1

Chapter 18: On the Nature of Argumentation Semantics: Existence and Uniqueness, Expressibility, and Replaceability

Ringo Baumann

ABSTRACT. This chapter is devoted to argumentation semantics which play the flagship role in Dung's abstract argumentation theory. Almost all of them are motivated by an easily understandable intuition of what should be acceptable in the light of conflicts. However, although these intuitions equip us with short and comprehensible formal definitions it turned out that their intrinsic properties such as *existence and uniqueness*, *expressibility*, and *replaceability* are not that easily accessible. The chapter reviews the mentioned properties for almost all semantics available in the literature. In doing so we include two main axes: namely first, the distinction between extension-based and labelling-based versions and secondly, the distinction of different kind of argumentation frameworks such as finite or unrestricted ones.

1 Introduction

Given the large variety of existing logical formalisms it is of utmost importance to select the most adequate one for a specific purpose, e.g. for representing the knowledge relevant for a particular application or for using the formalism as a modeling tool for problem solving. Awareness of the nature of a logical formalism, in other words, of its fundamental intrinsic properties, is indispensable and provides the basis of an informed choice. Apart from the deeper understanding of the considered formalism, the study of such intrinsic properties can help to identify interesting fragments or to develop useful extensions of a formalism. Moreover, the obtained insights can be used to refine existing algorithms, or even give rise to new ones.

Presumably, the best-known intrinsic property of logics is *monotonicity*. *Monotonic logics* like first order logic are perfectly suitable for the formalization of universal truths since in these logics, whenever a formula ϕ is a logical consequence of a set of axioms Σ, it remains true forever and without exception even if we add new axioms to Σ. Formalisms which do not satisfy monotonicity, commonly referred to as *nonmonotonic logics*, allow for defeasible reasoning, i.e. it is possible to withdraw former conclusions (cf. [Brewka, 1992; Gabbay *et al.*, 1994] for excellent overviews). Both kinds of logics have their traditional application domains and apart from this fundamental choice there are many other comparison criteria influencing the decision which logic or which specific semantics of a logic to use in a certain context.

One of the first intrinsic properties which comes to mind is *computational complexity*, i.e. how expensive is it to solve typical decision problems in the candidate formalism. A further related issue is *modularity* which is, among other things, engaged with the question whether it is possible to divide a given theory in subtheories, s.t. the formal semantics of the entire theory can be obtained by constructing the semantics of the subtheories. Both topics were studied in-depth for mainstream nonmonotonic formalisms like default logic [Gottlob, 1992; Turner, 1996], logic programming under certain semantics [Dantsin *et al.*, 1997; Lifschitz and Turner, 1994] as well as abstract argumentation frameworks under various argumentation semantics (cf. Chapters 14 and 19).

In this chapter we give an overview of three further intrinsic properties of abstract argumentation semantics.

1. *existence and uniqueness* Is it possible, and if so how, to *guarantee* the existence of at least one or exactly one extension/labelling by considering the structure of a given AF F only? (cf. Section 2)
2. *expressibility* Is it possible, and if so how, to *realize* a given candidate set of extensions/labellings within a single AF F? (cf. Section 3)
3. *replaceability* Is it possible, and if so how, to *simplify* parts of a given AF F, s.t. the modified version F' and F cannot be semantically distinguished by further information which might be added later to both simultaneously? (cf. Section 4)

The question whether a certain formalism always provides one with a formal meaning or even with a uniquely determined semantical answer is a crucial factor for its suitability for the application in mind. For instance, in contrast to problem solving where a plurality of solutions may possibly be desired, in decision making one might be interested in guaranteeing a single answer provided by a logical formalism. It is well-known that a given theory in propositional logic neither has to possess a model nor, in case of existence, has there to be exactly one. The same applies to logic programs under stable model semantics. In contrast, a propositional theory of positive formulae is always satisfiable and definite logic programs constitute a subclass of logic programs where even uniqueness is guaranteed. In Section 2 we will see that Dung's abstract argumentation semantics behave in a similar way, i.e. the existence or uniqueness of extensions/labellings depend on structural restrictions of argumentation frameworks.

Expressibility is concerned with the expressive power of logical formalisms. The question here is which kinds of model sets are realizable, that is, can be the set of models of a single knowledge base of the formalism. This is a decisive property from an application angle since potential necessary or sufficient properties of model sets may rule out a logic or make it perfectly appropriate for representing certain solutions. For instance, it is well-known that in case of propositional logic any finite set of two-valued interpretations is realizable. This means, given such a finite set $\mathcal{I}$, we always find a set of formulae T, s.t. $Mod(T) = \mathcal{I}$. In case of normal logic programs it is obvious that not all model sets can be expressed, since any set of stable models forms a $\subseteq$-antichain. Remarkably, being such an antichain is not only necessary but even sufficient for realizability w.r.t. stable model semantics [Eiter *et al.*, 2013; Strass, 2015]. In case of abstract argumentation we are equipped with a high number of semantics and in Section 3 we will see that characterizing properties are not that easy. Moreover, as expected, representational limits highly depend on the chosen semantics.

In case of propositional logic we have that – in contrast to all non-monotonic logics available in the literature – standard equivalence, i.e. sharing the same models, even guarantees intersubstitutability in any logical context without loss of information. As an aside, it is not the monotonicity of a certain logic but rather the so-called *intersection property* which guarantees this behavior (cf. [Baumann and Strass, 2016]). Substitutability is of great importance for dynamically evolving scenarios since it allows to simplify parts of a theory without looking at the rest. For this reason, much effort has been devoted to characterizing *strong equivalence* for nonmonotonic formalisms, such as logic programs [Lifschitz *et al.*, 2001], causal theories [Turner, 2004], default logic [Turner, 2001] and nonmonotonic logics in general [Truszczynski, 2006; Baumann and Strass, 2016]. In Section 4 we will see that characterization theorems in case of abstract argumentation are quite different from those for the aforementioned formalisms since being strongly equivalent can be decided by looking at the syntax only.

2 Existence and Uniqueness

Given a certain logical formalism $\mathcal{L}$ together with its semantics $\sigma_{\mathcal{L}}$. One central question is whether the semantics provides any $\mathcal{L}$-theory T with a formal meaning, i.e. $|\sigma_{\mathcal{L}}(T)| \geq 1$. A more demanding property than *existence* is *uniqueness*, i.e. $|\sigma_{\mathcal{L}}(T)| = 1$ for any $\mathcal{L}$-theory T. Clearly, these properties are interesting from several perspectives. For instance, in case of uniqueness, we observe a coincidence of sceptical and credulous reasoning modes. More precisely, if $\sigma_{\mathcal{L}}(T) = \{E\}$, then $\bigcap \sigma_{\mathcal{L}}(T) = \bigcup \sigma_{\mathcal{L}}(T) = E$. Furthermore, if a theory T is interpreted as *meaningful* if and only if $\sigma_{\mathcal{L}}(T) \neq \emptyset$, then existence might be a desired property. If the latter has to be neglected in the general case, then one further challenge is to identify sufficient properties of $\mathcal{L}$-theories guaranteeing their meaningfulness.

Let us turn to abstract argumentation frameworks [Dung, 1995]. Due to the practical nature of argumentation most work in the literature restricts itself to the case of finite AFs, i.e. any considered AF consists of finitely many arguments and attacks only. For this class of AFs a proof or disproof of existence or uniqueness is mostly straightforward. In the general infinite case however conducting such proofs is more intricate. It usually involves the proper use of set theoretic axioms, like the *axiom of choice* or equivalent statements. Dung already proposed the existence of preferred extensions in the case of infinite argumentation frameworks. It has later on (e.g. [Caminada and Verheij, 2010]) been pointed out that Dung has not been precise with respect to the use of principles. The existence of semi-stable extensions for finitary[1] argumentation frameworks was first shown by Weydert with the use of model-theoretic techniques [Weydert, 2011]. Later on, Baumann and Spanring presented a first comprehensive overview of results regarding existence and uniqueness for a whole bunch of semantics considered in the literature [Baumann and Spanring, 2015]. They provided complete or alternative proofs of already known results and contributed missing results for the infinite or finitary case. We mention two interesting results: Firstly, eager semantics is exceptional among the universally defined semantics since either there is exactly one or there are infinitely many eager extensions. Secondly, stage semantics behaves similarly to semi-stable in the sense that extensions are guaranteed as long as finitary AFs are considered. A further step forward in the systematic analysis of argumentation semantics in the infinite case was presented in [Spanring, 2015]. Spanring studied the relation between non-existence of extensions and the number of non-finitary arguments. It was shown that there are AFs where one single non-finitary argument causes a collapse[2] of semi-stable semantics. Interestingly, all known AFs which do not provide any stage extension possesses infinitely many non-finitary arguments. It is an open question whether this observation applies in general [Spanring, 2015, Conjecture 14].

2.1 Basic Definitions in Dung's Abstract Argumentation Theory

For the sake of self-containedness we review all relevant definitions (for more introductory comments we refer the reader to Chapter 4 in this handbook). The standard way of defining argumentation frameworks is to introduce a certain reference set $\mathcal{U}$, so-called *universe of arguments* and to require, that all arguments used in AFs are elements of this set. More formally, for any AF $F = (A, R)$ we have $A \subseteq \mathcal{U}$ and $R \subseteq A \times A$. In order to be able to consider AFs possessing an arbitrary finite number of arguments or even infinitely many we have to request that $|\mathcal{U}| \geq \aleph_0 = |\mathbb{N}|$. No further conditions are imposed. In the following we use $\mathcal{F}$ as an abbreviation for the set of all AFs (induced by $\mathcal{U}$). An AF F is called *finite* if it possesses finitely many arguments only. Furthermore, we say that F is *finitary* if every argument has only finitely many attackers.

Definition 2.1 *An AF $F = (A, R)$ is called*

1. *finite if* $|A| \in \mathbb{N}$,

[1] An argument is called *finitary* if it receives finitely many attacks only. Moreover, an AF is said to be *finitary* if and only if it consists of finitary arguments only (cf. Definition 2.1).

[2] The term *collapse* was firstly introduced in [Spanring, 2015] and it refers to a semantics not providing any extension/labelling for a given AF.

2. *finitary if for any* $a \in A$, $|\{b \in A \mid (b,a) \in R\}| \in \mathbb{N}$ *and*

3. *arbitrary or unrestricted if* $F \in \mathcal{F}$.

In order to formalize the notions of *existence* and *uniqueness* in the context of abstract argumentation theory we have to clarify what we precisely mean by a *semantics*. In the literature two main approaches to argumentation semantics can be found, namely so-called *extension-based* and *labelling-based* versions. The main difference is that extension-based versions return a set of sets of arguments (so-called *extensions*) for any given AF in contrast to a set of sets of n-tupels (so-called *labellings*) as in case of labelling-based approaches. However, from a mathematical point of view both kinds of semantics are instances of Definition 2.2. More precisely, extension-based versions are covered by $n = 1$ and labelling-based approaches can be obtained by setting $n \geq 2$. We use $\left(2^{\mathcal{U}}\right)^n$ to denote the n-ary cartesian power of $2^{\mathcal{U}}$, i.e. $\left(2^{\mathcal{U}}\right)^n = \underbrace{2^{\mathcal{U}} \times \cdots \times 2^{\mathcal{U}}}_{n-\text{times}}$.

Definition 2.2 *A semantics is a function* $\sigma : \mathcal{F} \to 2^{\left(2^{\mathcal{U}}\right)^n}$ *for some* $n \in \mathbb{N}$, *s.t.* $F = (A, R) \mapsto \sigma(F) \subseteq \left(2^A\right)^n$.

We now introduce the two different definedness statuses of argumentation semantics which capture the notions of existence and uniqueness, namely so-called *universal* and *unique definedness*. Both versions are relativized to a certain set of AFs. If clear from context, unimportant or if $\mathcal{C} = \mathcal{F}$ we will not mention explicitly the considered set of AFs.

Definition 2.3 *Given a semantics* σ *and a set* $\mathcal{C}$ *of AFs. We say that* σ *is*

1. *universally defined w.r.t.* $\mathcal{C}$ *if* $\forall F \in \mathcal{C}$, $|\sigma(F)| \geq 1$ *and*

2. *uniquely defined w.r.t.* $\mathcal{C}$ *if* $\forall F \in \mathcal{C}$, $|\sigma(F)| = 1$.

In this section we are interested in definedness statuses w.r.t. finite, finitary and arbitrary frameworks. Besides conflict-free and admissible sets (abbreviated by *cf* and *ad*) we consider a large number of mature semantics, namely naive, stage, stable, semi-stable, complete, preferred, grounded, ideal, eager semantics as well as the more exotic cf2 and stage2 semantics (abbreviated by *na*, *stg*, *stb*, *ss*, *co*, *pr*, *gr*, *il*, *eg*, *cf2* and *stg2* respectively). In the following we introduce the extension-based versions of these semantics (indicated by $\mathcal{E}_\sigma$). Any considered semantics possesses a 3-valued labelling-based version (denoted as $\mathcal{L}_\sigma$). It is important to note that for all considered semantics we do not observe any differences between the definedness statuses of their labelling-based and extension-based versions. For the mature semantics this is due the fact that there is a one-to-one correspondence between σ-extensions and σ-labellings implying that $|\mathcal{E}_\sigma(F)| = |\mathcal{L}_\sigma(F)|$ for any AF F (for more details confer Paragraph *Basic Properties and a Fundamental Relation* in Section 4 as well as Chapter 4 of this handbook).

Before presenting the definitions we have to introduce some notational conventions. Given an AF $F = (A, R)$ and a set $E \subseteq A$. We use E_F^+ or simply, E^+ for $\{b \mid (a, b) \in R, a \in E\}$. Moreover, $E_F^\oplus$ or simply, $E^\oplus$ is called the *range* of E and stands for $E^+ \cup E$. We say *a attacks b* (in F) if $(a, b) \in R$. An argument a is *defended* by E (in F) if for each $b \in A$ with $(b, a) \in R$, b is attacked by some $c \in E$. Finally, $\Gamma_F : 2^A \to 2^A$ with $I \mapsto \{a \in A \mid a \text{ is defended by } I\}$ denotes the so-called *characteristic function* (of F) [Dung, 1995].

Definition 2.4 *Let $F = (A, R)$ be an AF and $E \subseteq A$.*

1. $E \in \mathcal{E}_{cf}(F)$ *iff for no* $a, b \in E$, $(a, b) \in R$,
2. $E \in \mathcal{E}_{na}(F)$ *iff* $E \in \mathcal{E}_{cf}(F)$ *and for no* $I \in \mathcal{E}_{cf}(F)$, $E \subset I$,
3. $E \in \mathcal{E}_{stg}(F)$ *iff* $E \in \mathcal{E}_{cf}(F)$ *and there is no* $I \in \mathcal{E}_{cf}(F)$, *s.t.* $E^\oplus \subset I^\oplus$,
4. $E \in \mathcal{E}_{stb}(F)$ *iff* $E \in \mathcal{E}_{cf}(F)$ *and* $E^\oplus = A$,
5. $E \in \mathcal{E}_{ad}(F)$ *iff* $E \in \mathcal{E}_{cf}(F)$ *and* E *defends all its elements,*
6. $E \in \mathcal{E}_{ss}(F)$ *iff* $E \in \mathcal{E}_{ad}(F)$ *and there is no* $I \in \mathcal{E}_{ad}(F)$, *s.t.* $E^\oplus \subset I^\oplus$,
7. $E \in \mathcal{E}_{co}(F)$ *iff* $E \in \mathcal{E}_{ad}(F)$ *and for any* $a \in A$ *defended by* E *in* F, $a \in E$,
8. $E \in \mathcal{E}_{pr}(F)$ *iff* $E \in \mathcal{E}_{ad}(F)$ *and for no* $I \in \mathcal{E}_{co}(F)$, $E \subset I$,
9. $E \in \mathcal{E}_{gr}(F)$ *iff* E *is the* $\subseteq$*-least fixpoint of* Γ_F,
10. $E \in \mathcal{E}_{il}(F)$ *iff* $E \in \mathcal{E}_{ad}(F)$, $E \subseteq \bigcap \mathcal{E}_{pr}(F)$ *and there is no* $I \in \mathcal{E}_{co}(F)$ *satisfying* $E \subset I \subseteq \bigcap \mathcal{E}_{pr}(F)$,
11. $E \in \mathcal{E}_{eg}(F)$ *iff* $E \in \mathcal{E}_{ad}(F)$, $E \subseteq \bigcap \mathcal{E}_{ss}(F)$ *and there is no* $I \in \mathcal{E}_{co}(F)$ *satisfying* $E \subset I \subseteq \bigcap \mathcal{E}_{ss}(F)$.

Finally, we introduce the recursively defined cf2 and stage2 semantics [Baroni *et al.*, 2005; Dvořák and Gaggl, 2012].

Definition 2.5 *Let $F = (A, R)$ be an AF and $E \subseteq A$.*

1. $E \in \mathcal{E}_{cf2}(F)$ *iff*
   - $E \in \mathcal{E}_{na}(F)$ *if* $|SCCs_F = 1|$ *and*
   - $\forall S \in SCCs_F (E \cap S) \in \mathcal{E}_{cf2}\left(F|_{UP_F(S,E)}\right)$,
2. $E \in \mathcal{E}_{stg2}(F)$ *iff*
   - $E \in \mathcal{E}_{stg}(F)$ *if* $|SCCs_F = 1|$ *and*
   - $\forall S \in SCCs_F (E \cap S) \in \mathcal{E}_{stg2}\left(F|_{UP_F(S,E)}\right)$.

Here $SCCs_F$ denotes the set of all strongly connected components of F, and for any $E, S \subseteq A$, $UP_F(S,E) = \{a \in S \mid \nexists b \in E \setminus S : (b,a) \in R\}$.

The following proposition summarizes well-known subset relations between the considered semantics. For two semantics σ, τ and a certain set of AFs $\mathcal{C}$ we use $\sigma \subseteq_{\mathcal{C}} \tau$ as a shorthand for $\sigma(F) \subseteq \tau(F)$ for any AF $F \in \mathcal{C}$. The presented relations hold for both extension-based as well as labelling-based versions of the considered semantics. In the interest of readability we present the relations graphically.

Proposition 2.6 *For semantics σ and τ, $\sigma \subseteq_{\mathcal{F}} \tau$ iff there is a path of solid arrows from σ to τ in Figure 1. A dotted arrow indicates that the corresponding subset relation is guaranteed for finite frameworks only.*

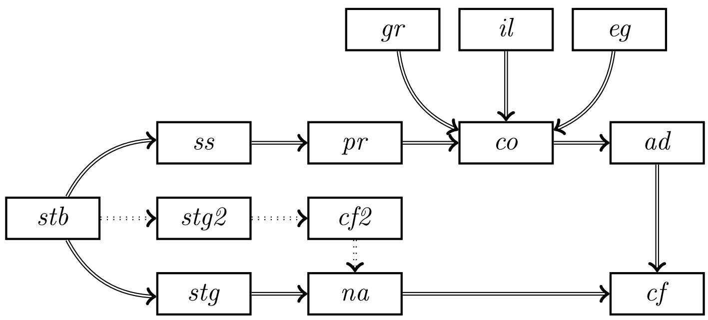

Figure 1: Subset Relations between Semantics

Detailed proofs can be found in [Baumann, 2014b, Proposition 2.7] as well as [Gaggl and Dvořák, 2016, Section 3.1]. Note that the shorthand $\sigma \subseteq_{\mathcal{C}} \tau$ requires that both semantics are total functions on $\mathcal{C}$ since a framework to which one of these semantics is undefined renders the subset shorthand undefined itself. The following simple example shows that Definition 2.5 does not always provide a definite answer on whether a certain candidate set is an *cf2*-extension or *stg2*-extension, respectively. This is due to the fact that the defined recursion does not terminate necessarily in case of non-finite AFs.[3] Consequently, *stg2* and *cf2* are not total functions regarding arbitrary frameworks.

Example 2.7 (Infinite Recursion [Baumann and Spanring, 2017])
Consider the following AF $F = (A \cup B, R)$ where

- $A = \{a_i \mid i \in \mathbb{N}\}$, $B = \{b_i \mid i \in \mathbb{N}\}$ *and*
- $R = \{(b_i, a_i), (a_{i+1}, a_i), (a_i, b_{i+1}) \mid i \in \mathbb{N}\}$

[3]We mention that the inventors of both semantics considered finite AFs only [Baroni *et al.*, 2005; Dvořák and Gaggl, 2012]. In case of finite AFs any recursion will terminate no matter which candidate set is considered.

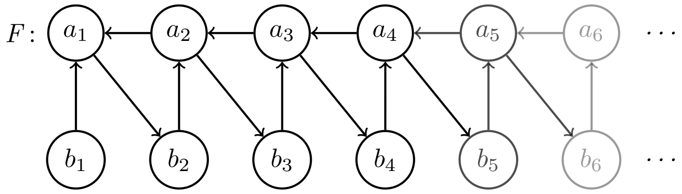

Let $\sigma \in \{cf2, stg2\}$. We want to check whether the candidate set $E = \{b_i \mid i \in \mathbb{N}\}$ is a σ-extension. Observe that the AF F possesses two SCCs, namely one consisting of the single argument b_1 and the other containing the remaining arguments, i.e. $S_1 = \{b_1\}$ and $S_2 = (A \cup B) \setminus \{b_1\}$. For S_1 we end up with the base case returning a positive answer. For S_2 we have to consider the AF $F' = F|_{UP_F(S_2,E)} = F|_{(A\cup B)\setminus\{a_1,b_1\}}$ (since a_1 is attacked by $b_1 \in E \setminus S_2$) and the set $S' = E \cap S_2 = \{b_i \mid i \in \mathbb{N}, i \geq 2\}$. Obviously, determining whether S' is an σ-extension w.r.t. F' is equivalent to decide whether S is an σ-extension w.r.t. F. This means, the consideration of the candidate set E leads to infinite recursion.

2.2 Finite AFs

As a matter of fact, in order to show that a certain semantics σ is not universally defined w.r.t. a certain set $\mathcal{C}$ it suffices to present an AF $F \in \mathcal{C}$, s.t. $\sigma(F) = \emptyset$. Contrastingly, an affirmative answer w.r.t. universal definedness requires a proof involving all AFs in $\mathcal{C}$. Let us consider finite AFs first. It is well-known that stable semantics does not warrant the existence of extensions/labellings even in the case of finite AFs. Witnessing examples are given by odd-cycles (cf. Example 2.8). Interestingly, in case of finite AFs we have that being odd-cycle free is sufficient for warranting at least one stable extension/labelling.[4]

Example 2.8 *The following minimalistic AFs cause a collapse of stable semantics, i.e. $stb(F_1) = stb(F_3) = \emptyset$.*

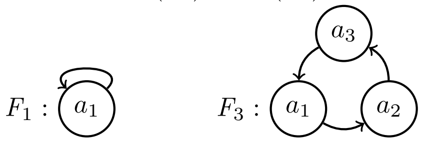

Observe that both frameworks do possess semi-stable, stage2 as well as stage extensions/labellings. The extensions are as follows: For any $\sigma \in \{ss, stg2, stg\}$, $\tau \in \{stg2, stg\}$, $\mathcal{E}_\sigma(F_1) = \{\emptyset\} = \mathcal{E}_{ss}(F_3)$ and $\mathcal{E}_\tau(F_3) = \{\{a_1\}, \{a_2\}, \{a_3\}\}$.

Let us consider now semi-stable semantics. Example 2.8 shows that AFs may possess semi-stable extensions even in the absence of stable extensions. Are semi-stable extensions possibly guaranteed in case of finite AFs? Consider

[4]This is due to the fact that firstly, in case of finite AFs, being odd-cycle free coincides with being *limited controversial* [Dung, 1995, Definition 32] and secondly, any limited controversial AFs warrants the existence of at least one stable extensions [Dung, 1995, Corollary 36].

the following explanations about the existence of semi-stable extensions taken from [Caminada, 2006]:

> For every argumentation framework there exists at least one semi-stable extension. This is because there exists at least one complete extension, and a semi-stable extension is simply a complete extension in which some property (the union of itself and the arguments it defeats) is maximal.

We would like to point out two issues. Firstly, the presented explanation should not be understood as: Since any semi-stable extension is a complete one and complete semantics is universally defined we conclude that semi-stable semantics is universally defined. Accepting this kind of (false) argumentation would imply the universal definedness of stable semantics since also any stable extension is a complete one. The second issue is that the presented explanation is not precise about *why* it is guaranteed that the non-empty set of complete extensions possesses at least one range-maximal member. The following statement gives a more precise explanation [Caminada *et al.*, 2012]:

> For every (finite) argumentation framework, there exists at least one semi-stable extension. This is because there exists at least one complete extension (the grounded) and the fact that the argumentation framework is finite implies that there exist at most a finite number of complete extensions. The semi-stable extensions are then simply those complete extensions in which some property (its range) is maximal.

This means, the additional argument that we have to compare finitely many complete extensions only justifies the universal definedness of semi-stable extensions in case of finite AFs. Obviously, in case of infinite AFs we cannot expect to have finitely many complete extensions implying that this kind of argumentation is no longer valid for finitary as well as infinite AFs in general.

In the rest of this subsection we want to argue why all considered semantics except the stable one are universally defined in case of finite AFs.[5] Remember that many semantics are looking for certain $\subseteq$-maximal elements. The main advantage in case of finiteness is that it is simply impossible to have infinite $\subseteq$-chains which guarantees the existence of $\subseteq$-maximal elements. Consider the following more detailed explanations. Given a finite AF $F = (A, R)$, i.e. $|A| = n \in \mathbb{N}$. Consequently, $1 \leq \left|2^A\right| = 2^n \in \mathbb{N}$. By definition of any extension-based semantics σ we derive $0 \leq |\mathcal{E}_\sigma(F)| \leq 2^n$ since $\mathcal{E}_\sigma(F) \subseteq 2^A$ (cf. Definition 2.2). This means, for any finite F and any semantics σ we have at least one candidate set for being a σ-extension (namely, the empty set) and at most finitely many σ-extensions. In any case, the empty set is conflict-free as well as admissible, i.e. $|\mathcal{E}_{cf}(F)|, |\mathcal{E}_{ad}(F)| \geq 1$. Furthermore, naive and

[5] We mention that grounded, ideal and eager semantics are even uniquely defined w.r.t. finite AFs. This will be a by-product of Theorem 2.23, Corollary 2.22 as well as Theorem 2.25.

preferred semantics are looking for $\subseteq$-maximal conflict-free or admissible sets, respectively. Since we have finitely many conflict-free as well as admissible sets only we derive the universal definedness of naive and preferred semantics in case of finite AFs. Combining $\mathcal{E}_{pr} \subseteq \mathcal{E}_{co}$ and $|\mathcal{E}_{pr}(F)| \geq 1$ yields the universal definedness of complete semantics in case of finite AFs. Moreover, since $1 \leq |\mathcal{E}_{cf}(F)|, |\mathcal{E}_{ad}(F)| \leq 2^n$ is given we obtain the universal definedness of stage and semi-stable semantics in case of finite AFs because the existence of $\subseteq$-range-maximal is guaranteed. Let us consider ideal and eager semantics. Candidate sets of both semantics are admissible sets being in the intersection of all preferred or semi-stable extensions, respectively. Note that there is at least one admissible set satisfying this property, namely the empty one since definitely $\emptyset \subseteq \bigcap \mathcal{E}_{pr}(F) \subseteq \mathcal{U}$ as well as $\emptyset \subseteq \bigcap \mathcal{E}_{ss}(F) \subseteq \mathcal{U}$. This means, the sets of candidates are non-empty and finite which guarantees the existence of $\subseteq$-maximal elements implying the universal definedness of ideal and eager semantics in case of finite AFs. The grounded extension, i.e. the $\subseteq$-least fixpoint of the characteristic function Γ_F, is guaranteed due to the monotonicity of Γ_F and the famous Knaster-Tarski theorem [Tarski, 1955]. Finally, even the more exotic stage2 as well as cf2 semantics are universally defined w.r.t. finite AFs. This can be seen as follows: Obviously, finitely many as well as initial SCCs are guaranteed due to finiteness. Consequently, one may start with computing stage/naive extension on these initial components and "propagate" the resulting extensions to the subsequent SCCs and so on. This procedure will definitely terminate and ends up with stage2/cf2 extensions. Apart from stable semantics we have argued that the extension-based versions of all considered semantics are universally defined w.r.t. finite AFs. In case of mature semantics, the result carry over to their labelling-based versions since any of these semantics possesses a one-to-one-correspondence between extensions and labellings. This property does not hold in case of admissible as well as conflict-free sets. However, since any admissible/conflict-free set induce at least one admissible/conflict-free labelling the result applies to their labelling versions too (cf. Chapter 4).

2.3 Arbitrary AFs

Non-well-defined Semantics

In contrast to all other semantics available in the literature, cf2 as well as stage2 semantics were originally defined recursively. The recursive schema is based on the decomposition of AFs along their strongly connected components (SCCs). Roughly speaking, the schema takes a base semantics σ and proceeds along the induced partial ordering and evaluates the SCCs according to σ while propagating relevant results to subsequent SCCs. This procedure defines a $\sigma 2$ semantics.[6] Given so-called *SCC-recursiveness* (cf. Chapter 19) we have to face some difficulties in drawing conclusions with respect to infinite AFs. Firstly,

[6]Following this terminology we have to rename *cf2* semantics to *na2* semantics since its base semantics is the naive semantics and not conflict-free sets.

arbitrary AFs need not to possess initial SCCs which is granted for finite AFs. This makes checking whether a certain set is an $\sigma 2$-extension more complicated and in particular, especially due to the recursive definitions not that easy to handle. Secondly, even worse, even if an AF as well as subsequent subframeworks of it possess initial SCCs there is no guarantee that any recursion will stop in finitely many steps. More precisely, as shown in Example 2.7 there might be candidate sets which lead to infinite recursion, i.e. the base case will never be considered. In [Gaggl and Dvořák, 2016, Propositions 2.12 and 3.2] the authors considered alternative non-recursive definitions of cf2 as well as stage2 semantics in case of finite AFs. It is an open question whether these definitions overcome the problem of undefinedness for arbitrary frameworks.

Collapsing Semantics

Dealing with finite AFs is a common as well as attractive and reasonable restriction, due to their computational nature. In the subsection before we have argued that apart from stable semantics all considered semantics are universally defined w.r.t. finite AFs. It is an important observation that warranting the existence of σ-extensions/labellings in case of finite AFs does not necessarily carry over to the infinite case, i.e. the semantics σ does not need to be universally defined w.r.t. arbitrary AFs. Take for instance semi-stable and stage semantics. To the best of our knowledge the first example showing that semi-stable as well as stage semantics does not guarantee extensions/labellings in case of non-finite AFs was given in [Verheij, 2003, Example 5.8.] and is picked up in the following example.

Example 2.9 (Collapse of Stage and Semi-stable Semantics) *Consider the following AF $F = (A \cup B \cup C, R)$ where*

- $A = \{a_i \mid i \in \mathbb{N}\}$, $B = \{b_i \mid i \in \mathbb{N}\}$, $C = \{c_i \mid i \in \mathbb{N}\}$ *and*
- $R = \{(a_i, b_i), (b_i, a_i), (b_i, c_i), (c_i, c_i) \mid i \in \mathbb{N}\} \cup$
  $\{(b_i, b_j), (b_i, c_j) \mid i, j \in \mathbb{N}, j < i\}$

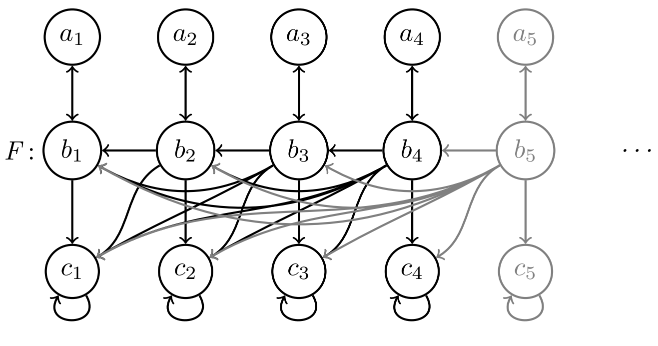

The set of preferred and naive extensions coincide, in particular $\mathcal{E}_{pr}(F) = \mathcal{E}_{na}(F) = \{A\} \cup \{E_i \mid i \in \mathbb{N}\}$ *where* $E_i = (A \setminus \{a_i\}) \cup \{b_i\}$. *Furthermore, none* *of these extensions is* $\subseteq$*-range-maximal since* $A^{\oplus} \subsetneq E_i^{\oplus} \subsetneq E_{i+1}^{\oplus}$ *for any* $i \in \mathbb{N}$. *In consideration of* $ss \subseteq pr$ *and* $stg \subseteq na$ *(cf. Figure 1) we conclude that this framework possesses neither semi-stable nor stage extensions/labellings.*

In Example 2.7 we have seen that cf2 as well as stage2 semantics are not well-defined in general. This means, there are infinite AFs and candidate sets leading to an infinite recursion implying that there is no definite answer on whether such a set is an extension. However, the following example shows that even if for any candidate set a definitive decision is possible there need not to be an extension in contrast to finite AFs.

Example 2.10 (Collapse of Cf2 and Stage2 Semantics) *Taking into account the AF* $F = (A \cup B \cup C, R)$ *from Example 2.9. Consider the AF* $G = F|_B$, *i.e. the restriction of* F *to* B.

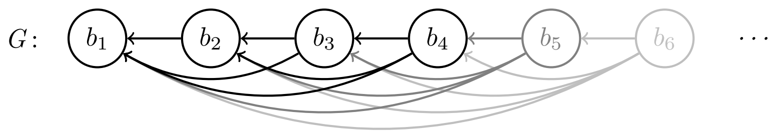

Let $\sigma \in \{cf2, stg2\}$. *Obviously, any argument* b_i *constitutes a SCC* $\{b_i\}$ *which is evaluated as* $\{b_i\}$ *by the base semantics of* σ. *Consequently,* $\emptyset$ *cannot be a* σ*-extension. Furthermore, a singleton* $\{b_j\}$ *cannot be a* σ*-extension either. The* b_i*'s for* $i > j$ *are not affected by* $\{b_j\}$ *and thus, the evaluation of* $G|_{UP_G(\{b_i\},\{b_j\})} = G|_{\{b_i\}} = (\{b_i\}, \emptyset)$ *do not return* $\emptyset$ *as required. Finally, any set containing more than two arguments would rule out at least one of them and thus, cannot be a* σ*-extension. Hence,* $|\mathcal{E}_\sigma(G)| = |\mathcal{L}_\sigma(G)| = 0$.

In Example 2.9 we have seen an AF F without any semi-stable and stage extensions/labellings. In [Baumann and Spanring, 2015] the authors studied the question of existence-dependency between both semantics in case of infinite AFs. More precisely, they studied whether it is possible that some AF does have semi-stable but no stage extensions or vice versa, there are stage but no semi-stable extensions. The following Example 2.11 shows that stage extensions might exist even if semi-stable semantics collapses.[7]

Example 2.11 (No Semi-Stable but Stage Extensions/Labellings) *Consider again the AF* F *depicted in Example 2.9. Using the components of* F *we define* $G = (A \cup B \cup C \cup D \cup E, R \cup R')$ *where*

- $D = \{d_i \mid i \in \mathbb{N}\}$ *and* $E = \{e_i \mid i \in \mathbb{N}\}$ *and*
- $R' = \{(a_i, d_i), (d_i, a_i), (b_i, d_i), (d_i, b_i), (d_i, c_i), (e_i, d_i), (e_i, e_i) \mid i \in \mathbb{N}\}$

[7] The AF $G = F|_B$ depicted in Example 2.10 witnesses the reverse case. It can be checked that $\mathcal{E}_{ss}(G) = \{\emptyset\}$ and $\mathcal{E}_{stg}(G) = \emptyset$ (cf. [Baumann and Spanring, 2015, Example 2] for further explanations).

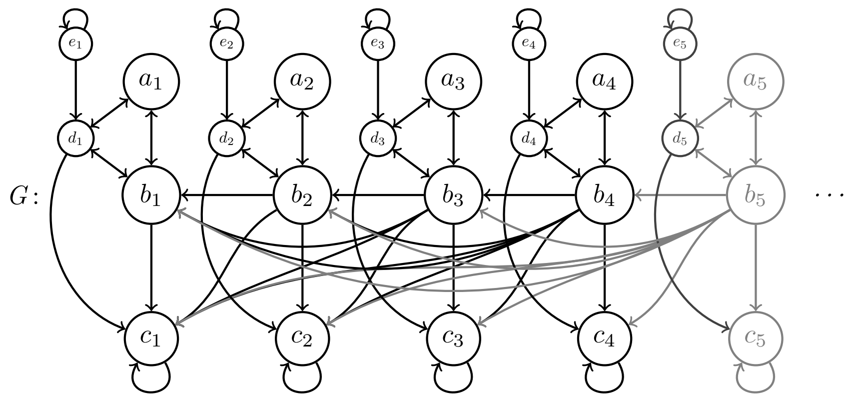

In comparison to Example 2.9 we do not observe any changes as far as preferred and semi-stable semantics are concerned. In particular, $\mathcal{E}_{pr}(G) = \{A\} \cup \{E_i \mid i \in \mathbb{N}\}$ *where* $E_i = (A \setminus \{a_i\}) \cup \{b_i\}$ *and again, none of these extensions is* $\subseteq$*-range-maximal. Hence,* $\mathcal{E}_{ss}(G) = \emptyset$*. Observe that we do have additional conflict-free as well as naive sets, especially the set* D*. Since any* $e \in E$ *is self-defeating and unattacked and furthermore,* $D^{\oplus} = A \cup B \cup C \cup D$ *we conclude,* $\mathcal{E}_{stg}(G) = \{D\}$*. Due to the one-to-one correspondence the collapse or non-collapse transfer to their labelling-based versions.*

Universally Defined Semantics

We now turn to semantics which are universally defined w.r.t. the whole class of AFs. The first non-trivial result in this line was already proven by Dung himself, namely the universal definedness of the extension-based version of preferred semantics [Dung, 1995, Corollary 12]. He argued that the Fundamental Lemma (cf. [Dung, 1995, Lemma 10] or Chapter 4 in this handbook) immediately implies that the set of all admissible sets is a complete partial order which means that any $\subseteq$-chain possesses a least upper bound. Then (and this was not explicitly stated in [Dung, 1995]), due to the famous *Zorn's lemma* [Zorn, 1935] the existence of $\subseteq$-maximal admissible sets, i.e. preferred extensions, is guaranteed.

In order to get an idea how things work in the general case we illustrate some proofs in more detail. We will see that a proof of universal definedness w.r.t. arbitrary AFs is completely different to the argumentation in case of finite ones. In order to keep this section self-contained we start with Zorn's lemma and an equivalent version of it.

Lemma 2.12 ([Zorn, 1935]) *Given a partially ordered set* $(P, \leq)$*. If any* $\leq$*-chain possesses an upper bound, then* $(P, \leq)$ *has a maximal element.*

Lemma 2.13 *Given a partially ordered set* $(P, \leq)$*. If any* $\leq$*-chain possesses an upper bound, then for any* $p \in P$ *there exists a maximal element* $m \in P$*, s.t.* $p \leq m$*.*

Having Lemma 2.13 at hand we may easily argue that any conflict-free/admissible set is bounded by a naive/preferred extension.

Lemma 2.14 *Given $F = (A, R)$ and $E \subseteq A$,*

1. *if $E \in \mathcal{E}_{cf}(F)$, then there exists $E' \in \mathcal{E}_{na}(F)$ s.t. $E \subseteq E'$ and*
2. *if $E \in \mathcal{E}_{ad}(F)$, then there exists $E' \in \mathcal{E}_{pr}(F)$ s.t. $E \subseteq E'$.*

Proof. For $F = (A, R)$ we have the associated power set lattice $(2^A, \subseteq)$. Consider now the partially ordered fragments $\mathcal{C} = (\mathcal{E}_{cf}(F), \subseteq)$ as well as $\mathcal{A} = (\mathcal{E}_{ad}(F), \subseteq)$. In accordance with Lemma 2.13 the existence of naive and preferred supersets is guaranteed if any $\subseteq$-chain possesses an upper bound in $\mathcal{C}$ or $\mathcal{A}$, respectively. Given a $\subseteq$-chain $\mathcal{E} \subseteq \mathcal{E}_{cf}(F)$ or $\mathcal{E} \subseteq \mathcal{E}_{ad}(F)$, respectively. Consider now $\bar{E} = \bigcup \mathcal{E}$. Obviously, $\bar{E}$ is an upper bound of $\mathcal{E}$, i.e. for any $E \in \mathcal{E}$, $E \subseteq \bar{E}$. It remains to show that $\bar{E}$ is conflict-free or admissible, respectively. Conflict-freeness is a finite condition. This means, if there were conflicting arguments $a, b \in \bar{E}$ there would have to be some conflict-free sets $E_a, E_b \in \bar{E}$, s.t. $a \in E_a$ and $b \in E_b$. Since $\mathcal{E}$ is a $\subseteq$-chain we have $E_a \subseteq E_b$ or $E_b \subseteq E_a$ which contradicts the conflict-freeness of at least one of them. Assume now $\bar{E}$ is not admissible. Consequently, there is some $a \in \bar{E}$ that is not defended by $\bar{E}$. Furthermore, there has to be an $E_a \in \mathcal{E}$, s.t. $a \in E_a$ contradicting the admissibility of $E_a \in \mathcal{E}_{ad}(F)$. ■

According to the last lemma, we may deduce the universal definedness of the extension-based versions of preferred as well as naive semantics as long as, for any AF F, the existence of at least one conflict-free or admissible set is guaranteed. This is an easy task since the empty set is conflict-free as well as admissible even in the case of arbitrary AFs. Consequently, universal definedness of both extension-based semantics is given and the same applies to their labelling-based versions due to their one-to-one correspondence.

Theorem 2.15 *Let $\sigma \in \{pr, na\}$. The semantics σ is universally defined.*

Remember that no matter which cardinality a considered AF possesses, we have that any preferred extension/labelling is a complete extension/labelling (Proposition 2.6). Thus, having the universal definedness of preferred semantics at hand we deduce that even complete semantics is universally defined w.r.t. the whole class of AFs .

Theorem 2.16 *The semantics co is universally defined.*

Let us consider now eager and ideal semantics. An eager extension is defined as the $\subseteq$-maximal admissible set that is a subset of each semi-stable extension. This is very similar to the definition of an ideal extension where the role of semi-stable extensions is taken over by preferred ones. On a more abstract level, both semantics are instantiations of the following schema.

Definition 2.17 *Let σ be a semantics (so-called base semantics). We define the σ-parametrized semantics ad^σ as follows. For any AF F,*

$$\mathcal{E}_{ad^\sigma} = \max_{\subseteq} \left\{ E \in \mathcal{E}_{ad}(F) \,\middle|\, E \subseteq \bigcap_{S \in \mathcal{E}_\sigma(F)} S \right\}.$$

These kind of semantics were firstly introduced in [Dvorák *et al.*, 2011]. The authors studied general properties of these semantics in case of finite AFs with the additional restriction that the base semantics σ has to be universally defined. The following general theorem requires neither finiteness of AFs, nor any assumption on the base semantics.

Theorem 2.18 *Any σ-parametrized semantics is universally defined.*

Proof. Given an AF $F = (A, R)$ and a σ-parametrized semantics ad^σ. Consider the set $\Sigma = \left\{ E \in \mathcal{E}_{ad}(F) \mid E \subseteq \bigcap_{S \in \mathcal{E}_\sigma(F)} S \right\}$. Note that in the collapsing case, i.e. $\mathcal{E}_\sigma(F) = \emptyset$, we have: $\bigcap_{S \in \mathcal{E}_\sigma(F)} S = \{x \in \mathcal{U} \mid \forall S \in \mathcal{E}_\sigma(F) : x \in S\} = \mathcal{U}$. However, in any case $\Sigma \neq \emptyset$ since for any F, $\emptyset \in \mathcal{E}_{ad}(F)$ and obviously, $\emptyset \subseteq \bigcap_{S \in \mathcal{E}_\sigma(F)} S \subseteq \mathcal{U}$. In order to show that $\mathcal{E}_{ad^\sigma}(F) \neq \emptyset$ it suffices to prove that $(\Sigma, \subseteq)$ possesses maximal elements. We will use Zorn's lemma. Given a $\subseteq$-chain $\mathcal{E} \in 2^\Sigma$. Consider now $\bar{E} = \bigcup \mathcal{E}$. Analogously to the proof of Lemma 2.14 we may easily show that $\bar{E}$ is conflict-free and even admissible. Moreover, since for any $E \in \mathcal{E}$, $E \subseteq \bigcap_{S \in \sigma(F)} S$ we deduce $\bar{E} \subseteq \bigcap_{S \in \sigma(F)} S$ guaranteeing $\bar{E} \in \Sigma$. Now, applying Lemma 2.12, we deduce the existence of $\subseteq$-maximal elements in Σ, i.e. $|\mathcal{E}_{ad^\sigma}(F)| \geq 1$ concluding the proof. ■

In particular, we obtain the result for the extension-based versions of eager and ideal semantics and thus, due to the one-to-one correspondence for both labelling-based versions too.

Corollary 2.19 *Let $\sigma \in \{eg, il\}$. The semantics σ is universally defined.*

One obvious question is whether the statement above can be strengthened in the sense that both semantics are even uniquely defined w.r.t. the whole class of AFs. The following proposition, in particular the second item, shows that the unique definedness of eager semantics w.r.t. finite frameworks does not carry over to the general unrestricted case.

Proposition 2.20 *For any F we have:*

1. $ss(F) = \emptyset \Rightarrow eg(F) = pr(F)$ *and*
2. $ss(F) = \emptyset \Rightarrow |eg(F)| \geq \aleph_0 = |\mathbb{N}|$.

Proof. We show both assertions for the extension-based versions.
1.) Given $F = (A, R)$ and let $\mathcal{E}_{ss}(F) = \emptyset$. Hence, $\bigcap_{S \in \mathcal{E}_{ss}(F)} S = \mathcal{U}$. Consequently, $\mathcal{E}_{ss}(F) = \max_{\subseteq} \{E \in \mathcal{E}_{ad}(F) \mid E \subseteq \mathcal{U}\}$. This means, $\mathcal{E}_{ss}(F) = \mathcal{E}_{pr}(F)$.
2.) We show the contrapositive. Assume $|\mathcal{E}_{eg}(F)| = n$ for some finite cardinal $n \in \mathbb{N}$. Due to the first statement we derive, $|\mathcal{E}_{pr}(F)| = n$. Since $ss \subseteq pr$ (cf. Proposition 2.6) we have finitely many candidates only. Furthermore, among these preferred extensions has to be at least one $\subseteq$-range-maximal set implying $\mathcal{E}_{ss}(F) \neq \emptyset$. ■

In a nutshell, if we observe a collapse of semi-stable semantics, then eager and preferred semantics coincide and moreover, we necessarily have infinitely many eager extensions/labellings. An AF witnessing such a behaviour can be found in Example 2.9.

Uniquely Defined Semantics

Although eager and ideal semantics are instances of σ-parametrized semantics we have shown the non-unique definedness (Proposition 2.20) for eager semantics only. This is no coincidence since preferred semantics, the base semantics of ideal semantics is universally defined in contrast to semi-stable semantics, the base semantics of the eager semantics. Moreover, the following theorem shows that any σ-parametrized semantics warrants the existence of exactly one extension if σ-extensions are conflict-free as well as guaranteed ([Dvorák *et al.*, 2011, Proposition 1]).

Theorem 2.21 *Given a σ-parametrized semantics ad^{σ}, s.t. $\sigma \subseteq cf$ and σ is universally defined w.r.t. a class $\mathcal{C}$, then ad^{σ} is uniquely defined w.r.t. $\mathcal{C}$.*

Proof. Given an AF $F = (A, R)$. We already know $|\mathcal{E}_{ad^{\sigma}}(F)| \geq 1$ (Theorem 2.18). Hence, it suffices to show $|\mathcal{E}_{ad^{\sigma}}(F)| \leq 1$. Suppose, to derive a contradiction, that for some $I_1 \neq I_2$ we have $I_1, I_2 \in \mathcal{E}_{ad^{\sigma}}(F)$. Consequently, by Definition 2.17, $I_1, I_2 \in \mathcal{E}_{ad}(F)$ and $I_1, I_2 \subseteq \bigcap_{S \in \mathcal{E}_{\sigma}(F)} S$ as well as neither $I_1 \subseteq I_2$, nor $I_2 \subseteq I_1$. Obviously, $I_1 \cup I_2 \subseteq \bigcap_{S \in \mathcal{E}_{\sigma}(F)} S$. Since $\mathcal{E}_{\sigma}(F) \neq \emptyset$ and I_1 as well as I_2 has to be subsets of any σ-extension (which are conflict-free by assumption) we deduce $I_1, I_2 \in \mathcal{E}_{cf}(F)$ and thus, $I_1 \cup I_2 \in \mathcal{E}_{cf}(F)$. Furthermore, since both sets are admissible in F we derive $I_1 \cup I_2 \in \mathcal{E}_{ad}(F)$ contradicting the $\subseteq$-maximality of at least one of the sets I_1 and I_2. ■

Corollary 2.22 *The semantics il is uniquely defined.*

A further prominent representative of uniquely defined semantics w.r.t. the whole class of AFs is the grounded semantics. Its unique definedness was already implicitly given in [Dung, 1995]. Unfortunately, this result was not explicitly stated in the paper. Nevertheless, in [Dung, 1995, Theorem 25] it was shown that firstly, the set of all complete extensions form a complete semilattice w.r.t. subset relation, i.e. the existence of a $\subseteq$-greatest lower bound for any non-empty subset S is implied. Secondly, it was proven that the grounded extension is the $\subseteq$-least complete extension. Consequently, the existence of such a $\subseteq$-least extension is justified via setting $S = \mathcal{E}_{co}(F)$ for any given F. Alternatively, one may stick to the original definition of the grounded extension, namely as $\subseteq$-least fixpoint of the characteristic function Γ_F and argue that the monotonicity of Γ_F as well as the Knaster-Tarski theorem [Tarski, 1955] imply its existence.

Theorem 2.23 *The semantics gr is uniquely defined.*

2.4 Finitary AFs

Let us consider now finitary AFs, i.e. AFs where each argument receives finitely many attacks only. It was already observed by Dung itself that finitary AFs possess useful properties. More precisely, if an AF is finitary, then the characteristic function Γ is not only monotonic, but even ω-continuous [Dung, 1995, Lemma 28] (which does not hold in case of arbitrary AFs [Baumann and Spanring, 2017, Example 1]). This implies that the least fixed point of Γ, i.e. the unique grounded extension, can be "computed" in at most ω steps by iterating Γ on the empty set (cf. [Rudin, 1976] for more details). A further advantage of finitary AFs is that for some semantics σ, the existence or even uniqueness of σ-extension is guaranteed which cannot be shown in general.

Consider again the AF F depicted in Example 2.9. In contrast to finite AFs where the existence of semi-stable as well stage extensions is guaranteed we observed a collapse of both semantics. Not that F is not finitary since, for example, the argument b_1 receives infinitely many attacks. A positive answer in case of semi-stable semantics, i.e. universal definedness w.r.t. finitary AFs was conjectured in [Caminada and Verheij, 2010, Conjecture 1] and firstly proven by Emil Weydert in [Weydert, 2011, Theorem 5.1]. Weydert proved his result in a first order logic setup using generalized argumentation frameworks. Later on, Baumann and Spanring provided an alternative proof using transfinite induction. Moreover, they showed that even stage semantics warrants the existence of at least one extension in case of finitary AFs [Baumann and Spanring, 2015, Theorem 14]. For detailed proofs we refer the reader to the mentioned scientific papers.

Theorem 2.24 *Let $\sigma \in \{ss, stg\}$. The semantics σ is universally defined w.r.t. finitary AFs.*

Applying Theorem 2.21 we derive that exactly one eager extension/labelling is guaranteed as long as the AF in question is finitary.

Theorem 2.25 *The semantics eg is uniquely defined w.r.t. finitary AFs.*

2.5 Summary of Results and Conclusion

In this section we gave an overview on the question whether certain semantics guarantee the existence or even unique determination of extensions/labellings. We have seen that these properties may vary from subclass to subclass. The following table gives a comprehensive overview over results presented in this section. The entry $''\exists''$ ($''\exists!''$) in row *certain* and column σ indicates that the semantics σ is universally (uniquely) defined w.r.t. the class of *certain* frameworks. No entry reflects the situation that a *certain* AF can be found which do not provide any σ-extension/labelling, i.e. σ collapses. The two question marks represent open problems. Note that we already observed that cf2 as well as stage2 semantics are not well-defined in case of finitary as well arbitrary AFs. This means, there are infinite AFs and candidate sets leading to an infinite recursion implying that there is no definite answer on whether such a set is an extension (Example 2.7). Nevertheless, even if for any candidate set a definitive decision is possible there are infinite (but non-finitary) AFs where both semantics collapse (Example 2.10). In [Baumann and Spanring, 2015, Conjecture 1] it is conjectured that this is impossible in case of finitary frameworks.

	stb	*ss*	*stg*	*cf2*	*stg2*	*pr*	*ad*	*co*	*gr*	*il*	*eg*	*na*	*cf*
finite		$\exists$	$\exists$	$\exists$	$\exists$	$\exists$	$\exists$	$\exists$	$\exists!$	$\exists!$	$\exists!$	$\exists$	$\exists$
finitary		$\exists$	$\exists$	?	?	$\exists$	$\exists$	$\exists$	$\exists!$	$\exists!$	$\exists!$	$\exists$	$\exists$
arbitrary						$\exists$	$\exists$	$\exists$	$\exists!$	$\exists!$	$\exists$	$\exists$	$\exists$

Table 1: Definedness Statuses of Semantics

For a detailed complexity analysis of the associated decision problems, i.e. *Given an AF F. Is* $|\sigma(F)| \geq 1$ *or even,* $|\sigma(F)| = 1$*?* we refer the reader to the complexity chapter of this handbook (Chapter 14, Table 1). The mentioned decisions problems are considered for finite AFs only since the input-length, i.e. the length of the formal encoding of an AF has to be finite (for finite representations of infinite AFs we refer the reader to [Baroni *et al.*, 2013]). Due to the table above some complexity results are immediately clear. For instance, the existence problem is trivial for all considered semantics except the stable one. An upper bound for the complexity of the uniqueness problem can be obtained via the complexity of the corresponding verification problem, i.e. *Given an AF F and a set E. Is* $E \in \mathcal{E}_\sigma(F)$*?*. More precisely, an algorithm which decides the uniqueness problem is the following two-step procedure: first, guessing a certain set E non-deterministically and second, verifying whether this set is an σ-extension.

As already mentioned, most of the literature concentrate on finite AFs for several reasons, especially due to their computational nature. However, allowing an infinite number of arguments is essential in applications where upper bounds on the number of available arguments cannot be established a priori, such as for example in dialogues [Belardinelli *et al.*, 2015] or modeling approaches including time or action sequences [Baumann and Strass, 2012].

Moreover, even actual infinite AFs frequently occur in the instantiation-based context. More precisely, the semantics of so-called *rule-based argumentation formalisms* (cf. Chapter 6 as well as [Prakken, 2010]) is given via the evaluation of induced Dung-style AFs. In this context, even a finite set of rules may lead to an infinite set of arguments as observed in (cf. [Caminada and Oren, 2014; Strass, 2015]).

In 2011, Baroni et al. wrote "As a matter of fact, we are not aware of any systematic literature analysis of argumentation semantics properties in the infinite case." [Baroni *et al.*, 2011, Section 4.4]. Since then only few works have contributed to a better understanding of infinite AFs. In [Baroni *et al.*, 2013] the authors studied to which extent infinite AFs can be finitely represented via formal languages and considered several decision problems within this context. In [Baumann and Spanring, 2015] a detailed study of the central properties of existence and uniqueness as presented in this chapter was given. Recently, the same authors addressed several central issues like *expressibility*, *intertranslatability* or *replaceability* (cf. Sections 3 and 4) in the general unrestricted case [Baumann and Spanring, 2017].

3 Expressibility

Given a certain logical formalism $\mathcal{L}$ used as knowledge representation language or modelling tool in general. Depending on the application in mind, it might be interesting to know which kinds of model sets are actually expressible in $\mathcal{L}$? More formally, if $\sigma_{\mathcal{L}}$ denotes the semantics of $\mathcal{L}$, we are interested in determining the set $\mathcal{R}_{\mathcal{L}} = \{\sigma_{\mathcal{L}}(T) \mid T \text{ is an } \mathcal{L}\text{-theory}\}$. This task, also known as *realizability* or *defineability*, highly depends on the considered formalism $\mathcal{L}$. Clearly, potential necessary or sufficient properties for being in $\mathcal{R}_{\mathcal{L}}$, i.e. being $\sigma_{\mathcal{L}}$*-realizable*, may rule out a logic or make it perfectly appropriate for a certain application. For instance, it is well-known that in case of propositional logic any finite set of two-valued interpretations is realizable. This means, given such a finite set $\mathcal{I}$, we always find a set of formulae T, s.t. $Mod(T) = \mathcal{I}$. Differently, in case of normal logic programs under stable model semantics we have that any finite candidate set is realizable if and only if it forms a $\subseteq$-antichain, i.e. any two sets of the candidate set have to be incomparable with respect to the subset relation. Remarkably, being such an $\subseteq$-antichain is not only necessary but even sufficient for realizability w.r.t. stable model semantics [Eiter *et al.*, 2013; Strass, 2015]. One major application of realizability issues are dynamic evolvements of $\mathcal{L}$-theories like in case of belief revision (cf. [Alchourrón *et al.*, 1985; Williams and Antoniou, 1998; Qi and Yang, 2008; Delgrande and Peppas, 2015; Delgrande *et al.*, 2008; Delgrande *et al.*, 2013; Baumann and Brewka, 2015a; Diller *et al.*, 2015] for several knowledge representation formalisms). Roughly speaking, belief revision deals with the problem of integrating new pieces of information to a current knowledge base which is represented by a certain $\mathcal{L}$-theory T. To this end, you are typically faced with the problem of modifying the given theory T in such a way that the revised version S satisfies $\sigma_{\mathcal{L}}(S) = M$ for some model set M. Now, before trying to do this revision in a certain minimal way it is essential to know whether M is realizable at all, i.e. $M \in \mathcal{R}_{\mathcal{L}}$.

The first formal treatment of realizability issues w.r.t. extension-based argumentation semantics was recently given by Dunne et al. [Dunne *et al.*, 2013; Dunne *et al.*, 2015]. They coined the term *signature* for the set of all realizable sets of extensions. The authors provided simple criteria for several mature semantics deciding whether a set of extensions is contained in the corresponding signature. For instance, two obvious necessary conditions in case of preferred semantics (as well as many other semantics) is that a candidate set $\mathbb{S}$ has to be non-empty, due to universal definedness of preferred semantics and second, $\mathbb{S}$ has to be a $\subseteq$-antichain, also known as *I-maximality criterion* [Baroni and Giacomin, 2007]. However, these conditions are not sufficient implying that further requirements has to hold. In case of preferred semantics it turned out that adding the requirement of so-called *conflict-sensitivity* indeed yield a set of characterizing properties. A $\subseteq$-antichain $\mathbb{S}$ is conflict-sensitive if for each pair of distinct sets A and B from $\mathbb{S}$ there are at least one $a \in A$ and one $b \in B$, s.t. a and b do not occur together in any set of $\mathbb{S}$. This implies that there exists an AF F in which the set of its preferred extension coincides with $\mathbb{S} = \{\{a,b\},\{a,c\},\{b,d\},\{c,d\}\}$. Furthermore, since $\{a,b\}$ and $\{b,d\}$ are already contained in $\mathbb{S}$ it is impossible to realize the set $\mathbb{T} = \mathbb{S} \cup \{\{a,d\}\}$ under preferred semantics. From a practical point of view, such realizability insights can be used to limit the search space when enumerating preferred extensions. More precisely, applying the mentioned characterization result we obtain that not only $\{a,d\}$, but also any other set $A \subseteq \{a,b,c,d\}$ can not be a further preferred extension of a certain AF given that we already computed all sets contained in $\mathbb{S}$. As a matter of fact, knowing that a certain set is realizable does not provide one automatically with a witnessing AF. Fortunately, there exist *canonical frameworks* showing realizability in a constructive fashion as shown in [Dunne *et al.*, 2013; Dunne *et al.*, 2015].

Later on, restricted versions of realizability were considered, namely *compact* as well as *analytic realizability* in case of extension-based semantics [Baumann *et al.*, 2014a; Baumann *et al.*, 2014b; Linsbichler *et al.*, 2015; Baumann *et al.*, 2016a]. Both versions are motivated by typical phenomena that can be observed for several semantics. First, there potentially exist arguments in a given AF that do not appear in any extension, so-called *rejected* arguments. Second, most of the argumentation semantics possess the feature of allowing *implicit conflicts*. An implicit conflict arises when two arguments are never jointly accepted although they do not attack each other. In order to understand in which way rejected arguments and implicit conflicts contribute to the expressive power of a certain semantics the notions of compact AFs as well as analytic AFs were introduced. The former kind disallows rejected arguments whereas the latter is free of implicit conflicts. It turned out that for many universally defined semantics the full range of expressiveness indeed relies on the use rejected arguments and implicit conflicts. This means, there are plenty of AFs which do not possess an equivalent AF which is in addition compact or analytic, respectively.

Recently, a first study of extension-based realizability w.r.t. arbitrary frameworks was presented in [Baumann and Spanring, 2017]. The authors compared the expressive power of several mature semantics in the unrestricted setting. Interestingly, the results reveal an intimate connection between arbitrary and finitely compact AFs in terms of expressiveness. Nevertheless, an in-depth analysis of realizability in the unrestricted setting is still missing. For instance, necessary and sufficient properties for being realizable are not considered so far.

There are only few works which have dealt with labelling-based realizability in the context of Dung-style argumentation frameworks. Dyrkolbotn showed that, as long as additional arguments are allowed any finite set of labellings is *realizable under projection* in case of preferred or semi-stable semantics [Dyrkolbotn, 2014]. In order to realize a set of labellings $\mathbb{S}$ under projection it suffices to come up with an AF F, s.t. its set of labellings modulo additional arguments coincide with $\mathbb{S}$. The second work by Linsbichler et al. deals with the standard notion of realizability adapted to labelling-based semantics [Linsbichler *et al.*, 2016]. The authors presented an algorithm which returns either "No" in case of non-realizability or a witnessing AF F in the positive case. Remarkably, the algorithm is not restricted to the formalism of abstract argumentation frameworks only. In fact, it can also be used to decide realizability in case of the more general abstract dialectical frameworks as well as various of its sub-classes [Brewka and Woltran, 2010; Brewka *et al.*, 2013].

3.1 Realizability and Signatures

Let us start with the two central concepts of this section, namely *realizability* as well as *signature*. In a nutshell, we say that a certain set $\mathbb{S}$ is realizable under the semantics σ, if there is an AF F such that its set of σ-labellings/σ-extensions coincides with $\mathbb{S}$. Collecting all realizable sets defines the concept of a signature. In accordance with the existing literature the main part of this section is devoted to finite realizability for extension-based semantics, i.e. signatures which contain set of σ-extensions of finite AFs only. Realizability w.r.t. labelling-based semantics as well as the consideration of infinite AFs will be briefly outlined only. Consider the following general definition of realizability in the context of abstract argumentation.

Definition 3.1 *Given a semantics $\sigma : \mathcal{F} \to 2^{\left(2^{\mathcal{U}}\right)^n}$ and a set $\mathcal{C} \subseteq \mathcal{F}$. A set $\mathbb{S} \subseteq \left(2^{\mathcal{U}}\right)^n$ is σ-realizable w.r.t. $\mathcal{C}$ if there is an AF $F \in \mathcal{C}$, s.t. $\sigma(F) = \mathbb{S}$.*

Definition 3.2 *Given a semantics σ and a set $\mathcal{C} \subseteq \mathcal{F}$. The σ-signature w.r.t. $\mathcal{C}$ is defined as $\Sigma_\sigma^{\mathcal{C}} = \{\sigma(F) \mid F \in \mathcal{C}\}$.*

If clear from context or unimportant we simply speak of *signatures* and write Σ without mentioning a semantics σ or set of AFs $\mathcal{C}$. Similarly, we say that a certain set is *realizable* instead of σ-realizable w.r.t. $\mathcal{C}$. Please observe that both concepts are intimately connected via the following relation: for any set $\mathbb{S}$ we have, $\mathbb{S}$ is realizable if and only if $\mathbb{S} \in \Sigma$. Consequently, if $\mathbb{S}$ is not contained in Σ, then there is no framework whose extensions/labellings are exactly $\mathbb{S}$. Hence, instead of searching for witnessing AFs (which might not exist) it is very attractive to find necessary as well as sufficient properties for the containment of a set $\mathbb{S}$ to a certain signature locally, i.e. by properties of $\mathbb{S}$ itself.

3.2 Signatures w.r.t. Finite AFs

We start with finite realizability. Instantiating Definitions 3.1 and 3.2 with $\mathcal{C} = \{F \in \mathcal{F} \mid F \text{ finite}\}$ formally capture the notions of realizability as well as signatures relativised to finite AFs. Consider the following definitions.

Definition 3.3 *Given a semantics* $\sigma : \mathcal{F} \to 2^{\left(2^{\mathcal{U}}\right)^n}$. *A set* $\mathbb{S} \subseteq \left(2^{\mathcal{U}}\right)^n$ *is finitely* σ*-realizable if there is an AF* $F \in \{F \in \mathcal{F} \mid F \text{ finite}\}$*, s.t.* $\sigma(F) = \mathbb{S}$.

Definition 3.4 *Given a semantics* σ*. The finite* σ*-signature is defined as* $\{\sigma(F) \mid F \in \mathcal{F}, F \text{ finite}\}$ *abbreviated by* Σ^f_σ.

We proceed with further notational shorthands (adjusted to the extension-based approach) which will be used throughout the whole section.

Definition 3.5 ([Dunne *et al.*, 2015]) *Given* $\mathbb{S} \subseteq 2^{\mathcal{U}}$*, we use*

- $Args_{\mathbb{S}}$ *to denote* $\bigcup_{S \in \mathbb{S}} S$ *and* $\|\mathbb{S}\|$ *for* $|Args_{\mathbb{S}}|$,
- $Pairs_{\mathbb{S}}$ *to denote* $\{(a,b) \mid \exists S \in \mathbb{S} : \{a,b\} \subseteq S\}$ *and*
- $dcl(\mathbb{S})$ *to denote (the so-called downward-closure)* $\{S' \subseteq S \mid S \in \mathbb{S}\}$

Furthermore, we say that $\mathbb{S}$ *is an extension-set if* $\|\mathbb{S}\|$ *is a finite cardinal.*

In order to familiarize the reader with the introduced definitions we give the following example.

Example 3.6 *Let* $\mathbb{S} = \{\{a\}, \{a,c\}, \{a,b,d\}\}$*. Then*

- $Args_{\mathbb{S}} = \{a,b,c,d\}$ *and* $\|\mathbb{S}\| = 4$*. This means,* $\mathbb{S}$ *is an extension-set.*
- $Pairs_{\mathbb{S}} = \{(a,a),(b,b),(c,c),(d,d),(a,b),(a,c),(a,d),(b,d)\} \cup \{(b,a),(c,a),(d,a),(d,b)\}$
- $dcl(\mathbb{S}) = \{\emptyset, \{a\}, \{b\}, \{c\}, \{d\}, \{a,b\}, \{a,c\}, \{a,d\}, \{b,d\}, \{a,b,d\}\}$

Furthermore, since naive extensions are defined as $\subseteq$*-maximal sets and obviously,* $\{a\} \subset \{a,c\}$ *we deduce that* $\mathbb{S}$ *is not na-realizable, i.e.* $\mathbb{S} \notin \Sigma^f_{\mathcal{E}_{na}}$*. Regarding complete semantics we obtain* $\mathbb{S} \in \Sigma^f_{\mathcal{E}_{co}}$ *witnessed by the following AF* F.

F : 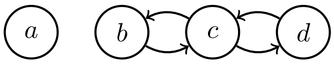

In the following we consider the signatures of the extension-based versions of stable, semi-stable, stage, naive, preferred, complete as well as grounded semantics [Dunne *et al.*, 2013; Dunne *et al.*, 2015]. We provide a bunch of properties where certain subsets of them exactly matches the containment conditions for certain signatures. All properties can be decided by looking on the set in question only.

Semantics based on Conflict-freeness

Our starting point are semantics based on conflict-free sets. Conflict-free sets by themselves inherited their conflict-freeness to any subset of them. More formally, the downward-closure does not vary the set of conflict-free sets for a given AF. A set possessing this property is called *downward-closed.* Clearly, downward-closedness does not hold in case of admissible sets as well as any other reasonable semantics σ where conflict-freeness is just one requirement among others for being a σ-extension. Take for instance naive semantics. Naive extension are defined as $\subseteq$-maximal conflict-free sets. Consequently, the set of all naive extensions is a $\subseteq$-antichain, i.e. any two naive extensions are *incomparable* w.r.t. subset relation. This property also applies to many other semantics, such as stable and stage semantics as well as any uniquely defined semantics. However, although incomparability is a necessary condition for many considered semantics it is certainly not sufficient. Consider therefore the following example taken from [Dunne *et al.*, 2015, Example 1].

Example 3.7 *Consider the $\subseteq$-antichain $\mathbb{S} = \{\{a,b\},\{a,c\},\{b,c\}\}$ and a semantics σ which selects its reasonable positions among the conflict-free sets, i.e. $\mathcal{E}_\sigma(F) \subseteq \mathcal{E}_{cf}(F)$ for any AF F. Now suppose there exists an AF F with $\mathcal{E}_\sigma(F) = \mathbb{S}$. Then F must not contain attacks between a and b, a and c, and respectively b and c. This means, $\{a,b,c\} \in \mathcal{E}_{cf}(F)$. But then $\mathcal{E}_\sigma(F)$ typically contains $\{a,b,c\}$.*

There are several ways to define the required property which excludes sets like $\mathbb{S}$ from above. It turned out that in order to characterize conflict-free based semantics like stable, stage and naive semantics a rather strong condition is required, so-called *tightness.* Roughly speaking, if an incomparable set is not tight, then there is a set $S \in \mathbb{S}$ and an argument a not belonging to S, s.t. for any $s \in S$ we find an other $S' \in \mathbb{S}$ with a and s being members of it. The idea behind the notion of being tight is simply that if an argument a does not occur in some extension S there must be a reason for that. The most simple reason one can think of is that there is a conflict between a and some $s \in S$, i.e. a and s do not occur jointly in any extension-set of $\mathbb{S}$ or, in other words, $(a,s) \notin \mathit{Pairs}_{\mathbb{S}}$. In a way, this limits the multitude of incomparable elements of an extension-set.

We proceed with the formal definitions.

Definition 3.8 ([Dunne *et al.*, 2013]) *Given* $\mathbb{S} \subseteq 2^{\mathcal{U}}$. *We call* $\mathbb{S}$

- *downward-closed if* $\mathbb{S} = dcl(\mathbb{S})$,
- *incomparable if* $\mathbb{S}$ *is a* $\subseteq$*-antichain and*
- *tight if for all* $S \in \mathbb{S}$ *and* $a \in Args_{\mathbb{S}}$ *it holds that if* $S \cup \{a\} \notin \mathbb{S}$ *then there exists an* $s \in S$ *such that* $(a, s) \notin Pairs_{\mathbb{S}}$.

Please observe that for incomparable $\mathbb{S}$, the premise of the tightness condition, i.e. $S \cup \{a\} \notin \mathbb{S}$, is always fulfilled. However, tightness and incomparability are independent of each other, i.e. neither tightness implies incomparability or comparability, nor incomparability implies tightness or non-tightness.

Example 3.9 *Consider again the extension-set* $\mathbb{S} = \{\{a, b\}, \{a, c\}, \{b, c\}\}$ *from Example 3.7. The set* $\mathbb{S}$ *is incomparable but not tight which can be seen as follows. If setting* $S = \{a, b\}$ *we observe* $S \cup \{c\} \notin \mathbb{S}$. *Moreover, for any* $s \in S$ *we find an* $S' \in \mathbb{S}$, *s.t.* $\{s, c\} = S'$ *implying that* $(s, c) \in Pairs_{\mathbb{S}}$. *More precisely, if* $s = a$, *then we have* $S' = \{a, c\}$ *and similarly, if* $s = b$ *we find* $S' = \{b, c\}$.

Furthermore, it can be checked that $\mathbb{S}' = \{\{a, b\}, \{a, c\}, \{b, d\}, \{c, d\}\}$ *or* $\mathbb{S}'' = \mathbb{S} \cup \{\{a, b, c\}\}$ *are witnessing examples for incomparability and tightness or tightness and comparability, respectively.*

Clearly, subsets of incomparable sets are incomparable. Such a kind of inheritance does not hold in case of tight sets (cf. $\mathbb{S}$ and $\mathbb{S}''$ as defined in Example 3.9). Nevertheless, there are non-trivial tight subsets of any tight set. For instance, in any case the set of all $\subseteq$-maximal elements is tight. Furthermore, if a tight set is even incomparable, then any subset of it is tight too.

In the following we present the main statements only. However, in many cases we provide some short comments indicating how to prove the statement in question. For full proofs we refer the reader to the referenced papers.

Lemma 3.10 ([Dunne *et al.*, 2015]) *For a tight extension-set* $\mathbb{S} \subseteq 2^{\mathcal{U}}$ *we have:*

1. *the* $\subseteq$*-maximal elements in* $\mathbb{S}$ *form a tight set, and*
2. *if* $\mathbb{S}$ *is incomparable then each* $\mathbb{S}' \subseteq \mathbb{S}$ *is tight.*

Note that the second statement of Lemma 3.10 implies that if the downward-closure of an incomparable extension-set $\mathbb{S}$ is tight, then $\mathbb{S}$ itself has to be tight too.

We proceed with a specific AF and check which properties apply to its different sets of extensions.

Example 3.11 *Consider the following AF* F.

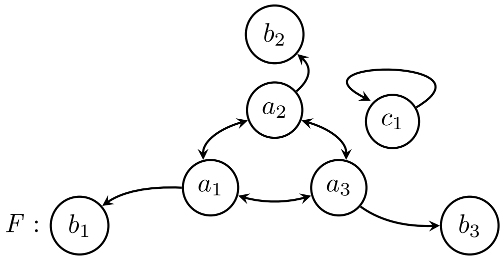

Since c_1 is self-defeating as well as unattacked we obtain $\mathcal{E}_{stb}(F) = \emptyset$. Furthermore, $\mathcal{E}_{stg}(F) = \{\{a_1, b_2, b_3\}, \{a_2, b_1, b_3\}, \{a_3, b_1, b_2\}\}$ and $\mathcal{E}_{na}(F) = \mathcal{E}_{stg}(F) \cup \{\{b_1, b_2, b_3\}\}$. We observe,

1. *$\mathcal{E}_{stb}(F), \mathcal{E}_{stg}(F)$ as well as $\mathcal{E}_{na}(F)$ are incomparable,*
2. *$\mathcal{E}_{stb}(F), \mathcal{E}_{stg}(F)$ as well as $\mathcal{E}_{na}(F)$ are tight and additionally,*
3. *$dcl(\mathcal{E}_{na}(F))$ and $dcl(\mathcal{E}_{stb}(F))$ are tight and obviously,*
4. *$\mathcal{E}_{stg}(F)$ and $\mathcal{E}_{na}(F)$ are non-empty.*

The first and the last items are not surprising since firstly, all considered semantics satisfy the I-maximality criterion which is just another name for incomparability and secondly, in Section 2 we have already seen that stage extensions are guaranteed for finitary (hence, for finite) frameworks and naive semantics is even universally defined w.r.t. the whole class of AFs. This means, incomparability or non-emptiness of the mentioned sets of σ-extensions do not depend on the specific AF F, but rather apply to any finite AF. Consequently, these properties represent necessary properties regarding realizability. The tightness statements of the second and third items can be checked in a straightforward manner. We now examine that $dcl(\mathcal{E}_{stg}(F))$ is non-tight. This can be seen as follows: Firstly, $\{b_2, b_3\} \in dcl(\mathcal{E}_{stg}(F))$. Now, for b_1 the premise of Definition 3.8 is satisfied, i.e. $\{b_1, b_2, b_3\} \notin dcl(\mathcal{E}_{stg}(F))$. Consequently, since $\{b_1, b_2\}, \{b_1, b_3\} \in dcl(\mathcal{E}_{stg}(F))$ and therefore, $(b_1, b_2), (b_1, b_3) \in Pairs_{dcl(\mathcal{E}_{stg}(F))}$ we deduce the non-tightness of $dcl(\mathcal{E}_{stg}(F))$. This means, tightness of the downward-closure of a given set can not be a necessary criterion for belonging to the stage signature.

We now present the characterization theorems for conflict-free, naive, stable as well as stage signatures. It is somehow surprising that only a few simple properties are sufficient to characterize these different signatures.

Theorem 3.12 ([Dunne *et al.*, 2015]) *Given a set $\mathbb{S} \subseteq 2^{\mathcal{U}}$, then*

1. *$\mathbb{S} \in \Sigma^f_{\mathcal{E}_{cf}} \Leftrightarrow \mathbb{S}$ is a non-empty, downward-closed, and tight extension-set,*
2. *$\mathbb{S} \in \Sigma^f_{\mathcal{E}_{na}} \Leftrightarrow \mathbb{S}$ is a non-empty, incomparable extension-set and $dcl(\mathbb{S})$ is tight,*
3. *$\mathbb{S} \in \Sigma^f_{\mathcal{E}_{stb}} \Leftrightarrow \mathbb{S}$ is a incomparable and tight extension-set,*

4. $\mathbb{S} \in \Sigma^f_{\mathcal{E}_{stg}} \Leftrightarrow \mathbb{S}$ *is a non-empty, incomparable and tight extension-set.*

We mention that a proof of the characterization theorem above requires two directions. Let us fix a certain semantics $\sigma \in \{cf, na, stb, stg\}$. The first part is to show that for any finite AF F, $\mathcal{E}_\sigma(F)$ satisfies the mentioned properties. Now, for the second part, if a certain extension-set $\mathbb{S}$ satisfies the properties in question, then we have to find a finite AF F, s.t. $\mathcal{E}_\sigma(F) = \mathbb{S}$.

Let us start with the first part. It suffices to consider tightness only since downward-closedness, non-emptiness and incomparability are clear (cf. some explanations given in Example 3.11). It is easy to see that $\mathcal{E}_{cf}(F)$ is tight because if augmenting a conflict-free set S with a non-conflicting argument a yields a conflicting set, then obviously there has to be at least one element in $s \in S$, s.t. $\{a, s\}$ is conflicting. In order to prove that $dcl(\mathcal{E}_{na}(F))$ is tight, it suffice to see that $dcl(\mathcal{E}_{na}(F)) = \mathcal{E}_{cf}(F)$. Consequently, applying Lemma 3.10 we obtain the tightness of $\mathcal{E}_{na}(F)$. Furthermore, with the same lemma, we get that every $\mathbb{S} \subseteq \mathcal{E}_{na}(F)$ is tight. In consideration of $stb \subseteq stg \subseteq na$ (Proposition 2.6) it follows that $\mathcal{E}_{stb}(F)$ as well as $\mathcal{E}_{stg}(F)$ are tight.

In order to show that the mentioned properties are not only necessary but even sufficient we have to come up with witnessing AFs. Consider therefore the following prototype.

Definition 3.13 ([Dunne *et al.*, 2015]) *Given an extension-set* $\mathbb{S}$*, we define the canonical argumentation framework for* $\mathbb{S}$ *as*

$$F^{cf}_{\mathbb{S}} = \big(Args_{\mathbb{S}}, (Args_{\mathbb{S}} \times Args_{\mathbb{S}}) \setminus Pairs_{\mathbb{S}}\big).$$

The idea behind the framework is simple: we draw a relation between two arguments iff they do not occur jointly in any set $S \in \mathbb{S}$. Consequently, for any $\mathbb{S}$, $F^{cf}_{\mathbb{S}}$ is symmetric. Moreover, in any case, it is self-loop-free since $a \in Args_{\mathbb{S}}$ implies $(a, a) \in Pairs_{\mathbb{S}}$. Let us consider the following example.

Example 3.14 *Let* $\mathbb{S} = \{\{a_1, b_2, b_3\}, \{a_2, b_1, b_3\}, \{a_3, b_1, b_2\}, \{b_1, b_2, b_3\}\}$ *and consider the corresponding canonical framework* $F^{cf}_{\mathbb{S}}$*.*

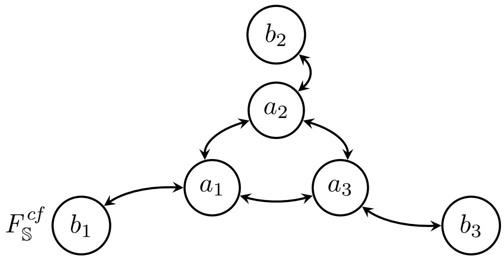

Please note that $\mathbb{S}$ *is non-empty, incomparable as well as possesses a tight downward-closure (cf. Example 3.11). Furthermore,* $F^{cf}_{\mathbb{S}}$ *realizes* $\mathbb{S}$ *under the naive semantics, i.e.* $\mathcal{E}_{na}\left(F^{cf}_{\mathbb{S}}\right) = \mathbb{S}$.

The following proposition shows that this is no coincidence.

Proposition 3.15 ([Dunne *et al.*, 2015]) *For each non-empty, incomparable extension-set* $\mathbb{S}$*, where* $dcl(\mathbb{S})$ *is tight,* $\mathcal{E}_{na}\left(F_{\mathbb{S}}^{cf}\right) = \mathbb{S}$.

Moreover, the canonical framework can also be used as witnessing framework in case of conflict-free sets as stated in the following proposition.

Proposition 3.16 ([Dunne *et al.*, 2015]) *For each non-empty, downward-closed and tight extension-set* $\mathbb{S}$*,* $\mathcal{E}_{cf}\left(F_{\mathbb{S}}^{cf}\right) = \mathbb{S}$.

We proceed with stable and stage semantics. In Theorem 3.12 the only difference between the characterizations of stable and stage signatures is the non-empty requirement for stage semantics. Remember that we are dealing with finite AFs and indeed in case of this restriction stable semantics is the only semantics which does not warrant the existence of extensions (cf. Table 1).[8] This means, stable semantics is the only semantics which may realize the empty extension-set (which is incomparable and tight too). The final step towards concluding Theorem 3.12 is to find witnessing frameworks for any non-empty, incomparable and tight extension-sets. At first we will show that the canonical framework does not do the job in case of these semantics. More precisely, given a non-empty, incomparable as well as tight extension-set $\mathbb{S}$, then the sets of stable as well as stage extensions of the canonical framework $F_{\mathbb{S}}^{cf}$ do not necessarily coincide with $\mathbb{S}$.

Example 3.17 *Consider again Example 3.14. We define* $\mathbb{T} = \mathbb{S} \setminus \{\{b_1, b_2, b_3\}\}$. *Please note that* $F_{\mathbb{T}}^{cf}$ *and* $F_{\mathbb{S}}^{cf}$ *are identical since* $Args_{\mathbb{S}} = Args_{\mathbb{T}}$ *and* $Pairs_{\mathbb{S}} = Pairs_{\mathbb{T}}$. *Furthermore, according to Example 3.11 we have that* $\mathbb{T}$ *is non-empty, incomparable and tight, but* $\mathcal{E}_{stb}\left(F_{\mathbb{T}}^{cf}\right) = \mathcal{E}_{stg}\left(F_{\mathbb{T}}^{cf}\right) = \mathcal{E}_{na}\left(F_{\mathbb{S}}^{cf}\right) = \mathbb{S} \neq \mathbb{T}$. *In order to get rid of the undesired stable as well as stage extension* $E = \{b_1, b_2, b_3\}$ *we may simply add a new self-defeating argument* e *to* $F_{\mathbb{S}}^{cf}$*, s.t.* e *is attacked by all other arguments excepting those stemming from* E*. The following framework* $F_{\mathbb{T}}^{stb}$ *illustrates this idea. Convince yourself that* $\mathcal{E}_{stb}\left(F_{\mathbb{T}}^{stb}\right) = \mathcal{E}_{stg}\left(F_{\mathbb{T}}^{stb}\right) = \mathbb{T}$.

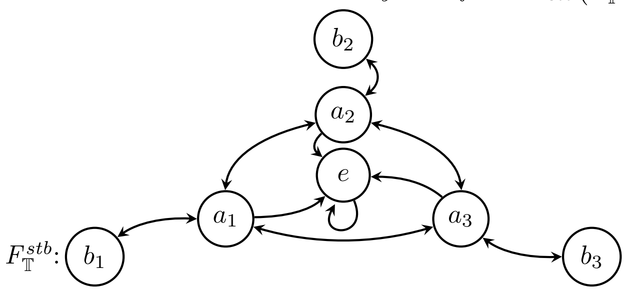

The following definition generalizes the construction idea from above to arbitrary many undesired sets. The subsequent proposition states that we

[8] For instance, $F = (\{a\}, \{(a, a)\})$ yields $\mathcal{E}_{stb}(F) = \emptyset$.

have indeed found witnessing examples for non-empty, incomparable and tight extension-sets as required for Theorem 3.12.

Definition 3.18 ([Dunne *et al.*, 2015]) *Given an extension-set* $\mathbb{S}$ *and its canonical framework* $F^{cf}_{\mathbb{S}} = (A^{cf}_{\mathbb{S}}, R^{cf}_{\mathbb{S}})$*. Let* $\mathbb{X} = \mathcal{E}_{stb}\left(F^{cf}_{\mathbb{S}}\right) \setminus \mathbb{S}$ *we define*
$$F^{stb}_{\mathbb{S}} = \left(A^{cf}_{\mathbb{S}} \cup \{\bar{E} \mid E \in \mathbb{X}\}, R^{cf}_{\mathbb{S}} \cup \{(\bar{E},\bar{E}), (a,\bar{E}) \mid E \in \mathbb{X}, a \in Args_{\mathbb{S}} \setminus E\}\right).$$

Proposition 3.19 ([Dunne *et al.*, 2015]) *For each non-empty, incomparable and tight extension-set* $\mathbb{S}$*,* $\mathcal{E}_{stb}\left(F^{stb}_{\mathbb{S}}\right) = \mathcal{E}_{stg}\left(F^{stb}_{\mathbb{S}}\right) = \mathbb{S}$*.*

Semantics based on Admissibility

Let us turn now to semantics based on admissible sets. In particular, we provide characterization theorems for the finite signatures w.r.t. admissible sets as well as preferred and semi-stable semantics. In contrast to semantics based on conflict-free sets where the notion of tightness played a decisive role (cf. Theorem 3.12) we have to introduce a new concept, so-called *conflict-sensitivity.* Conflict-sensitivity is a very basic property in the sense that it is fulfilled by almost all semantics σ (or rather, their corresponding sets of σ-extensions) available in the literature. Furthermore, it is strictly weaker than tightness, i.e. tight extension-sets are always conflict-sensitive, but not necessarily vice versa. To explain the difference between these two notions let us consider the following example taken from [Dunne *et al.*, 2015].

Example 3.20 *Consider the following framework* F*.*

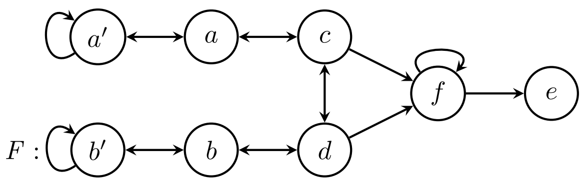

We have $\mathcal{E}_{pr}(F) = \mathcal{E}_{ss}(F) = \mathbb{S} = \{A, B, C\} = \{\{a,b\}, \{a,d,e\}, \{b,c,e\}\}$*. First, observe that* $\mathbb{S}$ *is not tight. This can be seen as follows: Obviously,* $A \cup \{e\} \notin \mathbb{S}$*, but both* (a,e) *and* (b,e) *are contained in* $Pairs_{\mathbb{S}}$ *since* $\{a,e\} \subseteq B$ *and* $\{b,e\} \subseteq C$*. This means, although* $A \cup \{e\}$ *is not a reasonable position w.r.t. preferred and semi-stable semantics we find witnessing extensions, namely* B *and* C*, showing that any argument in* A *is compatible with* e*, i.e. they can be accepted together. Please observe that this is not true for any two arguments in* A *and* B *or* A *and* C*, respectively. For instance,* $b,d \in A \cup B$*, but* $(b,d) \notin Pairs_{\mathbb{S}}$ *as well as* $a,c \in A \cup C$*, but* $(a,c) \notin Pairs_{\mathbb{S}}$*. Furthermore, the same applies to* B *and* C*, since* $c,d \in B \cup C$ *and* $(c,d) \notin Pairs_{\mathbb{S}}$*.*

The following definition precisely formalizes the observed property of the AF F presented in the example above.

Definition 3.21 ([Dunne *et al.*, 2015]) *A set $\mathbb{S} \subseteq 2^{\mathcal{U}}$ is called conflict-sensitive if for each $A, B \in \mathbb{S}$ such that $A \cup B \notin \mathbb{S}$ it holds that $\exists a, b \in A \cup B : (a, b) \notin Pairs_{\mathbb{S}}$.*

As the name suggests, the property checks whether the absence of the union of any pair of extensions in an extension-set $\mathbb{S}$ is justified by a conflict indicated by $\mathbb{S}$. Note that for $a, b \in A$ (likewise $a, b \in B$), $(a, b) \in Pairs_{\mathbb{S}}$ holds by definition. Thus the property of conflict-sensitivity is determined by arguments $a \in A \setminus B$, $b \in B \setminus A$, for $A, B \in \mathbb{S}$. As already indicated tightness implies conflict-sensitivity as stated in the following lemma.

Lemma 3.22 ([Dunne *et al.*, 2015]) *Every tight extension-set is also conflict-sensitive.*

Similarly to Lemma 3.10 one may show that the set of all $\subseteq$-maximal elements of a conflict-sensitive set is conflict-sensitive too. Moreover, if the initial set is incomparable in addition, then even any subset of it is conflict-sensitive. Furthermore, in contrast to tight extension-sets it is possible to add the empty set to a conflict-sensitive set without loosing conflict-sensitivity.[9]

Lemma 3.23 ([Dunne *et al.*, 2015]) *For a conflict-sensitive ext.-set $\mathbb{S} \subseteq 2^{\mathcal{U}}$,*

1. *the $\subseteq$-maximal elements in $\mathbb{S}$ form a conflict-sensitive set,*
2. *if $\mathbb{S}$ is incomparable then each $\mathbb{S}' \subseteq \mathbb{S}$ is conflict-sensitive, and*
3. *$\mathbb{S} \cup \{\emptyset\}$ is conflict-sensitive.*

Having conflict-sensitivity at hand, we are now ready to present characterization theorems for the signatures w.r.t. admissible sets as well as preferred and semi-stable semantics. Interestingly, it turns out that preferred and semi-stable semantics are equally expressive in case of finite AFs, i.e. $\Sigma^f_{\mathcal{E}_{pr}} = \Sigma^f_{\mathcal{E}_{ss}}$.

Theorem 3.24 ([Dunne *et al.*, 2015]) *Given a set $\mathbb{S} \subseteq 2^{\mathcal{U}}$, then*

1. $\mathbb{S} \in \Sigma^f_{\mathcal{E}_{ad}} \Leftrightarrow \mathbb{S}$ *is a conflict-sensitive ext.-set containing $\emptyset$,*
2. $\mathbb{S} \in \Sigma^f_{\mathcal{E}_{pr}} \Leftrightarrow \mathbb{S}$ *is a non-empty, incomparable and conflict-sensitive ext.-set,*
3. $\mathbb{S} \in \Sigma^f_{\mathcal{E}_{ss}} \Leftrightarrow \mathbb{S}$ *is a non-empty, incomparable and conflict-sensitive ext.-set.*

Let us first argue that the mentioned properties are necessary conditions for being in the corresponding signature. For admissible sets it suffices to recall the following two facts: First, the empty set is admissible by definition; and second, if the union of two admissible sets is conflict-free, then the union is admissible too. In other words, if the union fails to be admissible, then there

[9] Note that any one-element extension-set $\mathbb{S} \neq \{\emptyset\}$ is tight, whereas $\mathbb{S} \cup \{\emptyset\}$ is not.

has to be a conflict proving the conflict-sensitivity of any set of admissible sets. Now, for preferred and semi-stable semantics. Non-emptiness is due to the already shown universal definedness of both semantics in case of finite AFs (cf. Table 1). Moreover, incomparability is clear since both semantics satisfy the I-maximality criterion (cf. Chapter 17). Finally, conflict-sensitivity of sets of admissible sets transfer to sets of preferred extensions via statement 1 of Lemma 3.23 and therefore also to sets of semi-stable extensions via statement 2 of Lemma 3.23 and the fact that $ss \subseteq pr$ (Proposition 2.6).

In order to show that the mentioned properties are not only necessary but even sufficient we have to come up with witnessing AFs. In contrast to conflict-free based semantics we have to find AFs which encode the central notion of admissibility. Please note that the already introduced canonical frameworks $F^{cf}_{\mathbb{S}}$ as well as $F^{stb}_{\mathbb{S}}$ (cf. Definitions 3.13 and 3.18) do not comply with the requirements. Consider therefore the following example.

Example 3.25 *Let us consider again the non-empty, incomparable as well as tight set* $\mathbb{T} = \{\{a_1, b_2, b_3\}, \{a_2, b_1, b_3\}, \{a_3, b_1, b_2\}\}$ *together with its corresponding canonical framework* $F^{stb}_{\mathbb{T}}$ *as presented in Example 3.17. Due to Lemma 3.22 we have that any tight extension-set is even conflict-sensitive and thus,* $\mathbb{T}$ *satisfies the necessary requirements of Theorem 3.24. Inspecting the canonical framework reveals that* $\mathcal{E}_{pr}\left(F^{stb}_{\mathbb{T}}\right) = \mathbb{T} \cup \{\{b_1, b_2, b_3\}\} \neq \mathbb{T}$*. Although,* $\mathcal{E}_{ss}\left(F^{stb}_{\mathbb{T}}\right) = \mathbb{T}$ *one may easily check that non-empty, incomparable as well as conflict-sensitive set* $\mathbb{S} = \{\{a, b\}, \{a, d, e\}, \{b, c, e\}\}$ *mentioned in Example 3.20 shows that this equality does not hold in general. Likewise, one may prove that the framework* $F^{cf}_{\mathbb{S}}$ *is not appropriated as a witnessing prototype for semi-stable as well as preferred semantics.*

It turned out that suitable canonical AFs can be built by means of so-called *defense-formulae* as introduced in the following definition.

Definition 3.26 ([Dunne *et al.*, 2015]) *Given an extension-set* $\mathbb{S}$*, the defense-formula* $\mathcal{D}^{\mathbb{S}}_a$ *of an argument* $a \in Args_{\mathbb{S}}$ *in* $\mathbb{S}$ *is defined as:*

$$\mathcal{D}^{\mathbb{S}}_a = \bigvee_{S \in \mathbb{S}, s.t. a \in S} \ \bigwedge_{s \in S \setminus \{a\}} s.$$

$\mathcal{D}^{\mathbb{S}}_a$ *given as (a logically equivalent) CNF is called CNF-defense-formula* $\mathcal{CD}^{\mathbb{S}}_a$ *of* a *in* $\mathbb{S}$*.*

The main idea of the formula $\mathcal{D}^{\mathbb{S}}_a$ is to describe the conditions for the argument a being in an extension. Note that the variables coincide with the arguments. If $\mathbb{S}$ amounts to a set of admissible extensions, then each disjunct represents a set of arguments A which allows a to join in the sense that $A \cup \{a\}$ is a reasonable position w.r.t. admissible semantics. Put it differently, propositional models of $\mathcal{D}^{\mathbb{S}}_a \wedge a$ represent (if considered as set of atoms) supersets of certain reasonable position. Please not that a defense-formula $\mathcal{D}^{\mathbb{S}}_a$ is tautological if and only if $\{a\} \in \mathbb{S}$. We proceed with an example.

Example 3.27 *Consider again the non-empty, incomparable as well as conflict-sensitive set* $\mathbb{S} = \{\{a,b\},\{a,d,e\},\{b,c,e\}\}$ *stemming from Example 3.20. We obtain the following defense-formulae together with their corresponding CNF-defense-formulae (written in clause form).*

- $\mathcal{D}_a^{\mathbb{S}} = b \vee (d \wedge e) \equiv (b \vee d) \wedge (b \vee e)$ *and* $\mathcal{CD}_a^{\mathbb{S}} = \{\{b,d\},\{b,e\}\}$
- $\mathcal{D}_b^{\mathbb{S}} = a \vee (c \wedge e) \equiv (a \vee c) \wedge (a \vee e)$ *and* $\mathcal{CD}_b^{\mathbb{S}} = \{\{a,c\},\{a,e\}\}$
- $\mathcal{D}_c^{\mathbb{S}} = b \wedge e$ *and* $\mathcal{CD}_c^{\mathbb{S}} = \{\{b,e\}\}$
- $\mathcal{D}_d^{\mathbb{S}} = a \wedge e$ *and* $\mathcal{CD}_d^{\mathbb{S}} = \{\{a,e\}\}$
- $\mathcal{D}_e^{\mathbb{S}} = (a \wedge d) \vee (b \wedge c) \equiv (a \vee b) \wedge (d \vee b) \wedge (a \vee c) \wedge (d \vee c)$ *and* $\mathcal{CD}_d^{\mathbb{S}} = \{\{a,b\},\{a,c\},\{b,d\},\{c,d\}\}$

One simple idea for the realization of a certain set $\mathbb{S}$ under admissible semantics is the following two-step procedure. In the first step, we construct a framework F which maintains all elements of $\mathbb{S}$ as conflict-free sets. This can be done via the the canonical framework $F_{\mathbb{S}}^{cf}$. In the second step, we augment the initial framework $F_{\mathbb{S}}^{cf}$, s.t. only elements in $\mathbb{S}$ become admissible. The second step can be realized via adding a certain amount of additional arguments. More precisely, for any argument $a \in Args_{\mathbb{S}}$ we add n self-conflicting arguments $\alpha_{aC_1}, ..., \alpha_{aC_n}$ if $\left|\mathcal{CD}_a^{\mathbb{S}}\right| = |\{C_1,...,C_n\}| = n$. Then, for any $i \in \{1,..,n\}$, α_{aC_i} attacks a and is in turn attacked by any argument in C_i. Consider therefore the following example.

Example 3.28 *Again consider the extension-set* $\mathbb{S} = \{\{a,b\},\{a,d,e\},\{b,c,e\}\}$ *and its corresponding CNF-defense-formulae as presented in Example 3.27. In accordance with the above mentioned two-step procedure we obtain the dashed AF* $F_{\mathbb{S}}^{cf}$ *first. Then, in view of the CNF-defense-formulae we have to add 10 additional self-defeating arguments which attacks their corresponding argument. This intermediate step is depicted below.*

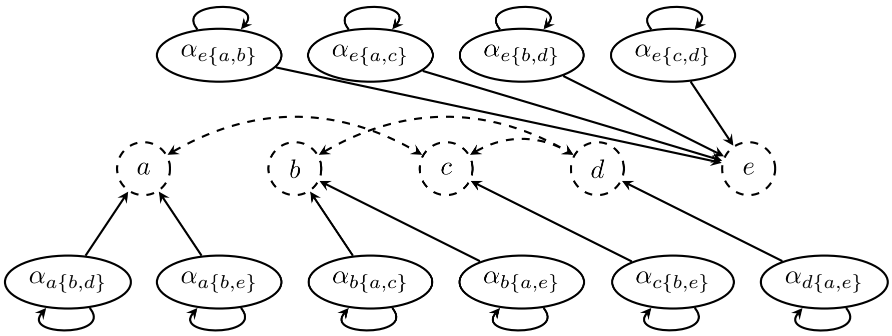

Let us consider the set $\{a,b\} \in \mathbb{S}$. *In order for* $\{a,b\}$ *to be admissible we have to add counter-attacks for the arguments* $\alpha_{a\{b,d\}}$, $\alpha_{a\{b,e\}}$, $\alpha_{b\{a,c\}}$ *and*

$\alpha_{b\{a,e\}}$. For instance, $\alpha_{a\{b,d\}}$ is attacked by b and d and so forth. The following figure (built on top of the previous one) depicts resulting counter-attacks for the mentioned 4 arguments highlighted as densely dotted edges. For the sake of clarity we do not perform this construction for the remaining arguments.

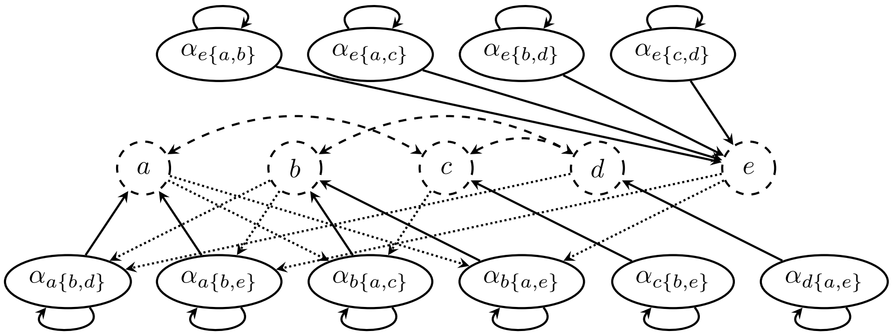

The following definition precisely formalizes the mentioned two-step procedure.

Definition 3.29 ([Dunne *et al.*, 2015]) *Given an extension-set $\mathbb{S}$, the canonical defense-argumentation-framework $F_{\mathbb{S}}^{def} = (A_{\mathbb{S}}^{def}, R_{\mathbb{S}}^{def})$ extends the canonical AF $F_{\mathbb{S}}^{cf} = (Args_{\mathbb{S}}, R_{\mathbb{S}}^{cf})$ as follows:*

$$A_{\mathbb{S}}^{def} = Args_{\mathbb{S}} \cup \bigcup_{a \in Args_{\mathbb{S}}} \left\{ \alpha_{a\gamma} \mid \gamma \in \mathcal{CD}_a^{\mathbb{S}} \right\}, \text{ and}$$

$$R_{\mathbb{S}}^{def} = R_{\mathbb{S}}^{cf} \cup \bigcup_{a \in Args_{\mathbb{S}}} \left\{ (b, \alpha_{a\gamma}), (\alpha_{a\gamma}, \alpha_{a\gamma}), (\alpha_{a\gamma}, a) \mid \gamma \in \mathcal{CD}_a^{\mathbb{S}}, b \in \gamma \right\}.$$

The subsequent proposition shows that not only all elements in $\mathbb{S}$ become admissible in the constructed AF $F_{\mathbb{S}}^{def}$, but rather that the set of admissible sets of $F_{\mathbb{S}}^{def}$ exactly coincides with $\mathbb{S}$ given that $\mathbb{S}$ is conflict-sensitive as well as contains the empty set.

Proposition 3.30 ([Dunne *et al.*, 2015]) *For each conflict-sensitive ext.-set $\mathbb{S}$ where $\emptyset \in \mathbb{S}$, it holds that $\mathcal{E}_{ad}\left(F_{\mathbb{S}}^{def}\right) = \mathbb{S}$.*

Interestingly, we may even use the canonical defense-AF to show that any non-empty, incomparable and conflict-sensitive extension-set $\mathbb{S}$ can be realized under the preferred semantics. This can be seen as follows: First, via Lemma 3.23 we obtain the conflict-sensitivity of $\mathbb{S} \cup \{\emptyset\}$ since $\mathbb{S}$ is assumed to be conflict-sensitive. Consequently, using Proposition 3.31 we obtain $\mathcal{E}_{ad}\left(F_{\mathbb{S}\cup\{\emptyset\}}^{def}\right) = \mathbb{S} \cup \{\emptyset\}$. Since $F_{\mathbb{S}}^{def} = F_{\mathbb{S}\cup\{\emptyset\}}^{def}$ and due to the incomparability of $\mathbb{S}$, we have $\mathcal{E}_{pr}\left(F_{\mathbb{S}}^{def}\right) = \mathbb{S}$ as stated in the following proposition.

Proposition 3.31 ([Dunne *et al.*, 2015]) *For each non-empty, incomparable and conflict-sensitive extension-set $\mathbb{S}$, it holds that $\mathcal{E}_{pr}\left(F_{\mathbb{S}}^{def}\right) = \mathbb{S}$.*

Furthermore, due to a translation result by Dvořák and Woltran we obtain that any non-empty, incomparable and conflict-sensitive extension-set $\mathbb{S}$ can be realized under semi-stable semantics too. More precisely, in [Dvořák and Woltran, 2011] it is shown that for any AF F exists an AF F', s.t. $\mathcal{E}_{pr}(F) = \mathcal{E}_{ss}(F')$.

Proposition 3.32 ([Dunne *et al.*, 2015]) *Each non-empty, incomparable and conflict-sensitive extension-set $\mathbb{S}$ is ss-realizable.*

Uniquely Defined Semantics

Let us finally turn to grounded, ideal and eager semantics. Remember that all mentioned semantics warrants the existence of exactly one extension given that the frameworks in question are finite (cf. Table 1). Furthermore, it is hardly surprising that this property is even sufficient for being in the corresponding signature, since any one-element extension-set $\mathbb{S} = \{E\}$ can be realized via $F_E = (E, \emptyset)$. In particular, we obtain that all three semantics are equally expressive.

Theorem 3.33 ([Dunne *et al.*, 2016]) *Given a set $\mathbb{S} \subseteq 2^{\mathcal{U}}$, then*

1. $\mathbb{S} \in \Sigma_{\mathcal{E}_{gr}}^{f} \Leftrightarrow \mathbb{S}$ *is an extension-set with* $|\mathbb{S}| = 1$,

2. $\mathbb{S} \in \Sigma_{\mathcal{E}_{il}}^{f} \Leftrightarrow \mathbb{S}$ *is an extension-set with* $|\mathbb{S}| = 1$ *and*

3. $\mathbb{S} \in \Sigma_{\mathcal{E}_{eg}}^{f} \Leftrightarrow \mathbb{S}$ *is an extension-set with* $|\mathbb{S}| = 1$.

Summary of Results and Further Remarks

In this subsection we provide a comprehensive overview of characterization results w.r.t. extension-based realizability in case of finite AFs. The following table collect and combine the results of the previous three subsections. The table has to be interpreted as follows: Consider a certain column σ. Then, the entries "×" in rows r_1,...,r_n indicate that for any extension-set $\mathbb{S}$, $\mathbb{S} \in \Sigma_{\mathcal{E}_\sigma}^{f} \Leftrightarrow r_1, ..., r_n$. Moreover, an entry "→" in row r reflects the fact that the collection of the properties $r_1, ..., r_n$ imply property r.

	cf	na	stb	stg	ad	pr	ss	gr	il	eg
$\mathbb{S} \neq \emptyset$	$\times$	$\times$		$\times$	$\rightarrow$	$\times$	$\times$	$\rightarrow$	$\rightarrow$	$\rightarrow$
$\emptyset \in \mathbb{S}$	$\rightarrow$				$\times$					
$\lvert\mathbb{S}\rvert = 1$								$\times$	$\times$	$\times$
$dcl(\mathbb{S})$ is tight		$\times$						$\rightarrow$	$\rightarrow$	$\rightarrow$
$\mathbb{S}$ is incomparable		$\times$	$\times$	$\times$		$\times$	$\times$	$\rightarrow$	$\rightarrow$	$\rightarrow$
$\mathbb{S}$ is tight	$\times$	$\rightarrow$	$\times$	$\times$				$\rightarrow$	$\rightarrow$	$\rightarrow$
$\mathbb{S}$ is conflict-sensitive	$\rightarrow$	$\rightarrow$	$\rightarrow$	$\rightarrow$	$\times$	$\times$	$\times$	$\rightarrow$	$\rightarrow$	$\rightarrow$
$dcl(\mathbb{S}) = \mathbb{S}$	$\times$									

Table 2: Characterizing Properties for Realizable Extension-sets

Remember that the decision whether a certain extension-set $\mathbb{S}$ is realizable can not be done via brute force (i.e., enumerating AFs and checking whether their extensions coincide with $\mathbb{S}$) since there are no a priori bounds on the number of required arguments. Consequently, the results depicted in Table 2 put us in a very good position since now, the question of realizability can be decided locally, i.e. by inspecting the set in question itself. Moreover, all mentioned properties can checked in polynomial time w.r.t. the size of the extensions as shown in [Dunne *et al.*, 2015, Theorem 6]. For the majority of the properties tractability is immediately apparent. The only exception is tightness of the downward-closure of a given extension-set $\mathbb{S}$ since its size is not polynomially bounded in the size of $\mathbb{S}$ (cf. [Dunne *et al.*, 2015, Proposition 12] for a way out of this problem).

By inspecting the respective properties as depicted in Table 2, we can immediately put the signatures of different semantics in relation to each other. The following theorem includes the signature w.r.t. complete semantics in addition. The reason why we did not included complete semantics in our considerations is simply that a precise characterization of the complete signature is still an open problem. Nevertheless, certain necessary properties are already found [Dunne *et al.*, 2015, Proposition 4] justifying items 3 and 4 of the following theorem.

Theorem 3.34 ([Dunne *et al.*, 2015]) *The following relations hold*

1. $\Sigma^f_{\mathcal{E}_{na}} \subset \Sigma^f_{\mathcal{E}_{stg}} \subset \Sigma^f_{\mathcal{E}_{ss}} = \Sigma^f_{\mathcal{E}_{pr}}$, *2.* $\Sigma^f_{\mathcal{E}_{stb}} = \Sigma^f_{\mathcal{E}_{stg}} \cup \{\emptyset\}$,

3. $\Sigma^f_{\mathcal{E}_{cf}} \subset \Sigma^f_{\mathcal{E}_{ad}} \subset \Sigma^f_{\mathcal{E}_{co}}$,

4. $\Sigma^f_{\mathcal{E}_\sigma} \subset \Sigma^f_{\mathcal{E}_\tau}$ *where* $\sigma \in \{gr, il, eg\}$, $\tau \in \{na, stb, stg, pr, ss, co\}$ *and*

5. $\left\{\mathbb{S} \cup \{\emptyset\} \mid \mathbb{S} \in \Sigma^f_{\mathcal{E}_{pr}}\right\} \subset \Sigma^f_{\mathcal{E}_{ad}}$.

The following Venn-diagram provides a compact overview of subset relations between the considered signatures. A bordered area represents a set of extension-sets. The outer ellipse $\mathcal{ES} = \{\mathbb{S} \subseteq 2^{\mathcal{U}} \mid \mathbb{S}$ is an ext.-set$\}$ stands for the set of all extension-sets over $\mathcal{U}$. Clearly, all other signatures are subsets of $\mathcal{ES}$ by definition. Furthermore, we use $\{\{\emptyset\}\}$ or $\{\emptyset\}$ the set consisting of the single extension-set $\{\emptyset\}$ (realizable by all considered semantics) or the set containing the empty extension-set (realizable by stable semantics only), respectively. The right side of Figure 2 shows signatures of semantics providing only incomparable extension-sets. The intersection of these signatures with $\Sigma^f_{\mathcal{E}_{co}}$ exactly coincides with $\Sigma^f_{\mathcal{E}_{gr}}$ as well as $\Sigma^f_{\mathcal{E}_{il}}$ and $\Sigma^f_{\mathcal{E}_{eg}}$ which contain all extension-sets $\mathbb{S}$ with $|\mathbb{S}| = 1$. Moreover, the only extension-set they have in common with the signatures of conflict-free and admissible sets is the extension-set containing the empty extension. This fact causes the "missing" intersection in the middle of Figure 2.

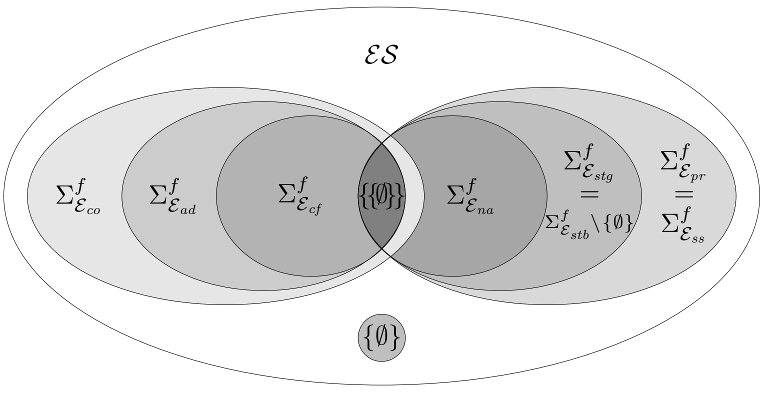

Figure 2: Subset Relations between Finite Signatures

Finally, we want to mention that all considered finite signatures, apart from the complete signature, are closed under non-empty intersections. More precisely, if two finitely σ-realizable sets $\mathbb{S}$ and $\mathbb{T}$ possess a non-empty intersection, then $\mathbb{S} \cap \mathbb{T}$ is finitely σ-realizable too. This feature is mainly due to the fact that subsets of incomparable and tight as well as incomparable and conflict-sensitive sets maintain these properties (cf. Lemmas 3.10 and 3.23).

Theorem 3.35 ([Dunne *et al.*, 2015]) *Let $\sigma \in \{cf, ad, na, stb, stg, pr, ss\}$. For any two finite AFs F_1, F_2 exists an finite AF F, s.t. $\mathcal{E}_\sigma(F) = \mathcal{E}_\sigma(F_1) \cap \mathcal{E}_\sigma(F_2)$ given that $\mathcal{E}_\sigma(F_1) \cap \mathcal{E}_\sigma(F_2) \neq \emptyset$.*[10]

3.3 Signatures w.r.t. Finite, Compact AFs

So far we considered realizibility without any restriction (apart from finiteness) for witnessing AFs. This means, realizing AFs may contain *rejected* arguments, i.e. arguments which do not appear in any extension. Rejected arguments are natural ingredients in typical argumentation scenarios and it is a priori completely unclear in which ways rejected arguments contribute to the expressibility of a particular semantics. In order to have a handle for analyzing the effect of rejected arguments, the class of *compact* AFs and its induced signatures were introduced and studied [Baumann *et al.*, 2014a; Baumann *et al.*, 2014b; Baumann *et al.*, 2016a]. An AF is compact with respect to a semantics σ, if it does not contain rejected arguments, i.e. each of its arguments appears in at least one σ-extension. Now, the main question is whether it is possible to get rid of rejected arguments without changing the outcome? or, in other words: Under which circumstances can AFs be transformed into equivalent compact ones? Note that studying compactness is far from being an academic exercise since there is a fundamental computational significance: When searching for extensions, arguments span the search space, since extensions are to be found among the subsets of the set of all arguments. Hence the more arguments, the larger the search space. Compact AFs are argument-minimal since none of the arguments can be removed without changing the outcome, thus leading to a minimal search space.

Let us first have a brief look on the naive semantics, which is defined as $\subseteq$-maximal conflict-free sets: Here, it is rather easy to see that any AF can be transformed into an equivalent compact AF by just removing all self-defeating arguments. In other words, the same outcome (in terms of the naive extensions) can be achieved by a simplified AF without rejected arguments. This means, naive semantics does not lose expressive power if we stick to compact AFs. However, it is not hard to find semantics where this coincidence does not hold implying that for such semantics the full range of expressiveness indeed relies on the concepts of rejected arguments. Consider therefore the following non-compact AF F.

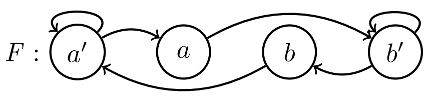

Let us consider admissible sets. We obtain $\mathbb{S} = \mathcal{E}_{ad}(F) = \{\emptyset, \{a, b\}\}$. Obviously, any attempt of realizing $\mathbb{S}$ with a compact AF $G = (\{a, b\}, R)$ is doomed to failure since if $\{a, b\}$ is admissible in G we necessarily obtain the admissibility of $\{a\}$ as well as $\{b\}$ proving $\mathbb{S} \neq \mathcal{E}_{ad}(G)$. It was one main result in [Baumann *et al.*, 2014a] to show that the finite, compact signatures w.r.t. stable, preferred,

[10] The prerequisite of a non-empty intersection can be dropped in case of stable semantics.

semi-stable, and stage semantics are strict subsets of their corresponding finite signatures. This means, in case of those semantics, sticking to finite, compact AFs implies a loss of expressive power.

Central Definitions and Preliminary Observations

In the following we formally introduce the central notions of *compact argumentation frameworks*, *compact realizibility* as well as *compact signatures*. As already stated, the main idea behind compact AFs is the absence of rejected arguments. For labelling-based semantics σ (i.e., a semantics returning n-tuples) we assume that the first component of their associated σ-labellings are interpreted as *acceptable sets of arguments* in analogy to σ-extensions in case of extension-based semantics. This means, if a certain argument occur in no first component of given σ-labellings we classify it as *rejected*. For a given labelling L we use L^{I} to refer to its first component.

Definition 3.36 *Given a semantics* $\sigma : \mathcal{F} \to 2^{\left(2^{\mathcal{U}}\right)^n}$. *An AF* $F = (A, R)$ *is compact for* σ *(or simply,* σ*-compact) if* $Args_{\mathcal{E}_\sigma(F)} = A$ *(in case of* $n = 1$*) or* $Args_{\{L^{\mathrm{I}} \mid L \in \mathcal{L}_\sigma(F)\}} = A$ *(for* $n \geq 2$*), respectively.*

Although extension-based and labelling-based semantics are formally different semantics (according to Definition 2.2) we often speak of the extension-based version or labelling-based version of a certain semantics. This can be formally justified for the considered semantics since there is a close relationship between both versions (cf. Section 2 in Chapter 4 for more explanatory comments as well as Facts 4.38 and 4.39 for some formal relations). The following fact shows that for all considered semantics σ there is no need to distinguish between σ-compactness w.r.t. the extension-based version of σ and σ-compactness w.r.t. the labelling-based version of σ. As an aside, such a coincidence does not require a one-to-one correspondence between the extension-based and labelling-based version of a semantics σ. It suffices that any σ-extension induces a σ-labelling and vice versa in such a way that accepted arguments are preserved (cf. statements 1 and 2 of Fact 4.38).

Fact 3.37 *For any* $\sigma \in \{stb, ss, stg, cf2, stg2, pr, ad, co, gr, il, eg, na, cf\}$ *and any*[11] *AF* F *we have:* F *is compact for* $\mathcal{E}_\sigma$ *iff* F *is compact for* $\mathcal{L}_\sigma$.

In the following we use CAF_σ for AFs compact for σ. Moreover, we use CAF_σ^f to indicate that the considered frameworks are finite in addition. It is intuitively clear that there are AFs F being σ-compact without being τ-compact for two different semantics σ and τ. The following example firstly presented in [Baumann *et al.*, 2014a, Figure 1] provides us with a witnessing framework.

[11] Indeed, no finiteness restriction is required here.

Example 3.38 *Consider the following AF F.*[12]

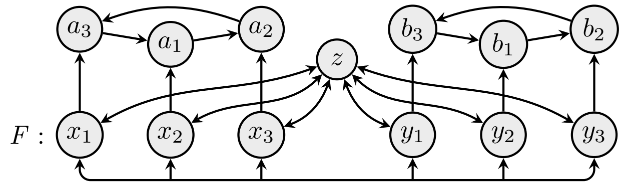

The preferred extensions of F are $\mathcal{E}_{pr}(F) = \{\{z\}, \{x_1, a_1\}, \{x_2, a_2\}, \{x_3, a_3\}, \{y_1, b_1\}, \{y_2, b_2\}, \{y_3, b_3\}\}$*, meaning that F is pr-compact (*$F \in CAF^f_{pr}$*) since each argument occurs in at least one preferred extension. On the other hand observe that* $\mathcal{E}_{ss}(F) = \mathcal{E}_{pr}(F) \setminus \{\{z\}\}$ *and* $\mathcal{E}_{stg}(F) = \{\{x_i, a_i, b_j\}, \{y_i, b_i, a_j\} \mid 1 \leq i, j \leq 3\}$*, i.e. z is not contained in any semi-stable or stage extension. Therefore F is neither compact for semi-stable nor compact for stage semantics (i.e.* $F \notin CAF^f_{ss}$ *and* $F \notin CAF^f_{stg}$*).*

How are the different sets of compact AFs related? We start with an easy observation.

Lemma 3.39 ([Baumann *et al.*, 2016a]) *For any two semantics σ and τ such that for each AF F, for every* $S \in \mathcal{E}_\sigma(F)$ *there is some* $S' \in \mathcal{E}_\tau(F)$ *with* $S \subseteq S'$*, we have* $CAF_\sigma \subseteq CAF_\tau$*.*

Note that $\sigma \subseteq \tau$ is a special case of the premise of Lemma 3.39. Thus, $CAF_\sigma \subseteq CAF_\tau$, whenever $\sigma \subseteq \tau$ (see Figure 1 for an overview). Strict subset relations have to be proven by providing a witnessing AF as presented in Example 3.38. Moreover, $CAF_{pr} = CAF_{co} = CAF_{ad}$ as well as $CAF_{na} = CAF_{cf}$ is justified by Lemma 2.14 and the fact that $pr \subseteq co \subseteq ad$ and $na \subseteq cf$. Finally, in case of the uniquely defined grounded and ideal semantics we have, $F = (A, R)$ is compact if and only if $R = \emptyset$. This in turn implies that F is compact for stable semantics. This means, $CAF_{gr} = CAF_{il} \subset CAF_{stb}$. Remember that eager semantics is uniquely defined w.r.t. finitary AFs only (Theorem 2.25, Example 2.9). Consequently, we may conclude $CAF^f_{gr} = CAF^f_{eg}$ only. Although, the majority of the results do not require the finiteness restriction we present the following theorem in terms of finite AFs. Detailed proofs for the relations between stable, semi-stable, preferred, stage and naive semantics can be found in [Baumann *et al.*, 2016a, Theorem 2].

Theorem 3.40 *The following relations hold:*

1. $CAF^f_{gr} = CAF^f_{il} = CAF^f_{eg}$,
2. $CAF^f_{pr} = CAF^f_{co} = CAF^f_{ad}$,
3. $CAF^f_{na} = CAF^f_{cf}$,

[12] The construct in the lower part of the figure represents symmetric attacks between each pair of distinct arguments. We will make use of this style in illustrations throughout the whole section.

4. $CAF^f_{gr} \subset CAF^f_{stb} \subset CAF^f_{ss} \subset CAF^f_{pr} \subset CAF^f_{na}$*,*

5. $CAF^f_{stb} \subset CAF^f_{stg} \subset CAF^f_{na}$ *and*

6. $CAF^f_{stg} \not\subseteq CAF^f_{\sigma}$ *as well as* $CAF^f_{\sigma} \not\subseteq CAF^f_{stg}$ *for any* $\sigma \in \{pr, ss\}$*.*

The following figure concisely summarizes all relations mentioned in the theorem above. Directed arrows between two boxes have to be interpreted as strict subset relations between the mentioned sets of compact AFs in these boxes.

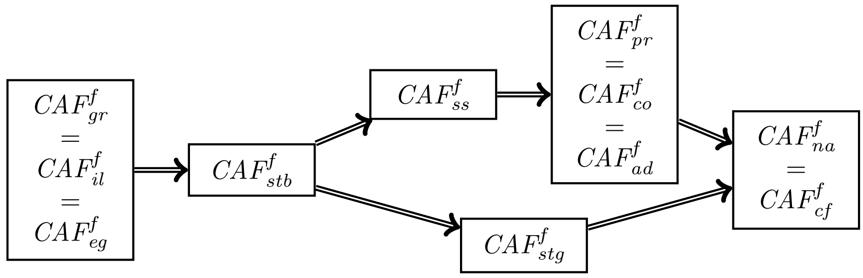

Figure 3: Subset Relations between Finite, Compact AFs

Instantiating Definitions 3.1 and 3.2 with $\mathcal{C} = CAF^f_{\sigma}$ formalize the notions of realizibility as well as signatures relativised to finite, compact AFs. Consider the following definitions.

Definition 3.41 *Given a semantics* $\sigma : \mathcal{F} \to 2^{\left(2^{\mathcal{U}}\right)^n}$*. A set* $\mathbb{S} \subseteq \left(2^{\mathcal{U}}\right)^n$ *is finitely, compactly* σ*-realizable if there is an AF* $F \in CAF^f_{\sigma}$*, s.t.* $\sigma(F) = \mathbb{S}$*.*

Definition 3.42 *Given a semantics* σ*. The finite, compact* σ*-signature is defined as* $\left\{\sigma(F) \mid F \in CAF^f_{\sigma}\right\}$ *abbreviated by* $\Sigma^{f,c}_{\sigma}$*.*

It is clear that $\Sigma^{c,f}_{\sigma} \subseteq \Sigma^{f}_{\sigma}$ holds for any semantics σ, i.e. finite, compact realizibility implies finite realizibility. In the following we shed light on the question whether the mentioned subset relation is strict for a given semantics? In other words, we answer the question whether we indeed lose expressive power if sticking to compact AFs.

The Loss or Stability of Expressive Power

Let us consider the uniquely defined grounded, ideal and eager semantics first. We already stated that a set $\mathbb{S}$ is realizable w.r.t. these semantics if and if only if $\mathbb{S}$ is an one-element extension-set if considering finite AFs (Theorem 3.33). Furthermore, it is immediate that an extension-set $\mathbb{S} = \{E\}$ can be compactly realized via $F_E = (E, \emptyset)$. This means, these semantics do not lose expressive power if we restrict ourselves to compact AFs. Furthermore, the attentive reader may have noticed that the canonical argumentation framework $F^{cf}_{\mathbb{S}}$, which was used as a witnessing framework for conflict-free sets and naive semantics (cf. Definition 3.13 as well as Propositions 3.15 and 3.16), does not involve further artificial arguments. Thus, it verifies finite, compact realizibility and shows that there is no expressive loss in case of conflict-free sets and naive semantics. For the other considered semantics, namely admissible, stable, stage, semi-stable, preferred as well as complete semantics we have to accept a strict weaker expressibility if we stick to compact AFs. In order to prove that in case of these semantics the full range of expressiveness indeed relies on the concept of rejected arguments we have to come up with witnessing extension-sets. Consider therefore the following example.

Example 3.43 *The extension-set* $\mathbb{S} = \{\{a,b\},\{a,d,e\},\{b,c,e\}\}$ *is realizable under preferred as well as semi-stable semantics (cf. Example 3.20 for a realizing non-compact framework). Let* $\sigma \in \{pr, ss\}$*. Now suppose there exists an AF* $F = (\{a,b,c,d,e\}, R)$*, s.t.* $\mathcal{E}_\sigma(F) = \mathbb{S}$*. Since* $\{a,d,e\},\{b,c,e\} \in \mathbb{S}$ *and* $\sigma \subseteq cf$ *we conclude that there is no attack in* R *involving* e*, i.e.* e *is an isolated argument in* F*. But then,* e *is contained in each* σ*-extension of* F *contradicting* $\{a,b\} \in \mathbb{S}$*. In Summary,* $\mathbb{S} \in \Sigma^{f}_{\mathcal{E}_\sigma} \setminus \Sigma^{f,c}_{\mathcal{E}_\sigma}$*.*

For further witnessing extension-sets we refer the reader to [Baumann *et al.*, 2016a, Propositions 35 and 57] and proceed with the main theorem.

Theorem 3.44 *It holds that*

1. $\Sigma^{f,c}_{\mathcal{E}_\sigma} = \Sigma^{f}_{\mathcal{E}_\sigma}$ *for* $\sigma \in \{cf, na, gr, il, eg\}$*, and*
2. $\Sigma^{f,c}_{\mathcal{E}_\sigma} \subset \Sigma^{f}_{\mathcal{E}_\sigma}$ *for* $\sigma \in \{ad, stb, stg, ss, pr, co\}$*.*

In both cases we may benefit of characterization theorems for finite signatures (cf. Theorems 3.12, 3.24 and 3.33). If both signatures are identical (first item), then necessary and sufficient properties for being finitely σ-realizable immediately carry over to finite, compact σ-realizibility. If we observe a strict subset relation (second item), then we obtain at least necessary properties for being in the finite, compact σ-signature.

Theorem 3.45 *Given a set* $\mathbb{S} \subseteq 2^{\mathcal{U}}$*, then*

1. $\mathbb{S} \in \Sigma^{f,c}_{\mathcal{E}_{cf}} \Leftrightarrow \mathbb{S}$ *is a non-empty, downward-closed and tight ext.-set,*
2. $\mathbb{S} \in \Sigma^{f,c}_{\mathcal{E}_{na}} \Leftrightarrow \mathbb{S}$ *is a non-empty, incomparable ext.-set and* $dcl(\mathbb{S})$ *is tight,*
3. $\mathbb{S} \in \Sigma^{f,c}_{\mathcal{E}_{gr}} \Leftrightarrow \mathbb{S}$ *is an ext.-set with* $|\mathbb{S}| = 1$*,*
4. $\mathbb{S} \in \Sigma^{f,c}_{\mathcal{E}_{il}} \Leftrightarrow \mathbb{S}$ *is an ext.-set with* $|\mathbb{S}| = 1$*,*
5. $\mathbb{S} \in \Sigma^{f,c}_{\mathcal{E}_{eg}} \Leftrightarrow \mathbb{S}$ *is an ext.-set with* $|\mathbb{S}| = 1$ *and*

6. $\mathbb{S} \in \Sigma^{f,c}_{\mathcal{E}_{stb}} \Rightarrow \mathbb{S}$ *is an incomparable and tight ext.-set,*

7. $\mathbb{S} \in \Sigma^{f,c}_{\mathcal{E}_{stg}} \Rightarrow \mathbb{S}$ *is a non-empty, incomparable and tight ext.-set,*

8. $\mathbb{S} \in \Sigma^{f,c}_{\mathcal{E}_{ad}} \Rightarrow \mathbb{S}$ *is a conflict-sensitive ext.-set containing* $\emptyset$,

9. $\mathbb{S} \in \Sigma^{f,c}_{\mathcal{E}_{pr}} \Rightarrow \mathbb{S}$ *is a non-empty, incomparable and conflict-sensitive ext.-set,*

10. $\mathbb{S} \in \Sigma^{f,c}_{\mathcal{E}_{ss}} \Rightarrow \mathbb{S}$ *is a non-empty, incomparable and conflict-sensitive ext.-set.*

Comparing Finite, Compact Signatures and Final Remarks
In the following we relate the finite, compact signatures of the semantics under consideration to each other. Recall that for finite signatures it holds that $\Sigma^f_{\mathcal{E}_{na}} \subset \Sigma^f_{\mathcal{E}_{stg}} = \left(\Sigma^f_{\mathcal{E}_{stb}} \setminus \{\emptyset\}\right) \subset \Sigma^f_{\mathcal{E}_{ss}} = \Sigma^f_{\mathcal{E}_{pr}}$ (cf. Figure 2). This picture changes dramatically when considering the relationships between finite, compact signatures as depicted in Figure 4 (incomparable semantics only) and formally stated in Theorem 3.46. The dashed areas represent particular intersections for which the question of existence of extension-sets is still an open question.

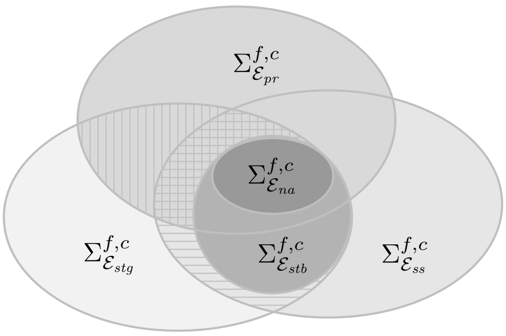

Figure 4: Subset Relations between Finite, Compact Signatures

We proceed with an enumeration of relationships between finite, compact signature including further semantics like conflict-free and admissible sets as well as grounded, ideal, eager and complete semantics. For formal proofs we refer the interested reader to [Baumann *et al.*, 2016a, Theorem 36, Proposition 58].

Theorem 3.46 *The following relations hold:*

1. $\Sigma^{f,c}_{\mathcal{E}_\sigma} \subset \Sigma^{f,c}_{\mathcal{E}_{na}} \subset \Sigma^{f,c}_{\mathcal{E}_\tau}$ *for* $\sigma \in \{gr, il, eg\}$ *and* $\tau \in \{stb, stg, ss, pr\}$,

2. $\Sigma^{f,c}_{\mathcal{E}_{stb}} \subset \Sigma^{f,c}_{\mathcal{E}_\sigma}$ *for* $\sigma \in \{stg, ss\}$,

3. $\Sigma^{f,c}_{\mathcal{E}_{cf}} \subset \Sigma^{f,c}_{\mathcal{E}_{ad}}$,

4. $\Sigma^{f,c}_{\mathcal{E}_{co}} \setminus \Sigma^{f,c}_{\mathcal{E}_{\sigma}} \neq \emptyset$ *and* $\Sigma^{f,c}_{\mathcal{E}_{\sigma}} \setminus \Sigma^{f,c}_{\mathcal{E}_{co}} \neq \emptyset$ *for* $\sigma \in \{cf, ad\}$,

5. $\Sigma^{f,c}_{\mathcal{E}_{pr}} \setminus (\Sigma^{f,c}_{\mathcal{E}_{stb}} \cup \Sigma^{f,c}_{\mathcal{E}_{ss}} \cup \Sigma^{f,c}_{\mathcal{E}_{stg}}) \neq \emptyset$,

6. $\Sigma^{f,c}_{\mathcal{E}_{stg}} \setminus (\Sigma^{f,c}_{\mathcal{E}_{stb}} \cup \Sigma^{f,c}_{\mathcal{E}_{pr}} \cup \Sigma^{f,c}_{\mathcal{E}_{ss}}) \neq \emptyset$,

7. $\Sigma^{f,c}_{\mathcal{E}_{stb}} \setminus \Sigma^{f,c}_{\mathcal{E}_{pr}} \neq \emptyset$,

8. $(\Sigma^{f,c}_{\mathcal{E}_{pr}} \cap \Sigma^{f,c}_{\mathcal{E}_{ss}}) \setminus (\Sigma^{f,c}_{\mathcal{E}_{stb}} \cup \Sigma^{f,c}_{\mathcal{E}_{stg}}) \neq \emptyset$ *and*

9. $\Sigma^{f,c}_{\mathcal{E}_{ss}} \setminus (\Sigma^{f,c}_{\mathcal{E}_{stb}} \cup \Sigma^{f,c}_{\mathcal{E}_{pr}} \cup \Sigma^{f,c}_{\mathcal{E}_{stg}}) \neq \emptyset$.

Comparing the results on expressiveness of the considered semantics as stated in Theorems 3.34 and 3.46 we observe notable differences. When allowing rejected arguments, preferred and semi-stable semantics are equally expressive and at the same time strictly more expressive than stable and stage semantics. As we have seen, this property does not carry over to the compact setting (with the exceptions $\Sigma^{f,c}_{\mathcal{E}_{stb}} \subset \Sigma^{f,c}_{\mathcal{E}_{ss}}$ and $\Sigma^{f,c}_{\mathcal{E}_{stb}} \subset \Sigma^{f,c}_{\mathcal{E}_{stg}}$) where signatures become incomparable.

Finally, regarding the open issues represented as dashed areas in Figure 4. More precisely, it is an open problem whether there are extension-sets lying in the intersection between $\Sigma^{f,c}_{\mathcal{E}_{pr}}$ (resp. $\Sigma^{f,c}_{\mathcal{E}_{ss}}$) and $\Sigma^{f,c}_{\mathcal{E}_{stg}}$ but outside of $\Sigma^{f,c}_{\mathcal{E}_{stb}}$. In [Baumann *et al.*, 2016a] it is conjectured that such extension-sets do not exist.

Conjecture 3.47 ([Baumann *et al.*, 2016a]) *It holds that* $\Sigma^{f,c}_{\mathcal{E}_{pr}} \cap \Sigma^{f,c}_{\mathcal{E}_{stg}} \subset \Sigma^{f,c}_{stb}$ *and* $\Sigma^{f,c}_{\mathcal{E}_{ss}} \cap \Sigma^{f,c}_{\mathcal{E}_{stg}} = \Sigma^{f,c}_{\mathcal{E}_{stb}}$.

3.4 Signatures w.r.t. Finite, Analytic AFs

We now turn to a further phenomenon, so-called *implicit conflicts* which can be frequently observed in typical argumentation scenarios. Consider therefore the following AF F.

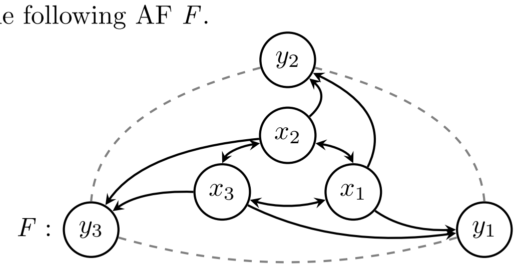

Let us consider stable semantics. Please note that any x_i is jointly acceptable with one specific y_j. More precisely, $\mathcal{E}_{stb}(F) = \{\{x_1, y_3\}, \{x_2, y_1\}, \{x_3, y_2\}\}$ implying that we do not have any rejected arguments, i.e. F is stable compact. What can be said about the two pairs of arguments x_1 and x_2 as well as y_1 and y_2? First of all, both pairs represent a semantical conflict in F since neither of those pairs occur together in any stable extension. In case of x_1

and x_2, the conflict is even a syntactical one since both arguments attack each other in contrast to the pair consisting of y_1 and y_2. This difference leads to the distinction between syntactically underlined *explicit conflicts* and syntactically unfounded *implicit* ones (indicated by dashed lines). In order to understand how implicit conflicts contribute to the expressiveness of a certain semantics, the set of *analytic* AFs and its induced signatures were introduced and studied [Linsbichler *et al.*, 2015; Baumann *et al.*, 2016a]. An analytic framework, i.e. a framework which is free of implicit conflicts maximizes the information on conflicts. One main question is: under which circumstances an arbitrary framework can be transformed into an equivalent analytic one? This question is interesting from a theoretical as well as practical point of view. On the one hand, analytic frameworks are natural candidates for normal forms of AFs, and on the other maximizing the number of explicit conflicts might help argumentation systems to evaluate AFs more efficiently.

Let us consider again the extension-set $\mathbb{S} = \{\{x_1, y_3\}, \{x_2, y_1\}, \{x_3, y_2\}\}$ stemming from the AF F depicted above. Replacing the dashed arrows with symmetric attacks in F shows that $\mathbb{S}$ can be analytically realized under stable semantics. Interestingly, this is no coincidence, since it was shown that in case of stable semantics any AF can be transformed into an equivalent analytical one. However, in general it is not that easy to make implicit conflicts explicit since there are frameworks where any suitable transformation requires the use of additional arguments as shown in [Linsbichler *et al.*, 2015].

Central Definitions and Preliminary Observations

In this section we consider the central notions of *analytic argumentation frameworks*, *analytic realizability* as well as *analytic signatures.* In order to define analytic AF we have to differentiate between the concept of an attack (as a syntactical element) and the concept of a conflict (with respect to the evaluation under a given semantics). More precisely, if two arguments cannot be accepted together, i.e. no reasonable position contain them jointly as elements, we say that these arguments are in conflict. If this conflict is syntactically underlined by an attack between them, we call this conflict explicit, otherwise implicit. Now, an analytic framework is an AF which simply does not contain any implicit conflicts. Consider the following definition.

Definition 3.48 *Given a semantics $\sigma : \mathcal{F} \to 2^{\left(2^{\mathcal{U}}\right)^n}$, an AF $F = (A, R)$ and two arguments $a, b \in A$. We say that*

1. *a and b are in conflict for σ if $(a, b) \notin Pairs_{\mathcal{E}_\sigma(F)}$ (in case of $n = 1$) or $(a, b) \notin Pairs_{\{L^I | L \in \mathcal{L}_\sigma(F)\}}$ (for $n \geq 2$), respectively,*
2. *the conflict is explicit w.r.t. σ if $(a, b) \in R$ or $(b, a) \in R$, otherwise implicit,*
3. *the AF F is analytic for σ (or σ-analytic) if all conflicts are explicit.*

Please notice that Definition 3.48 does not require a and b to be different arguments. In particular, an argument that is not contained in any reasonable position is in conflict with itself. This conflict is explicit if the argument is self-attacking and implicit otherwise. Furthermore, for all considered semantics σ we observe there is no need to distinguish between σ-analyticality w.r.t. the extension-based version of σ and σ-analyticality w.r.t. the labelling-based version of σ (similarly as in case of σ-compactness as stated in Fact 3.37). Please note that this coincidence is justified for any semantics σ whenever $\mathcal{E}_\sigma(F) = \{L^{\text{I}} \mid L \in \mathcal{L}_\sigma(F)\}$ is guaranteed.

Fact 3.49 *For any $\sigma \in \{stb, ss, stg, cf2, stg2, pr, ad, co, gr, il, eg, na, cf\}$ and any AF F we have: F is analytic for $\mathcal{E}_\sigma$ iff F is analytic for $\mathcal{L}_\sigma$.*

In the following we denote the set of all σ-analytic AFs as XAF_σ. To indicate that the frameworks under consideration are finite we use XAF^f_σ. We proceed with an example illustrating the new definitions.

Example 3.50 *As a simple example consider the following AF F depicted below.*

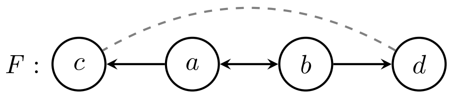

For $\sigma \in \{stb, pr, ss, stg\}$ we have $\mathcal{E}_\sigma(F) = \{\{a, d\}, \{b, c\}\}$. Observe that there is only one implicit conflict, namely the conflict between the arguments c and d, denoted by a dashed line. Hence, F is not σ-analytic, i.e. $F \notin XAF^f_\sigma$. However, since $\mathcal{E}_{na}(F) = \mathcal{E}_\sigma(F) \cup \{\{c, d\}\}$ we have that F is na-analytic, i.e. $F \in XAF^f_{na}$.

As indicated in Example 3.50 the sets of analytic AFs can differ for different semantics. Just like in case of compact AFs (cf. Lemma 3.39) one may easily verify the following lemma which allows to obtain a plenty of subset relations between sets of analytic AFs.

Lemma 3.51 ([Baumann *et al.*, 2016a]) *For any two semantics σ and τ such that for each AF F, for every $S \in \mathcal{E}_\sigma(F)$ there is some $S' \in \mathcal{E}_\tau(F)$ with $S \subseteq S'$, we have $XAF_\sigma \subseteq XAF_\tau$.*

In line with the existing literature we restrict our considerations to finite AFs. Regarding universal (but not uniquely) defined semantics we obtain the same relations as in case of compact AFs (see explanations below Lemma 3.39). In any case we have $XAF^f_{gr} \subseteq XAF^f_{il} \subseteq XAF^f_{eg}$ since ideal semantics accepts more arguments than grounded semantics and eager semantics is even more credulous than ideal semantics. Furthermore, $XAF^f_{eg} \subseteq XAF^f_{ss}$ because the unique eager extension is contained in all semi-stable extension by definition and moreover, semi-stable semantics guarantees reasonable positions in case of finite AFs. Now, let $F = (A, R)$ be analytic w.r.t. eager semantics and $\mathcal{E}_{eg}(F) = \{E\}$. We deduce that all arguments in $A \setminus E$ have to be self-defeating. Consequently, its corresponding (conflict-free) base semantics (cf. Definition 2.17) warrants exactly one extension for F. More precisely, $\mathcal{E}_{ss}(F) = \{E\}$. Finally, due to the self-conflicting arguments and the admissibility of E we obtain $\mathcal{E}_{pr}(F) = \{E\}$ and thus, $\mathcal{E}_{il}(F) = \{E\}$ showing that F is even analytic w.r.t. ideal semantics, i.e. $XAF^f_{il} = XAF^f_{eg}$. The AF $F = (\{a,b\}, \{(a,b),(b,a),(b,b)\})$ proves that a similar result in case of grounded and ideal semantics does not hold. Detailed proofs for the relations between stable, semi-stable, preferred, stage and naive semantics can be found in [Baumann *et al.*, 2016a, Theorem 4].

Theorem 3.52 *The following relations hold:*

1. $XAF^f_{gr} \subset XAF^f_{il} = XAF^f_{eg} \subset XAF^f_{ss}$,
2. $XAF^f_{pr} = XAF^f_{co} = XAF^f_{ad}$,
3. $XAF^f_{na} = XAF^f_{cf}$,
4. $XAF^f_{stb} \subset XAF^f_{ss} \subset XAF^f_{pr} \subset XAF^f_{na}$,
5. $XAF^f_{stb} \subset XAF^f_{stg} \subset XAF^f_{na}$,
6. $XAF^f_{stg} \not\subseteq XAF^f_{\sigma}$ *and* $XAF^f_{\sigma} \not\subseteq XAF^f_{stg}$ *for any* $\sigma \in \{pr, ss\}$,
7. $XAF^f_{\sigma} \not\subseteq XAF^f_{\tau}$ *and* $XAF^f_{\tau} \not\subseteq XAF^f_{\sigma}$ *for any* $\sigma \in \{gr, il, eg\}$, $\tau \in \{stb, stg\}$.

The following figure summarizes all relation in a compact way. Similarly to Figure 3, a directed arrow between two boxes has to be interpreted as strict subset relation between the mentioned sets of analytic AFs therein.

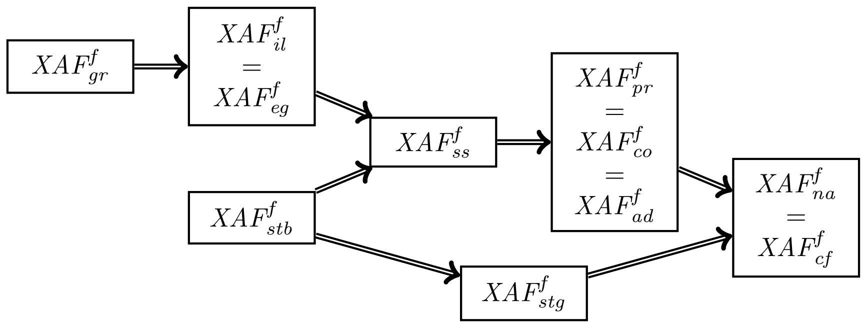

Figure 5: Subset Relations between Finite, Analytic AFs

At this point we want to mention that although Figures 3 and 5 look very similar we have that compactness and analyticality are sufficiently distinct properties. More precisely, apart from the uniquely defined semantics as well as naive semantics and conflict-free sets no subset relations between the sets of compact and analytic frameworks can be stated in general. Sticking to self-loop-free AFs allows one to draw further relations such as analyticality implies compactness for any considered semantics. The main reason for this general relation is that rejected arguments has to be self-defeating in case of analytic frameworks. A selection of proofs of relations listed below can be found in [Baumann *et al.*, 2016a, Proposition 5-8].

Proposition 3.53 *Given an AF F, then*

1. $CAF^f_\sigma \subset XAF^f_\sigma$ *for* $\sigma \in \{gr, il, eg, na, cf\}$,
2. $CAF^f_\sigma \not\subseteq XAF^f_\sigma$ *and* $XAF^f_\sigma \not\subseteq CAF^f_\sigma$ *for* $\sigma \in \{ad, stb, ss, pr, stg, co\}$.

If F is self-loop-free in addition, then

3. $F \in XAF^f_\sigma$ *and* $F \in CAF^f_\sigma$ *for* $\sigma \in \{na, cf\}$,
4. $F \in XAF^f_\sigma \Leftrightarrow F \in CAF^f_\sigma$ *for* $\sigma \in \{gr, il, eg\}$ *and*
5. $F \in XAF^f_\sigma \Rightarrow F \in CAF^f_\sigma$ *for* $\sigma \in \{ad, stb, ss, pr, stg, co\}$.

We now precisely formalize the notions of realizibility as well as signatures relativised to finite, analytic AFs. This can be formally done via instantiating Definitions 3.1 and 3.2 with $\mathcal{C} = XAF^f_\sigma$.

Definition 3.54 *Given a semantics* $\sigma : \mathcal{F} \to 2^{\left(2^{\mathcal{U}}\right)^n}$. *A set* $\mathbb{S} \subseteq \left(2^{\mathcal{U}}\right)^n$ *is finitely, analytically* σ*-realizable if there is an AF* $F \in XAF^f_\sigma$, *s.t.* $\sigma(F) = \mathbb{S}$.

Definition 3.55 *Given a semantics* σ. *The finite, analytic* σ*-signature is defined as* $\left\{\sigma(F) \mid F \in XAF^f_\sigma\right\}$ *abbreviated by* $\Sigma^{f,x}_\sigma$.

The Loss or Stability of Expressive Power

Clearly, every set in the finite, analytic signature of a semantics is also contained in the finite signature. Remember that in case of compact AFs we do not lose any expressive power if considering the uniquely defined grounded, ideal and eager semantics as well as naive semantics and conflict-free sets (Theorem 3.44). These equal expressiveness results carry over to analytic AFs and moreover, even stable and stage semantics may realize the same sets. For instance, consider again the non-analytic AF F as introduced in Example 3.50. One may easily verify that adding an attack from c to d or vice versa yields an AF F' analytic for stable semantics which does not change the set of stable extensions. However, in general it is not that easy to make implicit conflicts explicit but it was shown that the use of additional arguments indeed allows one to turn any finite framework in an analytical one without changing the set of stable or stage extensions, respectively [Baumann *et al.*, 2016a, Proposition 28, Theorem 29]. For the sake of completeness, we mention that it was an open question for a while, known as *Explicit Conflict Conjecture* [Baumann *et al.*, 2014a], whether it is possible, under stable semantics, to translate a given AF into an equivalent analytic one without adding further arguments. In [Baumann *et al.*, 2016a] the conjecture was refuted for stable and even stage semantics. For the remaining semantics, i.e. admissible, semi-stable, preferred and complete semantics the conjecture does not hold either since in case of these semantics we even have that the finite, analytic signature is a strict subset of the corresponding finite one. This means, the full range of expressiveness indeed relies on the use of implicit conflicts. Consider the following example firstly presented in [Baumann *et al.*, 2016a, Example 6].

Example 3.56 *Take into account the AF $F = (A, R)$ as depicted below.*

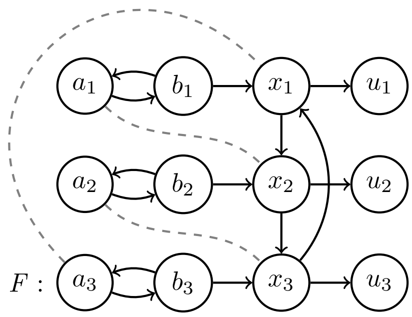

Formally, we have

$A = \{a_i, b_i, x_i, u_i \mid i \in \{1,2,3\}\}$ *and*
$R = \{(a_i, b_i), (b_i, a_i), (b_i, x_i), (x_i, u_i) \mid i \in \{1,2,3\}\} \cup \{(x_1, x_2), (x_2, x_3), (x_3, x_1)\}$.

Regarding the extension-based version of preferred semantics we obtain the set $\mathbb{S} = \mathcal{E}_{pr}(F) = \{S_a, S_b, A_1, A_2, A_3, B_1, B_2, B_3\}$ with

$$\begin{aligned} S_a &= \{a_1, a_2, a_3\} & S_b &= \{b_1, b_2, b_3, u_1, u_2, u_3\} \\ A_1 &= \{a_2, a_3, b_1, x_2, u_1, u_3\} & B_1 &= \{a_1, b_2, b_3, x_1, u_2, u_3\} \\ A_2 &= \{a_1, a_3, b_2, x_3, u_1, u_2\} & B_2 &= \{a_2, b_1, b_3, x_2, u_1, u_3\} \\ A_3 &= \{a_1, a_2, b_3, x_1, u_2, u_3\} & B_3 &= \{a_3, b_1, b_2, x_3, u_1, u_2\} \end{aligned}$$

We observe three implicit conflicts indicated by dashed lines. Consequently, F is not analytic w.r.t. preferred semantics. Moreover, we claim that $\mathbb{S}$ is not analytically pr-realizable at all. For a contradiction we assume that there exists an AF $G \in XAF^f_{pr}$, s.t. $\mathcal{E}_{pr}(G) = \mathbb{S}$. We now investigate this hypothetical AF G. The main idea is to show that if the conflict between a_1 and x_2 is made explicit, then $\mathbb{S} \neq \mathcal{E}_{pr}(G)$. First, note that G contains at least all arguments in A since $Args_{\mathbb{S}} = A$. Due to A_3 and B_3 we deduce that $S_a \cup \{u_2\}$ is conflict-free in G. Furthermore, due to A_1, the admissibility of S_a in G and the assumption that all conflicts has to be explicit, we infer that a_1 attacks x_2. Moreover, in consideration of $\mathbb{S}$, it is easy to see that x_2 is the only possible attacker of u_2

among $Args_{\mathbb{S}}$. *This implies that* S_a *defends* u_2 *against all arguments in* $Args_{\mathbb{S}}$. *Finally, any additional argument* $z \notin Args_{\mathbb{S}}$ *in* G *must be attacked by* S_a *since* G *is analytic w.r.t. preferred semantics and* S_a *must be admissible. This causes* $S_a \cup \{u_2\}$ *to be admissible in* G *and hence,* S_a *cannot be preferred in* G. *In summary, any AF realizing* $\mathbb{S}$ *has to be non-analytic for preferred semantics, i.e.* $\mathbb{S} \in \Sigma^{f}_{\mathcal{E}_{pr}} \setminus \Sigma^{f,x}_{\mathcal{E}_{pr}}$.

We proceed with the main theorem comparing finite signatures with their corresponding analytical ones.

Theorem 3.57 ([Baumann *et al.*, 2016a]) *It holds that*

1. $\Sigma^{f,x}_{\mathcal{E}_\sigma} = \Sigma^{f}_{\mathcal{E}_\sigma}$ *for* $\sigma \in \{cf, na, gr, il, eg, stb, stg\}$, *and*
2. $\Sigma^{f,x}_{\mathcal{E}_\sigma} \subset \Sigma^{f}_{\mathcal{E}_\sigma}$ *for* $\sigma \in \{ad, ss, pr, co\}$.

In the following we present characterization theorems for finite, analytic signatures or at least necessary properties for being finitely, analytically realizable. All results can be verified via combining the main theorem above as well as the already presented characterization theorems for finite signatures, namely Theorems 3.12, 3.24 and 3.33.

Theorem 3.58 *Given a set* $\mathbb{S} \subseteq 2^{\mathcal{U}}$, *then*

1. $\mathbb{S} \in \Sigma^{f,x}_{\mathcal{E}_{cf}} \Leftrightarrow \mathbb{S}$ *is a non-empty, downward-closed and tight ext.-set,*
2. $\mathbb{S} \in \Sigma^{f,x}_{\mathcal{E}_{na}} \Leftrightarrow \mathbb{S}$ *is a non-empty, incomparable ext.-set and* $dcl(\mathbb{S})$ *is tight,*
3. $\mathbb{S} \in \Sigma^{f,x}_{\mathcal{E}_{gr}} \Leftrightarrow \mathbb{S}$ *is an ext.-set with* $|\mathbb{S}| = 1$,
4. $\mathbb{S} \in \Sigma^{f,x}_{\mathcal{E}_{il}} \Leftrightarrow \mathbb{S}$ *is an ext.-set with* $|\mathbb{S}| = 1$,
5. $\mathbb{S} \in \Sigma^{f,x}_{\mathcal{E}_{eg}} \Leftrightarrow \mathbb{S}$ *is an ext.-set with* $|\mathbb{S}| = 1$,
6. $\mathbb{S} \in \Sigma^{f,x}_{\mathcal{E}_{stb}} \Leftrightarrow \mathbb{S}$ *is a incomparable and tight ext.-set,*
7. $\mathbb{S} \in \Sigma^{f,x}_{\mathcal{E}_{stg}} \Leftrightarrow \mathbb{S}$ *is a non-empty, incomparable and tight ext.-set and*
8. $\mathbb{S} \in \Sigma^{f,x}_{\mathcal{E}_{ad}} \Rightarrow \mathbb{S}$ *is a conflict-sensitive ext.-set containing* $\emptyset$,
9. $\mathbb{S} \in \Sigma^{f,x}_{\mathcal{E}_{pr}} \Rightarrow \mathbb{S}$ *is a non-empty, incomparable and conflict-sensitive ext.-set,*
10. $\mathbb{S} \in \Sigma^{f,x}_{\mathcal{E}_{ss}} \Rightarrow \mathbb{S}$ *is a non-empty, incomparable and conflict-sensitive ext.-set.*

Comparing Finite, Analytic Signatures and Final Remarks

So far we have compared finite signatures and finite, analytic signatures for the semantics under consideration. We have seen, for example, that preferred and semi-stable semantics can realize strictly more when allowing the use of implicit conflicts, while this is not the case for stable and stage semantics. In the following we relate the finite, analytic signatures of all considered semantics. Remember that we observed a considerable variety in the relations between incomparable semantics if sticking from finite to finite, compact signatures (cf. Figures 2 and 4). However, in the analytic case we have slight differences only as illustrated in Figure 8 (for incomparable semantics) and formally stated in Theorem 3.60. For instance, preferred and semi-stable signatures do not coincide anymore as shown by the following example taken from [Baumann *et al.*, 2016a, Figure 9, Proof of Theorem 34].

Example 3.59 *Consider the following AF F as depicted below.*

F :

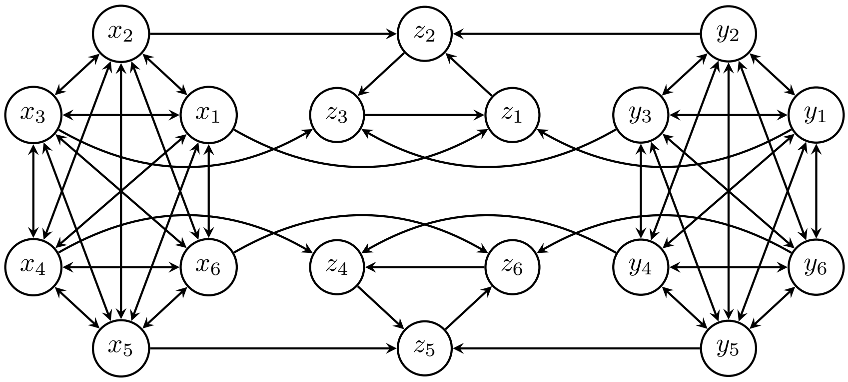

The preferred extension of F can be compactly presented via a cyclic successor functions. More precisely, if $\mathrm{s}(1) = 2, \mathrm{s}(2) = 3, \mathrm{s}(3) = 1$ *and* $\mathrm{s}(4) = 5, \mathrm{s}(5) = 6, \mathrm{s}(6) = 4$*, then* $\mathcal{E}_{pr}(F) = \mathbb{S} = \mathbb{S}_0 \cup \mathbb{S}_1 \cup \mathbb{S}_2$ *with*

$$\mathbb{S}_0 = \left\{\left\{x_i, y_j, z_{\mathrm{s}(i)}, z_{\mathrm{s}(j)}\right\} \mid i \in \{1,2,3\}, j \in \{4,5,6\} \text{ or } i \in \{4,5,6\}, j \in \{1,2,3\}\right\},$$
$$\mathbb{S}_1 = \left\{\left\{x_i, y_i, z_{\mathrm{s}(i)}\right\} \mid i \in \{1,2,3,4,5,6\}\right\} \text{ and}$$
$$\mathbb{S}_2 = \left\{\left\{x_i, y_{\mathrm{s}(i)}, z_{\mathrm{s}(\mathrm{s}(i))}\right\}, \left\{x_{\mathrm{s}(i)}, y_i, z_{\mathrm{s}(\mathrm{s}(i))}\right\} \mid i \in \{1,2,3,4,5,6\}\right\}.$$

This means, F is pr-analytic and therefore, $\mathbb{S} \in \Sigma^{f,x}_{\mathcal{E}_{pr}}$*. We show now that* $\mathbb{S} \notin \Sigma^{f,x}_{\mathcal{E}_{ss}}$*. Assume that there is some* $G = (B,S) \in XAF^f_{ss}$ *with* $\mathcal{E}_{pr}(G) = \mathbb{S}$*. We take a look at $\mathbb{S}_1$ and more specifically* $\{x_1, y_1, z_2\} \in \mathbb{S}_1$*. Now we need an explicit conflict between x_1 and x_4, but in the selected set only x_1 can possibly defend against this attack, hence* $(x_1, x_4) \in S$*. The same argument works for x_1 and x_3 as well as z_2 and z_3, meaning that also* $(x_1, x_3), (z_2, z_3) \in S$*. For symmetry reasons* $\{(x_i, x_j), (x_j, x_i), (y_i, y_j), (y_j, y_i) \mid i \in \{1,2,3\}, j \in \{4,5,6\}\} \subseteq S$ *and* $\{(x_{\mathrm{s}(i)}, x_i), (z_i, z_{\mathrm{s}(i)}) \mid i \in \{1,2 \dots 6\}\} \subseteq S$*.*

We take a look at $\mathbb{S}_2$ *and more specifically* $\{x_1, y_2, z_3\} \in \mathbb{S}_2$*. As there should be an explicit conflict between* x_1 *and* x_2 *with only* x_1 *possibly defending this extension against* x_2 *we need* $(x_1, x_2) \in S$*. Further as in this set only* y_2 *and* z_3 *can possibly attack* z_2 *we have the set* $\{y_2, z_3\}$ *attacking* z_2*. For symmetry reasons* $\{(x_i, x_{s(i)}), (y_i, y_{s(i)}) \mid i \in \{1, 2 \ldots 6\}\} \subseteq S$ *and each set* $\{x_i, z_{s(i)}\}, \{y_i, z_{s(i)}\}$ *for* $i \in \{1, 2 \ldots 6\}$ *attacks* z_i*.*

Finally we take a look at $\mathbb{S}_0$ *and specifically the set* $I = \{x_1, y_4, z_2, z_5\} \in \mathbb{S}_0$*. Since* I *necessarily is an admissible extension in an analytic AF we have that* I *attacks all rejected arguments. By the above observations we now have that* I *even attacks all arguments not being member of* I *in* G*, which means that* I *is a stable extension and stable semantics and semi-stable semantics thus coincide on* G*. But then, with* $J = \{x_1, y_1, z_2\} \in \mathbb{S}_1$ *not being in conflict with for instance* z_4 *we have that* J *can not be a stable or semi-stable extension in* G *concluding* $\mathbb{S} \notin \Sigma^{f,x}_{\mathcal{E}_{ss}}$*.*

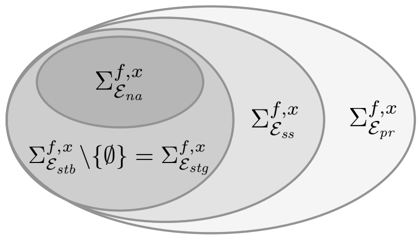

Figure 6: Subset Relations between Finite, Analytic Signatures

Theorem 3.60 ([Dunne *et al.*, 2015]) *The following relations hold:*

1. $\Sigma^{f,x}_{\mathcal{E}_\sigma} \subset \Sigma^{f,x}_{\mathcal{E}_{na}} \subset \Sigma^{f,x}_{\mathcal{E}_{stg}} \subset \Sigma^{f,x}_{\mathcal{E}_{ss}} \subset \Sigma^{f,x}_{\mathcal{E}_{pr}}$ *for* $\sigma \in \{gr, il, eg\}$*,*
2. $\Sigma^{f,x}_{\mathcal{E}_{stb}} = \Sigma^{f,x}_{\mathcal{E}_{stg}} \cup \{\emptyset\}$*,*
3. $\Sigma^{f,x}_{\mathcal{E}_{cf}} \subset \Sigma^{f,x}_{\mathcal{E}_{ad}}$ *and*
4. $\Sigma^{f,x}_{\mathcal{E}_{co}} \setminus \Sigma^{f,x}_{\mathcal{E}_\sigma} \neq \emptyset$ *and* $\Sigma^{f,x}_{\mathcal{E}_\sigma} \setminus \Sigma^{f,x}_{\mathcal{E}_{co}} \neq \emptyset$ *for* $\sigma \in \{cf, ad\}$*.*

3.5 Remarks on Unrestricted AFs and Intertranslatability

Recently, some first results regarding expressibility w.r.t. unrestricted frameworks were presented in [Baumann and Spanring, 2017]. Remember that the set of unrestricted frameworks, abbreviated by $\mathcal{F}$, contains all AFs $F = (A, R)$, s.t. $A \subseteq \mathcal{U}$ (cf. Section 2.1 for further information). This means, $\mathcal{F}$ contains finite as well as infinite AFs and especially, AFs possessing all available arguments. It is obvious that signatures w.r.t. unrestricted frameworks contain more realizable sets then their finite counterparts since finite AFs may realize finite as well as finitely many extensions only. The following definition formally captures all considered types of signatures (cf. Definitions 3.4, 3.42 and 3.55) without any finite assumption.

Definition 3.61 *Given a semantics σ. We call the set S the*

1. *(unrestricted) σ-signature if $S = \{\sigma(F) \mid F \in \mathcal{F}\}$ abbreviated by Σ_σ,*
2. *compact σ-signature if $S = \{\sigma(F) \mid F \in CAF\}$ abbreviated by Σ^c_σ and*
3. *analytic σ-signature if $S = \{\sigma(F) \mid F \in XAF\}$ abbreviated by Σ^x_σ.*

In [Baumann and Spanring, 2017] the authors were interested in a comparison of the expressive power of several mature semantics in the unrestricted setting. The following result shows that the relation between unrestricted signatures is intimately connected to their relation in case of finite, compact signatures. More precisely, non-empty relative complements in case of finite, compact signatures between two semantics carry over to their unrestricted versions. The main reason for this relation is the fact that unrestricted frameworks may contain any available argument of the universe $\mathcal{U}$.

Theorem 3.62 ([Baumann and Spanring, 2017]) *Given two semantics $\sigma, \tau \in \{na, stb, stg, ss, pr, co, gr, il, eg, cf, ad\}$ we have:*

1. *If $\Sigma^{f,c}_{\mathcal{E}_\sigma} \setminus \Sigma^{f,c}_{\mathcal{E}_\tau} \neq \emptyset$, then $\Sigma^c_{\mathcal{E}_\sigma} \setminus \Sigma^c_{\mathcal{E}_\tau} \neq \emptyset$ and*
2. *If $\Sigma^c_{\mathcal{E}_\sigma} \setminus \Sigma^c_{\mathcal{E}_\tau} \neq \emptyset$, then $\Sigma_{\mathcal{E}_\sigma} \setminus \Sigma_{\mathcal{E}_\tau} \neq \emptyset$.*

The following example illustrates the main proof idea.

Example 3.63 *Let $\mathcal{E} \in \Sigma^{f,c}_{\mathcal{E}_{pr}} \setminus \Sigma^{f,c}_{\mathcal{E}_{stb}}$ (cf. Figure 4) and $F = (A, R)$ a witnessing framework. This means, F is finite, $\mathcal{E}_{pr}(F) = \mathcal{E}$ and pr-compact, i.e. $\bigcup \mathcal{E} = A$. Consider now $H = (\mathcal{U}, R)$. Obviously, $\mathcal{E}' = \mathcal{E}_{pr}(H) = \{E \cup (\mathcal{U} \setminus A) : E \in \mathcal{E}\}$ and $\bigcup \mathcal{E}' = \mathcal{U}$ showing the σ-compactness of H. In particular, $\mathcal{E}' \in \Sigma^c_{\mathcal{E}_{pr}}$. Note that any stb-realization of $\mathcal{E}'$ has to be compact too since there are no additional arguments available. Assume $\mathcal{E}' \in \Sigma^c_{\mathcal{E}_{stb}}$, i.e. there is an AF $G' = (\mathcal{U}, R')$, s.t. $\mathcal{E}_{stb}(G') = \mathcal{E}'$. We observe that due to conflict-freeness there can not be attacks in G' between arguments from A and $\mathcal{U} \setminus A$ nor between any of the arguments from $\mathcal{U} \setminus A$. Consequently, $G = (A, R')$ is finite, $\mathcal{E}_{stb}(G) = \mathcal{E}$ and stb-compact implying that $\mathcal{E} \in \Sigma^{f,c}_{\mathcal{E}_{stb}}$ in contradiction to the initial assumption.*

Now we are prepared for a comparison in case of unrestricted frameworks. Ignoring the superscripts in Figure 4 provides you with a graphical representation for selected semantics.

Theorem 3.64 *For unrestricted signatures the following hold:*

1. *$\{\{E\} \mid E \subseteq \mathcal{U}\} = \Sigma_{\mathcal{E}_\sigma} \subset \Sigma_{\mathcal{E}_{na}} \subset \Sigma_{\mathcal{E}_\tau}$ for $\sigma \in \{gr, il\}$, $\tau \in \{stb, stg, ss, pr\}$,*

2. $\Sigma_{\mathcal{E}_{eg}} \subset \Sigma_{\mathcal{E}_{pr}}$,

3. $\Sigma_{\mathcal{E}_{stb}} \subset \Sigma_{\mathcal{E}_{\sigma}}$ *for* $\sigma \in \{stg, ss\}$,

4. $\Sigma_{\mathcal{E}_{pr}} \setminus (\Sigma_{\mathcal{E}_{stb}} \cup \Sigma_{\mathcal{E}_{ss}} \cup \Sigma_{\mathcal{E}_{stg}}) \neq \emptyset$,

5. $\Sigma_{\mathcal{E}_{stg}} \setminus (\Sigma_{\mathcal{E}_{stb}} \cup \Sigma_{\mathcal{E}_{pr}} \cup \Sigma_{\mathcal{E}_{ss}}) \neq \emptyset$,

6. $\Sigma_{\mathcal{E}_{stb}} \setminus \Sigma_{\mathcal{E}_{pr}} \neq \emptyset$,

7. $\Sigma_{\mathcal{E}_{ss}} \setminus (\Sigma_{\mathcal{E}_{stb}} \cup \Sigma_{\mathcal{E}_{pr}} \cup \Sigma_{\mathcal{E}_{stg}}) \neq \emptyset$,

8. $\Sigma_{\mathcal{E}_{co}} \setminus \Sigma_{\mathcal{E}_{\sigma}} \neq \emptyset$ *and* $\Sigma_{\mathcal{E}_{\sigma}} \setminus \Sigma_{\mathcal{E}_{co}} \neq \emptyset$ *for* $\sigma \in \{cf, ad\}$,

9. $\Sigma_{\mathcal{E}_{cf}} \subset \Sigma_{\mathcal{E}_{ad}}$.

Finally, we briefly consider the closely related topic of *intertranslatability*. Intertranslatability revolves around the idea of mapping one semantics to another. One main motivation for studying this issue is the possibility to reuse a solver for one semantics for another [Dvořák and Woltran, 2011]. The main tool for this endeavour are functions mapping AFs to AFs, so-called *translations* formally defined as follows.

Definition 3.65 *[Dvořák and Woltran, 2011] Given two semantics* σ, τ. *A function* $f : \mathcal{F} \to \mathcal{F}$ *is called an exact translation for* $\sigma \to \tau$, *if* $\sigma(F) = \tau(f(F))$ *for any AF* F. *It is called a faithful translation if for any AF* F *first* $|\sigma(F)| = |\tau(f(F))|$ *and second* $\sigma(F) = \{S \cap A(F) \mid S \in \tau(f(F))\}$.

Please note that for some semantics there are no exact translations available due to reasons inherent to those semantics. For instance, preferred semantics satisfies *I-maximality*, i.e. for any AF F, $\mathcal{E}_{pr}(F)$ forms a $\subseteq$-antichain [Baroni and Giacomin, 2007]. This implies that an exact translation $\mathcal{E}_{ad} \to \mathcal{E}_{pr}$ can not exist since for $F = (\{a\}, \emptyset)$ we observe $\{\emptyset, \{a\}\} = \mathcal{E}_{ad}(F)$. Sticking to faithful translations provides us with a positive answer if we consider finite AFs only [Spanring, 2012, Translation 3.1.85]. Interestingly, the considered translation does not serve in the general unrestricted case and interestingly, it was shown that a search for a suitable translation will never succeed (cf. [Baumann and Spanring, 2017, Example 6]).

The following theorem (a generalization of the finite version from [Dvořák and Spanring, 2016, Section 6.1]) establishes a close relation between realizability and intertranslatability as promised, namely: if τ is not less expressive than σ, then σ can be exactly translated to τ and vice versa.

Theorem 3.66 ([Baumann and Spanring, 2017]) *Given semantics* σ, τ. *We have:* $\Sigma_\sigma \subseteq \Sigma_\tau$ *if and only if there is an exact translation for* $\sigma \to \tau$.

The following Figure 7 illustrates translational (im)possibilities in an eye-catching way. Figure 7b summarizes known results regarding faithful translations in the finite case [Dvořák and Woltran, 2011; Spanring, 2012; Dvořák and

Spanring, 2016], augmented with obvious insights for unique status semantics *il* and *eg*. For semantics σ, τ, encirclement in the same component indicates bidirectional translations. An arrow from σ to τ means directional translations. If there is no directed path (for instance for *na* to *cf*, or for *cf* to *gr*), then there is no translation. Figure 7a features the same visualization for unrestricted AFs. Dropping the finiteness restriction has some further consequences for the considered semantics, namely exact and faithful intertranslatability coincide. It is an open question whether both forms of translations are essentially the same in the general unrestricted setting. In consideration of Theorem 3.66 we may interpret Figure 7a as a comparison of the expressiveness of the considered semantics. That is, $\Sigma_\sigma \subset \Sigma_\tau$ if and only if there is a directed path from σ to τ.

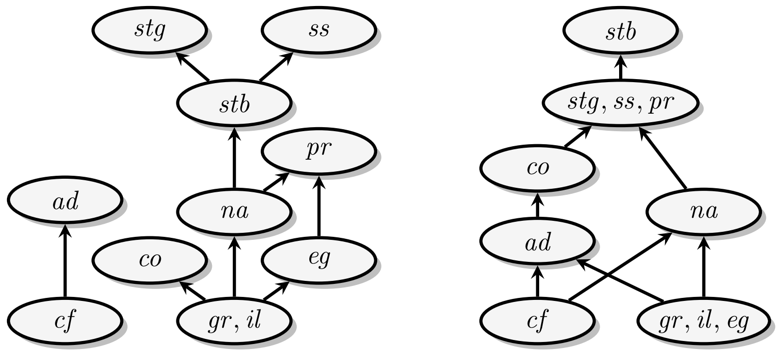

(a) Unrestricted AFs (faithful, exact) (b) Finite AFs (faithful)

Figure 7: Translational (Im)Possibilities

As a final note, in contrast to the unrestricted setting Baumann and Spanring observed that for slightly restricted AFs $F = (A, R)$, s.t. $|A| \leq |\mathcal{U} \setminus A|$ it is possible to provide exact and efficiently computable translations from preferred to semi-stable semantics via $f(F) = F' = (A', R')$ with $A' = A \cup \{a' \mid a \in A\}$ and $R' = R \cup \{(a, a'), (a', a') \mid a \in A\}$. It is an interesting question whether this restriction allows for similar translational possibilities as in case of finite AFs.

3.6 Realizability and Signatures for Labelling-based Versions

Although any considered semantics σ possesses an extension-based version (indicated by $\mathcal{E}_\sigma$) as well as a closely related 3-valued labelling-based version (indicated by $\mathcal{L}_\sigma$) we formally have that both versions are different semantics (or more precisely, functions) in the sense of Definition 2.2. This formal difference has some impact on realizability as well as signatures. Let us consider realizability in the realm of finiteness. As a matter of fact, for any considered 3-valued labelling-based version $\mathcal{L}_\sigma$ we have: if $F = (A, R)$ and $L = (L^{\mathrm{I}}, L^{\mathrm{O}}, L^{\mathrm{U}}) \in \mathcal{L}_\sigma(F)$, then $A = L^{\mathrm{I}} \cup L^{\mathrm{O}} \cup L^{\mathrm{U}}$. This means, σ-labellings assign a status to any argument in F. Now, in case of finite AFs we know that potentially realizable sets of labellings have to involve finitely many arguments only. Moreover, these finitely many arguments precisely determine the set A of witnessing AFs $F = (A, R)$.[13] Consider therefore the following example.

Example 3.67 *Consider the following set of 3-valued labellings* $\mathbb{S} = \{(\{a\}, \emptyset, \{b, c\}), (\{a, b\}, \{c\}, \emptyset)\}$. *Is* $\mathbb{S}$ *co-realizable? Since* $\{a\} \cup \emptyset \cup \{b, c\} = \{a, b, c\}$ *we deduce that candidates have to be members of* $\mathcal{C} = \{F = (A, R) \mid A = \{a, b, c\}\}$. *Note that* $|\mathcal{C}| = 2^{|\{a,b,c\}|^2} = 2^9 = 512$. *Clearly, this is a huge number, but it is a finite one. Consequently, the question of realizability can be decided by computing the* σ*-labellings of all AFs in* $\mathcal{C}$. *Of course, any intelligent search algorithm would involve further information like* $\{a, b\}$ *has to be conflict-free in a witnessing AF. Such an observation would decrease the number of candidates to* $2^5 = 32$. *However, in both cases one would find the unique witnessing framework* F, *i.e.* $\mathcal{L}_{co}(F) = \mathbb{S}$, *as depicted below.*

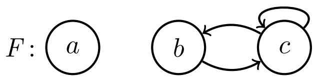

The example above shows that the search space can be very large even in case of small numbers of arguments. Consequently, locally verifiable necessary as well as sufficient properties for realizability just like in case of extension-based semantics are of high interest too. To the best of our knowledge only two papers have dealt with labelling-based realizability in the context of AFs. The first study was presented by Dyrkolbotn [Dyrkolbotn, 2014]. The author showed that, as long as additional arguments are allowed any finite set of labellings is realizable under preferred and semi-stable semantics. It is important to emphasize that Dyrkolbotn uses a more relaxed notion of realizibility, namely *realizibility under projection* (cf. Definition 3.72). The other work [Linsbichler *et al.*, 2016] deals with the standard notion of finite realizability (Definition 3.3). The authors presented an algorithm which returns either "No" in case of non-realizibility or a witnessing AF F in the positive case. The algorithm is not purely a guess-and-check method since it also includes a propagation step where certain necessary properties of witnessing AFs are processed. Remarkably, the algorithm is not restricted to the formalism of abstract argumentation frameworks only. In fact, it can also be used to decide realizability in case of the more general abstract dialectical frameworks as well as various of its sub-classes [Brewka and Woltran, 2010; Brewka *et al.*, 2013].

Preliminary Results for Labelling-based Signatures

In the following we shed light on general relations between the labelling-based and extension-based signatures of the considered semantics. Fortunately, due to former characterization results we will even achieve characterizing or at least necessary properties for finite realizability regarding labelling-based versions. We proceed with the definition of an *labelling-set* which is the n-valued analogon (for $n \geq 2$) to an extension-set as introduced in Definition 3.5. A labelling-set

[13] This is exactly the point which does not carry over to finite realizability in case of extension-based semantics (cf. statement 2 of Theorem 3.44).

is a finite set of n-tuples which are dealing with the same set of arguments and moreover, any n-tuple assigns exactly one status to each argument in question.

Definition 3.68 *Given* $\mathbb{S} \subseteq \left(2^{\mathcal{U}}\right)^n$. *Args*$_{\mathbb{S}}$ *denotes* $\bigcup_{L=(L_1,\ldots,L_n)\in\mathbb{S}} \bigcup_{i=1}^{n} L_i$ *and* $\|\mathbb{S}\|$ *stands for* $|Args_{\mathbb{S}}|$. *We say that* $\mathbb{S}$ *is a labelling-set if*

1. $\|\mathbb{S}\|$ *is a finite cardinal,*
2. *for any* $L = (L_1, \ldots, L_n) \in \mathbb{S}$, $Args_{\mathbb{S}} = \bigcup_{i=1}^{n} L_i$ *and*
3. *for any* $L = (L_1, \ldots, L_n) \in \mathbb{S}$, $L_1, \ldots, L_n$ *are pairwise disjoint.*

The following proposition establishes a connection between extension-based and labelling-based realizibility for any considered semantics. Roughly speaking, it states that labelling-based realizability requires extension-based realizability of the corresponding sets of in-labelled arguments. For a 3-tuple $L = (L_1, L_2, L_3)$ we also write $(L^{\text{I}}, L^{\text{O}}, L^{\text{U}})$ as usual.

Proposition 3.69 *Given a set of 3-tuples* $\mathbb{S} \subseteq \left(2^{\mathcal{U}}\right)^3$. *For any semantics* $\sigma \in \{stb, ss, stg, cf2, stg2, pr, ad, co, gr, il, eg, na, cf\}$ *we have,*

1. $\mathbb{S} \in \Sigma_{\mathcal{L}_\sigma} \Rightarrow \{L^{I} \mid L \in \mathbb{S}\} \in \Sigma_{\mathcal{E}_\sigma}$ *(unrestricted realizability)*
2. $\mathbb{S} \in \Sigma^{c}_{\mathcal{L}_\sigma} \Rightarrow \{L^{I} \mid L \in \mathbb{S}\} \in \Sigma^{c}_{\mathcal{E}_\sigma}$ *(compact realizability)*
3. $\mathbb{S} \in \Sigma^{x}_{\mathcal{L}_\sigma} \Rightarrow \{L^{I} \mid L \in \mathbb{S}\} \in \Sigma^{x}_{\mathcal{E}_\sigma}$ *(analytic realizability)*
4. $\mathbb{S} \in \Sigma^{f}_{\mathcal{L}_\sigma} \Rightarrow \{L^{I} \mid L \in \mathbb{S}\} \in \Sigma^{f}_{\mathcal{E}_\sigma}$ *(finite realizability)*
5. $\mathbb{S} \in \Sigma^{f,c}_{\mathcal{L}_\sigma} \Rightarrow \{L^{I} \mid L \in \mathbb{S}\} \in \Sigma^{f,c}_{\mathcal{E}_\sigma}$ *(finite, compact realizability)*
6. $\mathbb{S} \in \Sigma^{f,x}_{\mathcal{L}_\sigma} \Rightarrow \{L^{I} \mid L \in \mathbb{S}\} \in \Sigma^{f,x}_{\mathcal{E}_\sigma}$ *(finite, analytic realizability)*

Please note that the implications above are justified for any semantics σ whenever the different versions of it satisfy $\mathcal{E}_\sigma(F) = \{L^{\text{I}} \mid L \in \mathcal{L}_\sigma(F)\}$ for any relevant AF F. In the former sections we already presented characterization theorems or at least necessary properties for being finitely realizable regarding extension-based versions (cf. Theorems 3.12, 3.24 and 3.33). Combining these results with the proposition above yields the following necessary properties for finite realizability in the labelling-based case. Note that the mentioned implications apply to finite, compact as well as finite, analytic signatures too since $\Sigma^{f,c}_{\mathcal{L}_\sigma} \subseteq \Sigma^{f}_{\mathcal{L}_\sigma}$ as well as $\Sigma^{f,x}_{\mathcal{L}_\sigma} \subseteq \Sigma^{f}_{\mathcal{L}_\sigma}$ by definition. In case of grounded, ideal and eager semantics we have that being an one-element labelling-set is necessary and even sufficient for being finitely realizable. One may easily verify that the only-if-directions of these semantics are justified by the witnessing framework $F_L = (L^{\text{I}} \cup L^{\text{O}} \cup L^{\text{U}}, \{(i, o) \mid i \in L^{\text{I}}, o \in L^{\text{O}}\} \cup \{(u, u) \mid u \in L^{\text{U}}\})$ given that $\mathbb{S} = \{L\}$.

Theorem 3.70 *Given a set of 3-tuples* $\mathbb{S} \subseteq \left(2^{\mathcal{U}}\right)^3$*, then*

1. $\mathbb{S} \in \Sigma^f_{\mathcal{L}_{cf}} \Rightarrow \{L^I \mid L \in \mathbb{S}\}$ *is a non-empty, downward-closed and tight extension-set,*
2. $\mathbb{S} \in \Sigma^f_{\mathcal{L}_{na}} \Rightarrow \{L^I \mid L \in \mathbb{S}\}$ *is a non-empty, incomparable extension-set and* $dcl(\mathbb{S})$ *is tight,*
3. $\mathbb{S} \in \Sigma^f_{\mathcal{L}_{gr}} \Leftrightarrow \mathbb{S}$ *is a labelling-set with* $|\mathbb{S}| = 1$,
4. $\mathbb{S} \in \Sigma^f_{\mathcal{L}_{il}} \Leftrightarrow \mathbb{S}$ *is a labelling-set with* $|\mathbb{S}| = 1$,
5. $\mathbb{S} \in \Sigma^f_{\mathcal{L}_{eg}} \Leftrightarrow \mathbb{S}$ *is a labelling-set with* $|\mathbb{S}| = 1$,
6. $\mathbb{S} \in \Sigma^f_{\mathcal{L}_{stb}} \Rightarrow \{L^I \mid L \in \mathbb{S}\}$ *is a incomparable and tight extension-set,*
7. $\mathbb{S} \in \Sigma^f_{\mathcal{L}_{stg}} \Rightarrow \{L^I \mid L \in \mathbb{S}\}$ *is a non-empty, incomparable and tight extension-set,*
8. $\mathbb{S} \in \Sigma^f_{\mathcal{L}_{ad}} \Rightarrow \{L^I \mid L \in \mathbb{S}\}$ *is a conflict-sensitive ext.-set containing* $\emptyset$,
9. $\mathbb{S} \in \Sigma^f_{\mathcal{L}_{pr}} \Rightarrow \{L^I \mid L \in \mathbb{S}\}$ *is a non-empty, incomparable and conflict-sensitive extension-set,*
10. $\mathbb{S} \in \Sigma^f_{\mathcal{L}_{ss}} \Rightarrow \{L^I \mid L \in \mathbb{S}\}$ *is a non-empty, incomparable and conflict-sensitive extension-set.*

Realizibility under Projection

We turn now to *realizability under projection* which was firstly considered in [Dyrkolbotn, 2014]. In order to realize a set of labellings $\mathbb{S}$ under projection it suffices to come up with an AF F, s.t. its set of labellings restricted to the relevant arguments coincide with $\mathbb{S}$. Consider therefore the following illustrating example.

Example 3.71 *Given* $\mathbb{S} = \{(\{a\}, \{b\}, \emptyset), (\{b\}, \{a\}, \emptyset), (\emptyset, \{a, b\}, \emptyset)\}$. *We observe that the corresponding set of sets of in-labelled arguments* $\mathbb{S}^I = \{\emptyset, \{a\}, \{b\}\}$ *violates incomparability. Thus, applying statement 9 of Theorem 3.70 we derive that* $\mathbb{S}$ *is not finitely* pr*-realizable. Consider now the following AF* F.

F :
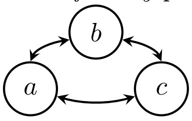

We obtain $\mathcal{L}_{pr}(F) = \{(\{a\}, \{b, c\}, \emptyset), (\{b\}, \{a, c\}, \emptyset), (\{c\}, \{a, b\}, \emptyset)\}$. *Now, if we restrict any labelling* $L = (L^I, L^O, L^U) \in \mathcal{L}_{pr}(F)$ *to the arguments* a *and* b, *i.e.* $L|_{\{a,b\}} = (L^I \cap \{a, b\}, L^O \cap \{a, b\}, L^U \cap \{a, b\})$ *we obtain exactly all labellings in* $\mathbb{S}$. *In this sense,* $\mathbb{S}$ *is* pr*-realizable under projection.*

We proceed with the formal definitions. For the sake of completeness we introduce realizability under projection and its corresponding signatures w.r.t. any kind of semantics as defined in Definition 2.2.

Definition 3.72 *Given a semantics* $\sigma : \mathcal{F} \to 2^{\left(2^{\mathcal{U}}\right)^n}$. *A set* $\mathbb{S} \subseteq \left(2^{\mathcal{U}}\right)^n$ *is* σ*-realizable under projection if there is an AF* F*, s.t.* $\sigma(F)|_{Args_{\mathbb{S}}} = \{E|_{Args_{\mathbb{S}}} \mid E \in \mathcal{E}_\sigma(F)\} = \mathbb{S}$ *(in case of* $n = 1$*) or* $\sigma(F)|_{Args_{\mathbb{S}}} = \{L|_{Args_{\mathbb{S}}} \mid L \in \mathcal{L}_\sigma(F)\} = \mathbb{S}$ *(for* $n \geq 2$*), respectively.*

Definition 3.73 *Given a semantics* σ*. The unrestricted as well as finite* σ*-projection-signatures are defined as follows:*

1. $\Sigma_\sigma^p = \{\sigma(F)|_B \mid F = (A, R) \in \mathcal{F}, B \subseteq A\}$ *and*
2. $\Sigma_\sigma^{f,p} = \{\sigma(F)|_B \mid F = (A, R) \in \mathcal{F}, F$ *is finite,* $B \subseteq A\}$

Analogously to Proposition 3.69 we state the following relation between labelling-based and extension-based versions of the considered semantics.

Proposition 3.74 *Given a set of 3-tuples* $\mathbb{S} \subseteq \left(2^{\mathcal{U}}\right)^3$*. For any semantics* $\sigma \in \{stb, ss, stg, cf2, stg2, pr, ad, co, gr, il, eg, na, cf\}$ *we have,*

1. $\mathbb{S} \in \Sigma_{\mathcal{L}_\sigma}^p \Rightarrow \{L^I \mid L \in \mathbb{S}\} \in \Sigma_{\mathcal{E}_\sigma}^p$ *(unrestr. realizability under projection)*
2. $\mathbb{S} \in \Sigma_{\mathcal{L}_\sigma}^{f,p} \Rightarrow \{L^I \mid L \in \mathbb{S}\} \in \Sigma_{\mathcal{E}_\sigma}^{f,p}$ *(finite realizability under projection)*

As a matter of fact, any projection signature is a superset of the corresponding signature. The following question then arises naturally: how much more sets can be generated if we stick to realizability under projection? For instance, we have already seen that even comparable sets are realizable under projection by semantics satisfying incomparability (Example 3.71). It was the main result in [Dyrkolbotn, 2014, Theorem 3.1] that in case of semi-stable and preferred semantics indeed any 3-valued labelling-set is finitely realizable under projection. The proof relies on two basic constructions. The first step *generates* an AF, consisting of so-called *circuits*, s.t. its set of preferred as well as semi-stable labellings restricted to the relevant arguments contains any possible labelling. The second construction *eliminates* undesired labellings step by step. Combining this realizability result with statement 2 of Proposition 3.74 yields the following theorem.

Theorem 3.75 *Let* $\sigma \in \{pr, ss\}$*. We have,*

1. $\Sigma_{\mathcal{L}_\sigma}^{f,p} = \left\{\mathbb{S} \subseteq \left(2^{\mathcal{U}}\right)^3 \mid \mathbb{S} \text{ is a labelling-set}\right\}$ *and*
2. $\Sigma_{\mathcal{E}_\sigma}^{f,p} = \left\{\mathbb{S} \subseteq 2^{\mathcal{U}} \mid \mathbb{S} \text{ is an extension-set}\right\}$.

3.7 Final Remarks and Conclusion

We have dealt with different forms of realizability in the context of abstract argumentation frameworks. In accordance with the existing literature the main part of this section was devoted to finite realizability for extension-based semantics. However, for any semantics σ we may state the following general subset relations depicted as Venn-diagram.

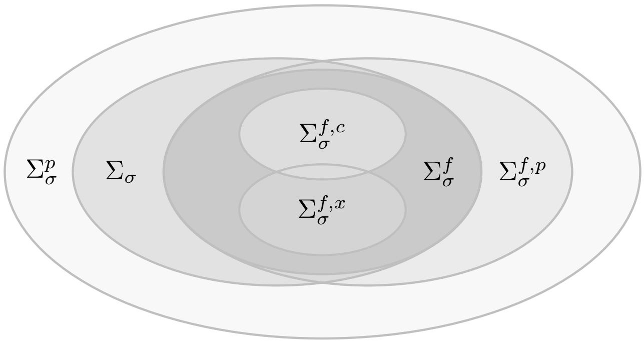

Figure 8: Subset Relations between Different Kinds of Signatures

In case of the extension-based versions of naive, grounded, ideal, eager, stable, stage, preferred and semi-stable semantics as well as conflict-free and admissible sets we provided exact characterizations for their corresponding general signatures. We have seen that for some semantics we do not lose any expressive power if sticking to compact or analytic AFs, i.e. $\Sigma_\sigma^f = \Sigma_\sigma^{f,c}$ or $\Sigma_\sigma^f = \Sigma_\sigma^{f,x}$, respectively. However, for certain prominent semantics, e.g. preferred semantics we have that the expressive power indeed relies on the use of rejected arguments or implicit conflicts. For such semantics, it remains an open problem to present exact characterizations for finite, compact or finite, analytic realizability, respectively. In case of labelling-based versions of semi-stable and preferred semantics we have seen that any labelling-set is realizable under projection. In [Dyrkolbotn, 2014] it was already noted that this equality does not hold for any semantics. For instance, the empty labelling is admissible for any AF F. Hence, in case of admissible semantics, no labelling-set is realizable under projection if it fails to include the empty labelling.

Finally, let us mention some computational issues not considered so far. It can be said that on the one hand, the classes of finite, compact and finite, analytic provide computational benefits both in practice and in terms of theoretical worst-case analysis. On the other hand testing for membership in one of the classes is, for most of the semantics, of rather high complexity and thus these classes cannot be directly used to improve systems. We refer the interested reader to [Baumann *et al.*, 2016a] for more details. Moreover, in general, given an extension-set $\mathbb{S}$, deciding whether $\mathbb{S}$ is compactly realizable is a hard problem, that is, by definition of the decision problem there are no good reasons to believe that we can do any better than guessing a compact AF and checking whether its extension-set coincides with $\mathbb{S}$. Nevertheless, for some semantics we have seen that finite, compact realizability can be characterized locally, i.e. by properties of $\mathbb{S}$ itself (as shown in Theorem 3.45). In this case, finite, compact realizability can be checked in polynomial time as for standard finite realizability [Dunne *et al.*, 2015, Theorem 6]. Moreover, in [Baumann *et al.*, 2016a] a huge number of shortcuts to detect non-compactness are provided. By shortcut we mean a property of the given extension-set $\mathbb{S}$ that is easily computable (preferably in polynomial time) which (sometimes) provides us with a definitive answer to the decision problem. These shortcuts are related to numerical aspects of argumentation frameworks like results concerning maximal number of extensions [Baumann and Strass, 2013].

4 Replaceability

Given a certain logical formalism $\mathcal{L}$ and two syntactically different $\mathcal{L}$-theories T_1 and T_2. One central question is whether, and if so, how to decide that these $\mathcal{L}$-theories represent the same information? Of course, in order to answer this question we have to clarify what we exactly mean by sharing the same information first. Note that there is neither a uniquely determined, nor a certain preferred interpretation by the formalism $\mathcal{L}$ itself. For instance, equating information with possessing the same semantics yields to the well-known notion of *ordinary* or *standard equivalence*. This means, assuming that $\sigma_{\mathcal{L}}$ is the semantics of $\mathcal{L}$ we might answer that T_1 and T_2 are equivalent if and only if $\sigma_{\mathcal{L}}(T_1) = \sigma_{\mathcal{L}}(T_2)$. A more demanding interpretation of sharing the same information is to require that T_1 and T_2 are semantically indistinguishable even if further $\mathcal{L}$-theories T are added to both simultaneously. More formally, we may state: T_1 and T_2 are considered to be equivalent if and only if $\sigma_{\mathcal{L}}(T_1 \cup T) = \sigma_{\mathcal{L}}(T_2 \cup T)$ for any theory T. This notion is known as *strong equivalence* and is of high interest for any logical formalism since it allows one to locally replace, and thus give rise for simplification, parts of a given theory without changing the semantics of the latter. In contrast to classical (monotone) logics where standard and strong equivalence coincide (cf. [Baumann and Strass, 2016] for more detailed information on this issue), it is possible to find ordinary but not strongly equivalent objects for any nonmonotonic formalism available in the literature. Consequently, much effort has been devoted to characterizing strong equivalence for nonmonotonic formalisms such as logic programs [Lifschitz *et al.*, 2001], causal theories [Turner, 2004], default logic [Turner, 2001] as well as nonmonotonic logics in general [Truszczynski, 2006].

In [Oikarinen and Woltran, 2011] the authors introduced the notion of strong equivalence for abstract AFs. They provided a series of characterization theorems for deciding strong equivalence of two AFs with respect to several semantics. In view of the fact that strong equivalence is defined semantically it is the main and quite surprisingly insight that being strongly equivalent can be decided syntactically. More precisely, they introduced the notion of a *kernel* of an AF F which is (informally speaking) a subgraph of F where certain attacks are deleted and showed that syntactical identity of suitably chosen kernels characterizes strong equivalence w.r.t. the considered semantics. Strong equivalence is, as its name suggests, a very (and often unnecessarily to) strong notion of equivalence if dynamic evolvements are considered. In many argumentation scenarios the type of modification which may potentially occur can be anticipated and furthermore, more importantly, does not range over *arbitrary expansions* as required for strong equivalence. Let us consider the instantiation-based context where AFs are built from an underlying knowledge base. Here, we typically observe that older arguments and their corresponding attacks survive and only new arguments which may interact with the previous ones arise given that a new piece of information is added to the underlying knowledge base. This type of dynamic evolvement is a so-called *normal expansion* and its corresponding equivalence notion were firstly studied in [Baumann, 2012a]. Over the last five years several equivalence notions taking into account specific types of evolvements reflecting the nature of various argumentation scenarios were defined and characterized. The considered dynamic scenarios range from the most general form, so-called *updates* [Baumann, 2014a] where arguments and attacks can be deleted and added to different types of *expansions* [Oikarinen and Woltran, 2011; Baumann, 2012a; Baumann and Brewka, 2010] and *deletions* [Baumann, 2014a] where arguments and/or attacks are allowed to be added or deleted in a certain way only.

Into the year 2015 all characterization theorems were stated in terms of extension-based semantics. Recently, Baumann presents their labelling-based counterparts and showed that, although labelling-based semantics contain more information then there extension-based counterpart, there is a majority of equivalence relations where labelling-based and extension-based versions coincide [Baumann, 2016]. Even more recently, a first consideration of strong equivalence regarding unrestricted frameworks were presented in [Baumann and Spanring, 2017]. It turned out that there are no characterizational differences compared to the finite case as long as the AFs in question are *jointly expandable*, i.e. that the existence of fresh arguments is guaranteed.

Another approach somehow complementary to the ones mentioned before is presented in [Baroni *et al.*, 2014] where sharing the same information is interpreted as possessing the same Input/Output behavior. Roughly speaking, the main idea is to consider an argumentation framework as a kind of black box which receives some input from the external world (i.e, a set of external arguments) via incoming attacks and produces an output to the external world via outgoing attacks. Such an interacting module is called an *argumentation multipole*. Two multipoles connected with the same external world are considered as *Input/Output equivalent* if the effects, i.e. the produced labellings for external arguments are the same for any reasonable input-labelling. This notion yields the possibility of replacing a multipole with another one embedded in a larger framework without affecting the labellings of the unmodified part of the initial framework. The interested reader is referred to Chapter 19 for further information. In the following we shed light on equivalence notions induced by certain dynamic scenarios.

4.1 Dynamic Scenarios and Corresponding Equivalence Notions

There are two main classes of dynamic scenarios, namely *expansions* and *deletions*. Both of them can be further divided in *normal* and *local* versions. These scenarios are motivated by real-world argumentation as well as instantiation-based argumentation [Caminada and Amgoud, 2007]. For instance, let us consider the dynamics of a discussion or dispute illustrated by the following citation [Besnard and Hunter, 2009]:

> How does argumentation usually take place? Argumentation starts when an initial argument is put forward, making some claim. An objection is raised, in the form of a counterargument. The latter is addressed in turn, eventually giving rise to a counter-counterargument, if any. And so on.

This means, in order to strengthen the own point of view or to rebut the opponents arguments it is natural that one tries to come up with *stronger* arguments, i.e. new arguments which are not attacked by the former arguments. This type of dynamics is formally captured by so-called *strong expansions* [Baumann and Brewka, 2010]. The formal counterpart of it, so-called *weak expansions* [Baumann and Brewka, 2010], where the new arguments do not attack (but may be attacked by) the old ones seem to be more an academic exercise than a task with practical relevance with regard to real-world argumentation.[14] Let us turn to instantiation-based argumentation where arguments and attacks stem from an underlying knowledge base (cf. Chapter 6 for detailed information as well as Figure 2 in Chapter 4 for an illustration). What happens on the abstract level if a new piece of information is added? It turns out that in almost all deductive argumentation systems older arguments and their corresponding attacks survive and only new arguments which may interact with the previous ones arise. This type of dynamic evolvement is formally captured by so-called *normal expansions*. *Local expansions* in contrast, i.e. expansions where new attacks are added only correspond to re-instantiations if we change to a less restrictive notion of attack (cf. [Besnard and Hunter, 2001] for different attack notions).

We start with the definition of the different types of expansions together with some introducing examples.

Definition 4.1 ([Baumann and Brewka, 2010]) *An AF G is an expansion of AF $F = (A, R)$ (for short, $F \preceq_E G$) iff $G = (A \dot\cup B, R \dot\cup S)$ for some (maybe empty) sets B and S. An expansion is called*

[14] We mention that they do play a decisive role w.r.t. computational issues, so-called *splitting methods* (cf. [Baumann, 2011; Baumann *et al.*, 2011; Baumann *et al.*, 2012]).

1. *normal* ($F \preceq_N G$) *iff* $\forall ab\ ((a,b) \in S \rightarrow a \in B \vee b \in B)$,
2. *strong* ($\mathcal{F} \preceq_S G$) *iff* $\mathcal{F} \preceq_N G$ *and* $\forall ab\ ((a,b) \in S \rightarrow \neg(a \in A \wedge b \in B))$,
3. *weak* ($\mathcal{F} \preceq_W G$) *iff* $\mathcal{F} \preceq_N G$ *and* $\forall ab\ ((a,b) \in S \rightarrow \neg(a \in B \wedge b \in A))$,
4. *local* ($F \preceq_L G$) *iff* $B = \emptyset$.

For short, being a normal expansion means that new attacks must involve at least one new argument in contrast to local expansions where new attacks involve old arguments only. Moreover, strong and weak expansions are normal and their names refer to properties of the additional arguments, namely arguments which are never attacked by former arguments (so-called *strong* arguments) and arguments which do not attack former arguments (so-called *weak* arguments).

Observe that any arbitrary expansion can be splitted up in a normal and a local part. This can be nicely seen in the following example.

Example 4.2 *The AF F is the initial framework. An arbitrary, normal, strong, weak or local expansion of it are given by F_E, F_N, F_S, F_W and F_L, respectively. Grey-highlighted arguments or attacks represent added information.*

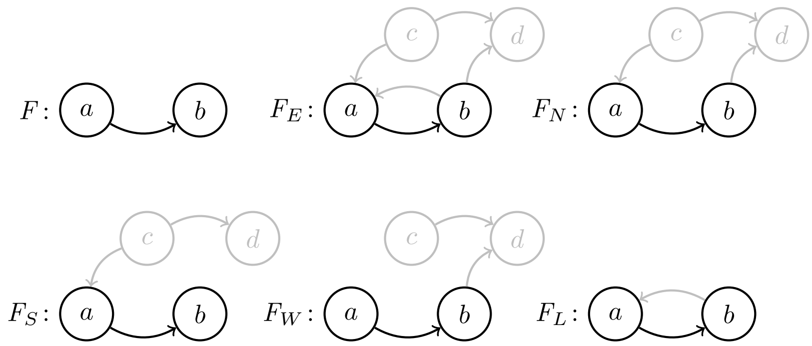

Figure 9: Different Kinds of Expansions

In 2014 the natural counter-parts (or more precisely, inverse operations) to arbitrary, normal and local expansions, so-called *deletions* were introduced [Baumann, 2014a]. Furthermore, the most general form of a dynamic scenario (where expansion and deletion can be combined) a so-called *update* were considered too. Analogously to expansions, any arbitrary deletion can be splitted in a normal and a local part. This means, a *normal deletion* retract arguments and their corresponding attacks. *Local deletions* in contrast delete attacks only.[15]

[15] We mention that *strong* as well as *weak deletions* are not introduced/considered so far. They could be easily defined as inverse operations of their expansion counterparts. Before doing so, it would be interesting to identify real-world situations or instantiation-based dynamics were such kind of evolvements naturally occur.

The main motivation behind these notions stems from instantiation-based context. More precisely, a normal deletion on the abstract level correspond to deleting information of a given knowledge base. Changing to a more restrictive notion of attack correspond to a local deletion and a combination of both of them give rise to an arbitrary deletion on the abstract level. We proceed with the formal definitions as well as introductory examples.

Definition 4.3 ([Baumann, 2014a]) *Given an AF $F = (A, R)$, a set of arguments B and a set of attacks S as well as a further AF H. The AF*

$$G = (F \setminus [B, S]) \cup H := \big((A, R \setminus S)|_{A \setminus B}\big) \cup H$$

is called an update of F (for short, $F \asymp_U G$). An update is called a

1. *deletion ($F \succeq_D G$) iff $H = (\emptyset, \emptyset)$,*
2. *normal deletion ($F \succeq_{ND} G$) iff ($F \succeq_D G$) and $S = \emptyset$,*
3. *local deletion ($F \succeq_{LD} G$) iff $F \succeq_D G$ and $B = \emptyset$.*

Let us take a closer look at the definition of $G = (F \setminus [B, S]) \cup H$. The AF H plays the role of added information, i.e. it contains new arguments and attacks. Consequently, for all kind of deletions we have $H = (\emptyset, \emptyset)$ which leaves us with $G = F \setminus [B, S]$. The set B contains arguments which have to deleted. Since attacks depend on arguments we have to delete the attacks which involve arguments from B too. This operation is formally captured by the restriction of F to $A \setminus B$. Furthermore, the set S contains particular attacks which have to be deleted. This means, the pair $[B, S]$ does not necessarily have to be an AF. Therefore we use $[B, S]$ instead of (B, S). If clear from context we use B and S instead of $[B, \emptyset]$ or $[\emptyset, S]$, i.e. we simply write $F \setminus B$ as well as $F \setminus S$ for normal or local deletions, respectively.

Example 4.4 *The AF F represents the initial situation. An update as well as arbitrary, normal or local deletion of it are given by F_U, F_D, F_{ND} and F_{LD}. Grey-highlighted arguments or attacks represent added information in contrast to dotted arguments and attacks which represent deleted objects.*[16] *More formally, in accordance with Definition 4.3 we have that $F_U = (F \setminus [B, S]) \cup H$, $F_D = F \setminus [B, S]$, $F_{ND} = F \setminus B$, $F_{LD} = F \setminus S$ where the set of arguments $B = \{c\}$, the set of attacks $S = \{(b, a)\}$ and the AF $H = (\{b, d, e, f\}, \{(d, b), (e, f), (f, d)\})$.*

[16] This convention will be used throughout the whole section.

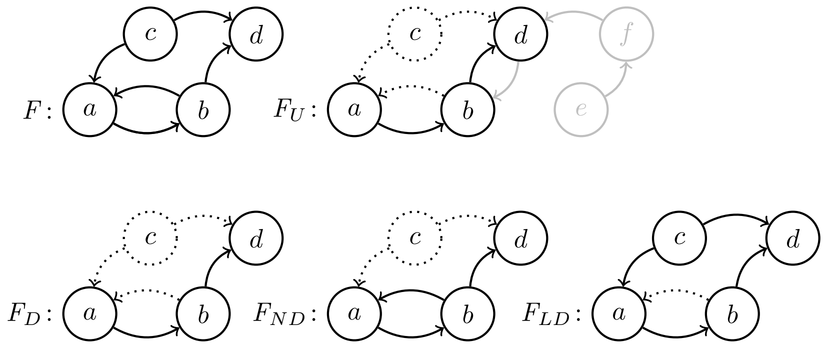

Figure 10: An Update and Different Kinds of Deletions

We now turn to the corresponding equivalence notions (cf. [Baumann and Strass, 2015, Section 3.8] for chronological order). Two AFs F and G are said to be *ordinarily equivalent* w.r.t. a semantics σ if they possess the same σ-extensions/labellings. In this case, we say that F and G possess the same *explicit* information. In contrast, sharing the same *implicit* information, i.e. being semantically indistinguishable w.r.t. any suitable future scenario is a much more demanding property which allows to replace F and G by each other without loss of semantical information.

Example 4.5 *Consider the following AFs F and G. We have $\mathcal{E}_{pr}(F) = \mathcal{E}_{pr}(G) = \{\{a\}\}$. This means, F and G possess the same explicit information w.r.t. preferred semantics or in other words, they are ordinarily equivalent.*

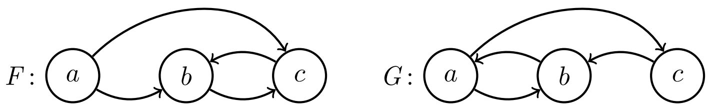

Assume that expansions as well deletions are the dynamic scenarios of interest. This means, we ask whether the AFs F and G even possess the same implicit information w.r.t. expansions or deletions, respectively? In order to give a negative answer one has to come up with one single dynamic scenario were the revised versions possess different preferred extensions. A positive answer in contrast is a statement about infinitely many dynamic scenarios (even in case of finite AFs). In this example, we give a negative answer for both modification types.

In case of expansions, we conjoin to both the AF $H = (\{a, b\}, \{(b, a)\})$. Consider the resulting frameworks below. We have $\mathcal{E}_{pr}(F \cup H) = \{\{a\}, \{b\}\}$ and since $G \cup H = G$ we obtain $\mathcal{E}_{pr}(G \cup H) = \{\{a\}\}$ without re-computing.

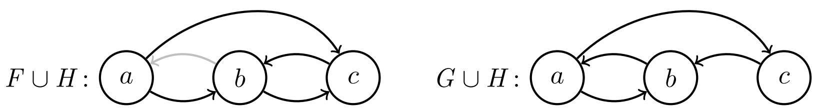

To reveal the inherent difference between F and G in case of deletions we may retract with the argument c. Consider the resulting (normal) deletions $F \setminus \{c\}$ and $G \setminus \{c\}$ of F or G, respectively. Now, $\{b\}$ becomes a preferred extension in $F \setminus \{c\}$ but still not in $G \setminus \{c\}$.

We now formally define what we precisely mean by possessing the same implicit information. As already stated, the first paper in this line of work was [Oikarinen and Woltran, 2011] engaged with characterizing *strong equivalence*. For the sake of clarity and comprehensibility we use the term *expansion equivalence* since strong equivalence [Oikarinen and Woltran, 2011, Definition 2] corresponds to semantical indistinguishability w.r.t. arbitrary expansions.

Definition 4.6 *Given a semantics σ. Two AFs F and G are*

1. *ordinarily equivalent w.r.t. σ (for short, $F \equiv^\sigma G$) iff $\sigma(F) = \sigma(G)$,*
2. *expansion equivalent w.r.t. σ (for short, $F \equiv^\sigma_E G$) iff for each AF H we have, $F \cup H \equiv^\sigma G \cup H$,*
3. *normal expansion equivalent w.r.t. σ (for short, $F \equiv^\sigma_N G$) iff for each AF H, such that $F \preceq_N F \cup H$ and $G \preceq_N G \cup H$ we have, $F \cup H \equiv^\sigma G \cup H$,*
4. *strong expansion equivalent w.r.t. σ (for short, $F \equiv^\sigma_S G$) iff for each AF H, such that $F \preceq_S F \cup H$ and $G \preceq_S G \cup H$ we have, $F \cup H \equiv^\sigma G \cup H$,*
5. *weak expansion equivalent w.r.t. σ (for short, $F \equiv^\sigma_W G$) iff for each AF H, such that $F \preceq_W F \cup H$ and $G \preceq_W G \cup H$ we have, $F \cup H \equiv^\sigma G \cup H$,*
6. *local expansion equivalent*[17] *w.r.t. σ (for short, $F \equiv^\sigma_L G$) iff for each AF H, such that $A(H) \subseteq A(F \cup G)$ we have, $F \cup H \equiv^\sigma G \cup H$.*
7. *update equivalent w.r.t. σ (for short, $F \equiv^\sigma_U G$) iff for any pair $[B, S]$ and any AF H we have, $(F \setminus [B, S]) \cup H \equiv^\sigma (G \setminus [B, S]) \cup H$,*
8. *deletion equivalent w.r.t. σ (for short, $F \equiv^\sigma_D G$) iff for any pair $[B, S]$ we have, $F \setminus [B, S] \equiv^\sigma G \setminus [B, S]$,*
9. *normal deletion equivalent w.r.t. σ (for short, $F \equiv^\sigma_{ND} G$) iff for any set of arguments B we have, $F \setminus B \equiv^\sigma G \setminus B$,*
10. *local deletion equivalent w.r.t. σ (for short, $F \equiv^\sigma_{LD} G$) iff for any set of attacks S we have, $F \setminus S \equiv^\sigma G \setminus S$,*

[17] Note that a suitable AF H is not necessarily a local expansion of F and G in the sense of Definition 4.1. Nevertheless, we may loosely speak about local expansions.

Remember that there are several relations between the considered dynamic scenarios. For instance, in accordance with Definitions 4.1 and 4.3, any normal expansion (deletion) is an arbitrary expansion (deletion). Furthermore, in the light of Definition 4.6, we certainly affirm that expansion equivalence is much more demanding then local expansion equivalence. In other words, local expansion equivalence of two AFs is an immediate and unavoidable consequence of being expansion equivalent. Finally, any considered equivalence notion is at least as demanding then ordinary equivalence.[18] Please note that these relations do not depend on certain properties of a considered semantics. Consequently, Figure 11 gives a preliminary overview for such interrelations (arising from the definitions) between the introduced equivalence notions for any possible semantics. For reasons, which will become clearer later, we also consider the identity relation. For two equivalence notion Φ and Ψ we have $\Phi \subseteq \Psi$ iff there is a link from Φ to Ψ.

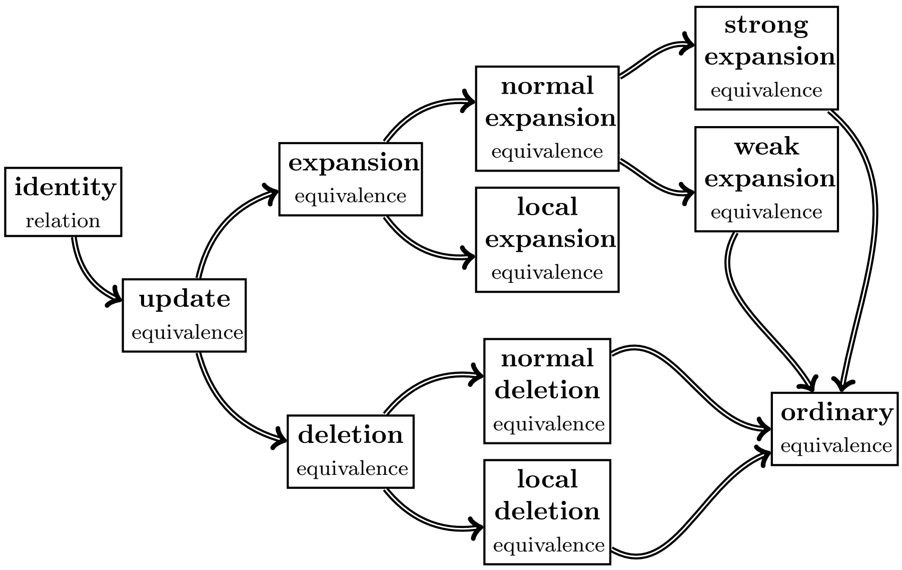

Figure 11: Preliminary Subset Relations between Equivalence Notions

In the remainder of this section we shed light on the question of *how* to determine whether two AFs are equivalent w.r.t. certain scenarios? As a by-product of these characterization results we will see that for many semantics the preliminary relations between the introduced equivalence notions depicted above can be delineated in a much more compact way. The majority of the presented characterization results is devoted to finite AFs as well as extension-

[18]The empty framework $(\emptyset, \emptyset)$ as well as the empty pair $[\emptyset, \emptyset]$ justifies this assertion for any type of expansions or deletions, respectively.

based semantics. We will see that there are some differences if sticking to unrestricted frameworks or the corresponding labelling-based versions.

4.2 Characterization Theorems for Extension-based Semantics

The Central Notion of Expansion Equivalence

In order to get an idea of how to find a characterization we start with some reflections. For this purpose we consider the most restrictive semantics, namely the stable one as well as the most prominent type of equivalence, namely expansion equivalence. What are necessary features of expansion equivalence w.r.t. stable semantics, i.e. which properties are implied if two AFs F and G are expansion equivalent? In consideration of Figure 11 we deduce their ordinary equivalence, i.e. $\mathcal{E}_{stb}(F) = \mathcal{E}_{stb}(G)$. Note that possessing the same set of extensions neither imply sharing the same arguments nor sharing the same self-loops as shown in the following example.

Example 4.7 *Consider the AFs F, G and H. Each two of them are ordinarily equivalent since $\mathcal{E}_{stb}(F) = \mathcal{E}_{stb}(G) = \mathcal{E}_{stb}(H) = \{\{a\}\}$.*

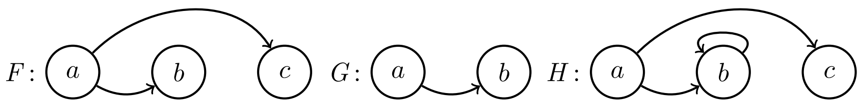

The AFs $I_1 = (\{c\}, \emptyset)$ and $I_2 = (\{a, b, c\}, \{(b, a), (b, c)\})$ witness that neither F and G, nor F and H are expansion equivalent w.r.t. stable semantics. Convince yourself that $\mathcal{E}_{stb}(F \cup I_1) = \{\{a\}\} \neq \{\{a, c\}\} = \mathcal{E}_{stb}(G \cup I_1)$ and $\mathcal{E}_{stb}(F \cup I_2) = \{\{a\}, \{b\}\} \neq \{\{a\}\} = \mathcal{E}_{stb}(G \cup I_2)$.

Restricting ourselves to finite AFs, it is not difficult to see that in case of expansion equivalence w.r.t. stable semantics the observed relation between non-sharing the same arguments/loops and non-equivalence does hold in general. In other words, possessing the same arguments as well as possessing the same loops are indeed necessary conditions for being expansion equivalent in the finite setting.

Let us summarize our observations in the following fact.

Fact 4.8 *Given two finite AFs F and G. If $F \equiv_E^{\mathcal{E}_{stb}} G$, then*

1. $\mathcal{E}_{stb}(F) = \mathcal{E}_{stb}(G)$,
2. $A(F) = A(G)$ *and*
3. $L(F) = L(G)$.

As already stated in Figure 11, being identical (i.e. $A(F) = A(G)$ and $R(F) = R(G)$) is sufficient for being expansion equivalent. Combining this undeniable fact together with the second and third items of Fact 4.8 encourages one to search for syntactical properties sufficient as well as necessary for being expansion equivalent. In order to guarantee the first item of Fact 4.8

we have to identify attacks which do not contribute anything when computing stable extensions. Moreover, these attacks which do not affect the evaluation of a given AF F have to be *redundant*, no matter how F is extended. Remember that being a stable extension can be simply verified by checking whether the set in question is conflict-free and possesses a full range.[19] This means, good candidates for "useless" attacks w.r.t. stable semantics should fulfill the following two properties: firstly, having or not having such an attack does not change the status of a set from being conflict-free to conflicting or vice versa and secondly, having or not having such an attack does not affect the range of a conflict-free set. Certainly, an attack (a, b) stemming from a self-defeating argument a does not change the conflict status of a certain set E. This can be seen as follows: If $a \in E$, then E was conflicting as well as remains conflicting after deleting or adding (a, b). Furthermore, if $a \notin E$, then E might be conflicting or not. In either case the conflict status of E does not change if (a,b) is added or removed since $\{a, b\} \not\subseteq E$. Finally, such an attack (a, b) might have an influence on the range of conflicting sets but it definitely has not in case of conflict-free sets since $a \notin E$ can not be questioned.

Example 4.9 *Consider the following AF F. We have, $\mathcal{E}_{stb}(F) = \{\{a\}\}$.*

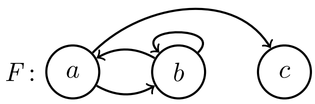

According to our considerations above adding or deleting an attack stemming from the self-defeating argument b does not change the semantics. Consider therefore the following three possible "manipulations".

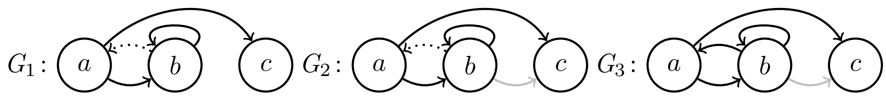

Indeed, $\mathcal{E}_{stb}(F) = \mathcal{E}_{stb}(G_1) = \mathcal{E}_{stb}(G_2) = \mathcal{E}_{stb}(G_3) = \{\{a\}\}$ support our claims for the static case. We encourage the reader to try to do the impossible, namely semantically distinguish the AFs F and its manipulations by an arbitrary expansion.

It was the main result in [Oikarinen and Woltran, 2011] that expansion equivalence can be indeed decided by looking at the syntax only. The authors introduced so-called *kernels* which are simply functions mapping each AF F to its redundancy-free version. This means, the kernel of an AF F does not possess any redundant attack. Put it differently, for any surviving attack exist at least one dynamic scenario were deleting this attack would cause a semantical difference. We proceed with the formal definition of the very first kernels

[19] The topic of *verifiability* of argumentation semantics σ was studied in [Baumann *et al.*, 2016b]. The main question is which (minimal amount of) information on top of conflict-free sets is exactly needed to determine whether a certain set is a σ-extension.

already introduced in [Oikarinen and Woltran, 2011]. We sometimes call them *classical*.

Definition 4.10 *Let $\sigma \in \{stb, ad, gr, co\}$. The σ-kernel $k(\sigma) : \mathcal{F} \to \mathcal{F}$ with $k(\sigma)(F) = F^{k(\sigma)} = \big(A, R^{k(\sigma)}\big)$ for a given AF $F = (A, R)$ is defined as:*

$$\begin{aligned}
R^{k(stb)} &= R \setminus \{(a,b) \mid a \neq b, (a,a) \in R\},\\
R^{k(ad)} &= R \setminus \{(a,b) \mid a \neq b, (a,a) \in R, \{(b,a),(b,b)\} \cap R \neq \emptyset\},\\
R^{k(gr)} &= R \setminus \{(a,b) \mid a \neq b, (b,b) \in R, \{(a,a),(b,a)\} \cap R \neq \emptyset\},\\
R^{k(co)} &= R \setminus \{(a,b) \mid a \neq b, (a,a),(b,b) \in R\}.
\end{aligned}$$

In order to get an idea of how the classical kernels work we proceed with an example.

Example 4.11 *Consider again the AF G_3 depicted in Example 4.9. We apply now all classical kernels.*

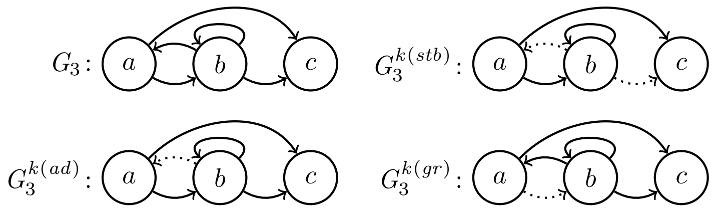

The stable kernel deletes all attacks (a,b) stemming from a self-defeating argument a. A deletion of (a,b) in case of the grounded kernel additionally requires that a is counter-attacked by b or b is self-defeating or both. Interchanging a and b yields the condition for deletion in case of the grounded kernel. Finally, $G_3^{k(co)} = G_3$ since deleting an attack (a,b) w.r.t. the complete kernel requires that both arguments a and b are self-defeating.

Before turning to characterization theorems, we collect some useful properties of the introduced kernels. The following fact contains intrinsic properties of the classical kernels.[20] More precisely, any classical kernel k is *node-preserving* and *loop-preserving*, i.e. the sets of arguments and self-defeating arguments do not change if applying k. Moreover, in the absence of self-loops, each AF coincides with its classical kernels. Furthermore, the decision whether an attack (a,b) has to be deleted does not depend on further arguments than a and b. Put differently, the reason of being redundant is *context-free*, i.e. it stems from the arguments themselves. The last two properties claim that equality of kernels is *robust* w.r.t. further compositions as well as deleting arguments and corresponding attacks. For a given AF $F = (B, S)$ we use $A(F), R(F)$ and $L(F)$

[20] Although most of the properties are immediately clear even in case of unrestricted frameworks we will state all of them for finite AFs only as done in the existing literature. The same applies to Fact 4.18. Some results regarding unrestricted frameworks can be found in Section 4.3.

to refer to its arguments, attacks and self-defeating arguments, i.e. $A(F) = B$, $R(F) = S$ and $L(F) = \{a \in A(F) \mid (a,a) \in R(F)\}$.

Fact 4.12 (cf. [Oikarinen and Woltran, 2011; Baumann, 2014a]) *Given $k \in \{k(stb), k(ad), k(gr), k(co)\}$. For any finite AF F we have:*

1. $A(F) = A\left(F^k\right)$, *(node-preserving)*
2. $L(F) = L\left(F^k\right)$, *(loop-preserving)*
3. $L(F) = \emptyset \;\Rightarrow\; F = F^k$ *and* *(sufficient condition for identity)*
4. $(a,b) \in R\left(F^k\right) \Leftrightarrow (a,b) \in R\left((F|_{\{a,b\}})^k\right)$. *(context-freeness)*

Furthermore, for finite AFs F and G we have:

4. *If $F^k = G^k$, then $(F \cup H)^k = (G \cup H)^k$ for any finite AF H and* *($\cup$-robustness)*
5. *If $F^k = G^k$, then $(F \setminus B)^k = (G \setminus B)^k$ for any finite set of args B.* *($\setminus$-robustness)*

We proceed with extrinsic properties, i.e. features of kernels in presence of semantics. More precisely, stable, admissible, grounded and complete semantics are insensitive w.r.t. the application of their corresponding classical σ-kernel, i.e. the set of σ-extensions remains unchanged. Furthermore, the admissible kernel neither effects semi-stable, eager, preferred and ideal semantics. Similarly in case of stable kernel and stage semantics.

Fact 4.13 ([Oikarinen and Woltran, 2011; Gaggl and Woltran, 2013]) *For any finite AF F we have:*

1. $\mathcal{E}_\sigma(F) = \mathcal{E}_\sigma\left(F^{k(\sigma)}\right)$ *for $\sigma \in \{stb, ad, gr, co\}$,*
2. $\mathcal{E}_\sigma(F) = \mathcal{E}_\sigma\left(F^{k(ad)}\right)$ *for $\sigma \in \{ss, eg, pr, il\}$ and*
3. $\mathcal{E}_{stg}(F) = \mathcal{E}_{stg}\left(F^{k(stb)}\right)$.

As already mentioned, kernels play a decisive role in deciding expansion equivalence. In general, we say that an equivalence notion $\equiv$ *is characterizable through k* or simply, *k is a characterizing kernel (of $\equiv$)* if for any two AFs F and G, $F \equiv G$ iff $F^k = G^k$. This means, proving whether two frameworks are equivalent can be done by simply checking whether the corresponding kernels are identical. Note that all classical kernels can be efficiently constructed from a given AF. The following main theorem states that for all nine considered semantics σ there is a certain classical kernel k, s.t. expansion equivalence w.r.t. σ is characterizable through k in the finite setting. This is a very remarkable result since expansion equivalence is defined semantically. For instance, two finite

AFs F and G are expansion equivalent w.r.t. stable semantics if and only if the associated stable kernels $F^{k(stb)}$ and $G^{k(stb)}$ are syntactically equal. Observe that there is no need to introduce further kernels since one single kernel may serve for different semantics.

Theorem 4.14 *[Oikarinen and Woltran, 2011; Gaggl and Woltran, 2013] For finite AFs F and G we have:*

1. $F \equiv_E^{\mathcal{E}_\sigma} G \Leftrightarrow F^{k(\sigma)} = G^{k(\sigma)}$ *for any* $\sigma \in \{stb, ad, co, gr\}$,
2. $F \equiv_E^{\mathcal{E}_\sigma} G \Leftrightarrow F^{k(ad)} = G^{k(ad)}$ *for any* $\sigma \in \{pr, il, ss, eg\}$ *and*
3. $F \equiv_E^{\mathcal{E}_{stg}} G \Leftrightarrow F^{k(stb)} = G^{k(stb)}$.

Having Theorem 4.14 at hand we can now formally verify that all AFs depicted in Example 4.9 are expansion equivalent w.r.t. stable semantics. This means, the recommended search for arbitrary expansions revealing semantical difference between them will never succeed. As an aside, one might get the impression that the syntactical characterization presented in Theorem 4.14 is somehow unique. This is not true. Consider therefore the equivalence class $[F]_E^{\mathcal{E}_{stb}} = \{G \mid F \equiv_E^{\mathcal{E}_{stb}} G\}$ induced by F. Mathematically speaking, the stable kernel $F^{k(stb)}$ represents the least (w.r.t. subgraph-relation) element in $[F]_E^{\mathcal{E}_{stb}}$. It is not difficult to prove that $[F]_E^{\mathcal{E}_{stb}}$ even possesses a greatest element, namely $F^{k'(stb)} = (A(F), R(F) \cup \{(a,b) \mid a \neq b, (a,a) \in R(F)\})$, i.e. the framework resulting from F by adding (instead of deleting) all redundant attacks. In case of finite AFs it can be shown with reasonable effort that expansion equivalence w.r.t. stable semantics is characterizable through $k'(stb)$ too. In the same manner, all other semantics considered in Theorem 4.14 possess alternative "greatest elements" characterizations. We will see that the so-called *naive kernel* (compare Definition 4.17) provides such a kind of characterization for naive semantics. The reason for this "choice" is simply that the induced equivalence classes do not necessarily possess a least element in case of naive semantics.

Finally, let us turn to the more exotic cf2 as well as stage2 semantics which are defined via a recursive schema based on the decomposition of AFs along their strongly connected components (SCCs). These semantics are exceptional regarding expansion equivalence since in contrast to all other semantics considered in this section we have that even attacks between two self-attacking arguments are *meaningful.* This means, the presence or absence of such attacks may change the outcome of an AF. Moreover, it turned out that any attack is non-redundant. In summary, expansion equivalence coincides with syntactical identity or more formally, for any finite AF F, $\left|[F]_E^{\mathcal{E}_{cf2}}\right| = \left|[F]_E^{\mathcal{E}_{stg2}}\right| = |\{F\}| = 1$.

Theorem 4.15 *[Gaggl and Woltran, 2013; Gaggl and Dvořák, 2016] Given* $\sigma \in \{cf2, stg2\}$. *For finite AFs F and G we have,*

$$F \equiv_E^{\mathcal{E}_\sigma} G \Leftrightarrow F = G.$$

Further Equivalence Notions Characterizable through Kernels

Let us turn to the remaining equivalence notions? Are there similar syntax-based characterization results?

Weaker Notions of Expansion Equivalence Let us consider less demanding notions than expansion equivalence, e.g. normal and local expansion equivalence. In consideration of Definition 4.1 we do not have good reasons to believe that two AFs could be semantically distinguished by normal or local expansions, given that we only have a witnessing arbitrary expansion showing their non-equivalence. It was one surprising result in this line of research, that for many semantics, expansion equivalence coincide with definitorially weaker notions of it. This implies that weaker notions than expansion equivalence can be characterized by classical kernels too. The first results in this respect were already given in [Oikarinen and Woltran, 2011, Theorem 8]. The authors showed that for some semantics expansion equivalence and local expansion equivalence coincide if considering finite AFs. It is worthwhile to gain a thorough understanding of this relation since it actually means that if there is an arbitrary expansion which semantically distinguish two finite AFs, than there has to be a local expansion doing likewise. Later it was shown that even normal expansion equivalence coincides with expansion equivalence for a whole bunch of semantics [Baumann, 2012a]. Interestingly, in contrast to local expansion equivalence, there are (to the best of our knowledge) no semantics together with witnessing AFs known which show that this coincidence does not hold in general.

Example 4.16 *Consider the following AFs F and G. According to Theorem 4.14 they are not expansion equivalent w.r.t. preferred semantics since $F^{k(ad)} = F \neq G = G^{k(ad)}$.*

As already stated (up to now) normal expansion equivalence coincides with expansion equivalence for any considered semantics. One possible scenario which makes the predicted different behaviour explicit is the following.

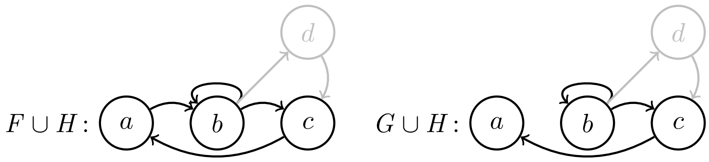

Formally, we define $H = (\{b, c, d\}, \{(b, d), (d, c)\})$ and we obtain $\{\{a, d\}\} = \mathcal{E}_{pr}(F \cup H) \neq \{\emptyset\} = \mathcal{E}_{pr}(G \cup H)$. We encourage the reader to try to find a witnessing example showing that F and G are not local expansion equivalent w.r.t. preferred semantics. Due to Theorem 4.20 there has to be at least one distinguishing local expansion.

How do the semantics behave in case of strong expansion equivalence? Remember, a special feature of strong expansions is that a former attack between old arguments will never become a counterattack to an added attack. In this sense, former attacks do not play a role with respect to being a potential defender of an added argument. Hence, in contrast to arbitrary expansions where such attacks might be relevant, we may delete them without changing the behavior with respect to further evaluations. To make this point clearer consider again the AF $F \cup H$ depicted in Example 4.16. Note that the already existing attack (a, b) in F becomes a defending attack of the newly added argument d. This means, such attacks in fact play an important role with respect to further evaluation in case of arbitrary expansions. It was one main result in [Baumann, 2012a] that for some semantics attacks like (a, b) in F are indeed redundant w.r.t. strong expansions. Even more surprising, strong expansion equivalence is characterizable through kernels. Therefore, more involved kernel definitions, so-called *σ-*-kernels* had to be introduced. These kernels allow more deletions than their classical counterparts for expansion equivalence. In contrast to them, σ-*-kernels are *context-sensitive*, i.e. the question whether an attack (a, b) is redundant can not be answered by considering the arguments a and b only [Baumann, 2014a].

The first three kernels presented in the definition below were firstly introduced in [Baumann, 2012a] with the objective to characterize strong expansion equivalence with respect to certain semantics. For the sake of completeness we also present the so-called *stg-*-kernel* as well as *na-kernel* [Baumann and Woltran, 2016; Baumann *et al.*, 2016b].[21]

Definition 4.17 *Let* $\sigma \in \{ad, gr, co, stg\}$*. The σ-*-kernel* $k^*(\sigma) : \mathcal{F} \to \mathcal{F}$ *with* $k^*(\sigma)(F) = F^{k^*(\sigma)} = \left(A, R^{k^*(\sigma)}\right)$ *for a given AF* $F = (A, R)$ *is defined as:*

$$
\begin{aligned}
R^{k^*(ad)} &= R \ \setminus \{(a,b) \mid a \neq b, ((a,a) \in R \wedge \{(b,a),(b,b)\} \cap \ R \neq \emptyset) \\
&\quad \vee ((b,b) \in R \wedge \forall c \ ((b,c) \in R \to \{(a,c),(c,a),(c,c),(c,b)\} \cap R \neq \emptyset))\}, \\
R^{k^*(gr)} &= R \setminus \{(a,b) \mid a \neq b, ((b,b) \in R \wedge \{(a,a),(b,a)\} \cap \ R \neq \emptyset) \\
&\quad \vee ((b,b) \in R \wedge \forall c \ ((b,c) \in R \to \{(a,c),(c,a),(c,c)\} \cap R \neq \emptyset))\}, \\
R^{k^*(co)} &= R \setminus \{(a,b) \mid a \neq b, ((a,a),(b,b) \in R) \vee ((b,b) \in R \wedge (b,a) \notin R \\
&\quad \wedge \forall c \ ((b,c) \in R \to \{(a,c),(c,a),(c,c),(c,b)\} \cap R \neq \emptyset))\}, \\
R^{k^*(stg)} &= R \ \setminus \{(a,b) \mid a \neq b, (a,a) \in R \vee \forall c \ (c \neq a \to (c,c) \in R)\} \\
R^{k(na)} &= R \cup \{(a,b) \mid a \neq b, \{(a,a),(b,a),(b,b)\} \cap R \neq \emptyset\}.
\end{aligned}
$$

The latter represents the so-called na-kernel $F^{k(na)} = \left(A, R^{k(na)}\right)$.

For an illustrating example we refer the reader to Example 4.19. Analogously to Fact 4.12 we collect some properties of the newly introduced kernels. The

[21] As an aside, we use the supplement "*", whenever the kernel in question is non-classical and expansion equivalence is already characterized by another kernel.

first three properties are immediately clear by definition.[22] The robustness w.r.t. deletions and corresponding attacks is less obvious but it is already shown for all considered kernels (except the *stg*-*-kernel) in case of finite AFs (cf. [Baumann, 2014a, Theorems 6 and 14]).

Fact 4.18 *Given* $k \in \{k^*(ad), k^*(gr), k^*(co), k^*(stg), k(na)\}$ *as well as* $k^* \in \{k^*(ad), k^*(gr), k^*(co), k^*(stg)\}$*. For two finite AFs* F *and* G *we have:*

1. $A(F) = A\left(F^k\right)$, *(node-preserving)*
2. $L(F) = L\left(F^k\right)$, *(loop-preserving)*
3. $L(F) = \emptyset \Rightarrow F = F^{k^*}$ *and* *(sufficient condition for identity)*
4. *If* $F^k = G^k$*, then* $(F \setminus B)^k = (G \setminus B)^k$ *for any finite set of arguments* B*.* *($\setminus$-robustness)*

Let us consider the *ad*-*-kernel (which, as we shall see, characterizes strong expansion equivalence for preferred semantics) in more detail. Consider the first disjunct. This first condition is exactly the same as in case of the *ad*-kernel (compare Definition 4.10), i.e. an attack (a, b) has to be deleted if a is self-attacking and at least one of the attacks (b, a) or (b, b) exist. The second disjunct provides one with further options to delete an attack (a, b), namely if b is self-defeating and furthermore, for all arguments c which are attacked by b at least one of the following conditions has to be fulfilled:

1. a attacks c,
2. c attacks a,
3. c attacks c,
4. c attacks b.

The motivation for the second disjunct is the following: At first observe that b cannot be an element of any conflict-free set. Consequently, in case of strong expansions the attack (a, b) may only be relevant with respect to the defense of c. In the first three cases this relevance becomes unimportant since $\{a, c\}$ is conflicting. In the fourth case the redundancy of (a, b) with respect to the defense of c is given by the fact that c already defends itself against b. Please note that the consideration of $c = a$ or $c = b$ is not excluded by Definition 4.17. The following frameworks exemplify different cases.

[22] The AF $F = (\{a, b\}, \{(a, b)\})$ shows that the naive kernel has to be excluded from item 3 of Fact 4.18 since $F^{k(na)} = (\{a, b\}, \{(a, b), (b, a)\}) \neq F$.

Example 4.19 *The following graphs show six frameworks and their corresponding ad-*-kernels. The dotted attacks represent initial attacks which have to be deleted if applying the ad-*-kernel.*

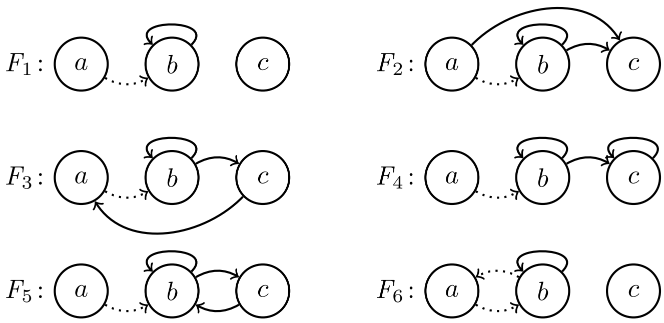

Consider the formal description of $R^{k^*(ad)}$ as given in Definition 4.17. The AF F_1 is somehow the base case since the only argument c, s.t. $(b,c) \in R(F_1)$ is b itself. Since $(b,b) \in R(F_1)$ we deduce that the considered intersection is non-empty and thus, the deletion of (a,b) is justified. The subsequent four frameworks F_2, F_3, F_4 and F_5 are the base case plus one further argument c different from a and b, s.t. for any $i \in \{2,3,4,5\}$, $(b,c) \in R(F_i)$. The last framework F_6 illustrates the case b counterattacks a. Note that the reason to delete (a,b) is somehow self-referential since (additionally to the base case) it is justified by $(a,b) \in R(F_6)$. Due to the first disjunct (i.e. just like in case of the classical ad-kernel) even the attack (b,a) has to be deleted.

We proceed with further characterization theorems.[23] An comprehensive overview of equivalence notion and their characterizing kernels in case of finite AFs and extension-based semantics is presented in Figure 12.

Theorem 4.20 *[Oikarinen and Woltran, 2011; Baumann, 2012a; Baumann and Woltran, 2016; Baumann* et al.*, 2016b] For finite AFs F and G we have the following coincidences.*

1. $F \equiv_E^{\mathcal{E}_\sigma} G \Leftrightarrow F \equiv_N^{\mathcal{E}_\sigma} G$ *for* $\sigma \in \{stg, stb, ss, eg, ad, pr, il, gr, co, na, cf2, stg2\}$,
2. $F \equiv_E^{\mathcal{E}_\sigma} G \Leftrightarrow F \equiv_L^{\mathcal{E}_\sigma} G$ *for* $\sigma \in \{ss, eg, ad, pr, il, na\}$ *and*
3. $F \equiv_E^{\mathcal{E}_\sigma} G \Leftrightarrow F \equiv_S^{\mathcal{E}_\sigma} G$ *for* $\sigma \in \{stg, stb, ss, eg, na\}$.

Furthermore, for any two finite AFs F and G we have the following non-classical characterizations.

4. $F \equiv_L^{\mathcal{E}_{stg}} G \Leftrightarrow F^{k^*(stg)} = G^{k^*(stg)}$,
5. $F \equiv_S^{\mathcal{E}_\sigma} G \Leftrightarrow F^{k^*(ad)} = G^{k^*(ad)}$ *for* $\sigma \in \{ad, pr, il\}$,

[23] Please note that the results in case of cf2 and stage2 semantics have never been published before.

6. $F \equiv_S^{\mathcal{E}_\sigma} G \Leftrightarrow F^{k^*(\sigma)} = G^{k^*(\sigma)}$ *for* $\sigma \in \{co, gr\}$ *and*

7. $F \equiv_E^{\mathcal{E}_{na}} G \Leftrightarrow F^{k(na)} = G^{k(na)}$.

At this point we want to highlight a very surprising relation. Remember that normal expansion equivalence and normal deletion equivalence are completely unrelated in the general picture (cf. Figure 11). The observation that the characterizing kernels (including the identity map in case of cf2 and stage2 semantics) of normal expansion equivalence w.r.t. all considered semantics in this section satisfy \-robustness (cf. Facts 4.12 and 4.18) reveals that normal expansion equivalence implies normal deletion equivalence for these semantics.

Corollary 4.21 *Given* $\sigma \in \{stg, stb, ss, eg, ad, pr, il, gr, co, na, cf2, stg2\}$ *and two finite AFs F and G. We have:* $F \equiv_N^{\mathcal{E}_\sigma} G \Rightarrow F \equiv_{ND}^{\mathcal{E}_\sigma} G$.

The attentive reader may have noticed that we do not have characterized local expansion equivalence w.r.t. stable, complete as well as grounded extension-based semantics. We mention that all three equivalence notions are already characterized but the characterization theorems are not purely kernel-based (cf. [Oikarinen and Woltran, 2011, Theorems 9,10,11]). Furthermore, it can be checked that none of the kernels presented in Definitions 4.10 and 4.17 serve as a characterizing kernel. Consider therefore the following example [Oikarinen and Woltran, 2011, Example 15].

Example 4.22 *The AFs F and G are local expansion equivalent w.r.t. stable semantics. This can be seen as follows. Given an AF H, s.t.* $A(H) \subseteq \{a, b\}$*. If* $(a, b) \in R(H)$ *and* $(a, a) \notin R(H)$*, we obtain* $\mathcal{E}_{stb}(F \cup H) = \mathcal{E}_{stb}(G \cup H) = \{\{a\}\}$*. Otherwise,* $\mathcal{E}_{stb}(F \cup H) = \mathcal{E}_{stb}(G \cup H) = \emptyset$.

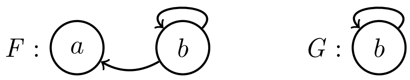

Remember that all introduced kernels are node-preserving (Facts 4.12 and 4.18). Consequently, none of them may serve as a characterizing kernel for local expansion equivalence w.r.t. stable semantics.

We mention that weak expansion equivalence is already characterized in case of stable semantics [Baumann, 2012a, Proposition 3] as well as admissible, preferred and complete semantics [Baumann and Brewka, 2015b, Theorem 1]. All characterization results are not kernel-based. For instance, two AFs are weak expansion equivalent w.r.t. stable semantics iff both do not possess stable extensions at all or if they share the same arguments and at the same time possess the same stable extensions. Consequently, $F = (\{a\}, \{(a, a)\})$ and $G = (\{a, b, c\}, \{(a, b), (b, c), (c, a)\})$ are weak expansion equivalent w.r.t. stable semantics. Both frameworks witness that any potential characterizing kernel k is necessarily neither node- nor loop-preserving.

As a final note, we are not aware of any study of weaker notions of expansion equivalence in case of cf2 as well as stage2 semantics.

Notions of Deletion Equivalence and Update Equivalence We start with local deletion equivalence. Remember that local deletion equivalent AFs cannot be semantically distinguished by deleting a certain set of attacks in both simultaneously. How "strong" is this notion? Are there redundant attacks or even redundant arguments?

Example 4.23 *Consider the following AFs F, G and H.*

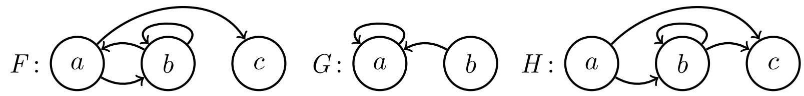

The AFs F and G do not possess the same arguments. Let us delete all occurring attacks, i.e. $S_A = R(F) \cup R(G)$. We obtain the following local deletions where $\{a, b, c\} \in \mathcal{E}_\sigma(F \setminus S_A) \setminus \mathcal{E}_\sigma(G \setminus S_A)$ for all semantics σ considered in this section.

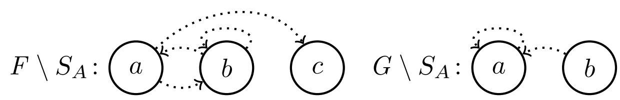

The AFs F and H possess the same arguments but differ in their attack-relation, e.g. $(b, c) \in R(H) \setminus R(F)$. This difference can be made more explicit if defining $S_R = (R(F) \cup R(H)) \setminus \{(b, c)\}$. Consider the resulting local deletions.

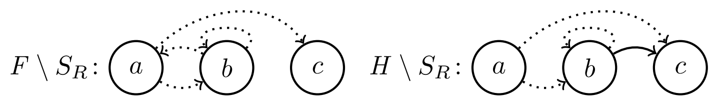

Once again we have $\{a, b, c\} \in \mathcal{E}_\sigma(F \setminus S_R)$ for all known semantics σ and $\{a, b, c\} \notin \mathcal{E}_\sigma(H \setminus S_R)$ if assuming conflict-freeness of the considered semantics.

The observations above indicate that there is not much space for redundancy in case of local expansion equivalence and indeed, it was one main result in [Baumann, 2014a] that local expansion equivalence collapse to identity for all semantics considered in this section. Moreover, instead of proving this one by one for any semantics the author followed the line in [Baroni and Giacomin, 2007] and provide abstract criteria guaranteeing the coincidence with syntactical identity. These criteria are very weak requirements, namely *conflict-freeness* ($\mathcal{CF}$) and the principle of *isolate-inclusion* ($\mathcal{II}$). The latter is fulfilled by a semantics σ iff for any AF F, the set of all isolated arguments is contained in at least one σ-extension. Observe that any considered semantics apart from stable semantics satisfy $\mathcal{II}$.[24]

[24]Note that only universally defined semantics σ, i.e. semantics which warrants the existence of at least one σ-extension (cf. Definition 2.3), may satisfy isolate-inclusion. A counter-example in case of stable semantics is given by $F = (\{a, b\}, \{(b, b)\})$. Obviously, a is isolated but $\mathcal{E}_{stb}(F) = \emptyset$. Nevertheless, local expansion equivalence in case of stable semantics collapse to identity too.

Theorem 4.24 ([Baumann, 2014a]) *Given a semantics σ satisfying $\mathcal{CF}$ and $\mathcal{II}$. For two finite AFs F and G we have:*

$$F \equiv^{\mathcal{E}_\sigma}_{LD} G \Leftrightarrow F = G.$$

Since being identical implies local deletion equivalence we deduce that all equivalence notion "inbetween" them collapse to identity too (cf. Figure 11).

Proposition 4.25 *Given a semantics σ satisfying $\mathcal{CF}$ and $\mathcal{II}$. For any two finite AFs F and G we have:*

$$F \equiv^{\mathcal{E}_\sigma}_{U} G \Leftrightarrow F \equiv^{\mathcal{E}_\sigma}_{D} G \Leftrightarrow F = G.$$

This means, for semantics satisfying conflict-freeness and isolate-inclusion any argument/attack may play a crucial role with respect to further evaluations if updates, deletions or local deletions are considered. Note that the results may apply to future semantics. In order to refine the general picture (as depicted in Figure 11) for the semantics considered in this section we state the following relations.[25]

Corollary 4.26 *Let $\sigma \in \{stg, stb, ss, eg, ad, pr, il, gr, co, na, cf2, stg2\}$. For any two finite AFs F and G we have:*

1. $F \equiv^{\mathcal{E}_\sigma}_{U} G \Leftrightarrow F \equiv^{\mathcal{E}_\sigma}_{D} G \Leftrightarrow F \equiv^{\mathcal{E}_\sigma}_{LD} G \Leftrightarrow F = G,$ *(k = id)*
2. $F \equiv^{\mathcal{E}_\sigma}_{D} G \Rightarrow F \equiv^{\mathcal{E}_\sigma}_{E} G,$ *(deletion vs. expansion)*
3. $F \equiv^{\mathcal{E}_\sigma}_{LD} G \Rightarrow F \equiv^{\mathcal{E}_\sigma}_{L} G.$ *(local versions)*

The Exceptional Case of Normal Deletion Equivalence

Normal deletion equivalence, where the retraction of arguments and corresponding attacks is considered, is exceptional in several regards. Firstly, the characterization theorems for admissible, complete and grounded semantics partially rely on σ-*-kernels. Remember that these kernels were originally introduced to characterize strong expansion (cf. Theorem 4.20). Secondly, normal deletion equivalent AFs do not even have to share the same arguments and thus give space for simplifications.

Example 4.27 *Consider the following AFs F and G. We have $\mathcal{E}_{ad}(F) = \mathcal{E}_{ad}(G) = \{\emptyset, \{a\}\}$. Even more, for any set of arguments B, $\mathcal{E}_{ad}(F \setminus B) = \mathcal{E}_{ad}(G \setminus B)$ showing their normal deletion equivalence, i.e. $F \equiv^{\mathcal{E}_{ad}}_{ND} G$.*

[25] The results in case of cf2 and stage2 semantics have never been published before. It can be checked that both semantics satisfy the preconditions of Theorem 4.25.

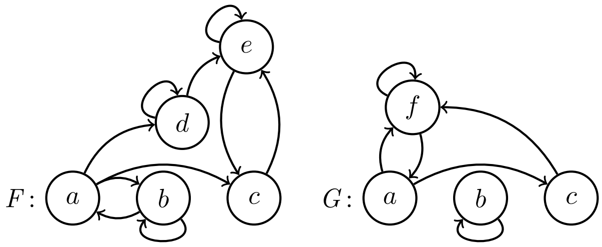

Observe that the non-shared arguments d, e and f do not play a role for the evaluation w.r.t. admissible semantics since firstly, they are self-defeating and thus cannot be part of an admissible set; and secondly, if they attack a non-looping argument shared by both arguments, e.g. e attacks c in F or f attacks a in G, then they are counter-attacked by the same argument, i.e. c attacks e in F and a attacks f in G. Consequently, they cannot influence potential admissible sets being a subset of $\{a, b\}$. *Finally, let us consider the ad-*-kernel of both frameworks (cf. Example 4.19 and the comments above for more details).*

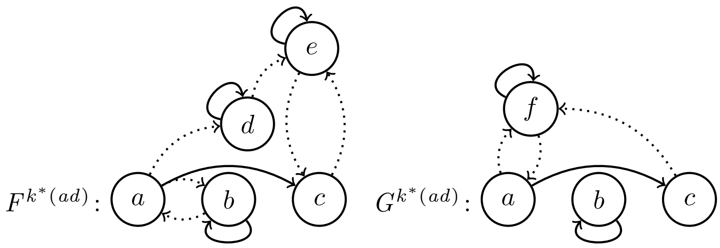

Obviously, F and G do not possess the same kernels but note that their restrictions to the shared arguments do, i.e. $\left(F|_{\{a,b,c\}}\right)^{k^*(ad)} = \left(G|_{\{a,b,c\}}\right)^{k^*(ad)}$.

It turned out that the issues raised in Example 4.27 are essential to characterize normal deletion equivalence w.r.t. admissible semantics. In case of complete and grounded semantics slightly different conditions have to be fulfilled, namely w.r.t. the non-shared arguments we have "it is forbidden to be attacked" instead of "counter-attack if attacked" like in case of admissible semantics and furthermore, instead of the ad-*-kernel the corresponding σ-*-kernels are used. Consider therefore the following definition and the characterization theorem. We use Δ to denote the symmetric difference, i.e. $A\Delta A' = A \setminus A' \cup A' \setminus A$. Moreover, $NL(F) = A(F) \setminus L(F)$, i.e. $NL(F)$ contains all arguments of F which are not self-defeating.

Definition 4.28 *Given* $F = (A, R)$ *and* $G = (A', R')$ *and let* $\sigma \in \{co, gr\}$.

1. *Loop*$(F,G) \Leftrightarrow_{def} L\left(F \cup G|_{A\Delta A'}\right) = A\Delta A'$,
   ("non-shared args are self-defeating")
2. $Att^{ad}(F,G) \Leftrightarrow_{def} \forall b \in A \setminus A' \; \forall a \in NL(F|_{A\cap A'}) : ((b,a) \in R \to (a,b) \in R)$
   $\wedge \forall b \in A' \setminus A \; \forall a \in NL(G|_{A\cap A'}) : ((b,a) \in R' \to (a,b) \in R')$,

("counter-attack if attacked")

3. $Att^{\sigma}(F,G) \Leftrightarrow_{def} \forall b \in A \setminus A' \; \forall a \in NL(F|_{A \cap A'}) : (b,a) \notin R$
$\wedge \forall b \in A' \setminus A \; \forall a \in NL(G|_{A \cap A'}) : (b,a) \notin R'$.
("it is forbidden to be attacked")

Theorem 4.29 ([Baumann, 2014a]) *Let $\sigma \in \{ad, co, gr\}$. Given two finite AFs $F = (A,R)$ and $G = (A',R')$ and let $I = A \cap A'$,*

$$F \equiv^{\mathcal{E}_\sigma}_{ND} G \Leftrightarrow Loop(F,G),\ \ Att^{\sigma}(F,G), (F|_I)^{k^*(\sigma)} = (G|_I)^{k^*(\sigma)}.$$

In contrast to admissible, complete and grounded semantics where normal deletion equivalence is indeed weaker than normal expansion equivalence we observe that these notions coincide in case of stable semantics. This means, normal deletion equivalence w.r.t. stable semantics is characterized by the classical stable kernel too.

The following theorem corrects the corresponding result in [Baumann, 2014a, Theorem 10] which did not take into account that an empty framework possess a stable extension, namely the empty one.[26]

Theorem 4.30 *For finite AFs F and G we have:*

$$F \equiv^{\mathcal{E}_{stb}}_{ND} G \Leftrightarrow F^{k(stb)} = G^{k(stb)}.$$

Proof. ($\Rightarrow$) We show the contrapositive, i.e. $F^{k(stb)} \neq G^{k(stb)} \Rightarrow F \not\equiv^{\mathcal{E}_{stb}}_{ND} G$.
1^{st} case: Assume $A\left(F^{k(stb)}\right) \neq A\left(G^{k(stb)}\right)$ and w.l.o.g. let $a \in A\left(F^{k(stb)}\right) \setminus A\left(G^{k(stb)}\right)$. Since the stable kernel is node-preserving (Fact 4.12) we obtain $G \setminus B = (\emptyset, \emptyset)$ and $F \setminus B \in \{(\{a\}, \emptyset), (\{a\}, \{(a,a)\})\}$ if $B = (A(F) \cup A(G)) \setminus \{a\}$. In either case, $\emptyset \in \mathcal{E}_{stb}(G) \setminus \mathcal{E}_{stb}(F)$ since $\mathcal{E}_{stb}(F) \in \{\emptyset, \{\{a\}\}\}$. From now on we assume $A\left(F^{k(stb)}\right) = A\left(G^{k(stb)}\right)$.
2^{nd} case: Consider $R\left(F^{k(stb)}\right) \neq R\left(G^{k(stb)}\right)$ and w.l.o.g. let $(a,b) \in R\left(F^{k(stb)}\right) \setminus R\left(G^{k(stb)}\right)$. Let $a = b$. Remember that the stable kernel is loop-preserving (Fact 4.12). Therefore, $(a,a) \in R(F) \setminus R(G)$. We obtain $G \setminus B = (\{a\}, \emptyset)$ and $F \setminus B = (\{a\}, \{(a,a)\})$ if $B = (A(F) \cup A(G)) \setminus \{a\}$. Hence, $\emptyset = \mathcal{E}_{stb}(F) \neq \mathcal{E}_{stb}(G) = \{\{a\}\}$. From now on we assume $L\left(F^{k(stb)}\right) = L\left(G^{k(stb)}\right)$. Consider now $a \neq b$. Consequently, $(a,b) \in R(F)$ and $(a,a) \notin R(F)$. Hence, $(a,a) \notin R(G)$ and furthermore, $(a,b) \notin R(G)$. Define $B = (A(F) \cup A(G)) \setminus \{a,b\}$. In any case, $\{a\} \in \mathcal{E}_{stb}(F \setminus B) \setminus \mathcal{E}_{stb}(G \setminus B)$ concluding the if-direction.
($\Leftarrow$) Given $F^{k(stb)} = G^{k(stb)}$. Applying Theorems 4.14 and 4.20 one after the other yields $F \equiv^{\mathcal{E}_{stb}}_{E} G$ and then $F \equiv^{\mathcal{E}_{stb}}_{N} G$. Finally, Corollary 4.21 justifies $F \equiv^{\mathcal{E}_{stb}}_{ND} G$ concluding the proof. ∎

[26] We mention that Theorem 10 in [Baumann, 2014a] hold, given that resulting AFs have to be non-empty. The claimed normal deletion equivalence of the AFs F and G depicted in [Baumann, 2014a, Example 4] can be disproved by setting $B = \{a, b, c, f\}$.

Characterization Theorems in Case of Self-loop-free AFs

We already observed that apart from naive kernel any mentioned kernel k does not change anything if the considered AF F is self-loop-free, i.e. $F = F^k$ (cf. Facts 4.12 and 4.18). Consequently, any equivalence relation characterizable through such a kernel collapses to identity if we restrict ourselves to self-loop-free AFs. This is stated in the following theorem.

Theorem 4.31 *Given a relation* $\equiv \subseteq \mathcal{F} \times \mathcal{F}$ *characterizable through* k *where* $k \in \{k(stb), k(ad), k(gr), k(co), k^*(ad), k^*(gr), k^*(co), k^*(stg)\}$. *For any self-loop-free AFs* F *and* G,

$$F \equiv G \Leftrightarrow F = G.$$

We will refrain from listing all combinations of semantics and equivalence notions characterizable through a kernel mentioned in the theorem above. Please confer Figures 12 and 15 for compact overviews. For all such combinations, self-loop-free AFs are redundancy-free, i.e. all attacks as well as arguments may play a crucial role w.r.t. further evaluations and thus, there is no space for simplification. In the introductory part of this section we noted that many equivalence notions, e.g. normal and local expansion equivalence are motivated by the instantiation-based context where AFs are built from an underlying knowledge base. However, we want to mention that there are some formalisms like classical logic-based argumentation where self-attacking arguments do not occur [Besnard and Hunter, 2001, Theorem 4.13], while for other systems, e.g. ASPIC self-defeating arguments indeed may arise [Prakken, 2010, Section 7].

Summary of Results and Conclusion

In the presented results the notion of a kernel played a crucial role. Indeed, kernels are interesting from several perspectives: First, they allow to decide the corresponding notion of equivalence by a simple check for topological (i.e. syntactical) equality. Moreover, all kernels we have obtained so far can be efficiently constructed from a given argumentation framework. This means, if a certain equivalence notion is characterizable through such a kernel, then we have tractability of the associated decision problem.

The following Figure 12 provides a comprehensive overview of the state of the art in case of extension-based semantics. The entry "k" in row M and column σ indicates that $\equiv^{\mathcal{E}_\sigma}_M$ is characterizable through k. The abbreviation "*id*" stands for identity map and the question mark represents an open problem. Further abbreviations like "L" and "Att^σ" refer to additional conditions relevant in case of normal deletion equivalence (cf. Theorem 4.29). The entry "$[m, n]$" indicates three facts. First, the characterization problem is already solved in Theorem/Proposition n in m.[27] Second, the characterization result is not (purely) kernel-based and third, it can be checked that none of the introduced kernels serve as a characterization.

[27] For m we use the following assignments: 1 = [Baumann, 2011], 2 =[Baumann and Brewka, 2015b], 3 =[Baumann and Brewka, 2013] and 4 =[Oikarinen and Woltran, 2011]

	stg	*stb*	*ss*	*eg*	*ad*	*pr*	*il*	*gr*	*co*	*na*	*cf2*	*stg2*
W	?	[1,3]	?	?	[2,1]	[3,1]	?	?	[2,1]	?	?	?
L	$k^*(stg)$	[4,9]	$k(ad)$	$k(ad)$	$k(ad)$	$k(ad)$	$k(ad)$	[4,10]	[4,11]	$k(na)$	?	?
E	$k(stb)$	$k(stb)$	$k(ad)$	$k(ad)$	$k(ad)$	$k(ad)$	$k(ad)$	$k(gr)$	$k(co)$	$k(na)$	id	id
N	$k(stb)$	$k(stb)$	$k(ad)$	$k(ad)$	$k(ad)$	$k(ad)$	$k(ad)$	$k(gr)$	$k(co)$	$k(na)$	id	id
S	$k(stb)$	$k(stb)$	$k(ad)$	$k(ad)$	$k^*(ad)$	$k^*(ad)$	$k^*(ad)$	$k^*(gr)$	$k^*(co)$	$k(na)$	?	?
ND	?	$k(stb)$	?	?	$k^*(ad)$ L, Att^{ad}	?	?	$k^*(gr)$ L, Att^{gr}	$k^*(co)$ L, Att^{co}	?	?	?
D	id	id	id	id	id	id	id	id	id	id	id	id
LD	id	id	id	id	id	id	id	id	id	id	id	id
U	id	id	id	id	id	id	id	id	id	id	id	id

Figure 12: Extension-based Characterizations for Finite AFs

Remember that any arbitrary expansion (deletion) can be split into a normal and local part. So one natural conjecture is that normal and local expansion (deletion) equivalence jointly imply expansion (deletion) equivalence. Using the results presented in this section we can not only verify the addressed conjecture but even give a significantly stronger result. In fact, the main and quite surprisingly relations for the considered semantics can be briefly and concisely stated in the following two equations, namely "normal expansion equivalence = expansion equivalence" and "local deletion equivalence = deletion equivalence".

The fact that different notions of equivalence might or might not coincide is interesting from a conceptual point of view. To illustrate this let us have a look at normal and strong expansion equivalence. Recall that normal expansions add new arguments and possibly new attacks which involve at least one of the fresh arguments, while strong expansions (a subclass of normal expansions) restrict the possible attacks between the new arguments and the old ones to a single direction. In dynamic settings, both concepts can be justified in the sense that new arguments might be raised but this will not influence the relation between already existing arguments. For strong expansions, only strong arguments will be raised, i.e. arguments which cannot be attacked by existing ones. The corresponding equivalence notions now check whether two AFs are "equally robust" to such new arguments, and indeed, normal expansion equivalence always implies strong expansion equivalence but the other direction is only true for some of the semantics, namely stage, stable, semi-stable, eager and naive semantics. One interpretation is that when two AFs are not normal expansion equivalent, then this can be made explicit by only posing strong arguments (not attacked by existing ones), while for the other semantics this is not the case. For this particular example, it seems that the notion of admissibility which is more "explicit" in the admissible, preferred, ideal, grounded and complete semantics is responsible for the fact that frameworks might be strong expansion equivalent but not normal expansion equivalent.

In Figure 11 we presented preliminary relations between several notions of equivalence which hold for any semantics. The refinement depicted in Figure 13 applies to any extension-based semantics considered in this section.

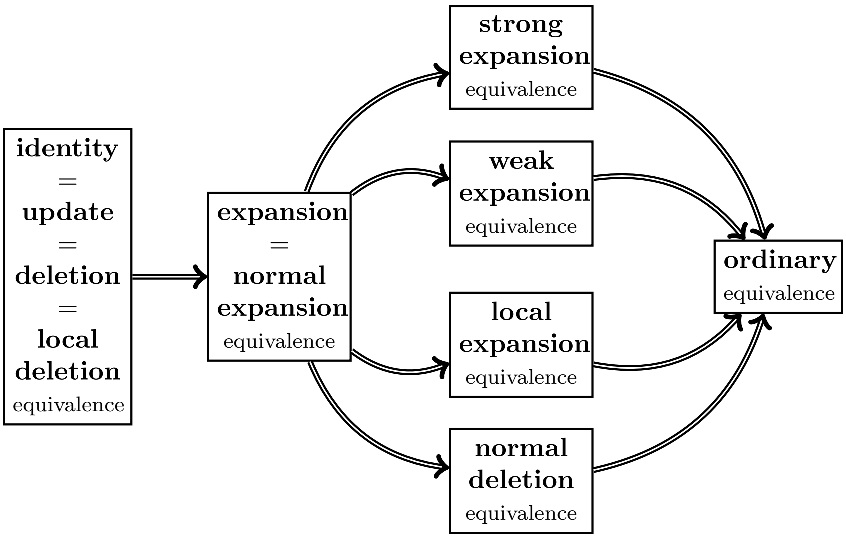

Figure 13: Relations for $\sigma \in \{stg, stb, ss, eg, ad, pr, il, gr, co, na, cf2, stg2\}$ - Extension-based Versions and Finite AFs

Finally, we present the overall picture for the most prominent semantics, namely the stable one. Interestingly, in contrast to Figure 13 all equivalence notions are comparable, i.e. they are totally ordered w.r.t. $\subseteq$. Comprehensive overviews for single semantics can be found in [Baumann, 2014b, Section 5.5.2] or [Baumann and Brewka, 2015b]. The latter also contains a comparison to different notions of *minimal change equivalence* firstly introduced in [Baumann, 2012b]. As an aside, very recently the authors of [Baumann *et al.*, 2017] introduced so-called *C-relativized equivalence* that subsumes ordinary and expansions equivalence as its extreme corner cases. The set C represents so-called *core* arguments which will not be directly touched by the possible expansions.

This means, for any set C we obtain a further intermediate notion between expansion and ordinary equivalence. However, due to its recency further relations are not studied so far.

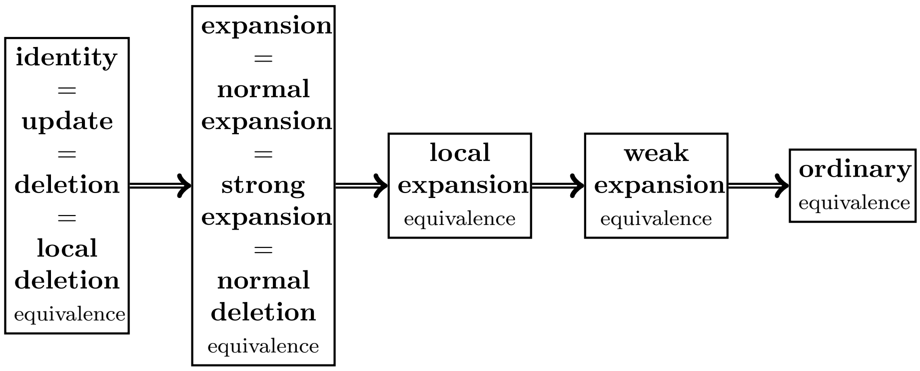

Figure 14: Stable Semantics - Extension-based Version and Finite AFs

4.3 Equivalence in the Light of Unrestricted Frameworks

Recently, a first study of several abstract properties in the unrestricted setting were presented in [Baumann and Spanring, 2017]. The main result regarding expansion equivalence can be summarized as follows: All characterization results carry over to the unrestricted setting as long as the AFs in question are *jointly expandable* (w.r.t. $\mathcal{U}$). Consider therefore the following definition and the corresponding characterization theorem.

Definition 4.32 *F and G are jointly expandable if $\mathcal{U} \setminus (A(F) \cup A(G)) \neq \emptyset$.*

Theorem 4.33 *[Baumann and Spanring, 2017] For jointly expandable AFs F and G we have:*

1. $F \equiv_E^{\mathcal{E}_\sigma} G \Leftrightarrow F^{k(\sigma)} = G^{k(\sigma)}$ *for any* $\sigma \in \{stb, ad, co, gr, na\}$,
2. $F \equiv_E^{\mathcal{E}_\sigma} G \Leftrightarrow F^{k(ad)} = G^{k(ad)}$ *for any* $\sigma \in \{pr, il, ss, eg\}$ *and*
3. $F \equiv_E^{\mathcal{E}_{stg}} G \Leftrightarrow F^{k(stb)} = G^{k(stb)}$.

The main proof strategies are straightforward extensions of those presented in [Oikarinen and Woltran, 2011]. However, finiteness assumptions are often used implicitly and one has to pay attention whether a certain reasoning step (e.g. subset relation between semantics, definedness statuses of semantics, finitely many extensions etc.) carry over to the infinite setting.

Interestingly, in case of the admissible as well as naive kernel we may even drop the restriction of joint expandability as stated in the following theorem.

Theorem 4.34 *[Baumann and Spanring, 2017] For unrestricted AFs F and G we have:*

1. $F \equiv_E^{\mathcal{E}_{na}} G \Leftrightarrow F^{k(na)} = G^{k(na)}$ *and*
2. $F \equiv_E^{\mathcal{E}_{\sigma}} G \Leftrightarrow F^{k(ad)} = G^{k(ad)}$ *for any* $\sigma \in \{ad, pr, il, ss, eg\}$.

The following two examples taken from [Baumann and Spanring, 2017] show that this assertion does not hold for all kernels considered in this section. The main reason for this different behaviour is that for some semantics it plays a decisive role whether AFs can be expanded by "fresh" arguments which is not given for unrestricted frameworks in general but guaranteed for jointly expandable AFs (cf. Definition 4.32).

Example 4.35 *Given* $c \in \mathcal{U}$ *and define* $F = (\mathcal{U} \setminus \{c\}, \{(a,a) \mid a \in \mathcal{U} \setminus \{c\}\})$ *and* $G = (\mathcal{U}, \{(a,a) \mid a \in \mathcal{U} \setminus \{c\}\})$. *For any H we observe* $\mathcal{E}_{stb}(F \cup H) = \mathcal{E}_{stb}(G \cup H)$. *In particular,*

$$\mathcal{E}_{stb}(F \cup H) = \begin{cases} \{\{c\}\}, & \text{if } \{(c,a) \mid a \in \mathcal{U} \setminus \{c\}\} \subseteq R(H) \text{ and } (c,c) \notin R(H) \\ \emptyset, & \text{otherwise} \end{cases}$$

Consequently, $F \equiv_E^{\mathcal{E}_{stb}} G$ *although* $A(F) \neq A(G)$ *(and thus,* $F^{k(stb)} \neq G^{k(stb)}$*).*

Example 4.36 *Consider the AFs* $F = (\mathcal{U}, \{(a,a) \mid a \in \mathcal{U}\})$ *as well as* $G = (\mathcal{U}, \{(a,b) \mid a,b \in \mathcal{U}, a \neq b)$. *Applying the grounded kernel does not change anything for either framework, i.e.* $F^{k(gr)} = F$ *and* $G = G^{k(gr)}$. *Due to the absence of unattacked arguments we deduce* $\mathcal{E}_{gr}(F \cup H) = \mathcal{E}_{gr}(G \cup H) = \{\emptyset\}$ *for any AF H. Consequently,* $F \equiv_E^{\mathcal{E}_{gr}} G$ *although* $F^{k(gr)} \neq G^{k(gr)}$.

4.4 Characterization Theorems for Labelling-Based Semantics

We now return to the finite setting and consider the second main approach used for evaluating argumentation scenarios, namely labelling-based semantics. As a matter of fact, the labelling-based versions of all considered semantics provides one with more information than their extension-based counter-parts. More precisely, the defined 3-valued labellings assign a status to any argument of the considered AF F, i.e. in addition to the information which arguments are *accepted* we also have labels for the remaining arguments indicating that they are either *rejected* or *undecided* with respect to F (cf. Chapter 4 for more background and precise definitions). It is well known that many semantics establish a one-to-one correspondence between their extension-based and labelling-based versions. This means, any labelling is associated with exactly one extension and vice versa. It is not immediately apparent whether this property guarantees that there is a coincidence of the extension-based and labelling-based equivalence notions. In [Baumann, 2016] a negative answer was given. The main reason for the invalidity is that AFs may possess the same extensions without sharing the same arguments which is impossible in case of labellings since any argument has to be labelled. Furthermore, even sharing the same arguments does not ensure the validity of the converse direction. Consider therefore the following example.

Example 4.37 *Consider the AFs F and G as depicted below. Although both frameworks possess the same unique preferred extension, they do not share the same preferred labellings. More precisely, $\mathcal{E}_{pr}(F) = \mathcal{E}_{pr}(G) = \{\{a\}\}$ but $\{(\{a\}, \{b\}, \emptyset)\} = \mathcal{L}_{pr}(F) \neq \mathcal{L}_{pr}(G) = \{(\{a\}, \emptyset, \{b\})\}$.*

Moreover, observe that $F^{k^*(ad)} = G = G^{k^*(ad)}$. Consequently, both frameworks are even strong expansion equivalent w.r.t. preferred extension-based semantics (Theorem 4.20). This means, equivalence notions may differ considerably if considered under the extension-based or labelling-based approach.

In contrast to extension-based semantics where characterization results are spread over a high number of publications there is only one reference, namely [Baumann, 2016] concerned with labelling-based semantics. The author considered 8 different equivalence notions w.r.t. 8 prominent labelling-based semantics in the finite setting. In effect, similarly to extension-based semantics, almost all labelling-based equivalence notions can be decided syntactically. Differently from the extension-based approach we observe a much more homogeneous picture. For instance, there is no need for the more sophisticated σ-*-kernels as we will see.

Basic Properties and a Fundamental Relation

Before turning to the main results we start with some preliminary facts relating σ-extensions and σ-labellings. In the following we restrict ourselves to the semantics considered in [Baumann, 2016]. For any 3-valued labelling $L = (L_1, L_2, L_3)$ we use $L = (L^{\text{I}}, L^{\text{O}}, L^{\text{U}})$ as usual.

Fact 4.38 *Given a finite AF $F = (A, R)$ and $E \subseteq A$. We write $E^{\mathcal{L}}$ for $(E, E^+, A \setminus E^{\oplus}))$. For all $\sigma \in \{stb, ss, eg, ad, pr, il, gr, co\}$ we have,*

1. *If $L \in \mathcal{L}_\sigma(F)$, then $L^{\text{I}} \in \mathcal{E}_\sigma(F)$,* *(extension induced by labelling)*
2. *If $E \in \mathcal{E}_\sigma(F)$, then $E^{\mathcal{L}} \in \mathcal{L}_\sigma(F)$ and* *(labelling induced by extension)*
3. *Obviously, $(E^{\mathcal{L}})^I = E$.* *($I \circ \mathcal{L} = id$)*

We point out that the first two properties mentioned in Fact 4.38 do not ensure that there is a one-to-one correspondence between σ-labellings and σ-extensions. This desirable feature (which would indeed justify the terms σ-labellings and σ-extensions) is given if additionally, labellings are uniquely determined by their in-labelled arguments.

Fact 4.39 *Given a finite AF $F = (A, R)$ and a set $E \subseteq A$. For all semantics $\sigma \in \{stb, ss, eg, pr, il, gr, co\}$ we have,*

1. *For any* $L, M \in \mathcal{L}_\sigma(F)$, $L^I = M^I$ *iff* $L = M$,
(uniquely determined by in-labels)

2. *Given* $L \in \mathcal{L}_\sigma(F)$, *then* $(L^I)^{\mathcal{L}} = L$ *and* *($\mathcal{L} \circ I = id$)*

3. $|\mathcal{L}_\sigma(F)| = |\mathcal{E}_\sigma(F)|$. *(same cardinality)*

As an aside, we mention that (although not immediately apparent) the first two items of Fact 4.39 are equivalent independently of any semantics definition. Please note that admissible labellings are excluded from Fact 4.39. The AF F depicted in Example 4.37 shows that this is no coincidence. It possesses two admissible labellings associated with one admissible extension. More precisely, the admissible labellings $(\{a\}, \{b\}, \emptyset)$ as well as $(\{a\}, \emptyset, \{b\})$ refer to the same admissible extension $\{a\}$.

We proceed with a general relation between labelling-based and extension-based versions of certain equivalence notion. More precisely, for any considered semantics and any equivalence notion presented in Definition 4.6 we have that being equivalent w.r.t. labellings implies being equivalent w.r.t. extensions. The main reason for this fundamental relation is the following lemma stating that possessing the same labellings implies sharing the same extensions. We mention that this property is already guaranteed if the semantics σ in question satisfies that any σ-extension induces an σ-labelling and vice versa (cf. statements 1 and 2 of Fact 4.38).

Lemma 4.40 ([Baumann, 2016]) *Given two finite AFs F and G. For any semantics $\sigma \in \{stb, ss, eg, ad, pr, il, gr, co\}$ we have,*

$$\mathcal{L}_\sigma(F) = \mathcal{L}_\sigma(G) \Rightarrow \mathcal{E}_\sigma(F) = \mathcal{E}_\sigma(G).$$

Proof. Reductio ad absurdum. Assume $\mathcal{E}_\sigma(F) \neq \mathcal{E}_\sigma(G)$. Then, w.l.o.g. exists $E \in \mathcal{E}_\sigma(F) \setminus \mathcal{E}_\sigma(G)$. Consequently, $E^{\mathcal{L}} \in \mathcal{L}_\sigma(F)$ (item 2 of Fact 4.38). Thus, $E^{\mathcal{L}} \in \mathcal{L}_\sigma(G)$ (assumption). Hence, $(E^{\mathcal{L}})^I \in \mathcal{E}_\sigma(G)$ (item 1 of Fact 4.38). Furthermore, $(E^{\mathcal{L}})^I = E \in \mathcal{E}_\sigma(G)$ (item 3 of Fact 4.38). Contradiction! ■

We now present the fundamental relation between labelling-based and extension-based equivalence notion.

Theorem 4.41 ([Baumann, 2016]) *Given two finite AFs F and G. For any $\sigma \in \{stb, ss, eg, ad, pr, il, gr, co\}$ and any $M \in \{W, L, E, N, S, ND, D, LD, U\}$ we have,*

$$F \equiv_M^{\mathcal{L}_\sigma} G \Rightarrow F \equiv_M^{\mathcal{E}_\sigma} G.$$

Proof. We show the contrapositive. Assume $F \not\equiv_M^{\mathcal{E}_\sigma} G$. This means, there is a certain scenario S according to M, s.t. $\mathcal{E}_\sigma(S(F)) \neq \mathcal{E}_\sigma(S(G))$.[28] Consequently, $\mathcal{L}_\sigma(S(F)) \neq \mathcal{L}_\sigma(S(G))$ (Lemma 4.40) proving $F \not\equiv_M^{\mathcal{L}_\sigma} G$. ■

In Example 4.37 we have seen that the converse direction does not hold in general. Nevertheless, there is huge number of equivalence notions where labelling-based and extension-based versions do indeed coincide (cf. Figure 15 for an overview).

Coincidences of Extension-based and Labelling-based Versions

Remember that the identity relation is the finest equivalence relation. Furthermore, it is already shown that deletion, local deletion as well as update equivalence w.r.t. $\mathcal{E}_\sigma$ collapse to identity (see Figure 13). Consequently, applying the fundamental relation stated in Theorem 4.41 we obtain the identical characterization results w.r.t. labelling-based semantics.

Theorem 4.42 ([Baumann, 2016]) *For finite AFs F and G, a scenario $M \in \{D, LD, U\}$ and a semantics $\sigma \in \{stb, ss, eg, ad, pr, il, gr, co\}$ we have,*

$$F \equiv_M^{\mathcal{L}_\sigma} G \Leftrightarrow F = G.$$

Analogously to extension-based semantics (cf. Fact 4.13) we have that there are combinations of kernels and semantics σ, s.t. the application of a kernel does not vary the set of σ-labellings.

Fact 4.43 *For any finite AF F,*

1. $\mathcal{L}_\sigma(F) = \mathcal{L}_\sigma\left(F^{k(\sigma)}\right)$ *for* $\sigma \in \{co, stb, gr\}$ *and*
2. $\mathcal{L}_\tau(F) = \mathcal{L}_\tau\left(F^{k(ad)}\right)$ *for* $\tau \in \{ss, eg, pr, il\}$.

The fact above is the decisive property which allows one to carry over further kernel-based characterization results for extension-based semantics to their labelling-based version. In order to show this result it was necessary to find a condition for equality of two complete labellings of different AFs. Remember that two complete labellings of the same framework are identical if and only if they possess the same in-labelled arguments (Fact 4.39). In case of different AFs we have to require additionally that both frameworks share the same arguments and the same range w.r.t. the set of in-labelled arguments.

Fact 4.44 *Given two finite AFs F and G as well as $L \in \mathcal{L}_{co}(F)$ and $M \in \mathcal{L}_{co}(G)$. We have $L = M$ iff simultaneously $A(F) = A(G)$, $L^I = M^I$ and $R_F^+(L^I) = R_G^+(M^I)$.*

[28]For instance, in case of expansion equivalence (i.e. M = E) a scenario S is simply the union with a further AF H, i.e. $S(F) = F \cup H$ and $S(G) = G \cup H$.

Please observe that admissible labellings do not fulfill Fact 4.44. Consider for instance again the AF F depicted in Example 4.37 and its two admissible labellings $(\{a\},\{b\},\emptyset)$ and $(\{a\},\emptyset,\{b\})$.

We proceed with the main coincidence theorem. It stipulates that several expansion equivalence relations as well as weaker notions do not distinguish between their labelling-based and extension-based version. This means, kernel-based characterization results (depicted in Figure 11) carry over to labelling-based semantics. Similarly to extension-based semantics we present an overview of characterizing kernels at the end of this section (cf. Figure 15).

Theorem 4.45 ([Baumann, 2016]) *Given finite AFs F and G. We have,*

1. $F \equiv^{\mathcal{E}_\sigma}_M G \Leftrightarrow F \equiv^{\mathcal{L}_\sigma}_M G$ *for* $\sigma \in \{stb, ss, eg, pr, il, gr, co\}, M \in \{E, N\}$,
2. $F \equiv^{\mathcal{E}_\sigma}_L G \Leftrightarrow F \equiv^{\mathcal{L}_\sigma}_L G$ *for* $\sigma \in \{ss, eg, pr, il\}$ *and*
3. $F \equiv^{\mathcal{E}_\sigma}_S G \Leftrightarrow F \equiv^{\mathcal{L}_\sigma}_S G$ *for* $\sigma \in \{stb, ss, eg\}$.

Non-Coincidence of Extension-based and Labelling-based Versions

We now leave the realm of uniformity of extension-based and labelling-based characterizations. This section is divided into three parts. We start with characterization theorems for admissible labellings. In particular, we will see that the admissible kernel (originally introduced to characterize equivalence notions w.r.t. admissible extension-based semantics) does not serve as characterizing kernel for admissible labellings. We then proceed with strong expansion equivalence w.r.t. labellings. We will see that the remaining notions are characterizable via traditional kernels instead of σ-*-kernels. In the third part we consider normal deletion equivalence w.r.t. labelling-based semantics. In contrast to their extension-based versions where many notions has defied any attempt of solving, we present characterization theorems based on traditional kernels for all eight considered semantics.

Expansion Equivalence w.r.t. Admissible Labellings Expansion equivalence as well as its local, normal and strong versions w.r.t. admissible extensions are characterizable through the admissible kernel. The following example shows that this assertion does not hold in case of admissible labellings.

Example 4.46 *The following two AFs possess the same admissible kernels, namely $F^{k(ad)} = G^{k(ad)} = F$. Consequently, applying characterization theorems for extension-based semantics we obtain $F \equiv^{\mathcal{E}_{ad}}_M G$ for $M \in \{L, E, N\}$ (cf. Figure 12).*

Observe that $(\{b\},\emptyset,\{a\}) \in \mathcal{L}_{ad}(G) \setminus \mathcal{L}_{ad}(F)$ because the argument a cannot be undecided in F since it attacks the in-labelled argument b. Thus $F \not\equiv^{\mathcal{L}_{ad}}_M G$ for $M \in \{L, E, N, S\}$.

Let us assume that the equivalence notions considered in the example above are characterizable through a certain kernel k. Due to the fundamental relation (Theorem 4.41) and the characterization results w.r.t. admissible extensions (Figure 12), we already know that the kernel k has to satisfy the following implication: $F^k = G^k \Rightarrow F^{k(ad)} = G^{k(ad)}$ for any two AFs F and G. This means, we are looking for a weaker kernel than the admissible one in the sense that first, everything which is redundant w.r.t. k has to be redundant w.r.t. the admissible kernel too; and second, an attack from a to b has to survive even if a is self-defeating and b counterattacks a. One candidate for k is the complete kernel since redundancy w.r.t. the complete kernel implies redundancy w.r.t. to the admissible one, and furthermore, it deletes an attack between two arguments if and only if both are self-defeating. And indeed, it was shown that expansion equivalence as well as its local, normal and strong variant w.r.t. admissible labellings are characterizable through the complete kernel as stated by the following theorem.

Theorem 4.47 ([Baumann, 2016]) *Given finite AFs F and G. We have,*

$$F \equiv_M^{\mathcal{L}_{ad}} G \Leftrightarrow F^{k(co)} = G^{k(co)} \text{ with } M \in \{L, E, N, S\}.$$

Strong Expansion Equivalence for Preferred, Ideal, Grounded and Complete Labellings In this subsection we present characterization theorems for strong expansion equivalence w.r.t. labelling-based preferred, ideal, grounded and complete semantics. Remember that in case of strong expansions a former attack between old arguments will never become a counterattack to an added attack. Consequently, in contrast to arbitrary expansions former attacks do not play a role with respect to being a potential defender of an added argument. The context-sensitive σ-*-kernels took these considerations into account and allow for more deletions than their classical counterparts.

Example 4.48 *According to Definition 4.17 we have, $F^{k^*(\sigma)} = G^{k^*(\sigma)}$ for any semantics $\sigma \in \{ad, gr, co\}$. More precisely, the attacks (a, b) in F as well as (c, b) in G are redundant w.r.t. all three σ-*-kernels. This means, in consideration of Figure 12 both frameworks are strong expansion equivalent w.r.t. the extension-based versions of preferred, ideal, grounded and complete semantics.*

Consider the following dynamic scenario where a stronger argument than the former ones is added. Formally, we conjoin the AF $H = (\{c, d\}, \{(d, c)\})$ to both frameworks F and G.

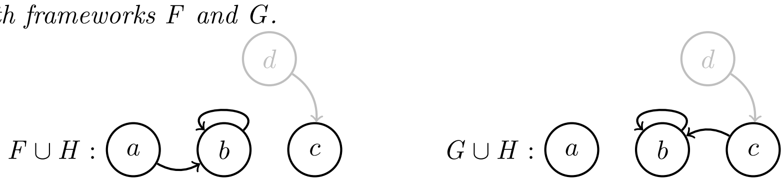

Note that both frameworks has to possess the same σ-extension since $G \equiv_S^{\mathcal{E}_\sigma} H$ for $\sigma \in \{pr, il, gr, co\}$ is already ensured. Furthermore, we observe $(\{a, d\}, \{b, c\}, \emptyset) \in \mathcal{L}_\sigma(F \cup H) \setminus \mathcal{L}_\sigma(G \cup H)$ since b cannot be out-labelled in $G \cup H$ because there is no in-labelled attacker. This means, $F \not\equiv_S^{\mathcal{L}_\sigma} G$ for $\sigma \in \{pr, il, gr, co\}$.

Analogously to the previous section let us assume that strong expansion equivalence w.r.t. the considered labelling-based semantics are characterizable through a certain kernel k. We immediately obtain, $F^k = G^k \Rightarrow F^{k^*(\sigma)} = G^{k^*(\sigma)}$ for any two AFs F and G. Possible candidates are the classical counterparts of the σ-*-kernels and indeed it was shown that these kernels guarantee the desired outcome. This means, in case of strong expansion equivalence w.r.t. preferred, ideal, grounded and complete semantics we have that the labelling-based version is characterizable through a classical σ-kernel if and only if the extension-based version is characterizable through the corresponding σ-*-kernel.

Theorem 4.49 ([Baumann, 2016]) *Given finite AFs F and G. We have,*

1. $F \equiv_S^{\mathcal{L}_\sigma} G \Leftrightarrow F^{k(ad)} = G^{k(ad)}$ *for* $\sigma \in \{pr, il\}$,
2. $F \equiv_S^{\mathcal{L}_{gr}} G \Leftrightarrow F^{k(gr)} = G^{k(gr)}$ *and*
3. $F \equiv_S^{\mathcal{L}_{co}} G \Leftrightarrow F^{k(co)} = G^{k(co)}$.

Normal Deletion Equivalence Characterizing normal deletion equivalence in case of extension-based semantics is exceptional in several regards. Remember that normal deletions retract arguments and their corresponding attacks. Firstly, only a few characterization results are achieved (cf. Figure 12). Furthermore, apart from stable semantics, none of the characterization results is purely kernel-based, i.e. beside the equality of kernels on certain parts of the frameworks further loop- as well as attack-conditions have to be satisfied. Finally, quite surprisingly, normal deletion equivalent AFs do not even have to share the same arguments enabling equivalence classes with an infinite number of elements. Being equivalent w.r.t. labellings and possessing different arguments at the same time is impossible in case of labellings since any argument has to be labelled. It turned out that any considered labelling-based semantics is characterizable through traditional kernels and thus, do not share any of the features mentioned above. Consider the following main theorem.

Theorem 4.50 ([Baumann, 2016]) *Given finite AFs F and G. We have,*

1. $F \equiv_{ND}^{\mathcal{L}_{stb}} G \Leftrightarrow F^{k(stb)} = G^{k(stb)}$,
2. $F \equiv_{ND}^{\mathcal{L}_\sigma} G \Leftrightarrow F^{k(ad)} = G^{k(ad)}$ *for* $\sigma \in \{ss, eg, pr, il\}$,
3. $F \equiv_{ND}^{\mathcal{L}_\sigma} G \Leftrightarrow F^{k(co)} = G^{k(co)}$ *for* $\sigma \in \{ad, co\}$ *and*
4. $F \equiv_{ND}^{\mathcal{L}_{gr}} G \Leftrightarrow F^{k(gr)} = G^{k(gr)}$.

Summary of Results and Conclusion

The following Figure 15 presents a comprehensive overview of the state of the art in case of labelling-based semantics. Analogously to Figure 12 the entry "k" in row M and column σ indicates that $\equiv^{\mathcal{L}_\sigma}_{M}$ is characterizable through k given the finiteness restriction. The abbreviation "id" stands for identity map and the question mark represents an open problem.[29] A grey-highlighted entry reflects the situation that extension-based and labelling-based version do not coincide.

	stb	ss	eg	ad	pr	il	gr	co
L	?	$k(ad)$	$k(ad)$	$k(co)$	$k(ad)$	$k(ad)$	?	?
E	$k(stb)$	$k(ad)$	$k(ad)$	$k(co)$	$k(ad)$	$k(ad)$	$k(gr)$	$k(co)$
N	$k(stb)$	$k(ad)$	$k(ad)$	$k(co)$	$k(ad)$	$k(ad)$	$k(gr)$	$k(co)$
S	$k(stb)$	$k(ad)$	$k(ad)$	$k(co)$	$k(ad)$	$k(ad)$	$k(gr)$	$k(co)$
ND	$k(stb)$	$k(ad)$	$k(ad)$	$k(co)$	$k(ad)$	$k(ad)$	$k(gr)$	$k(co)$
D	id	id	id	id	id	id	id	id
LD	id	id	id	id	id	id	id	id
U	id	id	id	id	id	id	id	id

Figure 15: Labelling-based Characterizations for Finite AFs

In contrast to extension-based semantics we observe a much more homogeneous picture. Firstly, there is no need for the more sophisticated σ-*-kernels. Secondly, normal deletion equivalence w.r.t. labelling-based semantics is naturally incorporated in the overall picture in the sense that it coincides with its corresponding expansion, normal expansion and strong expansion equivalence notions.

The following Figure 16 applies to each one of the eight labelling-based semantics considered in this section. In comparison to Figure 11 where pre-

[29] In contrast to extension-based semantics the labelling versions of conflict-free-based semantics like stage, naive, cf2 as well as stage2 semantics (cf. [Caminada, 2011; Gaggl and Dvořák, 2016]) as well as weak expansion equivalence at all were not considered so far and thus, represent open problems too.

liminary relations are depicted it illustrates (to a certain extent) a collapse of the diversity of the introduced equivalence notions in case of labelling-based semantics.

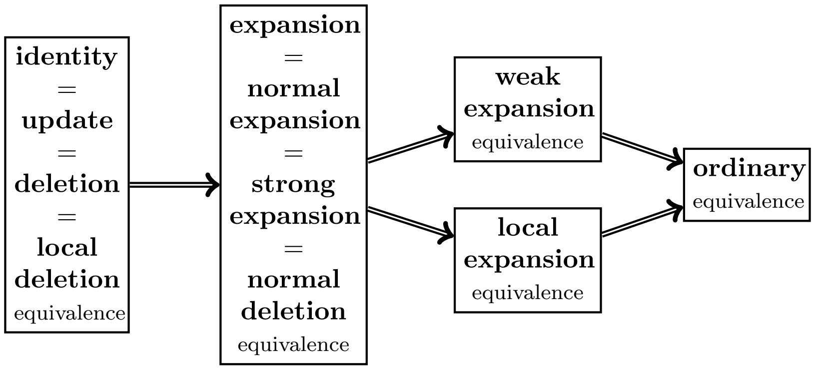

Figure 16: Relations for $\sigma \in \{stb, ss, eg, ad, pr, il, gr, co\}$ - Labelling-based Versions and Finite AFs

4.5 Final Remarks

In this section we motivated and discussed several notions of equivalence in the context of abstract argumentation and provided an exhaustive number of characterization theorems for extension-based as well as labelling-based semantics. In general we may state that Dung's abstract argumentation frameworks are a very compact formalism since the majority of the considered equivalence notion possess only little space for redundancy. Moreover, most of these notions collapse to identity if self-loop-free AFs are considered. This means, in this case any subframework of the AF in question may play a decisive role w.r.t. further evaluations and thus, cannot be locally replaced by another. This insight is sometimes used as an argument against the usefulness of the study of equivalence notions in the context of abstract argumentation. Obviously, we agree that if you are expecting much space for simplification, then the results are somehow disappointing but let us not lose sight of the fact that this is only clear *after* it has been proved. Furthermore, as already stated, the results underline that in case of abstract argumentation (almost) everything is meaningful similar to other non-monotonic formalisms available in the literature (cf. [Lifschitz *et al.*, 2001] for logic programs, [Turner, 2004] for causal theories, [Turner, 2001] for default logic and [Truszczynski, 2006] for nonmonotonic logics in general). However, one decisive difference to these formalisms is that equivalence notions in case of abstract argumentation can be decided syntactically. Indeed, kernels are interesting from several perspectives: First, they allow to decide the corresponding notion of equivalence by a simple check for topological equality and second, all kernels we have obtained so far can be efficiently constructed from a given argumentation framework. This means, if a certain equivalence notion is characterizable through such a kernel, then we have tractability of the associated decision problem.

Acknowledgements

We would like to thank Wolfgang Dvořák for fruitful discussions which helped to improve the article.

BIBLIOGRAPHY

[Alchourrón *et al.*, 1985] C. E. Alchourrón, P. Gärdenfors, and D. Makinson. On the logic of theory change: Partial meet contraction and revision functions. *Journal of Symbolic Logic*, 50:510–530, 1985.

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

[Baroni *et al.*, 2005] Pietro Baroni, Massimiliano Giacomin, and Giovanni Guida. SCC-recursiveness: a general schema for argumentation semantics. *Artificial Intelligence*, 168:162–210, 2005.

[Baroni *et al.*, 2011] Pietro Baroni, Martin Caminada, and Massimiliano Giacomin. An introduction to argumentation semantics. *The Knowledge Engineering Review*, 26:365–410, 2011.

[Baroni *et al.*, 2013] Pietro Baroni, Federico Cerutti, Paul E. Dunne, and Massimiliano Giacomin. Automata for infinite argumentation structures. *Artificial Intelligence*, 203:104–150, 2013.

[Baroni *et al.*, 2014] Pietro Baroni, Guido Boella, Federico Cerutti, Massimiliano Giacomin, Leendert W. N. van der Torre, and Serena Villata. On the input/output behavior of argumentation frameworks. *Artificial Intelligence*, 217:144–197, 2014.

[Baumann and Brewka, 2010] Ringo Baumann and Gerhard Brewka. Expanding argumentation frameworks: Enforcing and monotonicity results. In *Computational Models of Argument: Proceedings of COMMA 2010, Desenzano del Garda, Italy, September 8-10, 2010.*, pages 75–86, 2010.

[Baumann and Brewka, 2013] Ringo Baumann and Gerhard Brewka. Analyzing the equivalence zoo in abstract argumentation. In *Computational Logic in Multi-Agent Systems - 14th International Workshop, CLIMA XIV, Corunna, Spain, September 16-18, 2013. Proceedings*, pages 18–33, 2013.

[Baumann and Brewka, 2015a] Ringo Baumann and Gerhard Brewka. AGM meets abstract argumentation: Expansion and revision for dung frameworks. In *IJCAI, Proceedings of the Twenty-Fourth International Joint Conference on Artificial Intelligence, 2015, Buenos Aires, Argentina, July 25-31, 2015*, pages 2734–2740, 2015.

[Baumann and Brewka, 2015b] Ringo Baumann and Gerhard Brewka. The equivalence zoo for Dung-style semantics. *Journal of Logic and Computation*, 2015.

[Baumann and Spanring, 2015] Ringo Baumann and Christof Spanring. Infinite argumentation frameworks - On the existence and uniqueness of extensions. In *Advances in Knowledge Representation, Logic Programming, and Abstract Argumentation - Essays Dedicated to Gerhard Brewka on the Occasion of His 60th Birthday*, pages 281–295, 2015.

[Baumann and Spanring, 2017] Ringo Baumann and Christof Spanring. A study of unrestricted abstract argumentation frameworks. In *IJCAI, Proceedings of the 26th International Joint Conference on Artificial Intelligence, Melbourne, China, Australia, August 19-25 ,2017*, page to appear, 2017.

[Baumann and Strass, 2012] Ringo Baumann and Hannes Strass. Default reasoning about actions via abstract argumentation. In *Computational Models of Argument: Proceedings of COMMA 2012,, September 10-12, 2012, Vienna, Austria*, pages 297–309, 2012.

[Baumann and Strass, 2013] Ringo Baumann and Hannes Strass. On the maximal and average numbers of stable extensions. In *Theory and Applications of Formal Argumentation - Second International Workshop, TAFA 2013, Beijing, China, August 3-5, 2013, Revised Selected papers*, pages 111–126, 2013.

[Baumann and Strass, 2015] Ringo Baumann and Hannes Strass. Open problems in abstract argumentation. In *Advances in Knowledge Representation, Logic Programming, and Abstract Argumentation - Essays Dedicated to Gerhard Brewka on the Occasion of His 60th Birthday*, pages 325–339, 2015.

[Baumann and Strass, 2016] Ringo Baumann and Hannes Strass. An abstract logical approach to characterizing strong equivalence in logic-based knowledge representation formalisms. In *Principles of Knowledge Representation and Reasoning: Proceedings of the Fifteenth International Conference, KR 2016, Cape Town, South Africa, April 25-29, 2016.*, pages 525–528, 2016.

[Baumann and Woltran, 2016] Ringo Baumann and Stefan Woltran. The role of self-attacking arguments in characterizations of equivalence notions. *Journal of Logic and Computation*, 26(4):1293–1313, 2016.

[Baumann *et al.*, 2011] Ringo Baumann, Gerhard Brewka, and Renata Wong. Splitting argumentation frameworks: An empirical evaluation. In *Theorie and Applications of Formal Argumentation - First International Workshop, TAFA 2011. Barcelona, Spain, July 16-17, 2011, Revised Selected Papers*, pages 17–31, 2011.

[Baumann *et al.*, 2012] Ringo Baumann, Gerhard Brewka, Wolfgang Dvorák, and Stefan Woltran. Parameterized splitting: A simple modification-based approach. In *Correct Reasoning - Essays on Logic-Based AI in Honour of Vladimir Lifschitz*, pages 57–71, 2012.

[Baumann *et al.*, 2014a] Ringo Baumann, Wolfgang Dvorák, Thomas Linsbichler, Hannes Strass, and Stefan Woltran. Compact argumentation frameworks. In *ECAI 2014 - 21st European Conference on Artificial Intelligence, 18-22 August 2014, Prague, Czech Republic - Including Prestigious Applications of Intelligent Systems (PAIS 2014)*, pages 69–74, 2014.

[Baumann *et al.*, 2014b] Ringo Baumann, Wolfgang Dvorák, Thomas Linsbichler, Hannes Strass, and Stefan Woltran. Compact argumentation frameworks. *CoRR*, abs/1404.7734, 2014.

[Baumann *et al.*, 2016a] Ringo Baumann, Wolfgang Dvorák, Thomas Linsbichler, Christof Spanring, Hannes Strass, and Stefan Woltran. On rejected arguments and implicit conflicts: The hidden power of argumentation semantics. *Artificial Intelligence*, 241:244–284, 2016.

[Baumann *et al.*, 2016b] Ringo Baumann, Thomas Linsbichler, and Stefan Woltran. Verifiability of argumentation semantics. *CoRR*, abs/1603.09502, 2016.

[Baumann *et al.*, 2017] Ringo Baumann, Wolfgang Dvořák, Thomas Linsbichler, and Stefan Woltran. A general notion of equivalence for abstract argumentation. In *IJCAI, Proceedings of the 26th International Joint Conference on Artificial Intelligence, Melbourne, China, Australia, August 19-25 ,2017*, page to appear, 2017.

[Baumann, 2011] Ringo Baumann. Splitting an argumentation framework. In *Logic Programming and Nonmonotonic Reasoning - 11th International Conference, LPNMR 2011, Vancouver, Canada, May 16-19, 2011. Proceedings*, pages 40–53, 2011.

[Baumann, 2012a] Ringo Baumann. Normal and strong expansion equivalence for argumentation frameworks. *Artificial Intelligence*, 193:18–44, 2012.

[Baumann, 2012b] Ringo Baumann. What does it take to enforce an argument? minimal change in abstract argumentation. In *ECAI 2012 - 20th European Conference on Artificial Intelligence. Including Prestigious Applications of Artificial Intelligence (PAIS-2012) System Demonstrations Track, Montpellier, France, August 27-31 , 2012*, pages 127–132, 2012.

[Baumann, 2014a] Ringo Baumann. Context-free and context-sensitive kernels: Update and deletion equivalence in abstract argumentation. In *ECAI 2014 - 21st European Conference on Artificial Intelligence, 18-22 August 2014, Prague, Czech Republic - Including Prestigious Applications of Intelligent Systems (PAIS 2014)*, pages 63–68, 2014.

[Baumann, 2014b] Ringo Baumann. *Metalogical Contributions to the Nonmonotonic Theory of Abstract Argumentation.* College Publications - Studies in Logic, 2014.

[Baumann, 2016] Ringo Baumann. Characterizing equivalence notions for labelling-based semantics. In *Principles of Knowledge Representation and Reasoning: Proceedings of the Fifteenth International Conference, KR 2016, Cape Town, South Africa, April 25-29, 2016.*, pages 22–32, 2016.

[Belardinelli *et al.*, 2015] Francesco Belardinelli, Davide Grossi, and Nicolas Maudet. Formal analysis of dialogues on infinite argumentation frameworks. In *IJCAI, Proceedings of the 24th International Joint Conference on Artificial Intelligence, Buenos Aires, Argentina, July 25-31 ,2015*, pages 861–867, 2015.

[Besnard and Hunter, 2001] Philippe Besnard and Anthony Hunter. A logic-based theory of deductive arguments. *Artificial Intelligence*, 128:203–235, 2001.

[Besnard and Hunter, 2009] Philippe Besnard and Anthony Hunter. Argumentation based on classical logic. In *Argumentation in Artificial Intelligence*, pages 133–152. Springer, 2009.

[Brewka and Woltran, 2010] Gerhard Brewka and Stefan Woltran. Abstract dialectical frameworks. In *Principles of Knowledge Representation and Reasoning: Proceedings of the Twelfth International Conference, KR 2010, Toronto, Ontario, Canada, May 9-13, 2010*, 2010.

[Brewka *et al.*, 2013] Gerhard Brewka, Hannes Strass, Stefan Ellmauthaler, Johannes Peter Wallner, and Stefan Woltran. Abstract dialectical frameworks revisited. In *IJCAI, Proceedings of the 23rd International Joint Conference on Artificial Intelligence, Beijing, China, August 3-9, 2013*, pages 803–809, 2013.

[Brewka, 1992] Gerhard Brewka. Nonmonotonic reasoning: Logical foundations of commonsense. *SIGART Bull.*, 3(2):28–29, April 1992. Reviewer-Marek, V. W.

[Caminada and Amgoud, 2007] Martin Caminada and Leila Amgoud. On the evaluation of argumentation formalisms. *Artificial Intelligence*, 171:286–310, 2007.

[Caminada and Oren, 2014] Martin W. A. Caminada and Nir Oren. Grounded semantics and infinitary argumentation frameworks. In *Benelux Conference on Artificial Intelligence*, 2014.

[Caminada and Verheij, 2010] Martin W.A. Caminada and Bart Verheij. On the existence of semi-stable extensions. In *Benelux Conference on Artificial Intelligence*, 2010.

[Caminada *et al.*, 2012] Martin W. Caminada, Walter A. Carnielli, and Paul E. Dunne. Semi-stable semantics. *Journal of Logic and Computation*, 22:1207–1254, 2012.

[Caminada, 2006] Martin Caminada. Semi-stable semantics. In *Computational Models of Argument: Proceedings of COMMA 2006, September 11-12, 2006, Liverpool, UK*, pages 121–130, 2006.

[Caminada, 2011] Martin Caminada. A labelling approach for ideal and stage semantics. *Argument and Computation*, 2:1–21, 2011.

[Dantsin *et al.*, 1997] Evgeny Dantsin, Thomas Eiter, Georg Gottlob, and Andrei Voronkov. Complexity and expressive power of logic programming. In *Proceedings of the Twelfth Annual IEEE Conference on Computational Complexity, Ulm, Germany, June 24-27, 1997*, pages 82–101, 1997.

[Delgrande and Peppas, 2015] James P. Delgrande and Pavlos Peppas. Belief revision in horn theories. *Artificial Intelligence*, 218:1–22, 2015.

[Delgrande *et al.*, 2008] James P. Delgrande, Torsten Schaub, Hans Tompits, and Stefan Woltran. Belief revision of logic programs under answer set semantics. In *Principles of Knowledge Representation and Reasoning: Proceedings of the Eleventh International Conference, KR 2008, Sydney, Australia, September 16-19, 2008*, pages 411–421, 2008.

[Delgrande *et al.*, 2013] James P. Delgrande, Pavlos Peppas, and Stefan Woltran. Agm-style belief revision of logic programs under answer set semantics. In *Logic Programming and Nonmonotonic Reasoning, 12th International Conference, LPNMR 2013, Corunna, Spain, September 15-19, 2013. Proceedings*, pages 264–276, 2013.

[Diller *et al.*, 2015] Martin Diller, Adrian Haret, Thomas Linsbichler, Stefan Rümmele, and Stefan Woltran. An extension-based approach to belief revision in abstract argumentation. In *IJCAI, Proceedings of the Twenty-Fourth International Joint Conference on Artificial Intelligence, 2015, Buenos Aires, Argentina, July 25-31, 2015*, pages 2926–2932, 2015.

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

[Dunne *et al.*, 2013] Paul E. Dunne, Wolfgang Dvořák, Thomas Linsbichler, and Stefan Woltran. Characteristics of multiple viewpoints in abstract argumentation. In *4th Workshop on Dynamics of Knowledge and Belief*, pages 16–30, 2013.

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

[Dunne *et al.*, 2016] Paul E. Dunne, Christof Spanring, Thomas Linsbichler, and Stefan Woltran. Investigating the relationship between argumentation semantics via signatures. In *IJCAI, Proceedings of the Twenty-Fifth International Joint Conference on Artificial Intelligence, 2016, New York, NY, USA, 9-15 July 2016*, pages 1051–1057, 2016.

[Dvořák and Spanring, 2016] Wolfgang Dvořák and Christof Spanring. Comparing the expressiveness of argumentation semantics. *Journal of Logic and Computation*, 2016.

[Dvorák *et al.*, 2011] Wolfgang Dvorák, Paul E. Dunne, and Stefan Woltran. Parametric properties of ideal semantics. In *IJCAI, Proceedings of the 22nd International Joint Conference on Artificial Intelligence, Barcelona, Catalonia, Spain, July 16-22, 2011*, pages 851–856, 2011.

[Dvořák and Gaggl, 2012] Wolfgang Dvořák and Sarah Alice Gaggl. Incorporating stage semantics in the scc-recursive schema for argumentation semantics. In *Proceedings of the 14th International Workshop on Non-Monotonic Reasoning (NMR 2012)*, 2012.

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

[Dyrkolbotn, 2014] Sjur Kristoffer Dyrkolbotn. How to argue for anything: Enforcing arbitrary sets of labellings using afs. In *Principles of Knowledge Representation and Reasoning: Proceedings of the Fourteenth International Conference, KR 2014, Vienna, Austria, July 20-24, 2014*, 2014.

[Eiter *et al.*, 2013] Thomas Eiter, Michael Fink, Jörg Pührer, Hans Tompits, and Stefan Woltran. Model-based recasting in answer-set programming. *Journal of Applied Non-Classical Logics*, 23(1-2):75–104, 2013.

[Gabbay *et al.*, 1994] Dov M. Gabbay, C. J. Hogger, and J. A. Robinson, editors. *Handbook of Logic in Artificial Intelligence and Logic Programming: Nonmonotonic Reasoning and Uncertain Reasoning*, volume 3. Oxford University Press, Inc., New York, NY, USA, 1994.

[Gaggl and Dvořák, 2016] Sarah Alice Gaggl and Wolfgang Dvořák. Stage semantics and the SCC-recursive schema for argumentation semantics. *Journal of Logic and Computation*, 26(4):1149–1202, February 2016.

[Gaggl and Woltran, 2013] Sarah Alice Gaggl and Stefan Woltran. The cf2 argumentation semantics revisited. *Journal of Logic and Computation*, 23:925–949, 2013.

[Gottlob, 1992] Georg Gottlob. Complexity results for nonmonotonic logics. *Journal of Logic and Computation*, 2(3):397–425, 1992.

[Lifschitz and Turner, 1994] Vladimir Lifschitz and Hudson Turner. Splitting a logic program. In Pascal Van Hentenryck, editor, *International Conference on Logic Programming*, pages 23–37. MIT Press, 1994.

[Lifschitz *et al.*, 2001] Vladimir Lifschitz, David Pearce, and Agustín Valverde. Strongly equivalent logic programs. *ACM Transactions on Computational Logic*, 2:526–541, 2001.

[Linsbichler *et al.*, 2015] Thomas Linsbichler, Christof Spanring, and Stefan Woltran. The hidden power of abstract argumentation semantics. In *Theory and Applications of Formal Argumentation - Third International Workshop, TAFA 2015, Buenos Aires, Argentina, July 25-26, 2015, Revised Selected Papers*, pages 146–162, 2015.

[Linsbichler *et al.*, 2016] Thomas Linsbichler, Jörg Pührer, and Hannes Strass. A uniform account of realizability in abstract argumentation. In *ECAI 2016 - 22nd European Conference on Artificial Intelligence, 29 August-2 September 2016, The Hague, The Netherlands - Including Prestigious Applications of Artificial Intelligence (PAIS 2016)*, pages 252–260, 2016.

[Oikarinen and Woltran, 2011] Emilia Oikarinen and Stefan Woltran. Characterizing strong equivalence for argumentation frameworks. *Artificial Intelligence*, 175:1985–2009, 2011.

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

[Qi and Yang, 2008] Guilin Qi and Fangkai Yang. A survey of revision approaches in description logics. In *Proceedings of the 21st International Workshop on Description Logics (DL2008), Dresden, Germany, May 13-16, 2008*, 2008.

[Rudin, 1976] Walter Rudin. *Principles of Mathematical Analysis*. International series in pure and applied mathematics. McGraw-Hill, 1976.

[Spanring, 2012] Christof Spanring. Intertranslatability results for abstract argumentation semantics. Master's thesis, Uni Wien, 2012.

[Spanring, 2015] Christof Spanring. Hunt for the collapse of semantics in infinite abstract argumentation frameworks. In *2015 Imperial College Computing Student Workshop, ICCSW 2015, September 24-25, 2015, London, United Kingdom*, pages 70–77, 2015.

[Strass, 2015] Hannes Strass. The relative expressiveness of abstract argumentation and logic programming. In *Proceedings of the Twenty-Ninth AAAI Conference on Artificial Intelligence, January 25-30, 2015, Austin, Texas, USA.*, pages 1625–1631, 2015.

[Tarski, 1955] Alfred Tarski. A lattice-theoretical fixpoint theorem and its applications. *Pacific Journal of Mathematics*, 5(2):285–309, 1955.

[Truszczynski, 2006] Miroslaw Truszczynski. Strong and uniform equivalence of nonmonotonic theories - an algebraic approach. *Annals of Mathematics and Artificial Intelligence*, 48:245–265, 2006.

[Turner, 1996] Hudson Turner. Splitting a default theory. In William J. Clancey and Daniel S. Weld, editors, *Conference on Artificial Intelligence and Innovative Applications of Artificial Intelligence Conference*, pages 645–651. AAAI Press / The MIT Press, 1996.

[Turner, 2001] Hudson Turner. Strong equivalence for logic programs and default theories (made easy). In *Logic Programming and Nonmonotonic Reasoning, 6th International Conference, LPNMR 2001, Vienna, Austria, September 17-19, 2001, Proceedings*, pages 81–92, 2001.

[Turner, 2004] Hudson Turner. Strong equivalence for causal theories. In *Logic Programming and Nonmonotonic Reasoning, 7th International Conference, LPNMR 2004, Fort Lauderdale, FL, USA, January 6-8, 2004, Proceedings*, pages 289–301, 2004.

[Verheij, 2003] Bart Verheij. Deflog: On the logical interpretation of prima facie justified assumptions. *Journal of Logic and Computation*, 13(3):319–346, 2003.

[Weydert, 2011] Emil Weydert. Semi-stable extensions for infinite frameworks. In *Benelux Conference on Artificial Intelligence*, pages 336–343, 2011.

[Williams and Antoniou, 1998] Mary-Anne Williams and Grigoris Antoniou. A strategy for revising default theory Extensions. In *Principles of Knowledge Representation and Reasoning: Proceedings of the Sixth International Conference, KR 1998, Trento, Italy, June 2-5, 1998.*, pages 24–35, 1998.

[Zorn, 1935] Max Zorn. A remark on method in transfinite algebra. *Bulletin of the American Mathematical Society*, 41:667–670, 1935.

Ringo Baumann
Computer Science Institute, Intelligent Systems Group
Leipzig University, Germany
Email: baumann@informatik.uni-leipzig.de