A^\triangle)^S$ consists only of the unchallenged z_0 and is thus not vacuous. □

With respect to ud and $\mathbb{S} \subseteq 2^{\mathfrak{A}}$, $\mathbb{S} \mapsto F_{\bigcup \mathbb{S}}^\triangle$ is not a base framework since $\bigcup \mathbb{S} \notin \mathsf{ud}(F_{\bigcup \mathbb{S}}^\triangle)$, but we can easily complete the construction and turn it into a base framework for ud.

Definition 27. Given an extension set $\mathbb{S}$, the *cycle hub base framework* is obtained as follows.

1. Construct $F = (\mathcal{A}, \mathcal{R}) = F_{\bigcup \mathbb{S}}^\triangle$.

2. Create a new self-attacking argument $z_1 \notin \mathcal{A}$ and connect it to F via $z_1 \to z_0$.

We refer to the resulting framework as $F_{\mathbb{S}}^{\mathsf{ud}*}$.

Example 40. The right framework in Figure 25 shows $F_{\mathbb{S}}^{\mathsf{ud}*}$ constructed for an extension set consisting of the three arguments a, b, c.

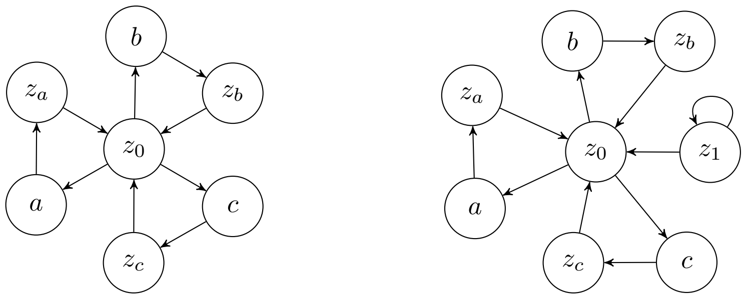

Figure 25: The frameworks $F_A^{\triangle}$ and $F_{\mathbb{S}}^{\mathsf{ud}*}$, for $A = \{a, b, c\} = \bigcup \mathbb{S}$.

The construction $\mathbb{S} \mapsto F_{\mathbb{S}}^{\mathsf{ud}*}$ is indeed a base framework.

Proposition 24. For any $\mathbb{S} \subseteq 2^{\mathfrak{A}}$, we have $\mathsf{ud}(F_{\mathbb{S}}^{\mathsf{ud}*}) = 2^{(\bigcup \mathbb{S})}$.

Proof. Let $S \in 2^{(\bigcup \mathbb{S})}$. If $S \subsetneq \bigcup \mathbb{S}$, $(F_{\mathbb{S}}^{\mathsf{ud}*})^S$ consists of a number of interlinked odd cycles around the central argument z_0, attacked by the self-attacking z_1, and there are obviously no non-empty admissible subsets. If $S = \bigcup \mathbb{S}$, only z_0, z_1 remain in the reduct; again, no non-empty subset is admissible. $\square$

We would like to investigate the expressiveness when using $F_{\mathbb{S}}^{\mathsf{ud}*}$ as a base framework in a similar manner as we did in the previous case studies. For the three-argument scenario we considered before, $F_{\mathbb{S}}^{\mathsf{ud}*}$ contains eight arguments and 11 attacks; this yields $2^{8^2-11-3} - 1 \approx 10^{15}$ possibilities to add further attacks, making it unfeasible to compute each single scenario. We resort to a randomized approach, in the hope that the relevant constructions are scattered densely enough across the search space.

Case Study 4. We apply attacks randomly to $F_{\mathbb{S}}^{\mathsf{ud}*}$ and compute the resulting extension set under undisputed semantics, collecting any realisations that we have not previously encountered in Case Studies 1, 2, and 3.

Result. Five additional realisation schemes have been discovered; representatives of their respective isomorphism classes are listed in Table 5. All in all, 41 different isomorphism classes were realised, including all realisations from the previous case studies, making this approach the one with the highest expressiveness so far.

4.8 Summary of constructions

In Subsections 4.5, 4.6, and 4.7 we attempted to derive construction methods from basic considerations; although our case studies were always limited to the concrete case of three-argument extension sets, the intention each time was to discover a systematic approach that also allows generalisations to arbitrary extension sets.

Tables 3, 4, and 5 summarise the construction recipes for three-argument undisputed extension sets that we have discovered so far. However, there are still extension sets in Σ_{ud} that cannot be realised by our construction methods; for some of these, we give concrete witnessing frameworks that prove their realisability in Table 6.

Finally, we list representatives of all three-argument extension sets from Σ_{ud} for which we do not have a realisation yet, in the open question raised below.

Realised extension set	Attacks, augmenting $F^{\mathsf{ud}*}_{\mathbb{S}}$
$\{\emptyset,\{a\},\{a,c\},\{b,c\},\{a,b,c\}\}$	$z_c \to z_a,\ z_c \to c,\ z_c \to b,\ z_b \to z_c,\ z_a \to z_b,\ z_a \to c,\ z_a \to b,\ z_a \to a,\ z_0 \to z_0,\ b \to z_a$
$\{\emptyset,\{a\},\{b\},\{b,c\},\{a,b,c\}\}$	$z_a \to z_a,\ z_0 \to z_1,\ z_1 \to c,\ z_1 \to a,\ c \to z_0,\ b \to z_c$
$\{\emptyset,\{a\},\{b\},\{c\},\{a,b\},\{a,b,c\}\}$	$z_c \to z_c,\ z_c \to b,\ z_b \to a,\ z_a \to z_b,\ z_a \to z_1,\ z_a \to c,\ z_0 \to z_1,\ z_1 \to c,\ z_1 \to a,\ c \to z_a,\ a \to z_0$
$\{\emptyset,\{a\},\{a,b\},\{a,c\},\{b,c\},\{a,b,c\}\}$	$z_c \to c,\ z_c \to b,\ z_c \to a,\ z_b \to c,\ z_b \to b,\ z_b \to a,\ z_a \to z_c,\ z_a \to z_b,\ c \to z_a,\ b \to z_a,\ a \to z_c,\ a \to z_b$
$\{\emptyset,\{a\},\{b\},\{a,b\},\{a,c\},\{b,c\},\{a,b,c\}\}$	$z_c \to z_a,\ z_a \to c,\ z_a \to b,\ z_a \to a,\ b \to z_a$

Table 5: Results from Case Study 4, showing only representatives of their respective isomorphism classes.

Extension set	Framework
$\{\emptyset,\{a\},\{b\},\{c\},\{a,b,c\}\}$	$(\{a,b,c,z_1,z_2,z_3,z_4,z_5\},\ \{(z_5,z_3),(z_5,b),(z_5,z_2),(a,z_5),(a,z_3),(z_4,z_5),(z_4,z_3),(z_4,z_2),(z_3,a),(z_3,b),(b,z_4),(z_2,b),(z_2,c),(z_1,a),(z_1,z_1),(z_1,c),(c,z_5),(c,z_2)\})$
$\{\emptyset,\{a\},\{b\},\{a,b\},\{a,c\},\{b,c\}\}$	$(\{a,b,c,z_1,z_2,z_3,z_4\},\ \{(z_4,b),(z_4,z_2),(a,z_4),(z_3,a),(z_3,c),(z_3,z_2),(b,z_3),(c,z_1),(z_2,a),(z_2,b),(z_2,c),(z_1,z_4),(z_1,z_2)\})$
$\{\emptyset,\{a\},\{b\},\{a,c\},\{b,c\},\{a,b,c\}\}$	$(\{a,b,c,z_1,z_2,z_3,z_4\},\ \{(z_4,b),(z_4,c),(a,z_4),(z_3,z_4),(z_3,z_1),(b,z_1),(z_2,a),(z_2,b),(z_2,z_2),(c,z_3),(z_1,a),(z_1,c)\})$

Table 6: Witnessing frameworks of some undisputed three-argument extension sets that were not realised by the constructions from Sections 4.5, 4.6, and 4.7.

Open Question 1. Are any of the following extension sets realisable by undisputed semantics, and what are their witnessing frameworks?

$$\{\emptyset,\{a,b\},\{a,c\},\{b,c\}\},$$
$$\{\emptyset,\{a\},\{a,b\},\{a,c\},\{b,c\}\},$$
$$\{\emptyset,\{a,b\},\{a,c\},\{b,c\},\{a,b,c\}\}.$$

4.9 A general solution for a special case

The final idea that we present in this section is a construction that is applicable to a wide range of non-admissible semantics, namely undisputed, strongly undisputed and weakly preferred semantics; its downside is that it works only for a limited set of cases, which are characterised as follows.

Definition 28. A set of sets $\mathbb{S} \subseteq 2^{\mathfrak{A}}$ is *uniquely indexed* if it is non-empty, and

$$\forall S \in \mathbb{S} : \exists a \in S : \forall S' \in \mathbb{S} : (a \in S' \iff S = S') .$$

For every such S and a, we say that a is a *unique index* of S.

Neither admissible nor weakly admissible semantics can realise uniquely indexed extension sets.

Corollary 2. If $\mathbb{S} \subseteq 2^{\mathfrak{A}}$ is uniquely indexed, then $\emptyset \notin \mathbb{S}$. □

We can show however that uniquely indexed extension sets exhibit the necessary required properties to be realised under ud, sud, and wpr.

Proposition 25. Every uniquely indexed $\mathbb{S} \subseteq 2^{\mathfrak{A}}$ is incomparable.

Proof. The unique indices of any $S_1, S_2 \in \mathbb{S}$ prohibit $S_1 \subseteq S_2$ if $S_1 \neq S_2$. □

Proposition 26. No uniquely indexed $\mathbb{S} \subseteq 2^{\mathfrak{A}}$ is disjointly supported.

Proof. If $\mathbb{S} = \{S\}$, then $S \neq \emptyset$ and $\bigcap \mathbb{S} = S \neq \emptyset$. Otherwise, every two sets $S_1, S_2 \in \mathbb{S}$ ($S_1 \neq S_2$) are in evident conflict because of the unique indices they contain. □

Proposition 27. No uniquely indexed $\mathbb{S} \subseteq 2^{\mathfrak{A}}$ is a distinct pairing.

Proof. For $S, S' \in \mathbb{S}$ with $S \neq S'$, we have $S \cup S' \notin \mathbb{S}$. □

Uniquely indexed extension sets are then realised under undisputed, strongly undisputed and weakly preferred semantics by Algorithm 1, which works as follows. First, mutual attacks are constructed between all unique indices. The remaining arguments are arranged in odd cycles; the unique indices then attack exactly those cycles whose arguments belong to their respective extensions. We call the resulting framework the *canonical uniquely indexed framework*.

Algorithm 1 Realise uniquely indexed extension sets.

Input: A uniquely indexed $\mathbb{S} = \{S_1, \dots, S_n\} \subseteq 2^{\mathfrak{A}}$ ($n \geq 1$).
let s_i be a unique index of S_i for $i \in \{1, \dots, n\}$.
let $S_i^* = S_i \setminus \{s_i\}$ for $i \in \{1, \dots, n\}$.
let $\mathcal{A} = \{s_1, \dots, s_n\}$, $\mathcal{R} = \{(s_i, s_j) \mid i, j \in \{1, \dots, n\}, i \neq j\}$.
for all $a \in \bigcup_i S_i^*$ **do** ▷ Construct cycles.
    let $z_1^a, z_2^a \in \mathfrak{A} \setminus \mathcal{A}$.
    $\mathcal{A} \leftarrow \mathcal{A} \cup \{z_1^a, z_2^a\}$.
    $\mathcal{R} \leftarrow \mathcal{R} \cup \{(z_1^a, a), (a, z_2^a), (z_2^a, z_1^a)\}$.
end for
for all $i \in \{1, \dots, n\}$ **do** ▷ Connect the cycles with the unique indices.
    for all $a \in S_i^*$ **do**
        $\mathcal{R} \leftarrow \mathcal{R} \cup \{(s_i, z_1^a)\}$.
    end for
end for
Output: $(\mathcal{A}, \mathcal{R})$.

Example 41. Figure 26 shows the resulting construction of Algorithm 1, applied to the input $\mathbb{S} = \{\{a, d\}, \{b, d, e\}, \{c, f\}\}$. The unique indices are a, b, and c.

Finally we show that the construction indeed produces the promised results.

Proposition 28. For any uniquely indexed $\mathbb{S} \subseteq 2^{\mathfrak{A}}$, let F be the canonical uniquely indexed framework. We have $\mathbb{S} = \mathsf{ud}(F) = \mathsf{sud}(F) = \mathsf{wpr}(F)$.

Proof. Let $\mathbb{S} = \{S_1, \dots, S_n\}$, and let s_i be a unique index of each S_i. We need to prove both of the following inclusions.

- "$\subseteq$": For any S_i, no indices s_i remain in the reduct F^{S_i}, and the odd cycles are either completely eliminated because they contain an argument $s \in S_i$ and are thus also attacked by s_i, or they are left intact and contain no non-empty, admissible subset; so we have $S_i \in \mathsf{ud}(F)$. The odd cycles can not contain an undisputed subset either, so we have $S_i \in \mathsf{sud}(F)$ as well. Lastly, none of the remaining odd cycles attack any arguments of S_i in F, so $S_i \in \mathsf{wadm}(F)$; no superset $S' \supsetneq S_i$ can be admissible in F, since in the reduct $F^{S'}$, there would be attackers $z_1^a \to a$ for each $a \in S' \setminus S_i$, so $S_i \in \mathsf{wpr}(F)$.

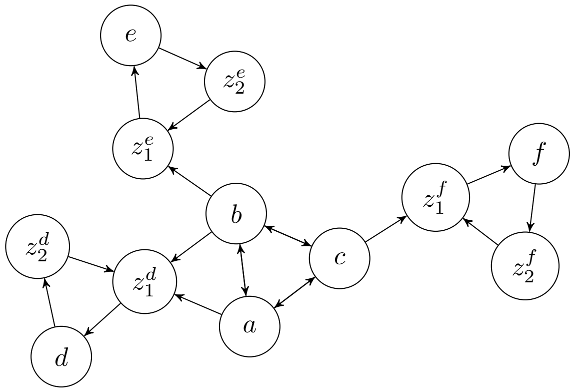

Figure 26: The canonical uniquely indexed framework, constructed for the set of argument-sets $\mathbb{S} = \{\{a, d\}, \{b, d, e\}, \{c, f\}\}$.

- "$\supseteq$": Every $S \in \mathsf{ud}(F)$ contains exactly one index s_i: if it contained none, then one of the s_i would be invariably part of the admissible extensions of the reduct with respect to S, since the s_i are not attacked from any other structures in F; S can not contain more than one s_i however, since that would introduce a conflict. Now consider the odd cycles that are not attacked by s_i: none of the arguments can be included in S, because the resulting reduct would contain non-empty admissible sets. In fact, the only possibility for a vacuous reduct is when all arguments are included in S which are contained in cycles that are attacked by s_i. Because of $\mathsf{adm}(F^S) \subseteq \mathsf{ud}(F^S)$, we can make an analogous argument for sud. Going further, the indices are uniquely included in each $S \in \mathsf{wadm}(F)$ because they are individually admissible, and their attacks break the corresponding odd cycles, making their contained arguments eligible for inclusion in S. Proposition 25 then allows the weakly preferred extensions to be realised. □

5 Discussion

We set out to explore the expressive power of various non-admissible semantics, in particular the weakly admissible, weakly preferred, undisputed, and strongly undisputed semantics. Comparing the signatures of weakly admissible and weakly preferred semantics to their classical counterparts, we found that, while they are not necessarily conflict-sensitive, they still retain the properties of non-emptiness (respectively, inclusion of the empty set) and, in the case of the weakly preferred semantics, incomparability. We then introduced the notions of disjointly supported extension sets as well as distinct pairings, and showed that the these classes of extension sets are unrealisable by the undisputed semantics. We estimated characterisations of the signatures of these semantics, which can be turned into precise characterisations if construction methods can be specified that realise extension sets under the respective semantics. In the search for such novel construction methods we orientated ourselves on the construction methods used in the classical cases. We identified four key features of classical constructions: the separation of extensions (i. e., the circumstance that the union of two extensions is not itself an extension); the suppression of auxiliary arguments introduced specifically for the construction so that they do not appear in the extensions; the presence of base frameworks; and the elimination of unwanted extensions as a result of a filter. While transferring these features to non-admissible cases, we found that different mechanisms for separating extensions were at work than in the classical scenarios, where the separation of extensions is always a consequence of evident conflict. We formalised these mechanisms and derived some necessary constructive properties of realising frameworks. We applied combinations of base frameworks and filters to examples, where we found that, while the combination of these concepts appears to have great expressive power, the resulting constructions themselves did not appear to be particularly comprehensible. In order to gain an understanding of possible construction patterns that realise non-admissible extension sets, we conducted a number of case studies that explored approaches of varying complexity; but even the large amount of constructions we obtained as a result did not appear to provide us with a comprehensive design idea, let alone a blueprint for a generic construction algorithm—we may marvel at the ingenuity of concrete realisations but fail to extract an underlying design principle. At least, however, we were able to present a solution for the special case of uniquely indexed extension sets, which proved to be applicable to undisputed, strongly undisputed, and weakly preferred semantics at the same time.

We briefly discuss some related works. We already mentioned the principle-based approach of Van Der Torre and Vesic [22] that evaluates semantics according to certain principles that they may satisfy; these principles typically relate structural features of the framework to properties of the extension set (the exceptions being the principles of *I-maximality*,[12] tightness, conflict-sensitivity, and com-closure, which we also considered), while we are only interested in properties of the extension set itself. We also mentioned the work of Dvořák and Woltran [15] as well as Dvořák and Spanring [13], who compare the expressiveness of semantics in terms of translatability of frameworks, but do not characterise semantics based on their signatures. Dyrkolbotn [16] constructs frameworks under labelling-based semantics with the help of auxiliary arguments and shows that, for the original arguments, arbitrary labellings can be realised under preferred and semi-stable[13] semantics (this does not contradict the results of Dunne *et al.* [12] that we build upon, since they also take into account the acceptability status of auxiliary arguments). Pührer [19] generalises the formalism from argumentation frameworks to abstract dialectical frameworks[14] and gives realisations of three-valued interpretations under certain classical semantics. Linsbichler, Pührer, and Strass [18] devise an algorithm that realises knowledge bases under a number of formalisms and semantics; so far, their implementation as well has been limited to several classical semantics. The question of expressiveness is also, in practical applications, typically linked to computational complexity; a representation is of little use if certain "canonical" problems, such as existence, credulous and sceptical acceptance, and verification [11] cannot be efficiently computed [20, 11].

To the best of our knowledge, no attempt in the way of the approach of Dunne *et al.* [12] to characterise non-admissible semantics in terms of their signatures has been carried out before. Our estimated characterisation of undisputed extension sets as non-empty and neither disjointly supported nor distinct pairings was prompted by our case studies, which hinted at possible non-realisable patterns; it should be instructive to conduct analogous studies for other non-admissible semantics and see whether the results provide indications that allow further refinements of our estimates. Perhaps an attempt to adapt the framework by Linsbichler, Pührer, and Strass [18] to deal with semantics that are of interest to us may shed more light on the fundamental reasons of the peculiar resistance that non-admissible semantics put up against their systematic realisation.

References

[1] Pietro Baroni, Martin Caminada, and Massimiliano Giacomin. Abstract argumentation frameworks and their semantics. In Pietro Baroni, Dov Gabbay, Massimiliano Giacomin, and Leendert van der Torre, editors, *Handbook of Formal Argumentation*, volume 1, pages 159–236. College Publications, 2018.

[2] Pietro Baroni and Massimiliano Giacomin. On principle-based evaluation of extension-based argumentation semantics. *Artificial Intelligence*, 171(10-15):675–700, 2007.

[3] Ringo Baumann, Gerhard Brewka, and Markus Ulbricht. Revisiting the foundations of abstract argumentation–semantics based on weak admissibility and weak defense. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 2742–2749, 2020.

[12] This term is synonymous to what we called *incomparability*.

[13] *Semi-stable* labellings [7] are complete labellings that minimise the set of *undecided* arguments.

[14] In an *abstract dialectical framework* (ADF) [6], the attack relation is replaced and generalised by an acceptance condition, i.e., a propositional formula for each argument, which expresses the argument's acceptance in terms other arguments' state of acceptance.

[4] Ringo Baumann, Gerhard Brewka, and Markus Ulbricht. Shedding new light on the foundations of abstract argumentation: Modularization and weak admissibility. *Artificial Intelligence*, 310:103742, 2022.

[5] Lydia Blümel and Matthias Thimm. Revisiting vacuous reduct semantics for abstract argumentation. In *Proceedings of the 27th European Conference on Artificial Intelligence (ECAI'24)*, October 2024.

[6] Gerhard Brewka and Stefan Woltran. Abstract dialectical frameworks. In *Proceedings of the Twelfth International Conference on Principles of Knowledge Representation and Reasoning*, pages 102–111, 2010.

[7] Martin WA Caminada and Dov M Gabbay. A logical account of formal argumentation. *Studia Logica*, 93:109–145, 2009.

[8] Marcos Cramer and Leon van der Torre. Scf2-an argumentation semantics for rational human judgments on argument acceptability. In *8th Workshop on Dynamics of Knowledge and Belief (DKB-2019) and 7th Workshop KI & Kognition (KIK-2019)*, 2019.

[9] Pierpaolo Dondio. Weakly-admissible semantics and the propagation of ambiguity in abstract argumentation semantics. 2019.

[10] Phan Minh Dung. On the acceptability of arguments and its fundamental role in nonmonotonic reasoning, logic programming and n-person games. *Artificial intelligence*, 77(2):321–357, 1995.

[11] Paul E Dunne. Well, to be honest, i wouldnt start from here at all. In *Computational Models of Argument*, pages 3–14. IOS Press, 2022.

[12] Paul E Dunne, Wolfgang Dvořák, Thomas Linsbichler, and Stefan Woltran. Characteristics of multiple viewpoints in abstract argumentation. *Artificial Intelligence*, 228:153–178, 2015.

[13] Wolfgang Dvořák and Christof Spanring. Comparing the expressiveness of argumentation semantics. *Journal of Logic and Computation*, 27(5):1489–1521, 2017.

[14] Wolfgang Dvořák, Markus Ulbricht, and Stefan Woltran. Recursion in abstract argumentation is hard—on the complexity of semantics based on weak admissibility. *Journal of Artificial Intelligence Research*, 74:1403–1447, 2022.

[15] Wolfgang Dvořák and Stefan Woltran. On the intertranslatability of argumentation semantics. *Journal of Artificial Intelligence Research*, 41:445–475, 2011.

[16] Sjur Kristoffer Dyrkolbotn. How to argue for anything: Enforcing arbitrary sets of labellings using afs. In *Fourteenth International Conference on the Principles of Knowledge Representation and Reasoning*, 2014.

[17] Antonis C Kakas and Paolo Mancarella. On the semantics of abstract argumentation. *Journal of Logic and Computation*, 23(5):991–1015, 2013.

[18] Thomas Linsbichler, Jörg Pührer, and Hannes Strass. A uniform account of realizability in abstract argumentation. In *ECAI 2016*, pages 252–260. IOS Press, 2016.

[19] Jörg Pührer. Realizability of three-valued semantics for abstract dialectical frameworks. In *Proceedings of the 24th International Conference on Artificial Intelligence*, pages 3171–3177, 2015.

[20] Hannes Strass. Expressiveness of two-valued semantics for abstract dialectical frameworks. *Journal of Artificial Intelligence Research*, 54:193–231, 2015.

[21] Matthias Thimm. On undisputed sets in abstract argumentation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 37, pages 6550–6557, 2023.

[22] Leon van der Torre and Srdjan Vesic. The principle-based approach to abstract argumentation semantics. *IfCoLog Journal of Logics and Their Applications*, 2017.