FernUniversität in Hagen
Faculty of Mathematics and Computer Science

AIG
Artificial Intelligence Group

Towards Characterisations of Non-Admissible Semantics in Abstract Argumentation

Master's Thesis

in partial fulfillment of the requirements for
the degree of Master of Science (M.Sc.)
in Practical Informatics (Praktische Informatik)

submitted by
Theo Doukas

First examiner: Prof. Dr. Matthias Thimm
Artificial Intelligence Group

Advisor: Prof. Dr. Matthias Thimm
Artificial Intelligence Group

Statement

Ich erkläre, dass ich die Masterarbeit selbstständig und ohne unzulässige Inanspruchnahme Dritter verfasst habe. Ich habe dabei nur die angegebenen Quellen und Hilfsmittel verwendet und die aus diesen wörtlich oder sinngemäß entnommenen Stellen als solche kenntlich gemacht. Die Versicherung selbstständiger Arbeit gilt auch für enthaltene Zeichnungen, Skizzen oder graphische Darstellungen. Die Arbeit wurde bisher in gleicher oder ähnlicher Form weder derselben noch einer anderen Prüfungsbehörde vorgelegt und auch nicht veröffentlicht. Mit der Abgabe der elektronischen Fassung der endgültigen Version der Arbeit nehme ich zur Kenntnis, dass diese mit Hilfe eines Plagiatserkennungsdienstes auf enthaltene Plagiate geprüft werden kann und ausschließlich für Prüfungszwecke gespeichert wird.

	Yes	No
I agree to have this thesis published in the library.	☑	☐
I agree to have this thesis published on the webpage of the artificial intelligence group.	☑	☐
The thesis text is available under a Creative Commons License (CC BY-SA 4.0).	☑	☐
The source code is available under a GNU General Public License (GPLv3).	☑	☐
The collected data is available under a Creative Commons License (CC BY-SA 4.0).	☑	☐

Stuttgart, 4 January 2025

(Place, Date) (Signature)

Abstract

An abstract argumentation framework is a formalism that uses a directed graph of interacting arguments to represent knowledge, thus modelling an exchange of opposing viewpoints. Semantics determine the sets of simultaneously acceptable arguments, the so-called extension-sets. We analyse representatives of the class of non-admissible semantics, particularly the undisputed, strongly undisputed, weakly admissible, and weakly preferred semantics, and determine properties of the extension-sets they produce, as well as structural features of frameworks that realise extension-sets under these semantics. We describe a class of extension-sets for which we show that it is not realisable under undisputed semantics. We also identify concepts 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.

Contents

1 Introduction

Abstract argumentation frameworks [Dun95] 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 abstract entities, but without internal structure, which attempt to invalidate each other through mutual attacks. Accordingly, an abstract argumentation 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. Taken on its own, this criterion can be seen as somewhat naive, as it does not consider possible attacks from arguments that lie outside the selected set. To safeguard against such attacks, one can further demand that all attacking arguments should themselves be counter-attacked from the selected set; one can go so far as to demand that the selected set attacks all arguments that are not part of it. It is not hard to see that, in the latter case, the existence of a set satisfying these conditions is no longer certain.

Thus, an essential parameter for finding a solution to an abstract argumentation problem is the specification of the exact conditions according to which the sets of valid arguments, in the following referred to as the *extensions* of the framework, are selected; these conditions are expressed by the *semantics* of the argumentation framework, namely in the form of a mapping from argumentation frameworks to their extensions. The focus of the present work lies on the expressive power of these mappings, i.e., on the question which sets of extensions the different semantics are able to generate at all.

1.1 Background and Relation to Previous Work

Abstract argumentation frameworks were introduced in the foundational work of Dung [Dun95]; the semantics described in there, which are now commonly referred to as *classical* semantics, are typically characterised by the fact that their extensions *defend* all attacks directed against any of their contained arguments, meaning that each argument attacking the extension is in turn attacked by another argument from the extension itself. Conflict-free sets of arguments exhibiting this property are called *admissible*.

One of the consequences arising from this concept of admissibility has already been questioned by Dung himself [Dun95, 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.

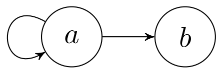

Figure 1: In this simple framework, b is attacked by a nonsensical attacker a.

The consideration of this example, as well as other similar examples, inspired the development of so called *non-admissible* semantics, which aim to weaken the concept of admissibility to such an extent that attacks by arguments that are in some way not to be regarded as serious do not have to be fended off by extensions. This new class of semantics includes the semantics based on the concept of *weak admissibility* by Baumann, Brewka and Ulbricht [BBU20, BBU22], as well as the semantics defined on the basis of *undisputed sets* by Thimm [Thi23]. Another approach is Dondio's also called "weakly admissible" semantics [Don19], which Baumann *et al.* compare along with other alternative approaches to their own [BBU22]. 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 [BG07] enumerated various principles and studied the fulfilment of these principles by the individual semantics; this consideration was extended by van der Torre and Vesic [vdTV17] to include further principles and semantics. Dvořák and Woltran [DW11] as well as Dvořák and Spanring [DS17] 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, Dvořák, Linsbichler, and Woltran [DDLW15] 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 extend the scope of consideration to non-admissible 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.

1.2 Contributions

The main idea of this thesis is to identify patterns and approaches found in the classical methods that determine properties of signatures and construct frameworks that realise extension-sets, and to investigate to what extent they are transferable to the non-admissible cases, particularly the weakly admissible, weakly preferred, undisputed and strongly undisputed semantics. To this end, we proceed as follows:

- We lay the necessary groundwork by recapitulating Dung's seminal theory of abstract argumentation and by introducing proponents of the novel class of non-admissible semantics in Section 2, where we aim to give a concise but thorough overview of the necessary prerequisites required to understand the present context as well as related work.
- We investigate the limits within which classical properties of signatures can be transferred to non-admissible semantics in Section 3. We introduce the class of *disjointly supported* extension-sets and show that they are not realisable by undisputed semantics.
- To characterise the signature of a semantics, one inevitably needs to specify a construction method that is able to realise all extension-sets in question. Thus in Section 4, we describe classical construction methods and identify common ideas and concepts: we observe that, typically, a construction method needs to separate extensions, make use of a base framework, subsequently apply filtering, and suppress any auxiliary arguments used in the construction.
- We transfer these ideas to non-admissible semantics starting with Section 5. We show that there are mechanisms at work that ensure that extensions remain separated from each other in the extension-set (in the sense that their union is not part of the extension-set despite no apparent conflicts exist), which differ from the analogous classical mechanisms, and go on to characterise structural properties of frameworks that implement these mechanisms. We generalise a proposition of Baumann *et al.* concerning the invariance of the order of reduct formation. We formally introduce the terms *base framework* and *filter*, and show that they play an important role in the classical construction schemes.
- Drawing on several case studies, we develop an intuition for the expressive power of the combination of base framework and filter in the non-admissible case in Section 6. There we introduce the *cycle hub base framework*, which has the highest expressive power among the constructions that we examine, and is the most promising generic candidate for the realisation of undisputed extension-sets. We introduce the class of *uniquely indexed* extension-sets and show that they can be realised by the undisputed, strongly undisputed, and the weakly preferred semantics.

We summarise and conclude our work in Section 7.

2 Preliminaries

This section explains notations and definitions, and reproduces results from existing literature insofar as they are relevant to further considerations.

2.1 Abstract Argumentation Frameworks and Classical Semantics

Following the presentation by Dung [Dun95] 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 $\mathfrak{F}_{\mathfrak{A}}$.

We now introduce basic notations. For $F = (\mathcal{A}, \mathcal{R})$ and $(a, b) \in \mathcal{R}$, we write $a \to_{\mathcal{R}} b$, or shorter $a \to b$ 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 $a \in \mathcal{A}$ (symbolically: $S \to a$) 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 $a \to S$. Extending the notation in the obvious way to two sets $S, T \subseteq \mathcal{A}$, we write $S \to T$ to denote that there are $s \in S, t \in T$ such that $s \to t$. Additionally, given a set $S \subseteq \mathcal{A}$, $S^+ = \{a \in \mathcal{A} \mid S \to a\}$ denotes the set of all arguments attacked by S, while $S^- = \{a \in \mathcal{A} \mid a \to S\}$ is the set of all attackers of S. Finally, we call $S^\oplus = S \cup S^+$ the *range of* S.

Definition 2 (Classical Defence)**.** 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 $a \to S$, we have $D \to a$.

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

Definition 3 (Conflict-Free and Admissible Sets)**.** Let $(\mathcal{A}, \mathcal{R}) \in \mathfrak{F}_{\mathfrak{A}}$ be a framework. A set $S \subseteq \mathcal{A}$ is *conflict-free* if it does not contain two arguments $a, b \in S$ with $a \to b$; furthermore, S is *admissible* if 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 (Semantics and Extensions)**.** A *semantics* is a mapping $\sigma : \mathfrak{F}_{\mathfrak{A}} \to 2^{2^{\mathfrak{A}}}$ 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 (in Definition 5) defined semantics were described by Dung [Dun95]; we will refer to them as the *classical* semantics.

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

- *preferred*, if it is a $\subseteq$-maximal admissible subset of $\mathcal{A}$;

- *complete*, if it is admissible and includes all arguments that it defends;
- *grounded*, if it is complete and $\subseteq$-minimal in $\mathcal{A}$;
- *stable*, if it is conflict-free and attacks all arguments (of $\mathcal{A}$) it does not contain.

These 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*).

In order to standardise the notation, we introduce the following designations for specific semantics.

Definition 6. Applied to a framework $F = (\mathcal{A}, \mathcal{R}) \in \mathfrak{F}_{\mathfrak{A}}$, we define:

- $\mathsf{cf} : F \mapsto \{S \subseteq \mathcal{A} \mid S \text{ is conflict-free in } F\}$;
- $\mathsf{naive} : F \mapsto \{S \subseteq \mathcal{A} \mid S \in \mathsf{cf}(F) \text{ and } S \text{ is } \subseteq\text{-maximal in } \mathcal{A}\}$;
- $\mathsf{adm} : F \mapsto \{S \subseteq \mathcal{A} \mid S \text{ is admissible in } F\}$;
- pr, com, grd, and stb designate the preferred, complete, grounded, and stable semantics.

We say that an argument of a framework F is *credulously accepted* with regard to a semantics σ if it appears in at least one extension of $\sigma(F)$, and that it is *sceptically accepted* if it appears in all extensions of $\sigma(F)$.

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 Semantics Based on Weak Admissibility and Weak Defence

Recalling the introductory example from Figure 1, we can generalize the pattern of a nonsensical or non-serious attack to attacks originating from *odd cycles* (Figure 2), as in the following example.

Example 1. Consider the odd cycle $S = \{a_1, \ldots, a_5\}$ of the rightmost framework in Figure 2. No subset of S is admissible in that framework, since every conflict-free subset of S (for example, $\{a_1, a_4\}$) has attackers (continuing the example, $\{a_3, a_5\}$) against whom it does not completely defend (the attack by a_3 is undefended). The same applies analogously to the other frameworks.

The frameworks in Figure 2 model situations where an argument is attacked from a point of view that itself is unsustainable. Yet, classical semantics demand that b does defend itself if it is to be accepted as valid. It appears that, in order to remedy the issue, the classical notion of admissibility needs to be adapted.

The intuition behind the following definition is to remove from consideration all arguments of a set and those that are attacked by the set itself, since these arguments cannot pose a threat any more.

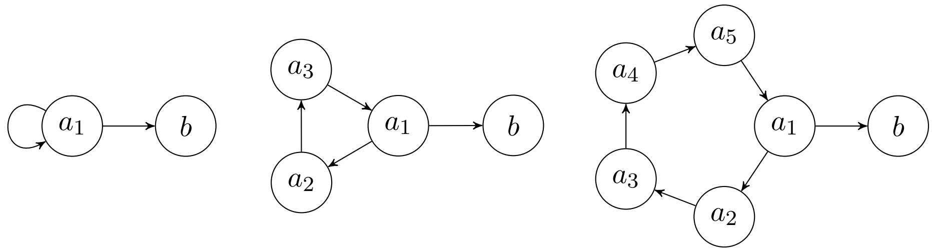

Figure 2: Odd cycles of length 1, 3, and 5 attacking an argument b; none of the attacks on b are considered serious.

Definition 7 (Reduct [BBU22])**.** 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 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.

Building on the concept of the reduct, we introduce the following modified notion of admissibility.

Definition 8 (Weak Admissibility [BBU22])**.** 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[1]

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

Admissible sets defend all of their elements, leaving no attackers in the reduct, so that the body of Definition 8 holds trivially for admissible sets as well.

Proposition 1 ([BBU22])**.** Admissibility implies weak admissibility. □

Example 2. Returning to the rightmost framework in Figure 2 (call it F), let us check for weak admissibility of the attacked set $\{b\}$. The reduct $F_1 = F^{\{b\}}$ consists of the odd cycle $S = \{a_1, \ldots, a_5\}$ 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 admissible indeed and does attack S' in F_1, so that S' cannot be weakly admissible in F_1 itself, rendering its attack on $\{b\}$ in F not serious. The cases for other conflict-free subsets of S are analogous and consequently, $\{b\}$ is weakly admissible.

The following weak notion of defence is slightly more involved than the classical one, since it needs to make sure that only serious attacks are defended against, and that the defenders themselves are sufficiently credible.

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

Definition 9 (Weak Defence [BBU22]). Let $F = (\mathcal{A}, \mathcal{R}) \in \mathfrak{F}_{\mathfrak{A}}$ be a framework, and let $D, S \subseteq \mathcal{A}$. D *weakly defends* S iff any attacker $a \to S$ ($a \in \mathcal{A}$) is either directly attacked by D; or (i) the attacker a does neither originate from D, nor from a weakly admissible set of the reduct F^D, and (ii) the defending set D is itself contained in a weakly admissible set of F.

Based on weak admissibility and defence, we can define counterparts to classical semantics from Definition 5.

Definition 10 ([BBU22]). Let $F = (\mathcal{A}, \mathcal{R}) \in \mathfrak{F}_{\mathfrak{A}}$. An extension is said to be:

- *weakly preferred*, if it is a $\subseteq$-maximal weakly admissible subset of $\mathcal{A}$;
- *weakly complete*, if it is weakly admissible and contains all of its supersets that it weakly defends;
- *weakly grounded*, if it is a $\subseteq$-minimal weakly complete subset of $\mathcal{A}$.

We denote the weakly preferred, the weakly complete, and the weakly grounded semantics with wpr, wcom, and wgrd, in accordance with Definition 6.

In addition to the above, Baumann *et al.* [BBU22] also investigate weak versions of stable semantics and find that they either coincide with weakly preferred or with classical stable semantics, concluding that a weak version of stable semantics is not obtained in an obvious fashion as was the case with the other weak versions of classical semantics.

2.3 Undisputed Sets

We have seen in Example 2 how the definition has to be applied recursively in order to determine the weak admissibility of a given set, suggesting that reasoning under weakly admissible semantics may be computationally challenging. Indeed, Dvořák, Ulbricht, and Woltran [DUW22] show that certain standard decision problems, such as the verification whether a given set is weakly admissible, are PSPACE-complete.[2]

An alternative approach to address the problems associated with semantics that build on the concept of admissibility and that even is designed to be more tractable computationally is presented by Thimm [Thi23]. We start with the following two preliminary definitions.

Definition 11 (Vacuity [Thi23]). A 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. $\sigma(F) \subseteq \{\emptyset\}$.

A framework that is vacuous with respect to a given semantics can be considered nonsensical, since none of its arguments are acceptable under that semantics.

[2] PSPACE-complete problems are among the computationally most challenging problems, and are at least as hard as NP-complete problems [Pap03].

Definition 12 (Vacuous Reduct Semantics [Thi23])**.** 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 with σ^τ.

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

Definition 13 ([Thi23])**.** We call

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

Example 3. We return to Figure 2 once more and want to assess undisputedness of $\{b\}$ in the rightmost framework, which we call F. $\{b\}$ is conflict-free, and we have seen earlier (in Example 1) that the reduct $F^{\{b\}}$ contains no admissible sets, so $\{b\}$ is undisputed.

It is worth noting that vacuous reduct semantics provide a template for describing other new as well as existing semantics, as in the following proposition.

Proposition 2 ([Thi23])**.** $\mathsf{pr} = \mathsf{adm}^{\mathsf{adm}}$. □

Corollary 1. An undisputed extension that is admissible is also preferred. □

2.4 Other Non-admissible Semantics

The semantics presented so far may harbour further possible inconsistencies, as the next example shows.

Example 4. Consider the framework in Figure 3 (call it F), where the arguments a and b both attack x, but mutually refute each other at the same time. We can imagine that a and b are incriminating statements that attack a claim of innocence in the form of x. However, since both statements contradict each other, a semantics that prioritises a presumption of innocence—for instance, in a legal context—in this case would have to dismiss both a and b, and accept x. Yet, we have $\mathsf{pr}(F) = \mathsf{wpr}(F) = \mathsf{ud}(F) = \mathsf{sud}(F) = \{\{a\}, \{b\}\}$, and $\mathsf{com}(F) = \mathsf{adm}(F) = \mathsf{wadm}(F) = \{\emptyset, \{a\}, \{b\}\}$.

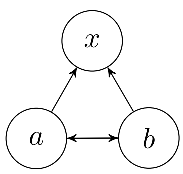

Figure 3: The two attackers of x refute each other. Should x be accepted?

The above example is due to Dondio [Don19]. He proposes a semantics that is based on Caminada and Gabbay's *labelling*-approach [CG09], which assigns labels to arguments in the following manner. Note that the term *weakly admissible* is also used below, as in Definition 8; both concepts are however unrelated.

Definition 14 ([Don19])**.** Let $F = (\mathcal{A}, \mathcal{R}) \in \mathfrak{F}_{\mathfrak{A}}$. A *weakly admissible labelling of F* is a total function $\mathcal{A} \to \{in, out, undecided\}$ so that

- when a is labelled *in*, then none of its attackers are labelled *in*;
- when a is labelled *out*, then at least one of its attackers is labelled *in*;
- when a is labelled *undecided*, then no attacker is labelled *in*, and at least one attacker is labelled *undecided* as well.

Arguments labelled *in* are considered explicitly accepted, arguments labelled *out* are explicitly rejected, and no judgement with regard to acceptance is made for the *undecided* arguments.

Example 5. With regard to the framework in Figure 3, we obtain the following weakly admissible labellings: $\{(a \mapsto in, b \mapsto out, x \mapsto out), (a \mapsto out, b \mapsto in, x \mapsto out), (a \mapsto undecided, b \mapsto undecided, x \mapsto in), (a \mapsto undecided, b \mapsto undecided, x \mapsto undecided)\}$, so that x is at least credulously accepted.

We mention in passing other semantics that do not make use of the concept of admissibility and therefore can be considered non-admissible semantics as well, for instance

- cf-based semantics, such as the previously defined naive semantics, as well as *stage* semantics, which maximize (with regard to set inclusion) the range $S^{\oplus}$ of conflict-free sets S, and which are closely related to stable semantics (all stage extensions are stable extensions as soon as one stable extension exists) [BCG11];
- and finally Dauphin, Rienstra and van der Torre's *qualified* and *semi-qualified semantics* [DRVDT20], which are based on labellings (similar to Definition 14) but allow for decomposition along the strongly connected components (SCC's) of a framework (SCC's are maximal sets of arguments which are mutually "reachable" through the attack relation) [BGG05].

2.5 Properties of Extension-Sets and Signatures of Semantics

Having introduced the various semantics that are of interest to us, we now turn to the question whether their extension-sets posses characteristic properties. For some of the semantics, this question has been answered affirmative. We list the respective properties below.

Definition 15 (Downward-Closure [DDLW15])**.** 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 $\mathbb{S}$, that is, $\text{dcl}(\mathbb{S}) = \{S \subseteq \bigcup \mathbb{S} \mid \exists S' \in \mathbb{S} : S \subseteq S'\}$.

Definition 16 (Properties of Extension-Sets [DDLW15])**.** 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 $\{a, b\} \not\subseteq S$ for all $S \in \mathbb{S}$; i.e., there are two arguments in $A \cup B$ that never occur jointly in any extension of $\mathbb{S}$;
- *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 $\{a, b\} \not\subseteq S$ 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 lemma reproduces two properties of tight sets given by Dunne *et al.* [DDLW15] that will be useful shortly.

Lemma 1 ([DDLW15])**.** Tightness implies conflict-sensitivity, and the $\subseteq$-maximal elements of a tight set form a tight set themselves. □

We now turn to the main topic of our consideration, which is the expressive power of semantics.

Definition 17 (Signature [DDLW15])**.** 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 σ can possibly produce. For some semantics, their signatures turn out to have exact characterisations; we reproduce some of the respective results of Dunne *et al.* [DDLW15] below.

Theorem 1 (Signatures of Semantics [DDLW15])**.** 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$\}$;
- $\Sigma_{\mathsf{naive}} = \{\mathbb{S} \subseteq 2^{\mathfrak{A}} \mid \mathbb{S}$ is non-empty, incomparable, and $\text{dcl}(\mathbb{S})$ is tight$\}$;
- $\Sigma_{\mathsf{stb}} = \{\mathbb{S} \subseteq 2^{\mathfrak{A}} \mid \mathbb{S}$ is incomparable 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$\}$;
- $\Sigma_{\mathsf{grd}} = \{\mathbb{S} \subseteq 2^{\mathfrak{A}} \mid |\mathbb{S}| = 1\}$. □

Note that we have not resorted to considering a framework's structural features (particularly, the attack relation) in order to characterise these semantics. In fact, one of the motivations behind this characterisation is to be able to assess whether it is feasible to represent modified or additional points of view within an existing argumentation scenario, before attempting to adapt the argumentation framework itself. Dunne *et al.* list further motivations and use cases [DDLW15].

Finally, between certain signatures of interest we have the following relations.

Theorem 2 ([DDLW15]). The following relations hold:

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

Proof. The implications follow from the characterisations of Theorem 1, combined with Lemma 1. The inequalities $\Sigma_{\mathsf{cf}} \neq \Sigma_{\mathsf{adm}}$ and $\Sigma_{\mathsf{naive}} \neq \Sigma_{\mathsf{pr}}$ are evidenced by $\mathbb{S}_1 = \{\{a,b,c\},\{c,d\},\{b,d,e\}\}$, which is incomparable and conflict-sensitive, but neither $\mathbb{S}_1$ or $\mathrm{dcl}(\mathbb{S}_1)$ are tight;[3] finally, $\mathbb{S}_2 = \{\{a\},\{a,b\}\}$ is conflict-sensitive but not incomparable, so $\mathbb{S}_2 \cup \{\emptyset\} \in \Sigma_{\mathsf{adm}}$, but $\mathbb{S}_2 \notin \Sigma_{\mathsf{pr}}$. □

[3] An even smaller example of a conflict-sensitive but not tight extension-set is $\{\emptyset, \{a,b\}\}$.

3 Properties of Semantics

The characterisations of Theorem 1 gave us necessary and sufficient properties of extension-sets under certain semantics. In this section, we will take a closer look at the necessary properties, that is, properties that apply commonly to all extension-sets of a certain 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.

Proposition 3 (Subset Relations [BBU20, Thi23])**.** For every 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 well-defined does transfer to undisputed semantics. In any case, we will investigate to what extent these relationships can provide indications regarding the common properties of the semantics being analysed here.

Another lead we will be following is the symmetry of the definitions of classical semantics and semantics based on the notions of weak admissibility and defense (compare Definitions 5 and 10); we will investigate whether some of the classical properties from Theorem 1 carry over to the corresponding non-admissible cases.

3.1 The Classical Properties

We start with the characteristic property of admissible semantics, namely conflict-sensitivity, and retrace the proof idea.

Proposition 4 ([DDLW15])**.** Every admissible extension-set of $F = (\mathcal{A}, \mathcal{R}) \in \mathfrak{F}_{\mathfrak{A}}$ is conflict-sensitive.

Proof-sketch. Let $A, B \in \mathsf{adm}(F)$ with $A \cup B \notin \mathsf{adm}(F)$. The only reason for $A \cup B$ not to be admissible is that there are conflicting arguments $a \in A$ and $b \in B$ with $a \rightarrow b$ or $b \rightarrow a$, which therefore can never appear jointly in one of F's extensions; this means that $\mathsf{adm}(F)$ is conflict-sensitive. □

Proposition 4 does not hold for weakly admissible semantics.

Proposition 5. Weakly admissible extension-sets are not necessarily conflict-sensitive.

Proof. The framework in Figure 4 is a counter-example, since $\mathsf{wadm}(F)$ is missing $\{a, b\} \cup \{b, c\} = \{a, b\} \cup \{a, c\} = \{a, b, c\}$. □

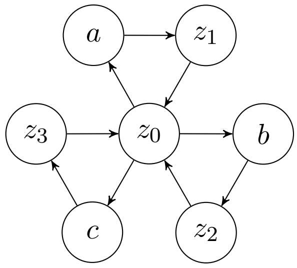

Figure 4: The weakly admissible extension-set of this framework (F) is not conflict-sensitive; we have $\mathsf{wadm}(F) = \{\emptyset, \{a\}, \{b\}, \{c\}, \{a, b\}, \{a, c\}, \{b, c\}\}$.

Dunne *et al.* also introduce the following relaxation of the concept of being conflict-sensitive.

Definition 18 ([DDLW15])**.** 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*).

Although the above property is tailored to apply to complete semantics [DDLW15], let us see whether it applies to weakly admissible semantics as well.

Proposition 6. Weakly admissible extension-sets are not necessarily com-closed.

Proof. The framework in Figure 4 serves again as a counter-example, because its extensions do not contain a completion-set that encompasses $\{a, b, c\}$. □

The property of being com-closed generalises conflict-sensitivity by requiring that, in the absence of evidence for conflict,[4] 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 4 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 holds for weak admissibility and state the only property common to all weakly admissible extension-sets remaining so far.

Proposition 7. 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 their common properties.

Proposition 8 ([DDLW15])**.** Every preferred extension-set of a framework $F \in \mathfrak{F}_{\mathfrak{A}}$ is non-empty, incomparable and conflict-sensitive.

[4] In the same sense as in the definition of conflict-sensitivity (Definition 16). The term ("evidently in conflict") will be formalised in Definition 19.

Proof. Preferred extensions are $\subseteq$-maximal admissible sets, so $\mathsf{pr}(F)$ is non-empty by definition; $\mathsf{pr}(F)$ does not contain a proper subset of any of its member sets and is thus incomparable. Furthermore, since every preferred extension-set contains only admissible extensions, it is also conflict-sensitive. □

Only the properties of non-emptiness and incomparability transfer to the weakly preferred semantics.

Proposition 9. Weakly preferred extension-sets are non-empty and incomparable, but not necessarily conflict-sensitive.

Proof. Non-emptiness follows from the definition, and incomparability is implied by $\subseteq$-maximality. Regarding conflict-sensitivity, the framework in Figure 4 gives a counterexample; its weakly preferred extensions are $\{a, b\}$, $\{a, c\}$ and $\{b, c\}$. □

We proceed to vacuous reduct semantics and repeat the following basic result.

Proposition 10 ([Thi23])**.** Undisputed extension-sets are non-empty, and strongly undisputed extension-sets are incomparable. □

For the sake of thoroughness we briefly examine whether the property of conflict-sensitivity applies.

Proposition 11. Undisputed and strongly undisputed extension-sets are neither necessarily conflict-sensitive nor necessarily com-closed.

Proof. A counterexample is again given by the framework (F) in Figure 4; we have $\mathsf{sud}(F) = \{\{a, b\}, \{a, c\}, \{b, c\}\}$ and $\mathsf{ud}(F) = \{\emptyset, \{a\}, \{b\}, \{c\}\} \cup \mathsf{sud}(F)$. □

Let us investigate other properties that may be induced by the subset relations of Proposition 3. The example from Figure 4 could suggest that at least the weakly admissible extension-sets may be downward-closed.

Proposition 12. The weakly admissible as well as the undisputed extension-sets of a framework are neither necessarily downward-closed nor necessarily tight.

Proof. Figure 5 gives a counterexample. Both extension-sets are obviously not downward-closed: $\{a\}$ is missing from $\mathsf{wadm}(F)$, and neither $\{a\}$ nor $\{b\}$ is present in $\mathsf{ud}(F)$. Tightness would require that the absence of the aforementioned sets is somehow reflected in $\mathsf{wadm}(F)$ and $\mathsf{ud}(F)$, but both contain $\{a, b\}$; and indeed, the framework shows no conflict between a and b. □

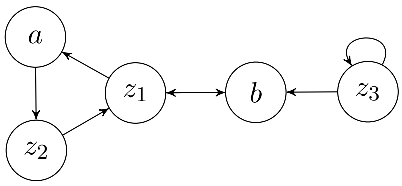

Figure 5: For this framework (F), we have $\mathsf{wadm}(F) = \{\emptyset, \{b\}, \{a,b\}\}$ as well as $\mathsf{ud}(F) = \{\emptyset, \{a,b\}\}$. These extension-sets are neither downward-closed nor tight.

3.2 Limits to the Expressiveness of Undisputed Semantics

At this point we may be led to suspect that the expressiveness of the undisputed semantics is only limited by the requirement to produce non-empty extension-sets; this suspicion will be proven wrong. We start with preliminaries.

Lemma 2. 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 3. If $\emptyset \in \mathsf{ud}(F)$, then $\mathsf{adm}(F^{\emptyset}) \subseteq \{\emptyset\}$ and $\mathsf{pr}(F) = \{\emptyset\}$. □

The following concept was already essential for the definition of conflict-sensitivity (Definition 16).

Definition 19. 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. We state this for the case of ud.

Lemma 3. For $F \in \mathfrak{F}_{\mathfrak{A}}$, let $X, Y \in \mathsf{ud}(F)$. If X and Y are not evidently in conflict with respect to $\mathsf{ud}(F)$, then $X \cap Y^{+} = \emptyset = Y \cap X^{+}$.

Proof. The negation of Definition 19 states that for all pairs $(x, y) \in X \times Y$ we have some $S \in \mathsf{ud}(F)$ 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 4. 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. 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 5. For a framework $F = (\mathcal{A}, \mathcal{R}) \in \mathfrak{F}_{\mathfrak{A}}$, let $S \in \mathsf{adm}(F)$, and let $S' \subseteq \mathcal{A}$. Then $S \setminus S' \in \mathsf{adm}(F^{S'})$.

Proof. All attackers of S that S' defended against are removed in the reduct $F^{S'}$. □

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

Definition 20. A set of sets $\mathbb{S} = \{A_1, \dots, A_n\} \subseteq 2^{\mathfrak{A}}$ $(n > 1)$ is *disjointly supported* if each $A_i \neq \emptyset$, $\bigcap_i A_i = \emptyset$, and the A_i are pairwise not in evident conflict.

Example 6. 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 13. 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 2 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 3; this leaves $S \cap A_i \neq \emptyset$ for all A_i. Lemma 4 then states that for at least one $A_j \in \mathbb{S}$ we have $S \setminus A_j \neq \emptyset$, and by Lemma 5 we have $S \setminus A_j \in \mathsf{adm}(F^{A_j})$, again contradicting $A_j \in \mathsf{ud}(F)$. □

3.3 Characteristics of Non-admissible Semantics

The results from the previous subsections suggest the following conjecture.

Conjecture 1. Weakly admissible, weakly preferred, as well as undisputed and strongly undisputed semantics are characterised as follows.

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

This conjecture can be proved by specifying construction methods for frameworks that realise arbitrary extension-sets which feature the required properties. The rest of the present work is dedicated to the study of such construction methods.

4 Classical Framework Construction

We would like to investigate to what extent the necessary properties of the extension-sets of various semantics determined in Section 3 are also sufficient, so that extension-sets to which these properties apply are indeed realisable by the respective semantics. One way to do this for a semantics is to specify a concrete method that constructs a framework for a given extension-set, so that the extension-set does indeed result from the semantics applied to the constructed framework. In this section we analyse existing construction methods for frameworks and attempt to isolate the concepts and strategies they contain.

4.1 The Canonical Frameworks

Non-emptiness, incomparability and conflict-sensitivity are not only necessary properties of preferred extension-sets (Proposition 8); they are also sufficient for their characterization, as the following construction shows.

Definition 21 (Canonical Argumentation Framework [DDLW15]). 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 other derived semantics) applied to $F^{\mathsf{cf}}_{\mathbb{S}}$ will encompass at least all elements of $\mathbb{S}$. More concretely:

Proposition 14. Let $\sigma \in \{\mathsf{cf}, \mathsf{adm}, \mathsf{pr}, \mathsf{stb}\}$, 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 $S \in \mathsf{adm}(F^{\mathsf{cf}}_{\mathbb{S}})$. If additionally each $S \in \mathbb{S}$ is $\subseteq$-maximal, then $S \in \mathsf{pr}(F^{\mathsf{cf}}_{\mathbb{S}})$. 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 $F^{\mathsf{cf}}_{\mathbb{S}}$, $s \rightarrow a$, so $S \in \mathsf{stb}(F^{\mathsf{cf}}_{\mathbb{S}})$.[5] □

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

[5] Dunne *et al.* [DDLW15] prove the partial statements of this proposition in various places; we have summarised them here.

Definition 22 (Canonical Defence Framework [DDLW15])**.** 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, $S^a_i = \{a\}$.
2. For each $D \in \mathcal{D}^a$, introduce a new argument z^a_D. This will create a total of $|\mathcal{D}^a| = (|S^a_1| - 1) \cdot \ldots \cdot (|S^a_n| - 1)$ new arguments.
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 [DDLW15]. To summarise:

Proposition 15 ([DDLW15])**.** 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 7. 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 $z^c_{\{b\}} = z_3$.

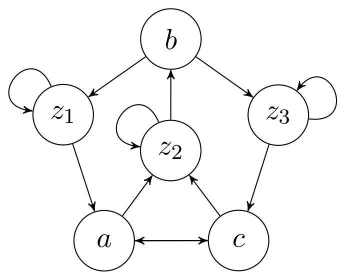

Figure 6: The defence framework $F^{\mathsf{adm}}_{\mathbb{S}}$ for $\mathbb{S} = \{\emptyset, \{a, b\}, \{b, c\}\}$.

In Definition 21 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 admissible semantics.

Proposition 16. 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 final construction method we want to investigate in this subsection aims at stable semantics. The method starts from the canonical argumentation framework and subsequently excludes unwanted extensions.

Definition 23 (Stable Canonical Framework [DDLW15]). 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 $a \rightarrow z_X$.

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.

Proposition 17 ([DDLW15]). 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, $\mathsf{stb}(F^{\mathsf{stb}}_{\mathbb{S}}) = \mathbb{S}$. □

4.2 Framework Translations

Further construction ideas are presented by Dvořák and Woltran [DW11] as well as Dvořák and Spanring [DS17], this time in the context of translations between frameworks, which (roughly speaking) are intended to preserve extensions across different semantics.

Definition 24 ([DW11]). A *framework translation* is a mapping $\mathfrak{F}_{\mathfrak{A}} \rightarrow \mathfrak{F}_{\mathfrak{A}}$. With respect to two semantics $\sigma, \sigma' : \mathfrak{F}_{\mathfrak{A}} \rightarrow 2^{2^{\mathfrak{A}}}$ and an arbitrary framework $F = (\mathcal{A}, \mathcal{R}) \in \mathfrak{F}_{\mathfrak{A}}$, a framework translation T is called

- *exact*, if $\sigma(F) = \sigma'(T(F))$;
- *faithful*, if $\sigma(F) = \{E \cap \mathcal{A} \mid E \in \sigma'(T(F))\}$ and $|\sigma(F)| = |\sigma(T(F))|$.

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 [DW11]: (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.

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

Example 8. The admissible extensions of the left framework in Figure 7 are identical to the complete extensions of framework to the right, namely $\{\emptyset, \{a\}, \{c\}, \{a, c\}\}$.

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 sets 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 will point out in the next subsection.

4.3 Construction Strategies

In the previous subsections we have studied various construction methods for frameworks. These methods appear to have certain commonalities, which we aim to identify below.

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.

5 Strategies for the Non-admissible Case

We would like to see how the features that were identified in Observation 1 translate to strategies for the construction of frameworks that attempt to realise extension-sets under non-admissible semantics.

5.1 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 below proposition.

Proposition 18 ([BBU20])**.** Let $F = (\mathcal{A}, \mathcal{R}) \in \mathfrak{F}_{\mathfrak{A}}$, let $\mathcal{A}^\circ = \{a \in \mathcal{A} \mid (a, a) \notin \mathcal{R}\}$, and let $F^\circ = F|_{\mathcal{A}^\circ}$. 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 9. In the framework from Figure 8 (F), we have $\mathsf{ud}(F) = \{\emptyset, \{c\}\}$ and $\mathsf{sud}(F) = \emptyset$, but $\mathsf{ud}(F^\circ) = \mathsf{sud}(F^\circ) = \{\emptyset\}$.

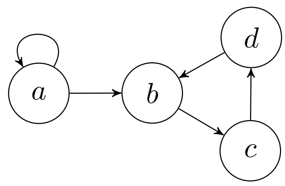

Figure 8: Undisputed and strongly undisputed semantics applied to this framework are not invariant with respect to removal of self-attacking arguments.

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 26) we will use this feature in a systematic way.

5.2 Separating Extensions

We turn to weakly admissible semantics and investigate mechanisms that can cause extensions to be separated from each other in an extension-set.

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 16). Non-admissible semantics introduce another mechanism that is able to separate extensions, namely odd cycles, as shown in the following example.

Example 10. Figure 9 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.

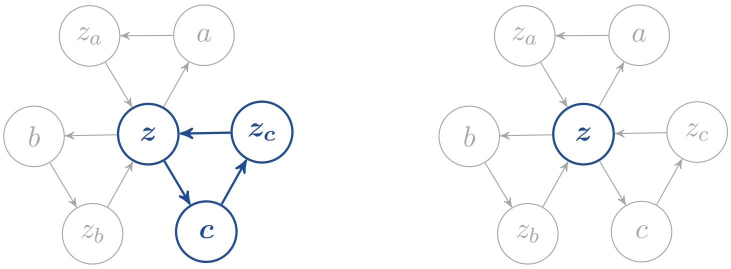

Figure 9: Mechanics of odd cycles in non-admissible semantics.

We would like to formalise the mechanics at work in Example 10. In order to do this, we need the following two results. The first is the *Modularisation Theorem* by Baumann *et al.*

Theorem 3 (*Modularisation* [BBU22]). Let $F \in \mathfrak{F}_\mathfrak{A}$ be a framework, 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 11. Consider the framework from Figure 10 (call it F), and let $X = \{a\}$ and $Y = \{b\}$. We have $F^X = (\{d\}, \emptyset) = (F^X)^Y \neq (F^Y)^X = (\emptyset, \emptyset) = F^{X \cup Y}$.

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.[6]

Lemma 6. Let $F = (\mathcal{A}, \mathcal{R}) \in \mathfrak{F}_\mathfrak{A}$ be a framework, and let $X, Y \subseteq \mathcal{A}$. If $X^+ \cap Y = \emptyset$, then $\left(F^X\right)^Y = F^{X \cup Y}$.

[6] Baumann *et al.* [BBU22, Proposition 3.3] give a similar proposition, but require $X \cap Y = \emptyset$ as well as $X \cup Y \in \mathsf{cf}(F)$.

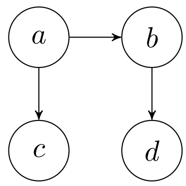

Figure 10: The reduct of this framework with respect to $\{a\}$ eliminates $\{a, b, c\}$, and the subsequent reduct with respect to $\{b\}$ cannot affect d any more.

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

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

while, for $F^{X \cup Y} = (\mathcal{A}_{XY}, \mathcal{R}_{XY})$, we have[7]

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

Finally, since $X^+_{\mathcal{R}} \cap Y = \emptyset$, Y attacks the same arguments in $\mathcal{R}$ as in $\mathcal{R}_X$, therefore $Y^+_{\mathcal{R}} = Y^+_{\mathcal{R}_X}$ and $\mathcal{A}_{X,Y} = \mathcal{A}_{XY}$. □

Corollary 2. 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 10.

Proposition 19. 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 and Y 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 then follows $X \cup Y \in \mathsf{wadm}(F)$, a contradiction. □

Example 12. Let F be the argumentation framework from Figure 9, 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 19 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.

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

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 19 when vacuous reduct semantics are concerned. Let us look at an example.

Example 13. Consider the framework (F) of Figure 11. 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 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).

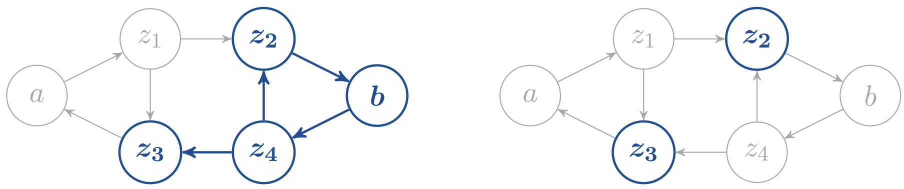

Figure 11: In this framework, $\{a, b\}$ is not among the undisputed extensions.

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

Proposition 20. 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.

Proof. The prerequisites demand $X \not\subseteq Y$ and $Y \not\subseteq X$, 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. $\square$

Example 14. Applying Proposition 20 to Figure 11, we have $X = \{a\}$, $Y = \{b\}$, and find $z_3 \to X$ and $z_2 \to Y$; neither X nor Y defend themselves.

Proposition 21. 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 non-empty admissible extension $S \in \mathsf{adm}(F^{X \cup Y})$ that is attacked by both X^+ and Y^+.

Proof. The existence of S 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 $(F^X)^Y = F^{X \cup Y}$, 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^+. For reasons of symmetry we have both $X^+ \to S$ and $Y^+ \to S$. □

Example 15. In the framework (F) of Figure 11, 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 $Y^+ = \{z_4\} \to S$.

5.3 Base Frameworks

We would like to find frameworks whose non-admissible extensions encompass a given extension-set, provided that the extension-set satisfies necessary properties, as in the following definition.

Definition 25 (Base Framework)**.** With respect to a semantics $\sigma : \mathfrak{F}_{\mathfrak{A}} \to 2^{2^{\mathfrak{A}}}$ and a set of extension-sets $\Sigma_\sigma \subseteq 2^{2^{\mathfrak{A}}}$ we call the mapping $\mathbb{S} \mapsto F_{\mathbb{S}}$ a *base framework* if for all $\mathbb{S} \in \Sigma_\sigma$ we have $\mathbb{S} \subseteq \sigma(F_{\mathbb{S}})$.

We have seen an example for a base framework in Proposition 14:

Example 16. For each $\sigma \in \{\mathsf{cf}, \mathsf{adm}, \mathsf{pr}, \mathsf{stb}\}$ 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 21) 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 realisation of a given extension-set. So let us see whether the canonical argumentation framework from Definition 21 can also function as a base framework for non-admissible semantics.[8] We start with weakly admissible semantics and assume that they are characterised according to Conjecture 1.

Proposition 22. Let $\Sigma_{\mathsf{wadm}} = \{\mathbb{S} \subseteq 2^{\mathfrak{A}} \mid \emptyset \in \mathbb{S}\}$. For all extension-sets $\mathbb{S} \in \Sigma_{\mathsf{wadm}}$ we have $\mathbb{S} \subseteq \mathsf{wadm}(F^{\mathsf{cf}}_{\mathbb{S}})$.

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. □

Thus, $\mathbb{S} \mapsto F^{\mathsf{cf}}_{\mathbb{S}}$ is a base framework with respect to weakly admissible semantics. For weakly preferred, undisputed and strongly undisputed semantics the situation is different.

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

Proposition 23. The canonical argumentation framework is not a base framework for $\sigma \in \{\mathsf{wpr}, \mathsf{ud}, \mathsf{sud}\}$ (with respect to Σ_σ from Conjecture 1).

Proof. Let $\mathbb{S} = \{\{a,b\},\{b,c\},\{a,c\}\}$; $\mathbb{S}$ is incomparable. We have $F^{\mathsf{cf}}_{\mathbb{S}} = (\{a,b,c\},\emptyset)$, and $\mathbb{S} \not\subseteq \mathsf{wpr}(F^{\mathsf{cf}}_{\mathbb{S}}) = \mathsf{ud}(F^{\mathsf{cf}}_{\mathbb{S}}) = \mathsf{sud}(F^{\mathsf{cf}}_{\mathbb{S}}) = \{\{a,b,c\}\}$. □

Next we investigate the canonical defence framework from Definition 22. Note that due to its construction, $(F^{\mathsf{adm}}_{\mathbb{S}})^\circ = F^{\mathsf{cf}}_{\mathbb{S}}$,[9] so the above results transfer in the case of weakly admissible semantics. Let us see however whether $\mathbb{S} \mapsto F^{\mathsf{adm}}_{\mathbb{S}}$ can be effective for undisputed semantics.

Proposition 24. The canonical defence framework is neither a base framework for ud nor for sud (with respect to $\Sigma_{\mathsf{ud}}, \Sigma_{\mathsf{sud}}$ from Conjecture 1).

Proof. Consider again the incomparable extension-set $\mathbb{S} = \{\{a,b\},\{b,c\},\{a,c\}\}$. Its canonical defence framework is depicted in Figure 12; for this framework, we have $\mathbb{S} \not\subseteq \mathsf{ud}(F^{\mathsf{adm}}_{\mathbb{S}}) = \mathsf{sud}(F^{\mathsf{adm}}_{\mathbb{S}}) = \{\{a,b,c\}\}$. □

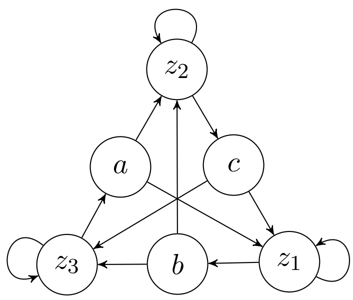

Figure 12: The canonical defence framework $F^{\mathsf{adm}}_{\mathbb{S}}$, constructed for the extension-set $\mathbb{S} = \{\{a,b\},\{b,c\},\{a,c\}\}$.

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

Proposition 25. $\mathbb{S} \mapsto F^{\emptyset}_{\mathbb{S}} = (\bigcup \mathbb{S}, \emptyset)$ is a base framework with respect to $\sigma = \mathsf{wadm}$ and $\mathbb{S} \in \Sigma_\sigma = \{\mathbb{S}' \subseteq 2^{\mathfrak{A}} \mid \emptyset \in \mathbb{S}'\}$.

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

A small modification yields a base framework for undisputed semantics.

Proposition 26. Let $z \notin \bigcup \mathbb{S}$. $\mathbb{S} \mapsto F^{\mathsf{ud}}_{\mathbb{S}} = (\bigcup \mathbb{S} \cup \{z\}, \{(z,z)\} \cup \{(z,s) \mid s \in \bigcup \mathbb{S}\})$ is a base framework with respect to $\sigma = \mathsf{ud}$ and $\Sigma_\sigma = \Sigma_{\mathsf{ud}}$ from Conjecture 1.

Proof. We have $\mathsf{ud}(F^{\mathsf{ud}}_{\mathbb{S}}) = 2^{(\bigcup \mathbb{S})}$, 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 17. Figure 13 shows the base framework $F^{\mathsf{ud}}_{\mathbb{S}}$ for any $\mathbb{S}$ with $\bigcup \mathbb{S} = \{a,b,c\}$.

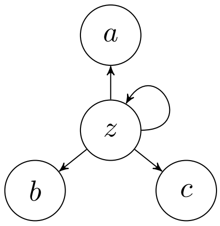

Figure 13: The undisputed extension-set of this framework is the set of all subsets of the original arguments $\{a, b, c\}$.

The base frameworks for wadm and ud described in Propositions 25 and 26 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 will only produce incomparable extension-sets, and the constructions that worked well for wadm and ud lack features that would separate extensions from another. We will revisit the underlying problem, namely the inclusion of mechanisms that separate extensions into the construction of base frameworks, in Subsection 6.5.

5.4 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 26. With respect to a semantics σ, a signature Σ_σ, and a set of 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 18. For $\sigma = \mathsf{stb}$ we used the following total filter (cf. Definition 23):

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

[9]The notation F° was introduced along with Proposition 18.

The concept of a filter shares similarities with framework translations which we described in Subsection 4.2; but instead of attempting to preserve extension-sets between different semantics, we seek to remove certain extensions under identical semantics. The following terminology is borrowed from the context of framework translations [DS17].

Definition 27. A filter f is called *covering* if for every $F = (\mathcal{A}, \mathcal{R})$ and every $\mathbb{X} \subseteq 2^{\mathfrak{A}}$ the resulting framework $f(F, \mathbb{X}) = (\mathcal{A}', \mathcal{R}')$ satisfies $\mathcal{A} \subseteq \mathcal{A}'$ and $\mathcal{R} \subseteq \mathcal{R}'$.

This means that 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 18 is covering, and the construction of the canonical defence framework (Definition 22) 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 19. Consider the following problem: for the framework (F) in Figure 14, 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 $\mathsf{wadm}(f(F, \{\{a\}\})) = \{\emptyset, \{a, b\}\}$?

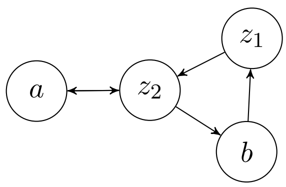

Figure 14: What modifications are needed to filter out one of this framework's weakly admissible extensions?

Possible solutions are shown in Figure 15; 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 utilize 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.

The solutions to the problem of Example 19 were found through an automated search that applied random augmentations to the given framework. The solutions

[10] Additionally, the framework translations featured in the work of Dvořák and Spanring [DS17] are generally covering; some are even *embedding*, meaning that no additional attacks between original arguments are introduced by the translation.

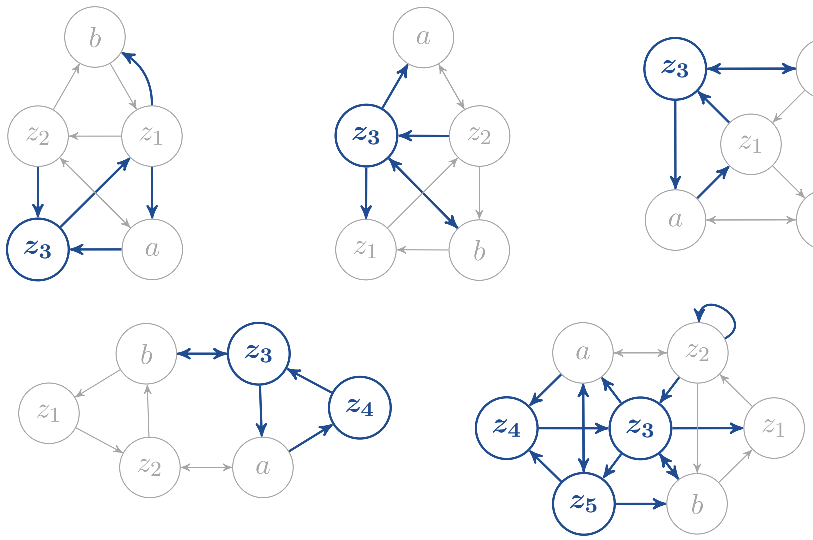

Figure 15: Solutions to the problem posed in Example 19. Added structures have been emphasized; all solutions are strict augmentations.

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 does appear to emerge.

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

Example 20. Figure 16 shows the counterpart of Figure 14, 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.

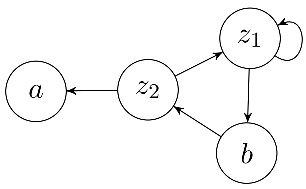

Figure 16: This framework has the undisputed extension-set $\{\emptyset, \{a\}, \{a, b\}\}$.

Several minimal solutions are presented in Figure 17. All frameworks shown there have the same set of undisputed extensions, namely $\{\emptyset, \{a, b\}\}$.

Much like in the previous Example 19, 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 power even in the non-admissible case. On the other hand,

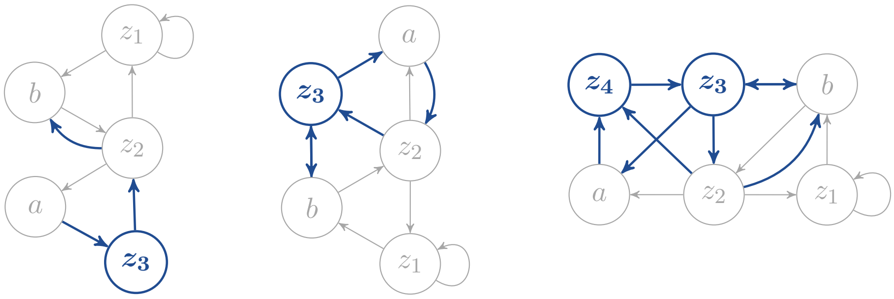

Figure 17: Some solutions to the problem from Example 20.

the filter algorithms appear to be much more complex than in the classical case, provided they exist and can be found at all. In the next section, we will further pursue the approach described in order to gain an impression of its possibilities and limitations.

6 Realisation of Non-admissible Extension-Sets

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

6.1 Minimal Scenarios

Examples 19 and 20 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 21. Let $\mathbb{S} \subseteq 2^{\mathfrak{A}}$ be an extension-set, let $F = F_{\mathbb{S}}^{\mathsf{ud}}$ be the base framework from Proposition 26, let $a, b \in \bigcup \mathbb{S}$, and let z be the auxiliary argument attacking all other $a \in \bigcup \mathbb{S}$. 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 21 apply only one attack each. There are also other, less obvious filter operations that use several attacks simultaneously.

Example 22. Figure 18 shows some filter operations on the base framework $F_{\mathbb{S}}^{\mathsf{ud}}$, instantiated for $\bigcup \mathbb{S} = \{a, b, c\}$. Each filter operation uses multiple attacks but does not introduce any new arguments.

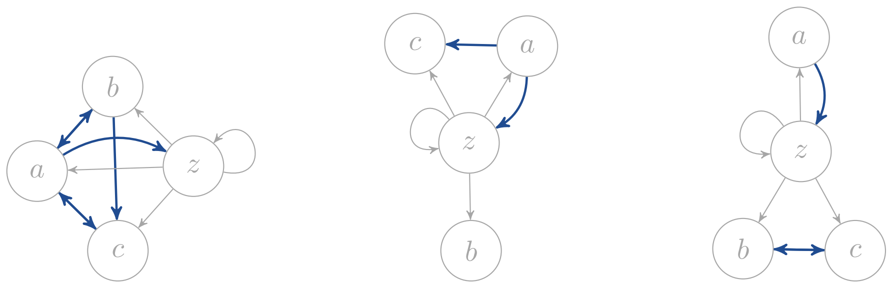

(a) $\mathsf{ud}(F) = \{\{a\}, \{b\}, \{c\}\}$ (b) $\mathsf{ud}(F) = \{\{a, b\}\}$ (c) $\mathsf{ud}(F) = \{\{a, b\}, \{a, c\}\}$

Figure 18: Filter operations that apply multiple attacks, and their outcomes.

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 $F_{\mathbb{S}}^{\mathsf{ud}}$ for a problem size of 3, i.e., $\bigcup \mathbb{S} = \{a, b, c\}$, and record their effect on $\mathsf{ud}(F_{\mathbb{S}}^{\mathsf{ud}})$. The number of scenarios to consider is $2^{4^2-4-3} - 1 = 511$, taking into account the four attacks that are invariably present in $F_{\mathbb{S}}^{\mathsf{ud}}$, as well as the fact that we do not need to consider self-attacking arguments a, b, c. From the 256 subsets of $2^{\{a,b,c\}}$, 38 do not contain all of the arguments $\{a, b, c\}$; from the remaining 218 sets, 53 are disjointly supported. Subtracting these, we arrive at 165 interesting subsets of $2^{\{a,b,c\}}$, and we would like to see how many can be realized by our construction.

Result. Only 30 out of the 165 interesting extension-sets were realized, although each of the 511 attack scenarios had an effect on the undisputed extension-set. The results are summarised in Table 1, where isomorphic[11] extension-sets have been omitted.

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

Table 1: 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 27. For a given $\mathbb{S} \in \Sigma_{\mathsf{ud}}$ (from Conjecture 1), let $A = \bigcup \mathbb{S}$, and for each $a \in A$, introduce a new argument z_a. The construction

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

is a base framework with respect to ud and Σ_{ud}, and we have $\mathsf{ud}(F_{\mathbb{S}}^{\mathsf{ud}+}) = 2^A$.

[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.

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. □

Example 23. Figure 19 shows the base framework $F_{\mathbb{S}}^{\mathsf{ud}+}$ for $\bigcup \mathbb{S} = \{a, b, c\}$.

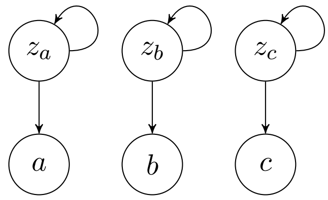

Figure 19: A base framework for ud, constructed from three arguments.

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

Example 24. 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 165 interesting extension-sets were realized; 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 realised extension-sets included all realised sets from Case Study 1; we summarise the results in Table 2.

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 25. The extension-set $\{\emptyset, \{a, b, c\}\}$ was among the unrealised extension-sets of Case Study 2. Figure 20 however shows that a realising framework exists, which is not an augmentation of the previously considered base framework.

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 have seen 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.

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 2: 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.

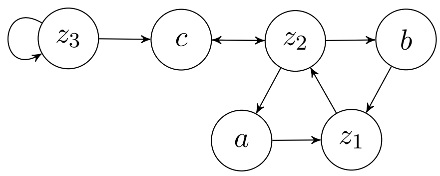

Figure 20: The undisputed extension-set of this framework is $\{\emptyset, \{a,b,c\}\}$.

6.2 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 different strategy in this subsection: instead of inserting arbitrary attacks into a given base framework, we define a number of mappings, which we will refer to as *transformations*, that each alter a given framework in a certain manner, and then proceed to combine these transformations.

The first transformation that we will be using 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

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

Example 26. Figure 21 illustrates an instance of the merge transformation μ.

Figure 21: 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

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

For a concrete application we continue to 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 22, which is not a base framework (its undisputed extension-set is $\{\emptyset\}$).

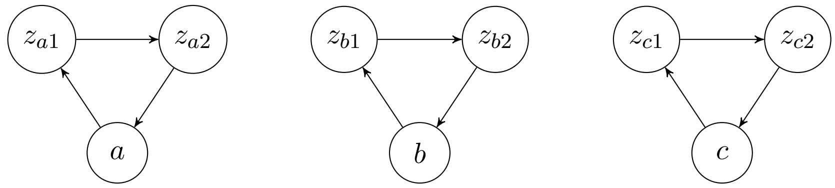

Figure 22: These three odd cycles serve as starting scenario to which subsequent transformations are applied.

Case Study 3. We apply the above defined transformations to the framework from Figure 22, 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 3. In summary, construction patterns for 35 different isomorphism classes representing extension-sets that contain three arguments were found.

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

6.3 The Cycle Hub Framework

We will now examine another base framework which, due to its complexity, is no longer suitable for an exhaustive search covering all possible attack configurations. Its construction draws inspiration 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 28 (Cycle Hub Framework)**.** For a given set of arguments A, the *cycle hub framework* $F_A^\triangle$ is constructed as follows. We start with the framework $F = (\mathcal{A}, \mathcal{R})$ 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 22 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 27. The left framework of Figure 23 shows the cycle hub framework $F_A^\triangle$ constructed for the three arguments $A = \{a, b, c\}$.

Proposition 28. 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 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 for a given $\mathbb{S} \in \Sigma_{\mathsf{ud}}$ (from Conjecture 1), $\mathbb{S} \mapsto F^\triangle_{\bigcup \mathbb{S}}$ is not a base framework since $\bigcup \mathbb{S}$ is missing from its extension-set, but we can easily complete the construction and turn it into a base framework for ud.

Definition 29 (Cycle Hub Base Framework)**.** Given an extension-set $\mathbb{S} \in \Sigma_{\mathsf{ud}}$ (from Conjecture 1), the *cycle hub base framework* is obtained as follows.

1. Construct $F = (\mathcal{A}, \mathcal{R}) = F^{\triangle}_{\bigcup \mathbb{S}}$.
2. Create a new self-attacking argument $z_1 \notin \mathcal{A}$ and connect it to F via $z_1 \to z_0$.

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

Example 28. The right framework in Figure 23 shows $F^{\mathsf{ud}*}_{\mathbb{S}}$ constructed for an extension-set consisting of the three arguments a, b, c.

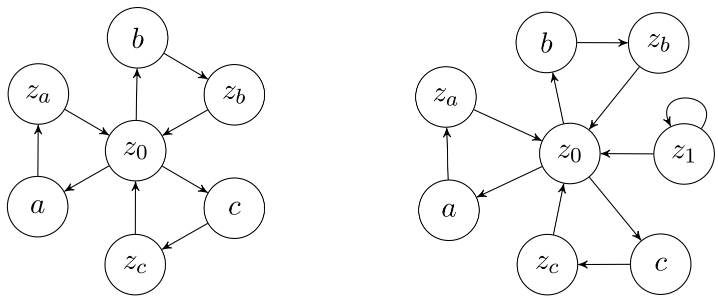

Figure 23: The frameworks $F^{\triangle}_{A}$ and $F^{\mathsf{ud}*}_{\mathbb{S}}$, for $A = \{a, b, c\} = \bigcup \mathbb{S}$.

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

Proposition 29. For any $\mathbb{S} \in \Sigma_{\mathsf{ud}}$ (from Conjecture 1), we have $\mathsf{ud}(F^{\mathsf{ud}*}_{\mathbb{S}}) = 2^{(\bigcup \mathbb{S})}$.

Proof. Let $S \in 2^{(\bigcup \mathbb{S})}$. If $S \subsetneq \bigcup \mathbb{S}$, $(F^{\mathsf{ud}*}_{\mathbb{S}})^S$ 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 admissible subsets. If $S = \bigcup \mathbb{S}$, only z_0, z_1 remain in the reduct, of which again no subset is admissible. □

We would like to investigate the expressiveness of $F^{\mathsf{ud}*}_{\mathbb{S}}$ as a base framework in a similar manner as we did in the previous case studies. For the three-argument scenario we considered before, $F^{\mathsf{ud}*}_{\mathbb{S}}$ 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 $F^{\mathsf{ud}*}_{\mathbb{S}}$ 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 4. All in all, 40 different isomorphism classes were realised, including all realisations from the previous case studies, making this approach the one with the highest expressiveness so far.

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

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

6.4 Summary of Constructions

In Subsections 6.1, 6.2, and 6.3 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 that also allows generalisations to arbitrary extension-sets.

Together, Tables 2, 3, and 4 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} (from Conjecture 1) that cannot be realised by our construction methods; for some of these, we can give concrete witnessing frameworks that prove their realisability. We do so in Table 5.

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 5: Witnessing frameworks of some undisputed three-argument extension-sets that were not realised by the constructions from Sections 6.1, 6.2, and 6.3.

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.

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

$$\begin{aligned}
&\{\{a,b\},\{a,c\},\{a,b,c\}\},\\
&\{\{a\},\{b\},\{a,b\},\{a,c\}\},\\
&\{\emptyset,\{a,b\},\{a,c\},\{b,c\}\},\\
&\{\{a\},\{b\},\{c\},\{a,b\},\{a,c\}\},\\
&\{\emptyset,\{a\},\{a,b\},\{a,c\},\{b,c\}\},\\
&\{\emptyset,\{a,b\},\{a,c\},\{b,c\},\{a,b,c\}\}.
\end{aligned}$$

6.5 Base Frameworks That Separate Extensions

So far we have used fairly generic base frameworks that produced the maximal set of extensions $2^{(\bigcup \mathbb{S})}$ for a given input extension-set $\mathbb{S}$. In the classical approach, by contrast, the base frameworks already employ strategies for separating extensions (cf. Definition 21 and Proposition 16). Similarly for the non-admissible case, we would like to incorporate our insights from Propositions 19, 20, and 21 into the construction of base frameworks, in the hope that this will result in a less complex filtering task and a broader range of realisable extension-sets. For the undisputed semantics then, we are presented with the following problem.

Problem 1. Given an extension-set $\mathbb{S} \in \Sigma_{\mathsf{ud}}$ (from Conjecture 1), construct a base framework $\mathbb{S} \mapsto F \in \mathfrak{F}_{\mathfrak{A}}$ with respect to ud and Σ_{ud}, so that for two extensions $X, Y \in \mathbb{S}$ where $X \cup Y \notin \mathbb{S}$, and in the absence of evident conflict between X and Y, F features the following:

1. There need to be undefended attacks on X and Y (Proposition 20);
2. X^+ and Y^+ need to attack a non-empty $S \in \mathsf{adm}(F^{X \cup Y})$ (Proposition 21).

In the above problem description, we consider only evidently non-conflicting sets since evident conflict can already be modelled with one-sided or mutual attacks; in this case, Propositions 20 and 21 do not apply any more. We further require that any construction method that intends to realise arbitrary undisputed extension-sets also respects the following.

Proposition 30 ([Thi23]). If $F \in \mathfrak{F}_{\mathfrak{A}}$ contains no odd cycles, then $\mathsf{ud}(F) = \mathsf{pr}(F)$. □

All of our proposed base frameworks for undisputed semantics (Propositions 26 and 27, and also Definition 29) indeed contain odd cycles. They all fulfil the first requirement of Problem 1, but since they make no effort to separate extensions, they do not fulfil the second; in the case studies we conducted, the realisation of the second requirement (Proposition 21) has always been a mere by-product of the subsequently applied filtering. Let us therefore entertain a somewhat naive, but straight-forward construction idea that incorporates separation mechanisms as a design principle. We should say upfront that the following algorithm is not able to produce adequate results in the general case, but at least highlights a characteristic problem one is confronted with when attempting to devise construction algorithms for non-admissible semantics.

Algorithm 1 Naively attempt to realise an undisputed extension-set.

Input: A set of sets $\mathbb{S} \subseteq 2^{\mathfrak{A}}$.
Output: *Construct*($\mathbb{S}$).

function CONSTRUCT($\mathbb{S} \subseteq 2^{\mathfrak{A}}$)
    let $\mathcal{A} = \emptyset, \mathcal{R} = \emptyset$. ▷ Start with an empty framework.
    let $P_{\text{ext}} = \{(X, Y) \in \mathbb{S} \times \mathbb{S} \mid X \cup Y \notin \mathbb{S}$ and X, Y are not in evident conflict$\}$.
    for all $C \in$ *Clusters*(P_{ext}) **do** ▷ Each cluster is a set of arguments.
        for all $a \in C$ **do**
            let $a_1, a_2 \in \mathfrak{A} \setminus \mathcal{A}$. ▷ Construct an odd cycle for a.
            $\mathcal{A} \leftarrow \mathcal{A} \cup \{a, a_1, a_2\}$.
            $\mathcal{R} \leftarrow \mathcal{R} \cup \{(a, a_1), (a_1, a_2), (a_2, a)\}$.
        end for
        let $z_C \in \mathfrak{A} \setminus \mathcal{A}$. ▷ Merge the cycles.
        $F \leftarrow \mu(\{a_2 \mid a \in C\}, z_C)$.
    end for
    let $P_{\text{arg}} = \{(a, b) \in \bigcup \mathbb{S} \times \bigcup \mathbb{S} \mid \forall S \in \mathbb{S} : \{a, b\} \nsubseteq S\}$.
    for all $(a, b) \in P_{\text{arg}}$ **do** ▷ Add attacks between conflicting arguments.
        $\mathcal{R} \leftarrow \mathcal{R} \cup \{(a, b), (b, a)\}$.
    end for
    return $(\mathcal{A}, \mathcal{R})$.
end function
function CLUSTERS($P_{\text{ext}} \subseteq 2^{\mathfrak{A}} \times 2^{\mathfrak{A}}$)
    let $S = \{X \cup Y \mid (X, Y) \in P_{\text{ext}}\}$.
    loop ▷ Combine sets that have common arguments.
        if $\exists S_1, S_2 \in S : (S_1 \neq S_2 \ \wedge \ S_1 \cap S_2 \neq \emptyset)$ **then**
            $S \leftarrow (S \setminus \{S_1, S_2\}) \cup \{S_1 \cup S_2\}$.
        else
            return S.
        end if
    end loop
end function

The interlinked odd cycles that we construct above are those of Definition 28. While Proposition 21 gives us some liberty in deciding the attackers X^+, Y^+ of the admissible extension of the reduct $S \in \mathsf{adm}(F^{X \cup Y})$, we choose to involve the maximal amount of arguments for the sake of simplicity. All in all, it would appear that we have fulfilled the requirements of Problem 1.

Example 29. We apply Algorithm 1 to the extension-set $\{\{a, b\}, \{b, c\}, \{a, c, d\}\}$. We find $P_{\text{ext}} = \{(\{a, b\}, \{b, c\})\}$ and $P_{\text{arg}} = \{(b, d)\}$, and construct the framework of Figure 24. The result contains the unwanted argument b_1 in one of its extensions.

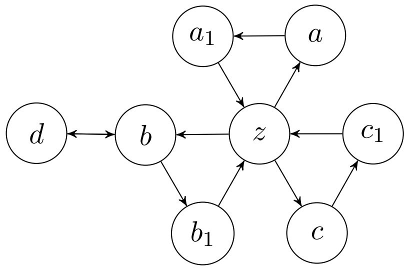

Figure 24: The result (F) of applying Algorithm 1 to the given extension-set $\{\{a, b\}, \{b, c\}, \{a, c, d\}\}$ produces $\mathsf{ud}(F) = \{\{b\}, \{a, b\}, \{b, c\}, \{a, c, d, b_1\}\}$.

Observation 4. It seems that combining the pattern of interlinked odd cycles with mutual attacks of evidently conflicting arguments has created some sort of an *a priori* unforeseen interference between both concepts, causing the auxiliary argument b_1 to emerge in an extension. This is indeed a typical instance of the difficulties that arise when one tries to construct frameworks that are supposed to realise non-admissible semantics from basic principles—it may well be possible to correct this particular problem *a posteriori*, e.g., by including additional structures that somehow suppress b_1 in the extension-set, but it is not clear that we actually need to do so until we have constructed our attempt at the solution and computed its outcome; and then it is not clear what the actual remedy would have to look like. Of course, any algorithmic treatment would need to have its outcome checked and possibly corrected as well, until we either arrive at an acceptable solution, or exhaust our possibilities. This approach then rather resembles a random testing of possibilities than a purposeful construction.

We leave the problem alone for the time being, but would like to remark that the concept of a base framework may still have some merit even in the non-admissible case. In the classical cases that we have explored, we have always found that a base framework laid the foundation for the design, which then had to be further refined. In this sense, the construction of the base framework was considerably less complex than the construction of the entire solution; if the situation behaves analogously for the non-admissible case, then the solution to Problem 1 could be a stepping stone on the way to the complete solution.

Open Question 2. For $\mathbb{S} \in \Sigma_{\mathsf{ud}}$, is there a general solution to Problem 1 that does not involve the realisation of $\mathbb{S}$ in its entirety?

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

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

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 3. 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 31. 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 32. 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. □

Uniquely indexed extension-sets are then realised under undisputed, strongly undisputed and weakly preferred semantics by Algorithm 2. We call the resulting framework the *canonical uniquely indexed framework*.

Example 30. Figure 25 shows the resulting construction of Algorithm 2, 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 33. 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 $\mathbb{S} = \{S_1, \ldots, S_n\}$, 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 $S_i \in \mathsf{wadm}(F)$.

Algorithm 2 Realise uniquely indexed extension-sets.

Input: A uniquely indexed $\mathbb{S} = \{S_1, \ldots, S_n\} \subseteq 2^{\mathfrak{A}}$ $(n \geq 1)$.

let s_i be a unique index of S_i for $i \in \{1, \ldots, n\}$.
let $S_i^* = S_i \setminus \{s_i\}$ for $i \in \{1, \ldots, n\}$.
let $\mathcal{A} = \{s_1, \ldots, s_n\}$, $\mathcal{R} = \{(s_i, s_j) \mid i, j \in \{1, \ldots, n\}, i \neq j\}$.
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 $i \in \{1, \ldots, n\}$ **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: $(\mathcal{A}, \mathcal{R})$.

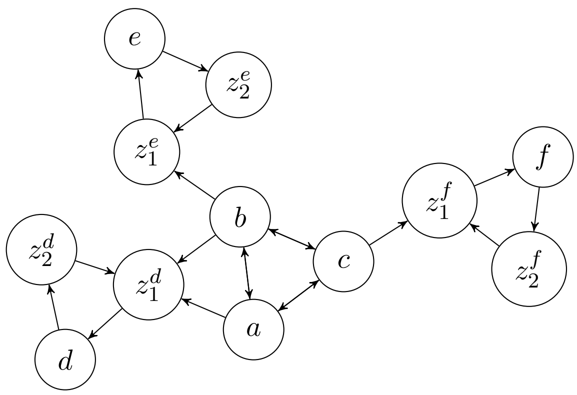

Figure 25: The canonical uniquely indexed framework, constructed for the set of sets $\mathbb{S} = \{\{a, d\}, \{b, d, e\}, \{c, f\}\}$.

- "$\supseteq$": 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 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 31 then allows the weakly preferred extensions to be realised. $\square$

7 Summary and Conclusion

We set out to explore the expressive power of various non-admissible semantics, and in doing so, we focused on 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 notion of disjointly supported extension-sets and showed that they are unrealisable by the undisputed semantics. We conjectured a characterisation of the signatures of these semantics, which can be proved 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.

In fact, at this point we are unable to specify a comprehensible construction method that is able to create base frameworks which separate evidently non-conflicting extensions under non-admissible semantics, something that we argued should be a helpful and less complex stepping stone towards a general solution to the construction problem. 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.

Related Work and Future Research We already mentioned the principle-based approach of Van Der Torre and Vesic [vdTV17] 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 [DW11] as well as Dvořák and Spanring [DS17], who compare the expressiveness of semantics in terms of translatability of frameworks, but do not characterise semantics based on their signatures.

Dyrkolbotn [Dyr14] 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.* [DDLW15] that we build upon, since they also take into account the acceptability status of auxiliary arguments). Pührer [Püh15] 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 [LPS16] 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 [Dun22] cannot be efficiently computed [Str15, Dun22].

To the author's knowledge, no attempt in the way of the approach of Dunne *et al.* [DDLW15] to characterise non-admissible semantics in terms of their signatures has been carried out before. Our conjectured characterisation of undisputed extension-sets as being non-empty and not disjointly supported 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 support or refute our conjecture. Perhaps an attempt to adapt the framework by Linsbichler, Pührer, and Strass [LPS16] 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.

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

[13] *Semi-stable* labellings [CG09] are complete labellings that minimise the set of *undecided* arguments.

[14] In an *abstract dialectical framework* (ADF) [BW10], 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.

References

[BBU20] 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.

[BBU22] 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.

[BCG11] Pietro Baroni, Martin Caminada, and Massimiliano Giacomin. An introduction to argumentation semantics. *The knowledge engineering review*, 26(4):365–410, 2011.

[BG07] Pietro Baroni and Massimiliano Giacomin. On principle-based evaluation of extension-based argumentation semantics. *Artificial Intelligence*, 171(10-15):675–700, 2007.

[BGG05] Pietro Baroni, Massimiliano Giacomin, and Giovanni Guida. Scc-recursiveness: a general schema for argumentation semantics. *Artificial Intelligence*, 168(1-2):162–210, 2005.

[BW10] 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.

[CG09] Martin WA Caminada and Dov M Gabbay. A logical account of formal argumentation. *Studia Logica*, 93:109–145, 2009.

[DDLW15] Paul E Dunne, Wolfgang Dvořák, Thomas Linsbichler, and Stefan Woltran. Characteristics of multiple viewpoints in abstract argumentation. *Artificial Intelligence*, 228:153–178, 2015.

[Don19] Pierpaolo Dondio. Weakly-admissible semantics and the propagation of ambiguity in abstract argumentation semantics. 2019.

[DRVDT20] Jeremie Dauphin, Tjitze Rienstra, and Leendert Van Der Torre. A principle-based analysis of weakly admissible semantics. In *Computational Models of Argument*, pages 167–178. IOS Press, 2020.

[DS17] Wolfgang Dvořák and Christof Spanring. Comparing the expressiveness of argumentation semantics. *Journal of Logic and Computation*, 27(5):1489–1521, 2017.

[Dun95] 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.

[Dun22] 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.

[DUW22] 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.

[DW11] Wolfgang Dvořák and Stefan Woltran. On the intertranslatability of argumentation semantics. *Journal of Artificial Intelligence Research*, 41:445–475, 2011.

[Dyr14] 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.

[LPS16] Thomas Linsbichler, Jörg Pührer, and Hannes Strass. A uniform account of realizability in abstract argumentation. In *ECAI 2016*, pages 252–260. IOS Press, 2016.

[Pap03] Christos H Papadimitriou. Computational complexity. In *Encyclopedia of computer science*, pages 260–265. 2003.

[Püh15] 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.

[Str15] Hannes Strass. Expressiveness of two-valued semantics for abstract dialectical frameworks. *Journal of Artificial Intelligence Research*, 54:193–231, 2015.

[Thi23] Matthias Thimm. On undisputed sets in abstract argumentation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 37, pages 6550–6557, 2023.

[vdTV17] Leon van der Torre and Srdjan Vesic. The principle-based approach to abstract argumentation semantics. *IfCoLog Journal of Logics and Their Applications*, 2017.