Artificial Intelligence 324 (2023) 104011

Contents lists available at ScienceDirect

Artificial Intelligence

journal homepage: www.elsevier.com/locate/artint

A claim-centric perspective on abstract argumentation semantics: Claim-defeat, principles, and expressiveness

Wolfgang Dvořák, Anna Rapberger *, Stefan Woltran

TU Wien, Institute of Logic and Computation, Austria

ARTICLE INFO

Article history:
Received 17 June 2022
Received in revised form 1 September 2023
Accepted 3 September 2023
Available online 7 September 2023

Keywords:
Argumentation semantics
Claim-based reasoning
Principle-based analysis
Expressiveness
Claim-augmented argumentation frameworks

ABSTRACT

Dung's *abstract argumentation frameworks (AFs)* are a key formalism in AI research nowadays. *Claims* are an inherent part of each argument; they substantially determine the structure of the abstract representation. Nevertheless, they are often not taken into account on the abstract level, which restricts the modeling capacities of AFs to problems that do not involve claims in the evaluation. In this work, we address this shortcoming and conduct a structural analysis of claim-based argumentation semantics utilizing *claim-augmented argumentation frameworks (CAFs)* which extend AFs by assigning a claim to each argument. Our main contributions are as follows: We first propose novel variants for preferred, naive, stable, semi-stable, and stage semantics based on claim-defeat and claim-set maximization, complementing existing CAF semantics. Among our findings is that for a certain subclass, namely *well-formed CAFs*, the different versions of preferred and stable semantics coincide, which is not the case for the other semantics. We then conduct a principle-based analysis of the semantics with respect to general and well-formed CAFs. Finally, we study the expressiveness of the semantics by characterizing their signatures. In summary, this paper provides a thorough analysis of fundamental properties of abstract argumentation semantics (along the lines of existing results for AFs) but from the perspective of the claims the arguments represent. This shift of perspective provides novel results which we deem relevant when abstract argumentation is used in an instantiation-based setting.

1. Introduction

Argumentation is a vibrant research area in the field of non-monotonic reasoning and knowledge representation. Originally studied predominantly under philosophical and linguistic aspects, argumentation theory has gained increasing attention in artificial intelligence research nowadays [1–3]. Argumentation theory, in its essence, is concerned with the justification of *defeasible statements (claims)*[1] through logical or evidence-based reasoning. In particular, in the light of an increasing amount of available data and worldwide information on demand, the trustworthy evaluation of defeasible statements is crucial. Thereby, an argumentative analysis provides means to assess the plausibility of claims, even in the light of inconsistent and conflicting information. The construction of arguments that support or challenge the defeasible state-

* Corresponding author.
E-mail addresses: dvorak@dbai.tuwien.ac.at (W. Dvořák), arapberg@dbai.tuwien.ac.at (A. Rapberger), woltran@dbai.tuwien.ac.at (S. Woltran).

[1] As pointed out by van Emeren and Verheij [4], the terms claim, conclusion, thesis, and debate proposition are used to refer to the same concept from different theoretical angles. Following their example, we will use these terms interchangeably; moreover, we add *defeasible statement* to this list.

https://doi.org/10.1016/j.artint.2023.104011

We should go to this restaurant, I heard their food is delicious.

We should go to this restaurant, their interior looks really great.

We should go to this restaurant, they have great reviews.

We should go to this restaurant, the service there is amazing.

(a) Natural language arguments with the same claim (highlighted in italic).

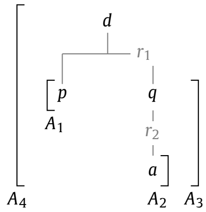

(b) ASPIC argument A_4 with claim d and sub-arguments A_1, A_2, A_3; constructed from rules $r_1 : p, q \to d$ and $r_2 : a \to q$ and premises p, a (cf. [10]).

$(\{a, a \to b\}, b)$

$(\{c, c \to b\}, b)$

$(\{a \wedge b\}, b)$

$(\{d \vee b, \neg d\}, b)$

(c) Logic-based arguments for the same claim b with different supports (consistent and $\subseteq$-minimal sets of formulae deriving b [15]).

Fig. 1. Three different settings depicting arguments and their claims. Fig. 1a depicts several arguments in natural language with the same claim; Figs. 1b and 1c give examples of rule-based and logic-based arguments, respectively.

ments in question can help to assess their plausibility. Arguments reveal the underlying conflicts between different claims and help to pinpoint their weaknesses. Systematic techniques to construct arguments and evaluate their claims—so-called *argumentation procedures* [3]—have proven useful in several different settings. They can model discourses, medical and legal cases [2]; moreover, they provide an orthogonal view on non-monotonic reasoning paradigms such as defeasible reasoning and logic programming [5–7]. The result of these procedures are *sets of jointly acceptable claims*. They represent the different viewpoints within a knowledge base and establish the correspondence to other non-monotonic formalisms. Interestingly, although these procedures are designed to evaluate the acceptability of claims, only little is known about the semantics of argumentation from a claim-centered perspective.

The argumentation process The process typically starts by constructing arguments from a potentially inconsistent knowledge base. An argument is generally understood as a complex structure based on assumptions, premises, or facts (common knowledge) which is used to support a particular claim. Logic-based approaches identify the claim of an argument with a logical formula and the support of the claim with a (defeasible) proof [8,9]. Likewise, the claim in rule-based approaches is a sentence of a formal language while the support consists of assumptions and facts that infer the claim based on a deductive system [10,11]. Other argumentation systems also consider evidence-based support of a defeasible statement. Examples include juridical argumentation [12,13] or decision-making procedures in medicine [14]. Fig. 1 gives examples of arguments in different formalisms.

Deciding the plausibility of a claim or a set of claims crucially depends on the acceptance of the arguments which support them: a set of claims is jointly acceptable iff there exists a set of jointly acceptable arguments with the desired claims (we say: such a set of arguments *realizes* the claim-set).

One of the most prominent approaches for the evaluation of argument acceptance in the area of computational argumentation are *abstract argumentation frameworks (AFs)* [5]. They provide a general schema for analyzing discourses by treating arguments as abstract entities while an attack relation encodes (potentially asymmetric) conflicts between them, thus giving rise to a graph-like representation of inconsistent information. In his seminal paper [5], Dung introduced fundamental argumentation semantics which are based on several crucial concepts, including *conflict-freeness*, which formalizes that arguments which are jointly acceptable should not attack each other; *defense*, stating that a set of jointly acceptable arguments should defend itself, i.e., it should counter-attack each attacker; and *maximality* of the outcome. In the final step, the acceptable claims are determined.

The last step in this procedure often receives only little attention. The acceptability of claims is identified by applying Dung's acceptability criteria to the constructed arguments and extract their claims in the last step. However, the particular way how argumentation semantics can be lifted to the level of claims is not unique. While conflict-freeness is linked to the realization of a particular claim-set since only arguments can attack each other, it is, for instance, less obvious how maximality-criteria translate to semantics on claim-level. Let us consider preferred semantics which returns $\subseteq$-maximal admissible (conflict-free and self-defending) sets. One way is to compute the preferred sets of arguments of the given AF and extract the claims of the acceptable arguments in a second step. Alternatively, $\subseteq$-maximality of admissible sets can be understood on claim-level and select those admissible sets which are $\subseteq$-maximal in terms of their claims. Note that both approaches formalize inherently different concepts, theoretically as well as practically: since several arguments can have the same claim (cf. Fig. 1) it is evident that the two approaches might yield different sets of jointly acceptable claims. Fig. 2 exemplifies these differences.

Claims in argumentation Claims are an integral part of argumentation formalisms; they are inextricably linked to arguments (cf. Fig. 1). Discussions on the difference between claim acceptance and argument acceptance can be traced back to the late eighties [16] where the acceptance of so-called *floating conclusions*, i.e., claims which are supported by different arguments, has been subject of debate. Claim acceptance plays a crucial role in the context of case-based reasoning [17] and has been studied in terms of different justification stages [18]. It has been also subject to a computational complexity analysis [19,20] x: Let's make *holidays at the sea*, hiking holidays are exhausting.

y: Let's make *holidays in the mountains*, swimming all the time is boring.

z: Making *holidays at the sea* can be combined with day trips in the mountains. Tenerife, for instance, has coasts and the highest mountain of Spain.

(a) Arguments supporting holidays at the sea (arguments x and z) resp. in the mountains (argument y).

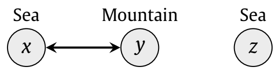

(b) Instantiated abstract framework with arguments x, y, and z. Arguments are labeled with the claim they support. The arguments x and y attack each other symmetrically due to the underlying conflict between hiking (argument x) and swimming (argument y).

Fig. 2. Natural language arguments (left) and their instantiated AF (right). The preferred argument extensions of the AF are $\{x, z\}$ and $\{y, z\}$, yielding the claim extensions {Sea} and {Sea,Mountain}. Maximizing over the admissible claim sets yields the unique claim-set {Sea,Mountain}.

where it has been shown that classical decision problems admit a higher complexity when considered in terms of claims than the analogous problem for AFs. Nevertheless, the role of claims in the evaluation procedure is often neglected. Dung's classical acceptability criteria are considered at the level of the arguments only.

The discrepancy between argument-level and claim-level maximization in, e.g., preferred semantics is then often circumvented by constructing frameworks under structural restrictions such that both variants coincide. However, for a certain class of semantics that take also the attacked elements into account, it is even impossible to capture the claim-level maximization with an according AF semantics using standard instantiation methods (e.g., L-stable semantics of logic programs, cf. Section 2 and [7]). These so-called *range-based semantics* maximize the union of all acceptable and defeated elements. Here, a genuine notion of defeated claims would be necessary to capture these semantics appropriately.

Our contributions In order to grasp the different approaches to maximization in claim-centered argumentation, we propose novel concepts of lifting argumentation semantics to claim-level. We identify two different classes of claim-based semantics: *inherited* and *hybrid semantics*. While inherited semantics (as introduced in [19]) evaluate a given AF and extract claims of the successful sets in the final step of the evaluation, the hybrid semantics give claims a more active role in the determination of the outcome: we consider *maximization on claim-level* (as outlined above regarding preferred semantics) and introduce a genuine notion of *claim-defeat*. Shifting maximization onto the claim-level yields hybrid variants of preferred and naive semantics. Moreover, our notion of claim-defeat gives rise to novel variants of (semi-)stable semantics. Combining both concepts gives rise to hybrid range-based semantics.

We present a systematic analysis of both classes of semantics. We adapt the *principle-based approach* to argumentation semantics [21,22] to the realm of claim-based reasoning; moreover, we consider the *relations* of all claim-based semantics and their expressiveness in terms of their *signature* [23], i.e., the set of all possible extension-sets a framework can possess under the given semantics. The principle-based methodology is well-suited for a systematic analysis of claim-based semantics: such a classification yields theoretical insights into the nature of the different semantics and can help to guide the search for suitable semantics appropriate in different scenarios. The claim-based analysis in terms of principles furthermore complements similar studies on classical argumentation semantics and sheds light on the different levels of arguments and claims (for instance, although preferred AF semantics satisfies the central principle of *I-maximality*, i.e., $\subseteq$-maximality of its extensions, it is not necessarily the case that inherited preferred semantics satisfies I-maximality, as we will see). The characterization of the signature of a semantics is key to understand its expressive power. Knowing which extensions can jointly be modeled within a single framework under a given semantics is, for instance, crucial in dynamic scenarios in which argumentation frameworks undergo certain changes [24,25]. Indeed, the design of dynamic operators that modify the outcome (for instance, removing one element from each extension) crucially depends on the signature of the semantics.

We base our work on *claim-augmented argumentation frameworks (CAFs)* [19] which extend AFs in a minimal way by assigning each argument a claim. We do not assume a particular structure of the claims; in the spirit of abstract argumentation, we consider them, analogous to arguments, as abstract entities. In our work, we furthermore investigate the behavior of claim-based semantics when restricted to a class of CAFs that appears in many instantiation procedures: in *well-formed CAFs*, arguments with the same claim attack the same arguments. Well-formedness conforms with a common behavior of attack construction: in many structured argumentation formalisms, an argument x attacks another argument y if the claim of x contradicts the support of y. In Fig. 1c, for instance, an argument with claim $\neg a$ attacks the first and the third argument; an argument with claim d attacks the last argument in this list. Nevertheless, there are several situations in which well-formedness is violated, for instance if preferences are taken into account [26,27]. Since we consider both classes of CAFs—unrestricted and well-formed CAFs—equally valuable, we take both of them into consideration.

Our main contributions are as follows:

- We introduce hybrid variants of preferred, naive, stable, semi-stable and stage semantics. By that, we provide argumentation semantics that shift maximization of extensions from argument-level to claim-level. We compare them to their inherited counter-parts.
- We provide a full picture of the relations between all considered inherited and hybrid semantics for both general and well-formed CAFs.

- We conduct a principle-based analysis of all claim-based (i.e., inherited and hybrid) argumentation semantics that we consider in this work. We introduce novel principles genuinely defined for claim-based semantics and study well-known properties of argumentation semantics such as e.g., I-maximality, naivety, and reinstatement in terms of claim-based reasoning. We compare hybrid semantics and inherited semantics as well as general CAFs and well-formed CAFs with respect to these properties.
- Finally, we study the expressiveness of claim-based semantics. We characterize the signatures of the considered semantics for both general CAFs and well-formed CAFs.

Parts of this paper have been presented at the 9th European Starting AI Researchers' Symposium (STAIRS), see [28], and at the 17th International Conference on Principles of Knowledge Representation and Reasoning (KR 2020), cf. [29]. The conference and workshop versions present hybrid semantics for CAFs and include results on I-maximality, relations between the semantics, and expressiveness of both variants of stable, semi-stable, stage, preferred, and naive semantics. The present paper extends the previous work significantly. Apart from providing full proofs in the present version, we improved the notation regarding hybrid and inherited semantics (cf. Remark 1.1) and claim-defeat. Furthermore, we significantly deepen the theoretical investigations of the claim-based semantics in the present paper: we conclude an in-depth principle-based analysis of the considered semantics with respect to general and well-formed CAFs (cf. Section 6); moreover, we extend our expressiveness results to admissible and complete semantics (cf. Section 7).

Remark 1.1. Following [30], we revised the notation and names of the CAF semantics in the present work. Originally, in [28, 29], we introduce hybrid semantics under the name *claim-level semantics* and denote hybrid (former: claim-level) semantics by *cl*-σ and inherited semantics by σ_c. Accordingly with the new naming, we have adapted the notation: in the present work, we write σ_i instead of σ_c to denote the inherited variant of the semantics σ; likewise, we write σ_h instead of *cl*-σ to denote the hybrid variant.

2. Logic programs and claim-based argumentation: a natural fit

The instantiation of logic programs (LPs) into AFs and generalizations thereof has been frequently discussed in the literature [5–7] and reveals the close connection between both formalisms. The correspondence of stable model semantics for LPs with stable semantics in AFs is probably the most fundamental example [5], but also other semantics of LP admit equivalent argumentation semantics [6]. In this section, we examine the close connection of logic programs and conclusion-driven argumentation formalisms. We reveal shortcomings of classical AF instantiations of logic programs regarding 3-valued model semantics and the maximization of range-based semantics. We propose an adaption of range-based semantics (in particular, for semi-stable semantics) for CAFs that naturally covers maximization on atom-level in LPs and thus gives rise to the missing argumentation-based counterpart of the so called L-stable model semantics.

We consider normal logic programs that consist of a set of rules of the form

$$c \leftarrow a_1, \ldots, a_m, \text{not } b_{m+1}, \ldots, \text{not } b_m$$

where c is the head and $\{a_1, \ldots, a_m, \text{not } b_{m+1}, \ldots, \text{not } b_m\}$ is called the body of the rule. By $\mathcal{L}(P)$ we denote the set of all atoms appearing in P. There are many different semantics for evaluating logic programs, most prominently stable model semantics. In the following, we consider *partially stable (p-stable) model semantics* based on 3-valued model semantics that generalize 2-valued model semantics by allowing for undefined atoms. A 3-valued model of a program P is a tuple (T, F) with minimal T and maximal F such that $T \cap F = \emptyset$, $T \cup F \subseteq \mathcal{L}(P)$, and (T, F) satisfies all (apart from certain undefined) rules in P. Atoms in T are considered *true* and atoms in F are set to *false*. Atoms that are neither contained in T nor F are considered *undefined* (in contrast to 2-valued model semantics it is not required that $T \cup F = \mathcal{L}(P)$). The model (T, F) or sometimes simply T is also called a *p-stable model* of P. We refer to Section 4.4 for the formal definition of these semantics.

In the following example, we will adapt an instantiation by [7].

Example 2.1. Consider the following logic program P:

$$\begin{array}{llllll} r_0: & a \leftarrow \text{not } d & r_2: & b \leftarrow \text{not } a & r_4: & e \leftarrow \text{not } e \\ r_1: & d \leftarrow \text{not } a & r_3: & c \leftarrow \text{not } a, \text{not } b & r_5: & e \leftarrow \text{not } a, \text{not } e \end{array}$$

We construct our corresponding AF as follows: First, each rule r_i yields an argument x_i. Second, attacks between arguments are obtained by considering the negative body elements of the associated rules as possible weak points: an argument x_i

attacks another argument x_j if the head of the corresponding rule r_i appears negated in the body of the rule r_j. We obtain the following AF (a formal definition of the translation can be found in Section 4.4):

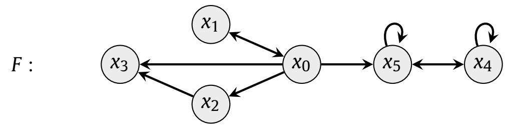

As the reader may verify, both the program P as well as its associated AF have no stable models (indeed, the atom e can neither be set to true nor to false in P, and the argument x_4 can never be member of or attacked by a conflict-free set in F). Under 3-valued model semantics that returns, roughly speaking, stable models that allow for undecided atoms, the program P returns the following answer sets (we consider only the atoms set to true in the 3-valued models): $\emptyset$, $\{a\}$, and $\{d, b\}$. As shown in [7], complete semantics preserves p-stable model semantics under the presented translation using an additional mapping that extracts the claims of the arguments: the complete extensions of F are given by $\emptyset$, $\{x_0\}$, and $\{x_1, x_2\}$; extracting the corresponding claims yields the p-stable models of P.

There is however a weakness in the translation: establishing the correspondence between AFs and LPs requires an intermediate step (i.e., the re-interpretation of the arguments in terms of their claims) and is thus not directly given. This issue can be circumvented by extending AFs in a minimal way: *claim-augmented argumentation frameworks (CAFs)* as introduced by [19] allow for assigning an abstract element to each argument that is considered the *claim* of the argument. This simple extension streamlines the correspondence between LPs and abstract argumentation in a natural way:

Example 2.2 *(Example 2.1 ctd.).* Consider again our LP from Example 2.1 and the AF instantiation, we obtain a CAF naturally by assigning each argument x_i the atom in the head of the corresponding rule. We obtain the following CAF:

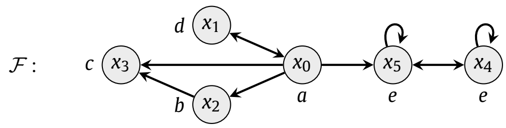

Utilizing the CAF representation of our LP P, we obtain the correspondence in a direct fashion: following [19], the accepted claim-sets of a CAF are obtained by evaluating the underlying AF and interpreting the conclusions in terms of the claims. The complete claim-sets of $\mathcal{F}$ are thus given by $\emptyset$, $\{a\}$, and $\{d, b\}$ which coincides with the p-stable models of our program P.

The representation of such instantiation procedures as CAFs handles the correspondence of conclusion-oriented formalisms such as logic programs without detours, i.e., no additional steps or mappings are needed. In this way, CAFs establish a closer relation between the two paradigms. Let us furthermore point out a conceptual advantage of CAFs that goes beyond their usage regarding instantiation procedures: with CAFs it is possible to capture situations in which two arguments represent the same conclusion, a scenario which cannot be formalized with standard argumentation frameworks without further assumptions (by definition of AFs, arguments are considered abstract, hence some additional function or modification of the definition would be required).

This observation reveals an even more powerful advantage of CAFs: they are flexible enough to capture semantics that make use of the conclusions in the evaluation. This advantage, however, has not been fully exploited so far. In [19], semantics fully depend on the argument structure of the AF. Claims come into play only in the very last step of the procedure when the accepted sets of arguments have already been identified; in this final step, the claim of each argument is inspected. These so-called *inherited semantics* yield the desired results in many situations; there are however semantics for which this simple evaluation procedure does not suffice.

Let us consider *semi-stable semantics* [31,32]. Semi-stable semantics yields admissible sets (i.e., sets of arguments that are conflict-free and defend themselves) with maximal range, i.e., it returns sets that contain as much accepted or explicitly rejected arguments as possible. This semantics constitutes a relaxation of stable semantics; it coincides with stable semantics if stable extensions exist and returns close-to-optimal solutions (models with as few undecided arguments as possible) otherwise. We note that semi-stable semantics is closely tied to the existence of odd cycles, i.e., it is well-known that semi-stable and stable semantics coincide if no odd cycles exist. *Least undefined stable (L-stable) model semantics* [33,34] can be considered as their LP-counterpart: here, the set of all atoms which are either considered true or false in a 3-valued stable model is maximized. However, when evaluating our running example with respect to this semantics, we observe an undesired discrepancy in the outcome.

Example 2.3 *(Example 2.1 ctd.).* Consider again our running example LP P and its associated CAF $\mathcal{F}$. We compute the L-stable models of P and obtain $\{a\}$ (atoms b, c, d are set to false) and $\{d, b\}$ (here, a, c are set to false). Evaluating the CAF $\mathcal{F}$ under semi-stable semantics however yields a unique extension: the argument x_0 attacks all remaining arguments except x_4 and is thus maximal in this respect. Therefore, $\{a\}$ is the only semi-stable claim-set of $\mathcal{F}$.[2]

This difference has been already observed by Caminada et al. [7] who proved that it is impossible to capture L-stable semantics on argument level under standard instantiation methods.

While it has been shown that inherited semantics often behave correctly, the example above reveals that, in some situations, results may deviate from the expected outcome of the original problem. A crucial observation is that semantics for LPs operate on conclusion (claim) level while abstract argumentation semantics as well as inherited CAF semantics are evaluated on argument level. We are thus interested in developing adequate semantics for CAFs which mimic the behavior of semantics performing maximization on conclusion-level of the original problem (e.g. L-stable model semantics for LPs).

We observe two sources that may lead to a different outcome of the respective evaluation methods:

- First, maximization is considered on different levels. In LPs, we maximize over sets of atoms while in the associated CAFs we maximize over arguments. This is however a mismatch since atoms in the LP correspond to *claims* in the CAF.
- The second issue is more subtle: while we successfully identify the claims of acceptable arguments with atoms that are set to true, we do not have a similar correspondence for atoms that are set to false. Coming back to our running example, we observe that the arguments with claim e (i.e., x_4, x_5) play a different role for the claim-sets $\{a\}$ and $\{b,d\}$: the realization $\{x_0\}$ of $\{a\}$ attacks one of them while the realization $\{x_1,x_2\}$ of $\{b,d\}$ does not. In the LP P on the other hand, the atom e is undecided with respect to both L-stable models. The underlying issue here is that evaluation methods for CAFs that have been considered so far do not take the *defeat of claims*, i.e., the successful attack of *all* arguments for a given claim, into account.

Inspired by these observations, we propose semantics that operate on the claim level (*hybrid semantics* or *h-semantics*, for short). With these adjustments, we are able to capture semantics of conclusion-oriented formalisms. Let us demonstrate the idea:

Example 2.4 *(Example 2.1 ctd.).* Let us consider again our CAF $\mathcal{F}$ and its complete argument-sets $\emptyset$, $\{x_0\}$, and $\{x_1,x_2\}$. We propose a new evaluation method for semi-stable semantics (co-called hybrid-semi-stable or *h-semi-stable semantics*) by maximizing accepted and defeated claims: The set $\{x_0\}$ defeats the claims b,c,d; the claim e is not defeated because x_0 does not attack all occurrences of e. The set of accepted and defeated claims with respect to the extension $\{x_0\}$ (the *claim-range* of $\{x_0\}$) is thus given by $\{a,b,c,d\}$. The set $\{x_1,x_2\}$ defeats the claims a,c. Therefore, $\{x_1,x_2\}$ has claim-range $\{a,b,c,d\}$ which coincides with the claim-range of $\{x_0\}$.

We can verify that these extensions are indeed our only h-semi-stable claim-sets: the claim-range of the empty set is empty; moreover, the set $\{x_1,x_3\}$ only defeats the claim a, therefore, its range $\{a,c,d\}$ is not $\subseteq$-maximal. We thus obtain that cl-semi-stable semantics in $\mathcal{F}$ yield the same outcome as L-stable model semantics for P.

Logic programming semantics is one out of many examples that adapt a conclusion-focused view in the evaluation; other examples are instances of default logic [35] or assumption-based argumentation [11]. In this work, we want to develop argumentation semantics that capture conclusion-oriented reasoning methods in its full expressiveness. In the scope of our studies, we will successfully identify argumentation semantics that matches L-stable model semantics (the definition of h-semi-stable semantics is given in Section 4.3). The result is provided in Section 4.4.

3. Preliminaries

Abstract argumentation We introduce abstract argumentation frameworks [5]; for a comprehensive introduction, see [3,36]. We fix U as the countably infinite domain of arguments.

Definition 3.1. An *argumentation framework (AF)* is a pair $F=(A,R)$ where $A\subseteq U$ is a finite set of arguments and $R\subseteq A\times A$ is the attack relation. Given an argument a, we say that *a attacks b* iff $(a,b)\in R$; a set of arguments $E\subseteq A$ *attacks b* iff $(a,b)\in R$ for some $a\in E$; E attacks another set of arguments $D\subseteq A$ iff E attacks some argument $b\in D$. A set of arguments $E\subseteq A$ *defends* an argument $a\in A$ iff E attacks each attacker of a. We use $a_F^+=\{b\mid(a,b)\in R\}$ and $a_F^-=\{b\mid(b,a)\in R\}$; we extend both notions to sets S as expected: $E_F^+=\bigcup_{a\in E}a_F^+$, $E_F^-=\bigcup_{a\in E}a_F^-$. We call $E_F^\oplus=E\cup E_F^+$ the *range* of E in F. If no ambiguity arises, we drop the subscript F.

[2] The attentive reader may have wondered whether the issue stems from the self-attacking arguments in the example. Let us point out that this is not the case. While the existence of odd-length cycles is crucial for distinguishing semi-stable and stable semantics, as mentioned in the previous paragraph, the actual length of the odd loops does not play a role. The present example can be modified such that the self-attackers are replaced with odd-length cycles greater than one; however, for the sake of simplicity, we decided to keep our examples as concise as possible.

Semantics for AFs are defined as functions σ which assign to each AF $F = (A, R)$ a set $\sigma(F) \subseteq 2^A$ of extensions. We consider for σ the functions *cf*, *ad*, *co*, *na*, *gr*, *stb*, *pr*, *ss* and *stg* which stand for conflict-free, admissible, complete, naive, grounded, stable, preferred, semi-stable and stage, respectively.

Definition 3.2. Let $F = (A, R)$ be an AF. A set $E \subseteq A$ is *conflict-free (in F)*, if there are no $a, b \in E$, such that $(a, b) \in R$. $cf(F)$ denotes the collection of sets being conflict-free in F. For $E \in cf(F)$, we define

- $E \in na(F)$, iff there is no $D \in cf(F)$ with $E \subset D$;
- $E \in ad(F)$, iff each $a \in E$ is defended by E in F;
- $E \in co(F)$, iff $E \in ad(F)$ and each $a \in A$ defended by E in F is contained in E;
- $E \in gr(F)$ iff E is a $\subseteq$-minimal complete extension;
- $E \in pr(F)$ iff E is a $\subseteq$-maximal complete extension;
- $E \in stb(F)$, iff $E^{\oplus}_F = A$;
- $E \in ss(F)$, if $E \in ad(F)$ and $\nexists D \in ad(F)$ with $E^{\oplus}_F \subset D^{\oplus}_F$;
- $E \in stg(F)$, iff $\nexists D \in cf(F)$, with $E^{\oplus}_F \subset D^{\oplus}_F$.

We recall that for each AF F, $stb(F) \subseteq stg(F) \subseteq na(F) \subseteq cf(F)$ and $stb(F) \subseteq ss(F) \subseteq pr(F) \subseteq ad(F)$; also $stb(F) = ss(F) = stg(F)$ in case $stb(F) \neq \emptyset$. Moreover, semantics $\sigma \in \{na, pr, stb, stg, ss\}$ deliver *incomparable* sets (anti-chains): for all $E, D \in \sigma(F)$, $E \subseteq D$ implies $E = D$. The property is also referred to as *I-maximal*.

Claim-based reasoning Next we define claim-augmented argumentation frameworks according to [19]. We fix C as the countably infinite domain of claims.

Definition 3.3. A *claim-augmented argumentation framework (CAF)* is a triple $\mathcal{F} = (A, R, cl)$ where $F = (A, R)$ is an AF and $cl : A \to C$ is a function which assigns a claim to each argument in A. The claim-function is extended to sets in the following way: For a set $E \subseteq A$, $cl(E) = \{cl(a) \mid a \in E\}$. We call an argument $a \in A$ an *occurrence* of claim $cl(a)$ in $\mathcal{F}$. Given a set of claims $S \subseteq cl(A)$, we call a set of arguments $E \subseteq A$ with $cl(E) = S$ a *realization of S in* $\mathcal{F}$.

Notation. For each CAF $\mathcal{F} = (A, R, cl)$ we denote by $F = (A, R)$ the underlying AF (for CAF $\mathcal{G}$, we denote the corresponding AF G). Also, we use subscript-notation $A_{\mathcal{F}}$, $R_{\mathcal{F}}$, and $cl_{\mathcal{F}}$ to indicate the affiliations.

In [19], semantics of CAFs are defined based on the standard semantics of the underlying AF. The extensions are interpreted in terms of the claims of the arguments. We call this variant *inherited semantics* (i-semantics).

Definition 3.4. For a CAF $\mathcal{F}$ and a semantics σ, we define the *inherited variant* of σ as $\sigma_i(\mathcal{F}) = \{cl(E) \mid E \in \sigma(F)\}$. We call a set $E \in \sigma(F)$ with $cl(E) = S$ a σ_i-realization of S in $\mathcal{F}$.

Example 3.5. Consider a CAF $\mathcal{F}$ given as follows:

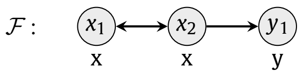

The CAF $\mathcal{F}$ has two arguments that support the same claim: both arguments x_1 and x_2 have the same claim x. The set $\{x\}$ has two realizations, namely $\{x_1\}$ and $\{x_2\}$. Since the sets are conflict-free and admissible, the set $\{x\}$ has two conflict-free (or admissible) realizations. It has, however, only one stb_i-realization: the set $\{x_2\}$ is stable in F as it attacks all remaining arguments whereas $\{x_1\}$ is not.

Basic relations between different semantics carry over from standard AFs, i.e. for any CAF $\mathcal{F}$, $stb_i(\mathcal{F}) \subseteq ss_i(\mathcal{F}) \subseteq pr_i(\mathcal{F}) \subseteq ad_i(\mathcal{F})$ and $stb_i(\mathcal{F}) \subseteq stg_i(\mathcal{F}) \subseteq na_i(\mathcal{F}) \subseteq cf_i(\mathcal{F})$; moreover, if $stb(\mathcal{F}) \neq \emptyset$ then $stb_i(\mathcal{F}) = ss_i(\mathcal{F}) = stg_i(\mathcal{F})$. On the other hand observe that we lose fundamental properties of semantics like I-maximality of preferred, naive, stable, semi-stable, and stage semantics:

Example 3.6. Let us consider again the CAF from Example 3.5. We observe that there are two stable extensions in F: $\{x_2\}$ and $\{x_1, y_1\}$. The resulting i-stable claim-sets are $\{x\}$ and $\{x, y\}$ which shows that i-stable semantics does not necessarily yield incomparable sets. Observe also that $na_i(\mathcal{F}) = stb_i(\mathcal{F}) = ss_i(\mathcal{F}) = stg_i(\mathcal{F}) = pr_i(\mathcal{F})$ in this case. Therefore, the observation furthermore holds for i-preferred, i-naive, i-stage, and i-semi-stable semantics.

We consider a class of CAFs that appears in many different contexts: *well-formed CAFs* incorporate the basic observation that attacks typically depend on the claim of the attacking argument.

Definition 3.7. A CAF (A, R, cl) is called *well-formed* if $a^+_F = b^+_F$ for all $a, b \in A$ such that $cl(a) = cl(b)$.

In well-formed CAFs we can speak of claims attacking arguments: we say that a claim $c \in cl(A)$ *attacks* an argument $a \in A$ if $(x, a) \in R$ for some (and thus for each) argument $x \in A$ having claim c. Likewise, we say that $S \subseteq cl(A)$ attacks $a \in A$ if there is a claim $c \in S$ that attacks a.

Observe that the instantiation procedure from [7] that has been adapted to CAFs in Example 2.1 returns well-formed CAFs since the outgoing attacks depend on the head of the corresponding rule. Indeed, it can be checked that the resulting CAF from Example 2.1 is well-formed. LP instantiations are one of the numerous examples of formalisms with well-formed attack relation [5,9–11]. Well-formed CAFs indeed appear in many different formalisms and applications. The classical instantiation of assumption-based argumentation, for instance, yields well-formed CAFs as well (indeed, it can be checked that the claims attack the same arguments, cf. [37,38]). Nevertheless, we observe that also non-well-formed CAFs play an important role in argumentation, e.g., when one considers preferences between arguments or other elements of the knowledge base (cf. [26, 27]). In ASPIC+, for example, only *successful attacks* are considered in the instantiation of the AF (a CAF instantiation can be easily obtained by assigning each argument its claim). Therefore, it can be the case that arguments with the same claim attack different arguments. As we consider the two classes as equally valuable, we conduct our analysis with respect to both classes.

4. Introducing hybrid semantics: maximization and defeat

In this section, we establish claim-based semantics that perform maximization on sets of acceptable claims as well as on the range on claim-level. Since the acceptability of a claim or a set of claims depends on the acceptability of the arguments that support them, we will consider a hybrid approach to lift these concepts to claims. We establish a defeat notion for claims *relative* to a set of arguments: intuitively, a claim is defeated by a set of arguments E if each occurrence of the claim is attacked by E. Our investigations give rise to novel versions of preferred and naive semantics (when considering maximization of claim-sets) which are discussed in Section 4.1; variants of stable semantics (using our novel notion of claim-defeat) which are introduced in Section 4.2; and semi-stable and stage semantics (when maximizing over sets of accepted and defeated claims) which combine both aspects and are discussed in Section 4.3. In Section 4.4, we discuss the connection to logic programming. We end this section with a summary of our investigations.

4.1. Maximization of claim-sets

Let us first consider two prominent semantics that involve maximization: preferred and naive semantics return $\subseteq$-maximal admissible and conflict-free sets, respectively. We introduce variants of preferred and naive semantics for CAFs by shifting maximization from argument- to claim-level.

Definition 4.1. Given a CAF $\mathcal{F}$ and a set of claims $S \subseteq cl(A_{\mathcal{F}})$. Then

- S is *h-preferred* ($S \in pr_h(\mathcal{F})$) iff S is $\subseteq$-maximal in $ad_i(\mathcal{F})$;
- S is *h-naive* ($S \in na_h(\mathcal{F})$) iff S is $\subseteq$-maximal in $cf_i(\mathcal{F})$.

For a set $S \in pr_h(\mathcal{F})$ ($S \in na_h(\mathcal{F})$), we call $E \in ad(F)$ ($E \in cf(F)$) with $cl(E) = S$ a *pr_h-realization* (*na_h-realization*, respectively) of S in $\mathcal{F}$.

Both variants of preferred (naive) semantics rely on admissibility (conflict-freeness) of the realization. While the inherited variant maximizes on the level of arguments, the newly introduced hybrid variant maximizes on the level of claims.

We consider the following example.

Example 4.2. Let us consider the following two CAFs $\mathcal{F}$ and $\mathcal{F}'$:

The CAF $\mathcal{F}$ already appears in Example 3.5; it is not well-formed and its i-preferred and i-naive claim-sets are $\{x\}$ and $\{x, y\}$: indeed, $\mathcal{F}$ has four conflict-free sets of arguments $\{x_1\}$, $\{y_1\}$, $\{x_1, y_1\}$, and $\{x_2\}$; all except $\{y_1\}$ are admissible; thus, the sets $\{x_1, y_1\}$ and $\{x_2\}$ are naive and preferred in the underlying AF. Extracting the claims of the sets yields $\{x, y\}$ and $\{x\}$.

Now, to compute the h-naive and h-preferred claim-sets of $\mathcal{F}$, we first compute the i-admissible and i-naive claim-sets of $\mathcal{F}$, which yields the conflict-free claim-sets $\{x\}$, $\{y\}$, and $\{x, y\}$; and the i-admissible claim-sets $\{x, y\}$ and $\{x\}$. Taking the $\subseteq$-maximal claim-sets, we obtain in both cases the unique claim-set $\{x, y\}$ as the h-preferred and h-naive outcome of $\mathcal{F}$.

The CAF $\mathcal{F}'$ yields the same claim-sets under inherited and hybrid preferred semantics, namely the sets $\{x, y\}$ and $\{x, z\}$. For naive semantics, the variants differ: inherited semantics yield the sets $\{x\}$, $\{x, y\}$ and $\{x, z\}$ while hybrid semantics return $\{x, y\}$ and $\{x, z\}$. Observe that $\mathcal{F}'$ is well-formed.

We first observe that maximization on claim-level constitutes a strengthening of their inherited counterparts that perform maximization on argument-level. That is, each h-preferred (h-naive) claim-set is also i-preferred (i-naive).

Proposition 4.3. *For each CAF $\mathcal{F}$, $\sigma_h(\mathcal{F}) \subseteq \sigma_i(\mathcal{F})$ for $\sigma \in \{pr, na\}$.*

Proof. We show that each $\subseteq$-maximal i-admissible (i-conflict-free) claim-set possesses a $\subseteq$-maximal admissible (conflict-free, respectively) realization: Consider a set $S \in \sigma_h(\mathcal{F})$ and let E denote an admissible (conflict-free) realization of S in $\mathcal{F}$ that is $\subseteq$-maximal among all admissible (conflict-free) realizations of S, i.e., E cannot be extended with further arguments with claims in S without violating admissibility (conflict-freeness, respectively). We observe that E is a $\subseteq$-maximal admissible (conflict-free) set in F: otherwise, there is an admissible (conflict-free) set of arguments $D \subseteq A$ such that $E \subset D$. By choice of E, D contains an argument a with claim $cl(a) \notin S$. Thus, we have found an i-admissible (i-conflict-free) set of claims $cl(D)$ that properly extends S, which contradicts $\subseteq$-maximality of S in $ad_i(\mathcal{F})$ ($cf_i(\mathcal{F})$, respectively). □

Note that the other direction does not hold: We have already seen that i-preferred as well as i-naive claim-sets are not necessarily I-maximal (cf. Example 3.5); h-preferred and h-naive semantics, on the other hand, yield I-maximal sets per definition.

The above proposition moreover reveals an alternative view on h-preferred and h-naive semantics: they can be equivalently defined by maximizing over i-preferred or i-naive claim-sets, respectively.

Proposition 4.4. *For a CAF $\mathcal{F}$ and a set of claims $S \subseteq cl(A_{\mathcal{F}})$, it holds that*

- *$S \in pr_h(\mathcal{F})$ iff S is $\subseteq$-maximal in $pr_i(\mathcal{F})$;*
- *$S \in na_h(\mathcal{F})$ iff S is $\subseteq$-maximal in $na_i(\mathcal{F})$.*

Proof. In Proposition 4.3, we have already seen that each h-σ claim-set is contained in $\sigma_i(\mathcal{F})$. We moreover observe that each set that is $\subseteq$-maximal in $\sigma_i(\mathcal{F})$ is also $\subseteq$-maximal in $ad_i(\mathcal{F})$ ($cf_i(\mathcal{F})$, respectively) by monotonicity of the claim-function. Furthermore, each $\subseteq$-maximal i-preferred (i-naive) claim-set has an admissible (conflict-free) realization. □

For well-formed CAFs, both variants of preferred semantics coincide. As we show next, i-preferred semantics yield claim-sets that are incomparable for this class.

Lemma 4.5. *For each well-formed CAF $\mathcal{F}$ and $E, D \in pr(F)$, $E \neq D$, it holds that $cl(E) \not\subseteq cl(D)$.*

Proof. First assume there exists an $a \in E$ attacking some $b \in D$ in F. It follows that $cl(a) \notin cl(D)$ (and thus $cl(E) \not\subseteq cl(D)$), otherwise the argument $c \in D$ with $cl(c) = cl(a)$ also attacks b due to well-formedness; since D is conflict-free, this cannot be the case.

Suppose now that there is no $a \in E$ that attacks some $b \in D$. If D does not attack E, then $E \cup D \in pr(F)$ which contradicts the assumption that E and D are preferred extensions of F. So D attacks E. By admissibility of E we obtain that E attacks D, which leads to a contradiction, as shown above. □

As a consequence we obtain that each i-preferred claim-set has a unique preferred realization in the underlying AF.

Corollary 4.6. *$|pr(F)| = |pr_i(\mathcal{F})|$ for every well-formed CAF $\mathcal{F}$.*

Moreover, inherited and hybrid preferred semantics coincide in well-formed CAFs, implying that i-preferred semantics also satisfy I-maximality in this case.

Proposition 4.7. *$pr_h(\mathcal{F}) = pr_i(\mathcal{F})$ for each well-formed CAF $\mathcal{F}$.*

For naive semantics, we cannot hope for an analogous result as the well-formed CAF $\mathcal{F}'$ from Example 4.2 demonstrates: Here, the two variants yield different claim-sets as outcome. The example furthermore shows that i-naive semantics violates I-maximality (even for well-formed CAFs).

4.2. Introducing claim-attacks - stable semantics

Having discussed maximization on claim-level, our next step is to establish the crucial notion of *defeat of claims*. As sketched in Section 2, inherited CAF semantics lack a notion of claim-defeat that indicates the difference between *defeated* (i.e., *false*) and *undecided* claims. Recall that in the CAF $\mathcal{F}$ associated to the LP in Example 2.1, the partial attack from set $\{x_0\}$ on the claim e (only one occurrence of e has been attacked) has led to accepting only the set $\{a\}$ as semi-stable claim-set, although the claim-range of $\{a\}$ and $\{b, d\}$ coincide. Our goal is to establish a definition of claim-defeat that renders e in this situation as undecided. The basic assumption is that a claim is defeated *if all occurrences are attacked*. Our choice is justified as such a behavior can be observed by LPs and other formalisms that evaluate on conclusion-level.

Let us furthermore point out that defeating a claim is achieved by a set of arguments rather than by a set of claims. In Example 2.1, another argument would be necessary that helps x_0 to attack all occurrences of e.

Definition 4.8. Given a CAF $\mathcal{F} = (A, R, cl)$, we say that a set of arguments $E \subseteq A$ *defeats a claim* $c \in cl(A)$ iff for all $x \in A$ with $cl(x) = c$, there is $y \in E$ such that $(y, x) \in R$, i.e., E attacks each occurrence of c in $\mathcal{F}$. We write $E^{*}_{\mathcal{F}} = \{c \in cl(A) \mid E$ defeats c in $\mathcal{F}\}$ to denote the set of claims that are defeated by E in $\mathcal{F}$.

Claim-defeat is complementary to claim-acceptance: while it is sufficient that a single occurrence of a given claim is accepted to render it acceptable, it is necessary to attack all occurrences in order to defeat it.

We note that the function $\cdot^{*}$ is monotone, that is, if $E \subseteq D$ then $E^{*}_{\mathcal{F}} \subseteq D^{*}_{\mathcal{F}}$ for any $E, D \subseteq A$.

Example 4.9. Consider the CAF $\mathcal{F}$ given as follows:

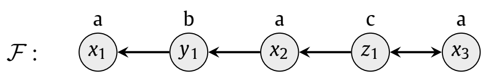

The set of arguments $\{y_1, z_1\}$ defeats the claim a (i.e., $\{y_1, z_1\}^{*}_{\mathcal{F}} = \{a\}$) because each occurrence of a is attacked: y_1 attacks x_1, and z_1 attacks x_2 and x_3. Moreover, the argument x_2 defeats claim b as it attacks the argument y_1 which is the unique argument carrying this claim.

Having established a notion for claim-defeat, we are ready to define the *claim-range* as a claim-based counterpart to the range in AFs. Again, the claim-range depends on a particular set of arguments. Intuitively, the claim-range of a set of arguments E contains all claims that are either *accepted* by E, i.e., all claims contained in E, or *rejected* by E, i.e., all claims that are defeated by E.

Definition 4.10. Given a CAF $\mathcal{F} = (A, R, cl)$ and a set of arguments $E \subseteq A$, by $E^{\circledast}_{\mathcal{F}} = cl(E) \cup E^{*}_{\mathcal{F}}$ we denote the *claim-range* of E in $\mathcal{F}$. If $E^{\circledast}_{\mathcal{F}} = cl(A)$ we say that E has *full claim-range* in $\mathcal{F}$.

Example 4.11. Let us consider again the CAF $\mathcal{F}$ from Example 4.9. The claim-range of $\{y_1, z_1\}$ is given by $\{a, b, c\}$ (i.e., $\{y_1, z_1\}^{\circledast}_{\mathcal{F}} = \{a, b, c\}$). Thus, the set has full claim-range, i.e., it holds that $cl(A) = \{y_1, z_1\}^{\circledast}_{\mathcal{F}}$. For $\{x_2\}$ we obtain $\{x_2\}^{\circledast}_{\mathcal{F}} = \{a, b\}$.

After having established the claim-range of a set of arguments, we could consider extending this definition to a set of *claims* by taking all realizations of the given set of claims into account. However, from Example 4.9, we learn that the claim-range of a set of claims is not necessarily unique: the maximal realization $\{x_1, x_2, x_3\}$ of claim a has full claim-range $\{a, b, c\}$, while the realization $\{x_1\}$ has claim-range $\{a\}$, and the realization $\{x_2\}$ has claim-range $\{a, b\}$. We see that the claim-range of a claim-set depends on the particular realization.

For well-formed CAFs, however, each claim-set admits a unique claim-range: recall that claims attack the same arguments in each well-formed CAF $\mathcal{F}$, i.e., $E^{+}_{F} = D^{+}_{F}$ for every two realizations E, D of a given claim-set S. It follows that each realization defeats the same claims.

Proposition 4.12. *Given a well-formed CAF $\mathcal{F}$ and a set of claims $S \subseteq cl(A_{\mathcal{F}})$, then $E^{*}_{\mathcal{F}} = D^{*}_{\mathcal{F}}$ and $E^{\circledast}_{\mathcal{F}} = D^{\circledast}_{\mathcal{F}}$ for every two realizations E, D of S in $\mathcal{F}$.*

Intuitively, we consider a set to be hybrid stable if it has full claim-range. As commonly observed for claim-based reasoning, the semantics depends on a particular realization. We thus consider a set of claims S to be h-stable in a given CAF $\mathcal{F}$ if it has a realization E that has full claim-range, i.e., $E^{\circledast}_{\mathcal{F}} = S \cup E^{*}_{\mathcal{F}} = cl(A)$. Following AFs, we furthermore require that the realization E is conflict-free in $\mathcal{F}$. While in AFs, a stable set of arguments is also admissible we observe that this is not generally the case for CAFs:

Example 4.13. Consider the CAF $\mathcal{F}$ from Example 4.9. Following our intuitive definition of hybrid stable semantics, we obtain that the set $\{a, b\}$ is h-stable in $\mathcal{F}$: indeed, the realization $E = \{y_1, x_3\}$ is conflict-free and defeats the claim c. Thus, the set has full claim-range: $E^{\circledast}_{\mathcal{F}} = \{a, b, c\}$. Observe that E is not admissible in $\mathcal{F}$ since the argument y_1 is not defended against the attack from x_2.

Inspired by this observation, we thus consider also an alternative variant of stable semantics that requires admissibility of the realization.

Definition 4.14. Given a CAF $\mathcal{F}$ and a set $S \subseteq cl(A_{\mathcal{F}})$. We say that

- S is a *h-cf-stable claim-set* ($S \in cf\text{-}stb_h(\mathcal{F})$) iff there exists a cf_i-realization E of S in $\mathcal{F}$ such that $E^{\circledast}_{\mathcal{F}} = cl(A_{\mathcal{F}})$;
- S is a *h-ad-stable claim-set* ($S \in ad\text{-}stb_h(\mathcal{F})$) iff there exists an ad_i-realization E of S in $\mathcal{F}$ such that $E^{\circledast}_{\mathcal{F}} = cl(A_{\mathcal{F}})$.

A set of arguments E *cf-stb_h-realizes* a claim-set S iff $cl(E) = S$, E is conflict-free in $\mathcal{F}$ and $E^{\circledast}_{\mathcal{F}} = cl(A_{\mathcal{F}})$; likewise, E *ad-stb_h-realizes* a claim-set S iff $cl(E) = S$, E is admissible in $\mathcal{F}$ and $E^{\circledast}_{\mathcal{F}} = cl(A_{\mathcal{F}})$.

The proposed variants relax inherited stable semantics. Indeed, a set of arguments E can have full claim-range without attacking all arguments that are not contained in E. For the hybrid variants it suffices that *some* argument with claim c is contained in E in order to accept c.

Example 4.15. Let us consider the following CAF $\mathcal{F}$:

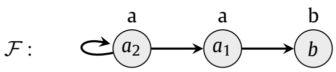

The framework has no stable extension since there is no argument that attacks the self-attacker a_2. Moreover, the only admissible set is $\emptyset$. Thus, there is no h-ad-stable claim-set either. We obtain however a h-cf-stable claim-set by considering the set $\{a_1\}$: the argument defeats claim b and carries claim a; thus $\{a_1\}^{\circledast}_{\mathcal{F}} = \{a, b\} = cl(A)$. We obtain that $cf\text{-}stb_h(\mathcal{F}) = \{\{a\}\}$. Observe that $\mathcal{F}$ is not well-formed.

Proposition 4.16. *$stb_i(\mathcal{F}) \subseteq ad\text{-}stb_h(\mathcal{F}) \subseteq cf\text{-}stb_h(\mathcal{F})$ for each CAF $\mathcal{F}$.*

Proof. To show that $stb_i(\mathcal{F}) \subseteq ad\text{-}stb_h(\mathcal{F})$, we observe that each stable extension E of the underlying AF F is admissible and attacks all remaining arguments. Thus, each claim is either accepted by E (i.e., E contains an occurrence of the claim in question) or defeated by E. We obtain $E^{\circledast}_{\mathcal{F}} = cl(A)$ for each stable extension of F. Moreover, we observe that each set of arguments E that realizes a h-ad-stable claim-set is also conflict-free. Consequently, we obtain that $ad\text{-}stb_h(\mathcal{F}) \subseteq cf\text{-}stb_h(\mathcal{F})$. □

The CAF $\mathcal{F}$ from Example 4.15 shows that $ad\text{-}stb_h(\mathcal{F}) \neq cf\text{-}stb_h(\mathcal{F})$ since $ad\text{-}stb_h(\mathcal{F}) = \emptyset$ but $cf\text{-}stb_h(\mathcal{F}) = \{\{a\}\}$. A small modification of the CAF $\mathcal{F}$ shows that $ad\text{-}stb_h(\mathcal{F}) \neq stb_i(\mathcal{F})$: If we delete the attack from a_2 to a_1 we obtain a single h-ad-stable claim-set $\{a\}$ (witnessed by the ad-realization $\{a_1\}$) but no i-stable claim-set exist. Observe that both considered CAFs are not well-formed. We will show next that for well-formed CAFs, all considered variants of stable semantics coincide.

Proposition 4.17. *$stb_i(\mathcal{F}) = ad\text{-}stb_h(\mathcal{F}) = cf\text{-}stb_h(\mathcal{F})$ for each well-formed CAF $\mathcal{F}$.*

Proof. Let $\mathcal{F} = (A, R, cl)$. We show that $cf\text{-}stb_h(\mathcal{F}) \subseteq stb_i(\mathcal{F})$: Consider a h-$cf$-stable claim-set S and a cf-stb_h-realization E of S in $\mathcal{F}$ that is $\subseteq$-maximal among all conflict-free realizations of S. We show that E is stable in the AF F. We show that E attacks all arguments that are not contained in E, i.e., $E^+_F = A \setminus E$. Let $x \in A \setminus E$ and let $cl(x) = c$. In case $c \notin S$, we have that all occurrences of c—including x—are attacked. Consider now the case $c \in S$, i.e., there is an argument $y \in E$ such that $cl(y) = c$. By maximality of E, we observe that $E \cup \{x\}$ is not conflict-free; thus, either (a) $(x, x) \in R$ or there is $z \in E$ such that either (b) $(z, x) \in R$ or (c) $(x, z) \in R$. In case (a) then also $(y, x) \in R$ by well-formedness; in case (b) we are done; in case (c) we have $(y, z) \in R$ by well-formedness and therefore E is not conflict-free, contradiction.

We obtain that $cf\text{-}stb_h(\mathcal{F}) \subseteq stb_i(\mathcal{F})$. By Proposition 4.16, $stb_i(\mathcal{F}) \subseteq ad\text{-}stb_h(\mathcal{F}) \subseteq cf\text{-}stb_h(\mathcal{F})$; thus, the statement follows. □

Finally, we show that both variants of hybrid stable semantics allow for alternative characterizations in terms of inherited complete and preferred semantics (for admissible-based h-stable semantics) and in terms of inherited naive semantics (for conflict-free-based stable semantics), respectively.

Proposition 4.18. *Given a CAF $\mathcal{F}$ and a set of claims $S \subseteq cl(A_{\mathcal{F}})$. Then the following statements are equivalent:*

(1) *$S \in ad\text{-}stb_h(\mathcal{F})$;*
(2) *there is a co_i-realization E of S in $\mathcal{F}$ with $E^{\circledast}_{\mathcal{F}} = cl(A_{\mathcal{F}})$;*
(3) *there is a pr_i-realization E of S in $\mathcal{F}$ with $E^{\circledast}_{\mathcal{F}} = cl(A_{\mathcal{F}})$.*

Moreover, the following two statements are equivalent:

(4) *$S \in cf\text{-}stb_h(\mathcal{F})$;*
(5) *there is an na_i-realization E of S in $\mathcal{F}$ with $E^{\circledast}_{\mathcal{F}} = cl(A_{\mathcal{F}})$.*

Proof. Let $\mathcal{F} = (A, R, cl)$. To prove (1) $\Leftrightarrow$ (2) $\Leftrightarrow$ (3), we first observe that (3) $\Rightarrow$ (2) $\Rightarrow$ (1) follows from the inclusions $pr(F) \subseteq co(F) \subseteq ad(F)$. To show (1) $\Rightarrow$ (3), consider a set $S \in ad\text{-}stb_h(\mathcal{F})$ and let E denote an ad_i-realization of S in $\mathcal{F}$ with $S \cup E^{*}_{\mathcal{F}} = cl(A)$. Then there is some $D \in pr(F)$ with $D \supseteq E$. We show that D is a pr_i-realization of S in $\mathcal{F}$, that is, $cl(D) = S$: Towards a contradiction, assume that there is some $c \in cl(A) \setminus S$ such that $c \in cl(D)$, that is, there is some $x \in D$ with $cl(x) = c$. By $S \cup E^{*}_{\mathcal{F}} = cl(A)$ we have $c \in E^{*}_{\mathcal{F}}$. Thus, there is some $y \in E \subseteq D$ that attacks x in F, which contradicts D being conflict-free. It follows that $cl(D) = S$. Moreover, D attacks each claim in $cl(A) \setminus S$ by monotonicity of $\cdot^{*}$; thus, the statement follows.

To prove (4) $\Leftrightarrow$ (5), it suffices to show (4) $\Rightarrow$ (5); the other direction is immediate since $cf(F) \subseteq na(F)$. Now, let $S \in ad\text{-}stb_h(\mathcal{F})$ and let E denote a cf_i-realization of S in $\mathcal{F}$ with $S \cup E^{*}_{\mathcal{F}} = cl(A)$. Similar as above, we consider a naive extension D in F with $E \subseteq D$ and show that $cl(D) = S$: In case there is some claim $c \in cl(A) \setminus S$ that is contained in $cl(D)$, there is some $y \in E \subseteq D$ that attacks an argument $x \in D$ with claim $cl(x) = c$, which contradicts D being conflict-free. We obtain that D is an na_i-realization of S in $\mathcal{F}$ that defeats all claims in $cl(A) \setminus S$. □

4.3. Bringing the two together - semi-stable and stage semantics

Semi-stable and stage semantics combine both methods that we have established in the preceding sections: they are designed to minimize undecidedness (starting from admissible or conflict-free sets, respectively). In terms of claims, semi-stable and stage semantics return $\subseteq$-maximal sets of claims that are either accepted or defeated with respect to a given extension.

Semi-stable and stage semantics weaken stable semantics by dropping the requirement that the claim-range has to contain all claims that are present in the framework.

Definition 4.19. Given a CAF $\mathcal{F}$ and a set of claims $S \subseteq cl(A_{\mathcal{F}})$. We say that

- S is a *h-stage claim-set* ($S \in stg_h(\mathcal{F})$) iff there exists a cf_i-realization E of S in $\mathcal{F}$ such that there is no $D \in cf(F)$ with $E^{\circledast}_{\mathcal{F}} \subset D^{\circledast}_{\mathcal{F}}$;
- S is a *h-semi-stable claim-set* ($S \in ss_h(\mathcal{F})$) iff there exists an ad_i-realization E of S in $\mathcal{F}$ such that there is no $D \in ad(F)$ with $E^{\circledast}_{\mathcal{F}} \subset D^{\circledast}_{\mathcal{F}}$.

A set of arguments E *stg_h-realizes* a claim-set S iff $cl(E) = S$, E is conflict-free in $\mathcal{F}$ and $E^{\circledast}_{\mathcal{F}})$ is $\subseteq$-maximal; likewise, E *ss_h-realizes* a claim-set S iff $cl(E) = S$, E is admissible in $\mathcal{F}$ and $E^{\circledast}_{\mathcal{F}})$ is $\subseteq$-maximal.

In contrast to the semantics we considered so far, we observe that the proposed variants of semi-stable and stage semantics neither constitute a strengthening nor a weakening of their inherited counterpart. The following example shows that even for well-formed CAFs, h-semi-stable and i-semi-stable (h-stage and i-stage, respectively) semantics potentially yield different claim-sets.

Example 4.20. Consider the following CAF $\mathcal{F}$:

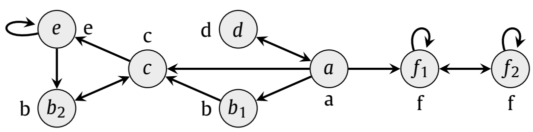

The i-admissible claim-sets of $\mathcal{F}$ are given by $S_1 = \{d\}$, $S_2 = \{b, d\}$ and $S_3 = \{a\}$. Let us now consider the claims they defeat: S_1 defeats claim a, S_2 defeats the claims c and a; and S_3 defeats claims c and d. Computing the claim-range of the sets yields the range $\{a, d\}$ for S_1; the range $\{a, b, c, d\}$ for S_2, and $\{a, c, d\}$ for S_3 (recall that for well-formed CAFs, each realization of a claim-set has the same range, it is thus possible to consider the claim-range of a set of claims). We obtain that $ss_h(\mathcal{F}) = \{\{b, d\}\}$. Observe that $\{a\}$ is the only i-semi-stable claim-set.

Regarding stage semantics, we furthermore consider the i-conflict-free claim-set $\{c\}$ that defeats claim e. We thus obtain two h-stage claim-sets: $\{c\}$ and $\{b, d\}$. We observe that $\{c\}$ together with $\{a\}$ are the i-stage extensions of $\mathcal{F}$.

We thus obtain that both semi-stable as well as stage semantics yield different extensions in both variants.

We consider alternative characterizations of both range-based semantics.

Proposition 4.21. *Given a CAF $\mathcal{F}$ and a set of claims $S \subseteq cl(A_{\mathcal{F}})$. The following statements are equivalent:*

(1) $S \in ss_h(\mathcal{F})$;
(2) *there is a co_i-realization E of S in $\mathcal{F}$ with $\subseteq$-maximal claim-range $E^{\circledast}_{\mathcal{F}}$ among complete extensions;*
(3) *there is a pr_i-realization E of S in $\mathcal{F}$ with $\subseteq$-maximal claim-range $E^{\circledast}_{\mathcal{F}}$ among preferred extensions.*

Moreover, the following two statements are equivalent:

(4) $S \in stg_h(\mathcal{F})$;
(5) *there is an na_i-realization E of S in $\mathcal{F}$ with $E^{\circledast}_{\mathcal{F}} = cl(A)$.*

Proof. The proof proceeds analogous to the proof of Proposition 4.18. To prove (1) ⇔ (2) ⇔ (3), we first observe that (3) ⇒ (2) ⇒ (1) follows from the inclusions $pr(F) \subseteq co(F) \subseteq ad(F)$. To show (1) ⇒ (3), consider a set $S \in ss_h(\mathcal{F})$ and let E denote a ss_h-realization of S in $\mathcal{F}$, that is, $E^{\circledast}_{\mathcal{F}}$ is $\subseteq$-maximal among i-admissible extensions. Then there is some $D \in pr(F)$ with $D \supseteq E$. As in the proof of Proposition 4.18, we obtain that D is a pr_i-realization of S in $\mathcal{F}$. Moreover, D defeats each claim that is defeated by E by monotonicity of $\cdot^*$, and thus, $E^{\circledast}_{\mathcal{F}} = D^{\circledast}_{\mathcal{F}}$ holds. Consequently, $D^{\circledast}_{\mathcal{F}}$ is $\subseteq$-maximal among preferred extensions: Assume otherwise, then there is a preferred extension T in F with $T^{\circledast}_{\mathcal{F}} \supset D^{\circledast}_{\mathcal{F}} = E^{\circledast}_{\mathcal{F}}$, which contradicts $\subseteq$-maximality of $E^{\circledast}_{\mathcal{F}}$ in F among i-admissible extensions. We have shown $D^{\circledast}_{\mathcal{F}}$ is $\subseteq$-maximal among preferred extensions; thus, the statement follows.

Likewise, we show (4) ⇒ (5) to prove the equivalence (4) ⇔ (5); the other direction is immediate since $cf(F) \subseteq na(F)$. Let $S \in stg_h(\mathcal{F})$ and let E denote a stg_h-realization of S in $\mathcal{F}$. As in the proof of Proposition 4.18, there exists a naive extension D in F with $E \subseteq D$ and $cl(D) = S$; similar as above, we obtain that $D^{\circledast}_{\mathcal{F}}$ is $\subseteq$-maximal among naive extensions. Thus, the statement follows. □

4.4. Hybrid semi-stable semantics: the missing counterpart for L-stable model semantics

In this section, we want to demonstrate that our novel hybrid semantics fill crucial gaps in the landscape of argumentation semantics. We come back to our motivating example and show that hybrid semi-stable semantics indeed capture L-stable model semantics for logic programs.

Logic programs in a nutshell We consider normal logic programs (LPs) [39] with default negation "not". Such programs consist of rules r of the form

$$r: \; c \leftarrow a_1, \ldots, a_n, \text{not}\, b_1, \ldots, \text{not}\, b_m$$

where $0 \leq n, m$ and the a_i, b_i and c are ordinary atoms. We let $head(r) = c$, $pos(r) = \{a_1, \ldots, a_n\}$ and $neg(r) = \{b_1, \ldots, b_m\}$; $\mathcal{L}(P)$ is the set of all atoms occurring in P.

We introduce 3-valued model semantics following [40] which generalizes stable model semantics [41] by allowing for undefined atoms.

Definition 4.22. A 3-valued Herbrand interpretation I of an LP P is a tuple $I = (T, F)$ with $T \cup F \subseteq \mathcal{L}(P)$ and $T \cap F = \emptyset$. We say $a \in \mathcal{L}(P)$ is true iff $a \in T$, false iff $a \in F$ and undefined otherwise.

Given a program P with Herbrand interpretation $I = (T, F)$ we define the *reduct* P/I of P w.r.t. I as follows: Starting from P,

(i) remove each rule r from P with $T \cap neg(r) \neq \emptyset$,
(ii) remove "not b" from each remaining rule whenever $b \in F$, and
(iii) for each $a \notin T \cup F$, replace each occurrence of "not a" by **u**.

Given two Herbrand interpretations $I = (T, F)$ and $I' = (T', F')$, we write $I \leq I'$ iff $T \subseteq T'$ and $F \supseteq F'$. A Herbrand interpretation $I = (T, F)$ is a 3-valued model of a program P iff I is a $\leq$-minimal model of P/I satisfying, for all atoms $a \in \mathcal{L}(P)$,

(a) $a \in T$ iff there is a rule $r \in P/I$ with $a = head(r)$ and $pos(r) \subseteq T$, and
(b) $a \in F$ iff for each rule $r \in P/I$ with $a = head(r)$ we have $pos(r) \cap F \neq \emptyset$.

As P/I is a positive program, such a model exists and is unique. These models I are also called *partially stable (p-stable)*. The undoubtedly most famous semantics for LPs is the so-called *stable model semantics*: a p-stable model is stable if it has no undecided atoms. It is a well-known result that stable model semantics correspond to stable AF semantics (when extracting the claims of the arguments) [5–7].

Similar in spirit to semi-stable semantics, *L-stable semantics* have been introduced to account for situations in which stable extensions do not exist. The idea is to minimize the undecided atoms. Below, we formally define L-stable model semantics:

Definition 4.23. A p-stable model $I = (T, F)$ of P is *L-stable* iff $T \cup F$ is $\subseteq$-maximal among all p-stable models of P.

Relating CAFs and LPs In [6,7], a correspondence between LPs and their associated AF has been established via appropriate mappings that assign each argument its own conclusion. Having incorporated this step in our formalism, we obtain the correspondence between CAFs and LP in a more direct fashion.

Below, we adapt the translation from LPs into AFs following the translation given in [7] to the realm of CAFs. First, we define the notion of an *argument* in an LP P.

Definition 4.24. For an LP P, A is an argument (in P) with

- $\textsc{Conc}(A) = c$,
- $\textsc{Rules}(A) = \bigcup_{i \le n} \textsc{Rules}(A_i) \cup \{r\}$, and
- $\textsc{Vul}(A) = \bigcup_{i \le n} \textsc{Vul}(A_i) \cup \{b_1, \dots, b_m\}$

iff there are arguments $A_1, \dots, A_n$ (in P) and a rule $r \in P$ with $r = c \leftarrow \textsc{Conc}(A_1), \dots, \textsc{Conc}(A_n)$, not $b_1, \dots,$ not b_m, and $r \notin \textsc{Rules}(A_i)$ for all $i \le n$. The rule r is called the *toprule* of A.

Given two arguments A and B, we say A attacks B if $\textsc{Conc}(A) \in \textsc{Vul}(B)$. Moreover, the claim of an argument A is its conclusion $\textsc{Conc}(A)$. The corresponding AF is denoted by $F_P = (A_P, R_P)$.

Definition 4.25. For an LP P, we define the corresponding CAF $\mathcal{F}_P = (A_P, R_P, cl_P)$ where A_P is the set of all arguments in P, $R_P = \{(A, B) \mid A, B \in A_P, \textsc{Conc}(A) \in \textsc{Vul}(B)\}$, and $cl_P(A) = \textsc{Conc}(A)$ for each $A \in A_P$.

In [7], it has been shown that p-stable and complete semantics correspond to each other when extracting the conclusions of the arguments. In terms of our adapted translation, we get the following more direct formulation:

Proposition 4.26. *Let P be a logic program. The model $I = (T, F)$ is P-stable in P iff $T \in co_i(\mathcal{F}_P)$.*

In a similar fashion, it can be shown that stable, well-founded, and regular LP semantics correspond to stable, grounded, and preferred (C)AF semantics; we recall these results in Appendix A.

As discussed in Section 2, L-stable semantics cannot be captured via established AF semantics that operate exclusively on argument-level. Having formally defined our hybrid version of semi-stable semantics, it can be shown that L-stable semantics indeed correspond to h-semi-stable semantics, as the following result demonstrates.

Theorem 4.27. *Let P be a logic program, $\mathcal{F}_P$ the associated CAF, and $I = (T, F)$ be a 3-valued interpretation. Then I is L-stable in P iff $T \in ss_h(\mathcal{F}_P)$.*

Hence, we have successfully established an argumentative counterpart of L-stable model semantics of LPs. The proof of this theorem can be found in Appendix A.

4.5. Summary

In the preceding subsections, we introduced novel variants of claim-based argumentation semantics by lifting certain evaluation-steps onto claim-level. Performing maximization on claim-level gave rise to alternative variants of preferred and naive semantics. We discussed claim-defeat which led to two novel hybrid variants of stable semantics; finally, bringing the two together gave rise to hybrid semi-stable and stage semantics.

Interestingly, it turned out that h-preferred and i-preferred as well as all stable variants collapse when we consider them on well-formed CAFs. This means that if arguments with the same claim have the same outgoing attacks, it holds that argument-level and claim-level maximization of i-admissible sets yield the same outcome. Also, if hybrid stable claim-sets in well-formed CAFs defeat all claims it follows that all arguments are attacked as well; i.e., for stable semantics in well-formed CAFs, claim-defeat and argument-attacks are interchangeable concepts. However, as we have seen, the notions do not coincide, even if the CAF is well-formed: range-based semantics potentially yield a different outcome as Example 4.20 demonstrates. This means as soon as we relax the condition and move to ⊆-maximality instead of universal quantification over the set of all arguments/claims not contained in the extension we observe fundamental differences between claim-defeat and argument-attack. Likewise, claim-set and argument-set maximization on arbitrary sets does not necessarily yield the same outcome in well-formed CAFs. As we have seen, i-naive and h-naive extensions potentially differ (cf. Example 4.2). It turns out that admissibility plays an important role for the concurrence of i- and h-preferred semantics.

In Section 4.4, we have shown that hybrid semi-stable semantics provide an argumentative counterpart for L-stable model semantics for LPs. With this, we demonstrate that hybrid semantics extend the applicability of argumentation semantics to semantics that could not be captured before. Hybrid semantics integrate central elements of a conclusion-oriented evaluation of argumentative scenarios that can be found in several different formalisms like assumption-based argumentation [11] and default logic [35]; hence, their advantages are by no means limited to LP semantics.

The case of claim-conflicts and claim-defense Let us end this section with a brief discussion about our focus on hybrid variants of maximization and defeat. We do not develop a notion of 'claim-conflict' because conflicts are inextricably linked to arguments. While it can be the case that claims are conflicting, for instance, if they are the logical negation of each other, or because they are inherently incompatible on an ethical, philosophical, or moral level, it is usually the task of the arguments constructed in favor of these claims to make their dissension explicit. We have seen this behavior in our logic programming instantiation (cf. Definition 4.25). Here, the claims of the constructed arguments correspond to the (derivable) atoms in the logic program. It is clear that the atoms of the program are not in conflict with each other per se. Indeed, we must first construct the arguments to reveal the conflicts between them.

Now, let us turn to the case of claim-defense and why we did not provide a hybrid variant of this fundamental notion (explaining the lack of hybrid variants of admissible, complete, and grounded semantics). Generally speaking, the reason is that claim-defense coincides with their traditional argument-based counter-part. Let us take a closer look on the notion. Intuitively, defense obeys the following logic: an entity (e.g., an argument, a claim) is *defended* iff each attacking unit is counter-attacked. Now, with our notion of claim-defeat at hand, this abstract view gives rise to the following notion of claim-defense:

a set of arguments E claim-defends *a claim c in a given CAF $\mathcal{F}$ iff E attacks each set of arguments D that claim-defeats c.*

That is, E must attack some argument $b \in D$ for each attacking set D of c. This means that there must be some argument x with claim c that is defended by E (in the underlying AF); otherwise, we can find a set of arguments that claim-defeats c but is not attacked by E. With these combinatorial considerations, claim-defense can be reformulated as follows: a set of arguments E claim-defends a claim c in $\mathcal{F}$ iff there exists an argument x with claim c that is defended by E in F. Thus, claim-defense coincides with classical defense on argument-level.

5. Relations between semantics

We first state a general observation which clarifies the relation between inherited and hybrid semantics in case every argument possesses a unique claim. If the claim furthermore coincides with the name of the argument name, we obtain that both variants coincide with the standard AF semantics.

Lemma 5.1. *For any $\sigma \in \{pr, na, stb, ss, stg\}$ and CAF $\mathcal{F} = (A, R, cl)$ with $cl(a) = a$ for all $a \in A$, we have $\sigma_h(\mathcal{F}) = \sigma_i(\mathcal{F}) = \sigma(F)$.*

It follows that negative results (via counter-examples) showing that two AF semantics are not in a subset-relation immediate apply to (well-formed) CAFs. We refer to [42, Figure 13] for an overview of the respective results for AF semantics.

Theorem 5.2. *The relations between the semantics depicted in Fig. 3 hold.*

As already discussed in Section 3 the relations between inherited semantics follow from the corresponding relations for AFs. Moreover, in Section 4 the relations between semantics that are based on the same Dung semantics have been settled: For arbitrary CAFs we have

$$stb_i(\mathcal{F}) \subseteq ad\text{-}stb_h(\mathcal{F}) \subseteq cf\text{-}stb_h(\mathcal{F})$$

by Proposition 4.16; moreover, by Proposition 4.3, it holds that

$$pr_h(\mathcal{F}) \subseteq pr_i \text{ and } na_h(\mathcal{F}) \subseteq na_i.$$

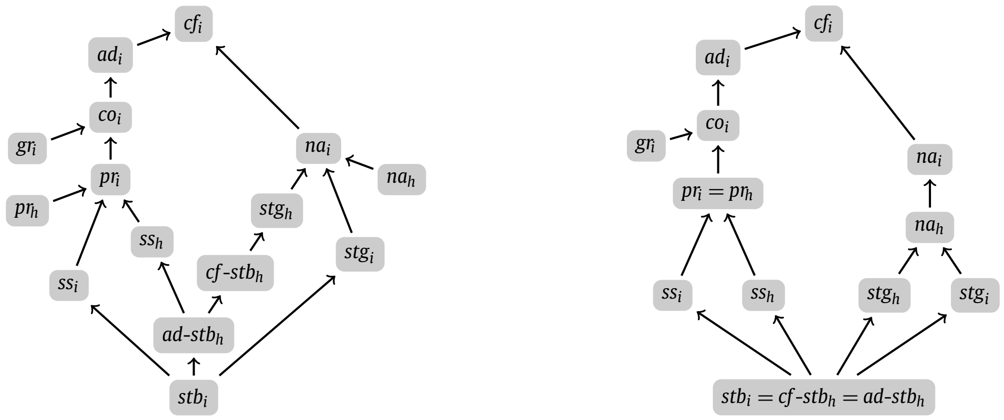

(a) Relations between CAF semantics.   (b) Relations for well-formed CAFs.

Fig. 3. Relations between semantics for general CAFs (a) and well-formed CAFs (b). An arrow from σ to τ indicates that $\sigma(\mathcal{F}) \subseteq \tau(\mathcal{F})$ for each (well-formed) CAF $\mathcal{F}$.

For well-formed CAFs, all stable variants coincide (by Proposition 4.17), also, i-preferred and h-preferred semantics yield the same outcome (by Proposition 4.7). Finally, semi-stable and stage semantics are incomparable, even in the well-formed case (cf. Example 4.20).

Next we discuss the remaining $\subseteq$-relations. First, we notice that each h-*ad*-stable claim-set is h-semi-stable, since each such set has full (and thus $\subseteq$-maximal) claim-range; likewise, each h-*cf*-stable set is h-stage.

Proposition 5.3. *$ad\text{-}stb_h(\mathcal{F}) \subseteq ss_h(\mathcal{F})$ and $cf\text{-}stb_h(\mathcal{F}) \subseteq stg_h(\mathcal{F})$ for any CAF $\mathcal{F}$.*

Furthermore, recall that h-semi-stable and h-stage semantics can be equivalently defined via preferred and naive AF semantics, respectively (cf. Proposition 4.21). We thus obtain that each h-semi-stable (h-stage) claim-set is h-preferred (h-naive, respectively).

Proposition 5.4. *$ss_h(\mathcal{F}) \subseteq pr_i(\mathcal{F})$ and $stg_h(\mathcal{F}) \subseteq na_i(\mathcal{F})$ for any CAF $\mathcal{F}$.*

This concludes the proofs for all $\subseteq$-relations for admissible-based semantics as shown in Fig. 3 for both well-formed and general CAFs.

Although h-naive semantics does not coincide with i-naive semantics in the well-formed case, we observe that h-naive semantics joins in the $\subseteq$-chain of conflict-free-based semantics: for well-formed CAFs, h-naive semantics are a superset of both inherited and hybrid stage semantics.

Lemma 5.5. *$stg_h(\mathcal{F}) \subseteq na_h(\mathcal{F})$ and $stg_i(\mathcal{F}) \subseteq na_h(\mathcal{F})$ for each well-formed CAF $\mathcal{F}$.*

Proof. First, consider a h-stage set $S \in stg_h(\mathcal{F})$. Towards a contradiction, assume $S \notin na_h(\mathcal{F})$. That is, there is some $T \in cf_i(\mathcal{F})$ with $T \supset S$. Now, since $\mathcal{F}$ is well-formed, each realization of S and T attack the same claim. By monotonicity of the range-function, we obtain that $D^{\circledast}_{\mathcal{F}} \supset E^{\circledast}_{\mathcal{F}}$ for each realization D of T and E of S; which contradicts $S \in stg_h(\mathcal{F})$.

Now, consider an i-stage claim-set $S \in stg_i(\mathcal{F})$, i.e., there is a set $E \subseteq A$ with $cl(E) = S$ such that $E \cup E^+_F$ is maximal wrt. subset-relation. Now, assume that $S \notin na_h(\mathcal{F})$, i.e. there exists a set $T \in cf_i(\mathcal{F})$ such that $T \supset S$. Consider a cf_i-realization D of T in $\mathcal{F}$. Now, since E is stage in F, there is some $x \in E \cup E^+_F$ such that $x \notin D \cup D^+_F$. By well-formedness, $D^+_F \supseteq E^+_F$; thus, we have $x \in E$ and $x \notin D$. We can assume that x and D are conflicting; otherwise consider $D' = D \cup \{x\}$ instead. Since x and D are conflicting and since $x \notin D^+_F$, there exists $y \in D$ such that $(x, y) \in R$. Since $T \subset S$, there is $z \in D$ such that $cl(x) = cl(z)$. By well-formedness, $(z, y) \in R$, which contradicts D being conflict-free. $\square$

We discuss counter-examples for the remaining cases: First, we use Lemma 5.1 to transfer known results for relations for AF semantics to CAF semantics.

Proposition 5.6. *Let* `Sem` *be the set of all semantics under consideration. There is a well-formed CAF $\mathcal{F}$ such that $\alpha(\mathcal{F}) \not\subseteq \beta(\mathcal{F})$ for*

1. $\alpha = cf_i$, $\beta \in \texttt{Sem} \setminus \{cf_i\}$;
2. $\alpha = ad_i$, $\beta \in \texttt{Sem} \setminus \{cf_i, ad_i\}$;
3. $\alpha = co_i$, $\beta \in \texttt{Sem} \setminus \{cf_i, ad_i, co_i\}$;

4. $\alpha = gr_i$, $\beta \in \mathtt{Sem} \setminus \{cf_i, ad_i, co_i, gr_i\}$;
5. $\alpha \in \{pr_h, pr_i\}$, $\beta \in \mathtt{Sem} \setminus \{cf_i, ad_i, co_i, pr_h, pr_i\}$;
6. $\alpha \in \{na_h, na_i\}$, $\beta \in \mathtt{Sem} \setminus \{cf_i, na_h, na_i\}$;
7. $\alpha \in \{ss_h, ss_i\}$, $\beta \in \{stg_h, stg_i, na_h, na_i, cf\text{-}stb_h, ad\text{-}stb_h, stb_i\}$ and
8. $\alpha \in \{stg_h, stg_i\}$, $\beta \in \{ad_i, ss_h, ss_i, pr_h, pr_i, cf\text{-}stb_h, ad\text{-}stb_h, stb_i\}$.

It remains to provide a counter-example for the absence of $\subseteq$-relations between ss_i, ss_h and pr_h (stg_i, stg_h and na_h respectively) for general CAFs.

Example 5.7. Consider the following (non-well-formed) CAF $\mathcal{F}$:

$\mathcal{F}$:

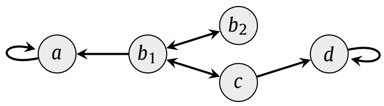

Let us first note that in $\mathcal{F}$, the set of conflict-free and admissible sets coincides; thus, all admissible-based and conflict-free based semantics coincide, in particular: $pr_h(\mathcal{F}) = na_h(\mathcal{F})$, $ss_h(\mathcal{F}) = stg_h(\mathcal{F})$, and $ss_i(\mathcal{F}) = stg_i(\mathcal{F})$.

The sets $E_1 = \{b_1\}$ and $E_2 = \{b_2, c\}$ are $\subseteq$-maximal conflict-free sets in F and have $\subseteq$-maximal (claim-)range: E_1 attacks the arguments a, b_2, and c; thus, it has argument-range $\{a, b_1, b_2, c\}$ and claim-range $\{a, b, c\}$; the set E_2 attacks arguments b_1 and d, yielding argument-range $\{b_1, b_2, c, d\}$ and claim-range $\{b, c, d\}$.

We obtain that $\{b\}$ and $\{b, c\}$ are inherited and hybrid semi-stable and stage in $\mathcal{F}$. On the other hand, the set $\{b, c\}$ is the unique h-naive and h-preferred claim-extension of $\mathcal{F}$.

The crucial observation in the above example is that h-naive and h-preferred semantics are I-maximal while the others are not; i.e., it might be the case that semi-stable and stage variants yield claim-sets S, T that are in $\subseteq$-relation to each other ($S \subset T$). Among other principles, we will discuss this property in depth in Section 6.

Finally, let us discuss the connection between h-stable, h-semi-stable, and h-stage semantics. Recall that for inherited semantics, $stb_i(\mathcal{F}) = ss_i(\mathcal{F}) = stg_i(\mathcal{F})$ in case $stb_i(\mathcal{F}) \neq \emptyset$. We observe that this does not extend to h-stable semantics.

Example 5.8. Let us consider the following CAF $\mathcal{F}$:

$\mathcal{F}$:

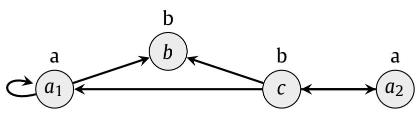

In $\mathcal{F}$, we have $ad\text{-}stb_h(\mathcal{F}) = ss_h(\mathcal{F}) = \{\{c\}\}$ and $cf\text{-}stb_h(\mathcal{F}) = stg_h(\mathcal{F}) = \{\{c\}, \{a, d\}\}$.

However, we can obtain the following weaker version.

Lemma 5.9. *For any CAF $\mathcal{F}$, (a) $cf\text{-}stb_h(\mathcal{F}) \neq \emptyset$ implies $cf\text{-}stb_h(\mathcal{F}) = stg_h(\mathcal{F})$ and (b) $ad\text{-}stb_h(\mathcal{F}) \neq \emptyset$ implies $ad\text{-}stb_h(\mathcal{F}) = ss_h(\mathcal{F})$.*

Proof. In case $cf\text{-}stb_h(\mathcal{F})$ is non-empty, we have that each $S \in stb_h(\mathcal{F})$ has $\subseteq$-maximal range (full range in fact), i.e., there is a cf_i-realization E of S in $\mathcal{F}$ such that $E^{\circledast}_{\mathcal{F}} = cl(A)$. We obtain $stb_h(\mathcal{F}) = stg_h(\mathcal{F})$. Similar arguments hold for the respective admissible-based semantics. □

6. Principles

In this section, we conduct a *principle-based analysis* of CAF semantics. The goal of our studies is to identify differences between inherited and hybrid semantics on the one hand and to analyze the different behavior of the semantics when restricted to well-formed CAFs when compared to the general case on the other hand. We have already experienced in Section 4 that differences between inherited and hybrid semantics vanish when restricting them to well-formed CAFs (cf. Proposition 4.7 and 4.17). Our principle-based analysis aims to work out such specific differences in greater detail. We consider principles restricted to the class $\mathfrak{C}_u$ of all CAFs as well as to the class $\mathfrak{C}_{wf}$ of all well-formed CAFs.

In this section, we identify not only principles that are genuine for CAF semantics, but consider also principles that extend well-known principles for AF semantics like conflict-freeness or reinstatement to claim-based reasoning. In this aspect, let us recall that AFs can be seen as a special case of CAFs by taking the identity function as claim-function. By Lemma 5.1, we obtain that negative results (via counter-examples) carry over to CAFs for those principles that are a faithful generalization of AF principles. To compare our principles with the corresponding AF case, it will be useful to consider the CAF-class $\mathfrak{C}_{id} = \{(F, id) \mid F \text{ is an AF}\}$ that contains each AF as equivalent CAF representation.

We subdivide our principles in three different groups: in Section 6.1, we consider principles that address properties of the underlying structure of the framework with respect to specific semantics; in Section 6.2, we consider basic properties like conflict-freeness and admissibility inspired by similar principles for AF semantics; and in Section 6.3, we study set-theoretical principles that give insight into the expressiveness of the considered semantics.

6.1. Meta-principles

In this section, we consider principles that formalize fundamental properties of the relation between claim-sets and their realization. We furthermore discuss principles that deal with independence of the formal language.

Realizability of claim-sets The first three principles deal with realizations of claim-sets. The so-called *realizability principle* is at the core of argumentative claim justification: a claim cannot be accepted if there is no argument for it.

We furthermore consider the *unique realization principle* and the *maximal realization principle* that give insights about the nature of the realizations. For a semantics that satisfies unique realizability, each claim-set has exactly one witnessing realization. Maximal realizability is closely related: the principle is satisfied if the union of all witnessing realizations of a claim-set is a witnessing realization itself. Satisfaction of these principle is advantageous for computational aspects. If a claim-set is uniquely realized it holds that argument-extensions and claim-extensions are in one-to-one correspondence to each other. This indicates that the computational complexity of standard reasoning tasks is comparable to the complexity of the corresponding AF semantics (cf. [19,20]). Maximal realizability, on the other hand, can guide the search for a suitable realization. Intuitively, the property can be exploited by over-approximating the feasible solution and iteratively narrow it down until a suitable candidate has been found. We refer to [19,20] for an in-depth discussion on this matter.

Principles 6.1. *Let ρ be a CAF semantics and let $\mathfrak{C}$ be a class of CAFs.*

(Realizability) *ρ satisfies the* realizability principle *in $\mathfrak{C}$ iff for every CAF $\mathcal{F} \in \mathfrak{C}$, for every claim-set S, $S \in \rho(\mathcal{F})$ only if there is a set of arguments $E \subseteq A_{\mathcal{F}}$ that realizes S in $\mathcal{F}$.*

(Unique Realizability) *ρ satisfies the* unique realizability principle *in $\mathfrak{C}$ iff for every CAF $\mathcal{F} \in \mathfrak{C}$, for every $S \in \rho(\mathcal{F})$ there is a unique set of arguments $E \subseteq A_{\mathcal{F}}$ that ρ-realizes S in $\mathcal{F}$.*

(Maximal Realizability) *ρ satisfies the* maximal realizability principle *in $\mathfrak{C}$ iff for every CAF $\mathcal{F} \in \mathfrak{C}$, for every $S \in \rho(\mathcal{F})$, the set $E^{\max} = \bigcup_{E\ \rho\text{-real. } S} E$ is a ρ-realization of S in $\mathcal{F}$.*

By definition, each semantics under consideration satisfies the realizability principle.

Proposition 6.2. *I-grounded, i-complete, i-preferred, i-semi-stable, i-naive, i-stage, and i-stable semantics satisfy unique realizability for well-formed CAFs.*

Proof. Since each realization of an i-complete claim-set in a well-formed CAF attacks—and thus defends—the same arguments, we obtain that each i-complete claim-set admits a unique co_i-realization. This property extends to all complete-based inherited semantics, i.e., to i-grounded, i-preferred, i-semi-stable, and i-stable semantics.

Let us next discuss the case for naive semantics: Consider a well-formed CAF $\mathcal{F}$ and let S denote an i-naive claim-set of $\mathcal{F}$. Now, assume that S has two na-realizations $E \neq D$ in $\mathcal{F}$. Since both E and D are naive in F, there must be a conflict between them. w.l.o.g., assume that there is some argument $x \in E$ that attacks some $y \in D$. On the other hand, it holds that $E^+ = D^+$ by well-formedness. Thus, there is some $z \in D$ that attacks y, which contradicts D being conflict-free.

Since each stage extension is naive, the statement follows for i-stage semantics. □

If a semantics satisfies unique realizability it also satisfies maximal realizability. We show that all hybrid semantics as well as i-conflict-free and i-admissible semantics satisfy this principle in $\mathfrak{C}_{wf}$.

Proposition 6.3. *All semantics under consideration satisfy maximal realizability for well-formed CAFs.*

Proof. Starting with inherited conflict-free and admissible semantics, we first observe that two cf_i-realizations E, D of a claim-set S are conflict-free since they attack the same arguments. Hence, $E \cup D$ cf_i-realizes S as well. Moreover, if E and D are ad_i-realizations of S, it holds that both defend the same arguments. Thus, $E \cup D$ ad_i-realizes S. We thus obtain that i-conflict-free and i-admissible semantics satisfy maximal realizability. The inherited semantics in question satisfy the principle since they build on either i-conflict-free or i-admissible semantics (and since they already satisfy unique realizability).

For h-preferred and all variants of stable semantics, the statement follows since they coincide with their respective inherited counter-parts. For the remaining semantics, it suffices to consider the i-preferred (for h-semi-stable semantics) respectively the i-naive (for h-naive and h-stage semantics) realization of the claim-set in question: Consider a well-formed

CAF $\mathcal{F}$ and let S denote a h-semi-stable claim-set of $\mathcal{F}$. By our results from Section 4, S has a pr_i-realization E in $\mathcal{F}$. This realization contains all ss_h-realizations of S in $\mathcal{F}$, i.e., $E = E^{\max}$.

The proof for h-naive and h-stage semantics is analogous. □

Interestingly, hybrid semantics are not uniquely realized as they do not require ⊆-maximality of their admissible (or conflict-free) realizations. Consider the following trivial example with only two arguments both having the same claim c.

Example 6.4. Consider the well-formed CAF $\mathcal{F} = (\{x, y\}, \emptyset, cl)$ with $cl(x) = cl(y) = c$. In $\mathcal{F}$, all hybrid semantics return the same claim-set $\{c\}$. However, the extension $\{c\}$ has three possible realizations: $\{x\}$, $\{y\}$, and $\{x, y\}$, all witnessing the acceptance of c.

We note that the alternative definitions of h-semantics that consider complete, preferred, or naive semantics (cf. Propositions 4.4, 4.18, and 4.21) as their base sets indeed satisfy unique realizability since the inherited variants transfer this property to the respective semantics.

Apart from i-grounded semantics, all remaining semantics considered in this paper violate unique and maximal realizability in the general case. It suffices to extend Example 6.4 in a minimal way:

Example 6.5. Consider the CAF $\mathcal{F} = (\{x, y\}, \{(x, y), (y, x)\}, cl)$ with $cl(x) = cl(y) = c$. In $\mathcal{F}$, all semantics return the claim-set $\{c\}$. However, the extension $\{c\}$ has two possible realizations $\{x\}$ and $\{y\}$ which shows that $\{c\}$ is neither uniquely realizable nor possesses a maximal realization.

Abstraction principles Next, we consider two closely related principles that both deal with aspects of abstraction. The *argument-names independence principle* states that the specific names of the arguments should not play a role when evaluating a given framework with respect to the claims. The *language independence principle* [21,22], also referred to as *abstraction principle* [43,44], formalizes that a semantics is independent of the specific names of the elements that occur in a framework. In contrast to argument-name independence, which states that two isomorphic frameworks yield *identical* claim-extensions independently of the considered argument-names, the language independence principle states that the evaluation process does not depend on the names of the abstract objects (i.e., arguments and claims) in the frameworks. Both principles address fundamental elements of argumentative reasoning.

In order to formalize these principles, we introduce two notions of CAF isomorphisms.

Definition 6.6. A bijective function $f : A_F \to A_G$ between two CAFs $\mathcal{F}$ and $\mathcal{G}$ is an *isomorphism* iff the following two conditions are satisfied:

- f is attack-preserving, i.e., for all $x, y \in A_F$, $(x, y) \in R_F$ iff $(f(x), f(y)) \in R_G$; and
- f is claim-preserving, i.e., for all $x, y \in A_F$, $cl(x) = cl(f(y))$.

$\mathcal{F}$ and $\mathcal{G}$ are isomorphic to each other iff there is an isomorphism $f : A_F \to A_G$.

A bijective function $f : A_F \to A_G$ between two CAFs $\mathcal{F}$ and $\mathcal{G}$ is a *generalized isomorphism* iff the following two conditions are satisfied:

- f is attack-preserving, i.e., for all $x, y \in A_F$, $(x, y) \in R_F$ iff $(f(x), f(y)) \in R_G$; and
- f is claim-preserving, i.e., for all $x, y \in A_F$, preserves the claim-structure, i.e., $cl(x) = cl(y)$ iff $cl(f(x)) = cl(f(y))$.

We say that $\mathcal{F}$ and $\mathcal{G}$ are generalized isomorphic to each other iff there is a generalized isomorphism $f : A_F \to A_G$. We call the function $f_c : cl(A_F) \to cl(A_G)$ with $f_c(cl(x)) = cl(f(x))$ *f-induced claim-isomorphism.*

Example 6.7. Let us consider our CAF $\mathcal{F}$ from Example 3.5 and another CAF $\mathcal{G}$ also having three arguments. Both $\mathcal{F}$ and $\mathcal{G}$ are depicted below:

The CAFs $\mathcal{F}$ and $\mathcal{G}$ are not isomorphic to each other as the claims which appear in the CAFs do not coincide. They are, however, generalized isomorphic to each other: indeed, the function f with $x_1 \mapsto a$, $x_2 \mapsto b$, and $y_1 \mapsto c$ satisfies $(x, y) \in R_{\mathcal{F}}$ iff $(f(x), f(y)) \in R_{\mathcal{G}}$ and preserves the claim-structure by associating claim x in $\mathcal{F}$ with claim α in $\mathcal{G}$ and claim y with claim β. The induced claim-isomorphism f_c behaves accordingly and maps x to α and y to β.

Principles 6.8. *Let ρ be a CAF semantics and let $\mathfrak{C}$ be a class of CAFs.*

Table 1
Meta-principles w.r.t. general CAFs.

	Realizability	Arg-name Ind.	Language Ind.	Unique Realizability	Maximal Realizability
cf_i	✓	✓	✓	✗	✗
ad_i	✓	✓	✓	✗	✗
gr_i	✓	✓	✓	✓	✓
co_i	✓	✓	✓	✗	✗
pr_i	✓	✓	✓	✗	✗
pr_h	✓	✓	✓	✗	✗
stb_i	✓	✓	✓	✗	✗
$cf\text{-}stb_h$	✓	✓	✓	✗	✗
$ad\text{-}stb_h$	✓	✓	✓	✗	✗
ss_i	✓	✓	✓	✗	✗
ss_h	✓	✓	✓	✗	✗
na_i	✓	✓	✓	✗	✗
na_h	✓	✓	✓	✗	✗
stg_i	✓	✓	✓	✗	✗
stg_h	✓	✓	✓	✗	✗

(Argument-names Independence) *ρ satisfies the* argument-names independence principle *in $\mathfrak{C}$ iff for every two CAFs $\mathcal{F}$ and $\mathcal{G}$ in $\mathfrak{C}$ which are isomorphic to each other, it holds that $\rho(\mathcal{F}) = \rho(\mathcal{G})$.*

(Language Independence) *ρ satisfies the* language independence principle *in $\mathfrak{C}$ iff for every two CAFs $\mathcal{F}$ and $\mathcal{G}$ in $\mathfrak{C}$ which are generalized isomorphic to each other (via isomorphism f), it holds that $\rho(\mathcal{F}) = \{f_c(S) \mid S \in \rho(\mathcal{G})\}$ for the f-induced claim-isomorphism $f_c : cl(A_F) \to cl(A_G)$.*

It is easy to see that all considered semantics satisfy both principles. We note that the argument-names independence principle plays a crucial role in the context of dynamic scenarios. In a nutshell, it makes a difference whether changes of a framework depend on argument names or on the graph topology. This issue has been discussed in [45] in the context of framework expansions.

Remark 6.9 *(Relation to AFs).* The language independence principle is a faithful adaption of the corresponding AF principle: restricting the principle to $\mathfrak{C}_{id}$ yields precisely the desired principle since each generalized isomorphism between $\mathcal{F}, \mathcal{G} \in \mathfrak{C}_{id}$ corresponds to an AF-isomorphism between F and G.

The adaption of argument-name independence to AFs by restricting it to the class $\mathfrak{C}_{id}$ yields a principle that allows to compare only identical AFs (due to the definition of CAF-isomorphism) and is thus trivially satisfied by all possible semantics. The alternative adaption of the principle by considering native AF-isomorphisms (graph-theoretically speaking, an arc-preserving bijection), on the other hand, results in a principle that is not satisfied by any non-trivial argumentation semantics considered in the literature since the names of the arguments trivially matter when evaluating AFs: Indeed, a simple counter-example are given by AFs $F = (\{a\}, \emptyset)$ and $G = (\{b\}, \emptyset)$ which are AF-isomorphic to each other but yield different extensions ($\{a\}$ and $\{b\}$, respectively).

Remark 6.10 *(Floating conclusions).* We note that in some situations, it might be desirable to consider semantics that enforce a stronger connection between argument names and claims. This is reflected in the discussion about *floating conclusions* in formal argumentation and related areas [46–48]. Floating conclusions are claims that appear as conclusions in all extensions, but no common argument supports them. The question was whether these conclusions can be considered skeptically acceptable. In the current work, we adapt the more generous approach in the tradition of Makinson and Schlechta [46] and consider them acceptable. This choice is reflected in the satisfaction of the argument-names independence principle that states that the semantics are not affected when changing the names of the arguments. The more skeptical approach to floating conclusions could be realized, for instance, by developing semantics that return claims in combination with arguments. However, this would cause a violation of the argument-names independence principle.

Summary Table 1 and 2 summarize our results from this section. Table 1 presents all considered principles for general CAFs while Table 2 contains all principles for well-formed CAFs. The realizability principle as well as the argument-name and language independence principle are satisfied by all considered semantics, which confirms that these principles formalize fundamental properties of claim-based reasoning.

Apart from formalizing general principles of claim-based reasoning, our results of this section revealed central differences between general and well-formed CAFs. We observe that the desirable unique and maximal realizability principles are not satisfied by any (except i-grounded) semantics in the general case. For well-formed CAFs, the picture is more diverse, in particular, due to the difference between inherited and hybrid semantics for the unique realizability principle. The closely related maximal realizability principle is satisfied by all semantics considered in this work for well-formed CAFs. At first

Table 2
Meta-principles w.r.t. well-formed CAFs.

	Realizability	Arg-name Ind.	Language Ind.	Unique Realizability	Maximal Realizability
cf_i	✓	✓	✓	✗	✓
ad_i	✓	✓	✓	✗	✓
gr_i	✓	✓	✓	✓	✓
co_i	✓	✓	✓	✓	✓
pr_i	✓	✓	✓	✓	✓
pr_h	✓	✓	✓	✗	✓
stb_i	✓	✓	✓	✓	✓
$cf\text{-}stb_h$	✓	✓	✓	✗	✓
$ad\text{-}stb_h$	✓	✓	✓	✗	✓
ss_i	✓	✓	✓	✓	✓
ss_h	✓	✓	✓	✗	✓
na_i	✓	✓	✓	✓	✓
na_h	✓	✓	✓	✗	✓
stg_i	✓	✓	✓	✓	✓
stg_h	✓	✓	✓	✗	✓

sight, these results deem inherited semantics preferable over hybrid semantics since no hybrid semantics is uniquely realizable. We note that from a computational perspective, this difference is negligible in the well-formed case since maximal realizability is often sufficient to compute a realization efficiently.

6.2. Basic principles

In this section, we deepen the study of claim-based semantics by investigating fundamental properties of argumentation semantics in the context of claim-based reasoning. The objective is to identify the characteristics of the claim-based semantics under consideration. We analyze principles of the underlying realization of claim-sets and identify principles for maximality-based semantics, semantics that are closed under defense, and range-based semantics.

To begin with, we study claim-based semantics on argument level by analyzing the corresponding realizations.

Realization-based analysis As discussed in Section 4.5, central notions like conflict, defense, or admissibility (being the combination of the former two) are inherent properties of arguments. The fundamental importance of these properties is well-known; they are at the core of argumentative reasoning. Nevertheless, they play a fundamental role in the context of claim-based argumentation because they are crucial properties of the underlying realizations of claim-sets. In order to grasp these concepts in the claim-based setting, we adapt well-known principles for argumentation semantics: the general idea to lift a given AF inspired principle is to require existence of a realization that satisfies the principle on argument level.

Principles 6.11. *Let ρ be a CAF semantics and let $\mathfrak{C}$ be a class of CAFs.*

(Conflict-freeness) *ρ satisfies* conflict-freeness *in $\mathfrak{C}$ iff for every CAF $\mathcal{F} \in \mathfrak{C}$, for every $S \in \rho(\mathcal{F})$, there is a conflict-free realization E of S in $\mathcal{F}$.*
(Defense) *ρ satisfies the* defense principle *in $\mathfrak{C}$ iff for every CAF $\mathcal{F} \in \mathfrak{C}$, for every $S \in \rho(\mathcal{F})$, there is a realization E of S in $\mathcal{F}$ that defends itself.*
(Admissibility) *ρ satisfies the* admissibility principle *in $\mathfrak{C}$ iff for every CAF $\mathcal{F} \in \mathfrak{C}$, for every $S \in \rho(\mathcal{F})$, there is an admissible realization E of S in $\mathcal{F}$.*

We note that this idea serves as blue-print to lift several other AF principles. In this work, we focus on the present principles as we consider them of particular importance for a principled study of claim-based semantics.

By definition, each semantics considered in this paper satisfies conflict-freeness. Defense and admissibility are satisfied by all admissible-based semantics.

Proposition 6.12. *I-admissible, i-complete, i-grounded, h-ad-stable, i-stable and both variants of semi-stable and preferred semantics satisfy defense and admissibility.*

We refer to Example 4.15 which shows that h-*cf*-stable semantics does not satisfy admissibility. By Lemma 5.1, we obtain counter-examples for the remaining semantics from the corresponding AF case.

Maximality-based principles We consider principles that are based on maximality requirements, namely the *naivety principle* [22] as well as the famous *I-maximality principle* [21].

The naivety principle states that each extension is a $\subseteq$-maximal conflict-free set. In the context of claims, this principle can be extended in two ways: First, by requiring the existence of a realization that is maximal with respect to set-inclusion, and second, by requiring that the claim-set itself is ⊆-maximal. Notice that this two natural choices reflect the different approaches that underlie inherited and hybrid semantics, respectively.

The I-maximality principle requires that extensions are pairwise incomparable, i.e., no extension is a subset of another one. This fundamental principle underlies several AF semantics, in particular, preferred and naive semantics. In the context of claims, it plays a central role for the hybrid variants of these semantics.

Principles 6.13. *Let ρ be a CAF semantics and let $\mathfrak{C}$ be a class of CAFs.*

(i-Naivety) *ρ satisfies the* inherited naivety principle *in $\mathfrak{C}$ iff for every CAF $\mathcal{F} \in \mathfrak{C}$, for every $S \in \rho(\mathcal{F})$, there is a conflict-free realization E of S in $\mathcal{F}$ which is ⊆-maximal in $cf(F)$.*

(h-Naivety) *ρ satisfies the* hybrid naivety principle *in $\mathfrak{C}$ iff for every CAF $\mathcal{F} \in \mathfrak{C}$, for every $S \in \rho(\mathcal{F})$, it holds that S is ⊆-maximal in $cf_{\mathrm{i}}(\mathcal{F})$.*

(I-Maximality) *ρ satisfies* I-maximality *in $\mathfrak{C}$ iff for every CAF $\mathcal{F} \in \mathfrak{C}$, for every $S, T \in \rho(\mathcal{F})$, if $S \subseteq T$ then $S = T$.*

Semantics that satisfy these principles are well-suited for applications in which as many claims as possible should be deemed acceptable. The naivety principles target situations in which agreement and disagreement is decided along conflicts between arguments and is satisfied by conflict-free-based semantics. In contrast, I-maximality can be tied to different acceptance conditions.

We first discuss satisfaction of the naivety principles. Restricting these principles to $\mathfrak{C}_{id}$ results in the naivety principle for AF semantics. By Lemma 5.1, we thus obtain counter-examples for i-admissible, i-complete, i-grounded, and all variants of preferred and semi-stable semantics.

By definition, argument-dependent naivety is satisfied by inherited naive semantics. It follows that all semantics ρ with $\rho(\mathcal{F}) \subseteq na_i(\mathcal{F})$ for all CAFs $\mathcal{F}$ satisfy this principle too. Apart from our versions of stable semantics which satisfy both principles, i-naivety can be seen as complementary to the admissibility principle.

Proposition 6.14. *All variants of naive, stage, and stable semantics satisfy inherited naivety.*

Claim-dependent naivety, on the other hand, is not satisfied by any of the considered semantics in the general case, except for h-naive semantics. As we will see, h-naive semantics is one of the few semantics that satisfy I-maximality in $\mathfrak{C}_u$. For h-naive semantics, the principle is satisfied by definition. By results from Section 5, we obtain the following result.

Proposition 6.15. *h-naive semantics satisfy hybrid naivety. Moreover, all variants of stage and stable semantics satisfy hybrid naivety in $\mathfrak{C}_{wf}$.*

By Example 4.2, i-naive semantics can realize claim-sets that are in subset-relation, even if the CAF is well-formed, showing that i-naive semantics violate hybrid naivety even in the well-formed case.

For I-maximality, is easy to see that h-preferred and h-naive semantics satisfy this principle. Moreover, i-grounded semantics yields a unique extension and thus satisfies this principle as well.

Proposition 6.16. *I-grounded, h-naive, and h-preferred semantics satisfy I-maximality.*

The principle is not satisfied by any of the remaining semantics under consideration for general CAFs. The CAF from Example 3.5 possesses the claim-extensions $\{x\}$, $\{x, y\}$ which are accepted under all except i-grounded, h-naive, and h-preferred semantics.

We obtain more positive results on well-formed CAFs: using our ⊆-inclusion results from Section 5, we obtain that all variants of preferred, stable, semi-stable, and stage semantics satisfy I-maximality in $\mathfrak{C}_{wf}$.

Proposition 6.17. *I-grounded, h-naive, all variants of preferred, semi-stable, stage, and stable semantics satisfy I-maximality in $\mathfrak{C}_{wf}$.*

Counter-examples for the remaining semantics are by the respective counter-examples for AFs (using Lemma 5.1).

Reinstating claims The *reinstatement principle* first studied in [21] states that an extension should contain all arguments it defends. CF-reinstatement additionally requires that the extension is not in conflict with the arguments it defends. We transfer this principles to the context of claims. These principles capture the characteristics of complete semantics.

We adapt reinstatement and CF-reinstatement as follows: In order to reinstate a claim c, it suffices that some argument that supports this claim is defended. We require that a given claim-set S contains all claims that are defended by some realization of S. For CF-reinstatement, we consider only realizations of S and occurrences of c that are conflict-free.

Principles 6.18. *Let ρ be a CAF semantics and let $\mathfrak{C}$ be a class of CAFs.*

(Reinstatement) *ρ satisfies* reinstatement *in $\mathfrak{C}$ iff for every CAF $\mathcal{F} \in \mathfrak{C}$, for every $S \in \rho(\mathcal{F})$, if there is a realization E of S in $\mathcal{F}$ that defends an argument $a \in A_{\mathcal{F}}$ then $cl(a) \in S$.*

(CF-Reinstatement) *ρ satisfies* CF-reinstatement *in $\mathfrak{C}$ iff for every CAF $\mathcal{F} \in \mathfrak{C}$, for every $S \in \rho(\mathcal{F})$, if there is a realization E of S in $\mathcal{F}$ that defends an argument $a \in A_{\mathcal{F}}$ and $E \cup \{a\}$ is conflict-free then $cl(a) \in S$.*

We prove the following relation between CF-reinstatement and h-naivety.

Proposition 6.19. *H-naivety implies CF-reinstatement.*

Proof. Consider a semantics ρ that satisfies h-naivety. Consider a CAF $\mathcal{F} \in \mathfrak{C}$ and let $S \in \rho(\mathcal{F})$. Furthermore assume that there is a realization E of S in $\mathcal{F}$ that defends an argument $a \in A_{\mathcal{F}}$ and $E \cup \{a\}$ is conflict-free. By the h-naivety principle, it holds that S is $\subseteq$-maximal in $cf_i(\mathcal{F})$.

Towards a contradiction, assume $cl(a) \notin S$. Then $S \cup \{cl(a)\} \in cf_i(\mathcal{F})$ and $S \subseteq S \cup \{cl(a)\}$. This is a contradiction to $\subseteq$-maximality of S in $cf_i(\mathcal{F})$. Hence, we obtain $cl(a) \in S$, as desired. □

Hence, we obtain that CF-reinstatement is satisfied for h-naive semantics and all variants of stable and stage semantics in $\mathfrak{C}_{wf}$. Moreover, since each realization of a claim-set S attacks—and thus defends—the same arguments in well-formed CAFs, we obtain that both principles are satisfied by each semantics which yields complete extensions in $\mathfrak{C}_{wf}$ (by definition, each semantics satisfies conflict-freeness as stated above).

Also, h-naive semantics satisfies CF-reinstatement also in the general case. This gives h-naive semantics an exclusive status as it is the only semantics under consideration that retains this fundamental property for general CAFs.

Proposition 6.20. *H-naive semantics satisfies CF-reinstatement. Moreover, i-complete, i-grounded, and all variants of semi-stable, preferred, and stable semantics satisfy reinstatement and CF-reinstatement in $\mathfrak{C}_{wf}$. Both variants of stage semantics satisfy CF-reinstatement in $\mathfrak{C}_{wf}$.*

I-admissible, i-conflict-free, and both variants of stage and naive semantics violate reinstatement, even for well-formed CAFs—the corresponding counter-examples coincide with those for AFs. Likewise, we obtain counter-examples for i-admissible and i-conflict-free semantics for CF-reinstatement.

For i-naive semantics, we obtain the following counter-example:

Example 6.21. Consider the CAF $\mathcal{F}$ given as follows:

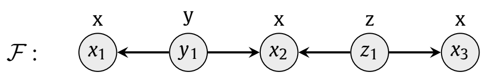

The i-naive extensions of $\mathcal{F}$ are $\{x\}$, $\{x, y\}$, $\{x, z\}$, and $\{y, z\}$. For $S = \{x\}$, we can find a conflict-free realization E of x, namely $E = \{x_3\}$, that defends y_1 (the argument has no attacker) and $E \cup \{y_1\}$ is conflict-free. Nevertheless, $cl(y_1) = y$ is not contained in S. Note that $\mathcal{F}$ is indeed well-formed. Therefore, i-naive semantics does not satisfy CF-reinstatement, not even on well-formed CAFs.

The following counter-examples show that none of the considered semantics satisfies reinstatement for general CAFs; moreover, CF-reinstatement is satisfied by h-naive semantics only:

Example 6.22. Let us consider the following three CAFs $\mathcal{F}$, $\mathcal{F}'$, and $\mathcal{F}''$, where the latter two are small adaptions of $\mathcal{F}$:

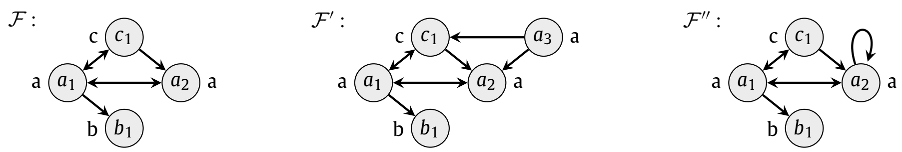

First, we consider the CAF $\mathcal{F}$ and observe that the claim-set $S = \{a\}$, witnessed by realization $\{a_1\}$, is a ρ-extension of $\mathcal{F}$ for all except i-grounded and h-naive semantics. The realization $E = \{a_2\}$ of S defends the argument b_1 against the attack from a_1; moreover, $E \cup \{b_1\}$ is conflict-free, nevertheless, $cl(b_1) = b$ is not contained in S.

For i-grounded semantics, we adapt $\mathcal{F}$ by adding another argument a_3 with claim a that attacks c_1 and a_2—the resulting CAF is called $\mathcal{F}'$ and is depicted above. This argument defends a_1; thus, $\{a\}$, witnessed by $\{a_1, a_3\}$, is i-grounded in the modified CAF. The realization $E = \{a_2\}$ of $\{a\}$ serves as counter-example also in this case. It follows that all except h-naive semantics fail to satisfy reinstatement and CF-reinstatement for general CAFs.

The third CAF $\mathcal{F}''$ shows that h-naive semantics fail to satisfy reinstatement for general CAFs: The realization $E = \{a_2\}$ of $S = \{a\}$ defends b_1 although b is not contained in S.

Range and rejection Finally, we turn towards principles that are characteristic for range-based semantics. We consider two different principles. The *definite claims principle* formalizes the intrinsic motivation of semantics that maximize over the range in the context of claim-based reasoning. The principle is satisfied by semantics that maximize the set of claims that are either accepted or rejected (with respect to a given set of arguments). Given a set of arguments E, we call claims that are either accepted by E (that is, contained in $cl(E)$) or defeated by E *defined* with respect to E. Having as little ambiguity (i.e., claims that are not defined) as possible approximates two-valued semantics. The *justified rejection principle* requires that a claim is contained in some extension if it is not defeated by any claim-extension. We consider only claims that are *cf-realizable*, that is, there is some argument with this claim that is not self-attacking. Intuitively, the principle formalizes that a claim must be either credulously acceptable (i.e., contained in some extension) or credulously defeated (i.e., defeated by some extension).

Principles 6.23. *Let ρ be a CAF semantics and let $\mathfrak{C}$ be a class of CAFs.*

(Definite Claims) *ρ satisfies the* definite claims principle *in $\mathfrak{C}$ iff for every CAF $\mathcal{F} \in \mathfrak{C}$, for every $S \in \rho(\mathcal{F})$, there is a ρ-realization E of S such that the set of defined claims w.r.t. E is $\subseteq$-maximal in $\rho(\mathcal{F})$.*

(Justified Rejection) *ρ satisfies* justified rejection *in $\mathfrak{C}$ iff for every CAF $\mathcal{F} \in \mathfrak{C}$, for every cf-realizable claim $c \in cl(A_{\mathcal{F}})$, if there is no $S \in \rho(\mathcal{F})$ with $c \in S$ then there is some ρ-realization E of a claim-set $S' \in \rho(\mathcal{F})$ that defeats c in $\mathcal{F}$.*

The definite claims principle is satisfied by all variants of stable semantics as well as by h-semi-stable and h-stage semantics.

Proposition 6.24. *I-grounded, h-semi-stable, h-stage, and all variants of stable semantics satisfy the definite claims principle.*

The well-formed CAF from Example 2.1 provides a counter-example for all remaining semantics. Indeed, $\{b, d\}$ and $\{a\}$ are acceptable claim-sets but the range of the unique realization of $\{b, d\}$ is not $\subseteq$-maximal.

We next discuss satisfaction of the justified rejection principle.

Proposition 6.25. *I-conflict-free, h-stage and all variants of naive and stable semantics satisfy justified rejection.*

Proof. I-conflict-free, i-naive, and h-naive semantics satisfy the justified rejection principle because, by definition, if a claim c has an occurrence that is not self-attacking, then there is an extension that contains this claim; thus, the premise is never satisfied. Moreover, all stable variants satisfy justified rejection: if an extension does not contain a given claim c then c is defeated by it.

Finally, also h-stage semantics satisfy justified rejection: Since there is some conflict-free set E in the underlying AF that contains the given claim c, either $cl(E)$ extends to a set with $\subseteq$-maximal range (thus, the premise is not satisfied) or there is some other set D that defeats c. □

Proposition 6.26. *I-stage semantics satisfy justified rejection in $\mathfrak{C}_{wf}$.*

Proof. Consider a CAF $\mathcal{F}$ and a claim $c \in cl(A)$. Assume that c is not contained in any i-stable set. Observe that the set E_c of all cf-realizable arguments with claim c attacks all remaining occurrences of it., i.e., all occurrences of c are contained in the range of E_c. By our assumption, there must be a stage set $E \subseteq A$ such that $x \in E^+$ for all arguments with $cl(x) = c$ (otherwise, E_c is incomparable with all other stage sets and is thus contained in some stage set). Hence, $cl(E)$ defeats c. □

In general, i-stage and i-semi-stable semantics do not satisfy this principle:

Example 6.27. Let us consider the following CAF $\mathcal{F}$:

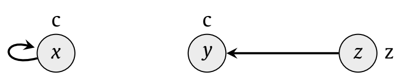

$\{z\}$ is the unique stage and semi-stable extension in the underlying AF. However, the extension does not defeat claim c.

For the remaining admissible-based semantics, we consider the following counter-example:

Table 3
Basic principles w.r.t. general CAFs.

	Confl.-free	Defense/Adm.	i-Na.	h-Na.	I-Max.	Reinst.	CF-Reinst.	Def. Claims	Just. Reject.
cf_i	✓	✗	✗	✗	✗	✗	✗	✗	✓
ad_i	✓	✓	✗	✗	✗	✗	✗	✗	✗
gr_i	✓	✓	✗	✗	✓	✗	✗	✓	✗
co_i	✓	✓	✗	✗	✗	✗	✗	✗	✗
pr_i	✓	✓	✗	✗	✗	✗	✗	✗	✗
pr_h	✓	✓	✗	✗	✓	✗	✗	✗	✗
stb_i	✓	✓	✓	✗	✗	✗	✗	✓	✓
$cf\text{-}stb_h$	✓	✗	✓	✗	✗	✗	✗	✓	✓
$ad\text{-}stb_h$	✓	✓	✓	✗	✗	✗	✗	✓	✓
ss_i	✓	✓	✗	✗	✗	✗	✗	✗	✗
ss_h	✓	✓	✗	✗	✗	✗	✗	✓	✗
na_i	✓	✗	✓	✗	✗	✗	✗	✗	✓
na_h	✓	✗	✓	✓	✓	✗	✓	✗	✓
stg_i	✓	✗	✓	✗	✗	✗	✗	✗	✗
stg_h	✓	✗	✓	✗	✗	✗	✗	✓	✓

Table 4
Basic principles w.r.t. well-formed CAFs.

	Confl.-free	Defense/Adm.	i-Na.	h-Na.	I-Max.	Reinst.	CF-Reinst.	Def. Claims	Just. Reject.
cf_i	✓	✗	✗	✗	✗	✗	✗	✗	✓
ad_i	✓	✓	✗	✗	✗	✗	✗	✗	✗
gr_i	✓	✓	✗	✗	✓	✓	✓	✓	✗
co_i	✓	✓	✗	✗	✗	✓	✓	✗	✗
pr_i	✓	✓	✗	✗	✓	✓	✓	✗	✗
pr_h	✓	✓	✗	✗	✓	✓	✓	✗	✗
stb_i	✓	✓	✓	✓	✓	✓	✓	✓	✓
$cf\text{-}stb_h$	✓	✓	✓	✓	✓	✓	✓	✓	✓
$ad\text{-}stb_h$	✓	✓	✓	✓	✓	✓	✓	✓	✓
ss_i	✓	✓	✗	✗	✓	✓	✓	✗	✗
ss_h	✓	✓	✗	✗	✓	✓	✓	✓	✗
na_i	✓	✗	✓	✗	✗	✗	✗	✗	✓
na_h	✓	✗	✓	✓	✓	✗	✓	✗	✓
stg_i	✓	✗	✓	✓	✓	✗	✓	✗	✓
stg_h	✓	✗	✓	✓	✓	✗	✓	✓	✓

Example 6.28. Let us consider the following well-formed CAF $\mathcal{F}$:

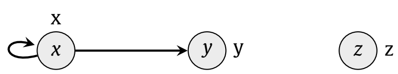

$\{z\}$ is the only (non-empty) admissible set; thus, it is the unique candidate for all admissible-based realizations. Nevertheless, z does not defeat y.

Summary Table 3 and Table 4 summarize our results for general and well-formed CAFs, respectively.

Our first observation is that several principles behave similar in both CAF classes. Since our primary goal of this section was to identify characteristic features of the different semantics, it is not surprising that the differences can be rather found along the semantics. Consider, for instance, the conflict-freeness, defense, and admissibility principle that characterize the underlying properties of the realizations. Apart from h-*cf*-stable semantics which violates defense and admissibility in the general case but satisfies both principles with respect to well-formed CAFs, the semantics behave the same irrespective of the considered class. Similar, the defeat-based principles definite claims and justified rejection are satisfied by the same semantics in both cases. Both principles are satisfied by stable variants of the semantics. Justified rejection is additionally satisfied by principles that are based on conflict-freeness while the definite claims principle is satisfied by h-semi-stable and h-stage semantics.

For h-naivety, I-maximality, reinstatement, and CF-reinstatement, the picture looks different: in the general case, reinstatement is not satisfied by any semantics while h-naivety and CF-reinstatement are both only satisfied by h-naive semantics. I-maximality is additionally satisfied by h-preferred and i-grounded semantics. As all of these properties are considered characteristic for naive semantics, our results indicate that the hybrid naive variant can be seen as reasonable generalization of naive semantics to claim-based semantics. This theory is underlined by the fact that i-naive semantics does not satisfy any of the aforementioned principles, even in the well-formed case.

All of these principles are designed to capture characteristics of claim-based semantics that are relevant in different situations. If one aims to satisfy as much principles as possible, a strong candidate would be i-stable, h-*ad*-stable, or h-

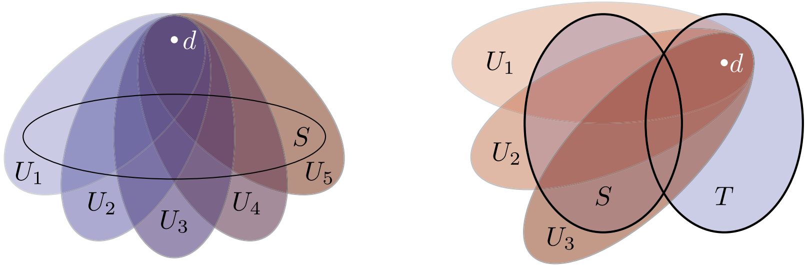

(a) Tightness-criteria visualized. (b) Conflict-sensitivity-criteria visualized.

Fig. 4. Graphical representation of the required conditions of tightness (4a) and conflict-sensitivity (4b): In Fig. 4a, the set S is covered by the upper union $\bigcup_{i\leq 5} U_i$ of d. If tightness is satisfied by semantics ρ, then $S\cup\{d\}$ is contained in $\rho(\mathcal{F})$ for each $\mathcal{F}$. Fig. 4b depicts the upper union $\bigcup_{i\leq 3} U_i$ of an element $d\in T$ which contains S. If S is contained in the upper union of each element of T, then $S\cup T$ is a claim-extension with respect to a semantics ρ that satisfies conflict-sensitivity.

naive semantics (although the latter is considered less preferred since it fails to satisfy defense and admissibility). In the well-formed case, we have a clear winner: stable semantics satisfy all considered principles. However, depending on the application, some principles might be of lower priority. For instance, in order to ensure that as much claims as possible are accepted, it is reasonable to choose a semantics that satisfies I-maximality. In such a situation, it might be negligible whether the remaining claims are defined (i.e., rejected) or not. Overall, the choice of the right semantics is highly context-dependent. Here, our results can help to guide the search for a suitable semantics.

6.3. Extension-set principles

In this section, our object of interest is the structure of so-called *extension-sets*, i.e., sets of claims or, to be more precise, the set of all claim-extensions that are acceptable with respect to a given semantics. The goal of this section is to identify properties of the semantics that characterize the expressiveness of a semantics. We recall classical set-theoretical principles and introduce novel principles in order to identify subtle differences between extension-sets for claim-based semantics.

First, we recall principles from [23,49] that play a central role for the expressiveness of AF semantics. Roughly speaking, these properties explain why particular sets (of arguments or, in our case, of claims) are not jointly acceptable with respect to a particular semantics. The following notion will be useful.

Definition 6.29. Given $\mathbb{S}\subseteq 2^{\mathcal{C}}$ and a set $S\subseteq\bigcup_{T\in\mathbb{S}} T$, we define the *upper union* of S in $\mathbb{S}$ as

$$\mathrm{up}_{\mathbb{S}}(S)=\bigcup_{S\subseteq T\in\mathbb{S}} T.$$

If we consider an I-maximal extension-set $\mathbb{S}$, we observe that the upper union becomes the identity function on $\mathbb{S}$. The upper union contains in this case only the input-set.

Lemma 6.30. *Given a semantics ρ that satisfies I-maximality and a CAF $\mathcal{F}$, it holds that $S=\mathrm{up}_{\rho(\mathcal{F})}(S)$ for each $S\in\rho(\mathcal{F})$.*

Principles 6.31. *Let ρ be a CAF semantics and let $\mathfrak{C}$ be a class of CAFs.*

(Downward Closure) *ρ is* downward closed *in $\mathfrak{C}$ iff for every CAF $\mathcal{F}\in\mathfrak{C}$, for every $S\in\rho(\mathcal{F})$, if $T\subseteq S$ then $T\in\rho(\mathcal{F})$.*

(Tightness) *ρ satisfies* tightness *in $\mathfrak{C}$ iff for every CAF $\mathcal{F}\in\mathfrak{C}$, for every $S\in\rho(\mathcal{F})$ and for every claim $d\in cl(A_{\mathcal{F}})$, if $S\in\mathrm{up}_{\rho(\mathcal{F})}(\{d\})$ then $S\cup\{d\}\in\rho(\mathcal{F})$.*

(Conflict-sensitivity) *ρ satisfies* conflict-sensitivity *in $\mathfrak{C}$ iff for every CAF $\mathcal{F}\in\mathfrak{C}$, for every $S,T\in\rho(\mathcal{F})$, if $S\in\mathrm{up}_{\rho(\mathcal{F})}(\{d\})$ for all $d\in T$ then $S\cup T\in\rho(\mathcal{F})$.*

Fig. 4 gives a graphical visualization of tightness and conflict-sensitivity. If tightness is satisfied by a semantics ρ, then $S\subseteq\bigcup_{i\leq 5} U_i=\mathrm{up}_{\rho(\mathcal{F})}(\{d\})$ (as shown in Fig. 4a) implies $S\cup\{d\}\in\rho(\mathcal{F})$ for all CAFs $\mathcal{F}$. Conflict-sensitivity is satisfied by a semantics ρ, if $S\subseteq\bigcup_{i\leq 3} U_i=\mathrm{up}_{\rho(\mathcal{F})}(\{d\})$ as depicted in Fig. 4b for all $d\in T$ implies $S\cup T\in\rho(\mathcal{F})$ for each CAF $\mathcal{F}$.

Remark 6.32. In [23], conflict-sensitivity and tightness has been introduced via so-called *pairs*: a couple c,d forms a pair if there is an extension that contains both a and b. A semantics satisfies conflict-sensitivity iff for every two extensions S,T, if every couple c,d forms a pair then the union of S and T is an extension itself. A semantics satisfies tightness if for every extension S, for every claim d, if each couple c,d is a pair for every $c\in S$, then $S\cup\{d\}$ is an extension. Our formulation is indeed equivalent to the original formulation: S is contained in the upper union of a claim d iff c,d form a pair for all $c\in S$; conflict-sensitivity generalizes this concept to each claim $d\in T$.

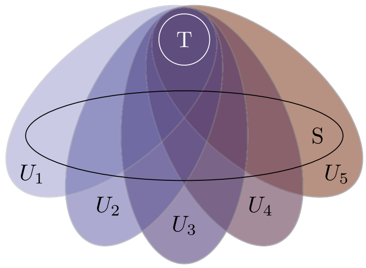

Fig. 5. Graphical representation of the required conditions of cautious closure: the set S is covered by the upper union $\bigcup_{i \leq 5} U_i$ of T. If cautious closure is satisfied by semantics ρ, then this implies that $S \cup T$ is contained in $\rho(\mathcal{F})$. We have replaced the single claim d in Fig. 4a by a set of claims T.

It turns out that these principles are too strong when it comes to claim-based semantics, even for well-formed CAFs. It is easy to see that downward closure is satisfied only by i-conflict-free semantics. Moreover, i-grounded semantics satisfies conflict-sensitivity and tightness since it is a single-status (also called uniquely defined [50]) semantics, i.e., it returns exactly one set. The remaining semantics fail to satisfy these principles, even in the well-formed case, as the following example demonstrates.

Example 6.33. We consider the extension-set $\mathbb{S} = \{\{a, b\}, \{b, c\}, \{a, c\}\}$ which is neither tight nor conflict-sensitive. We generate the following well-formed CAF $\mathcal{F}$:

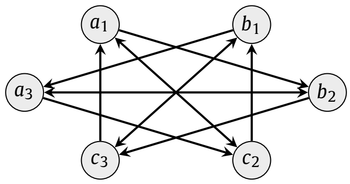

For each claim c in a claim-set $S_i \in \mathbb{S}$, we introduce an argument c_i in $\mathcal{F}$. Each claim-set S is attacked by claims not appearing in S, for example, the set $\{a, b\}$ is attacked by claim c. In this way, we ensure that $\mathcal{F}$ is well-formed. It can be checked that $\rho(\mathcal{F}) = \mathbb{S}$ for h-naive semantics and for all variants of preferred, stable, semi-stable, and stage semantics. Moreover, $\mathbb{S} \cup \{\emptyset\}$ corresponds to $ad_i(\mathcal{F})$ and $co_i(\mathcal{F})$, while $\mathbb{S} \cup \{\{a\}, \{b\}, \{c\}\} = na_i(\mathcal{F})$ and $\mathbb{S} \cup \{\emptyset, \{a\}, \{b\}, \{c\}\} = cf_i(\mathcal{F})$.

We consider a novel principle that can be seen as a relaxation of tightness and conflict-sensitivity.

Principle 6.34 *(Cautious closure). A semantics ρ is* cautiously closed *in $\mathfrak{C}$ iff for every CAF $\mathcal{F} \in \mathfrak{C}$, for every $S, T \in \rho(\mathcal{F})$, if $S \subseteq \mathrm{up}_{\rho(\mathcal{F})}(T)$ then $S \cup T \in \rho(\mathcal{F})$.*

We observe that instead of single claims $c \in cl(A)$, we consider claim-sets that are contained in $\rho(\mathcal{F})$. Fig. 5 provides a graphical representation of this generalized criteria. Next we show that each semantics that satisfies conflict-sensitivity also satisfies cautious closure. It follows that each AF semantics that satisfies conflict-sensitivity (e.g., admissible, grounded, preferred, stable, semi-stable, and stage semantics) satisfies the generalized principle as well.

Proposition 6.35. *Conflict-sensitivity implies cautious closure.*

Proof. Given a CAF $\mathcal{F}$ and two sets $S, T \in \rho(\mathcal{F})$ with $S \subseteq \mathrm{up}_{\rho(\mathcal{F})}(T)$. Hence, S is contained in the upper union of each single claim $d \in T$, i.e., $S \in \mathrm{up}_{\rho(\mathcal{F})}(\{d\})$ for all $d \in T$. If $\rho(\mathcal{F})$ is conflict-sensitive, we obtain $S \cup T \in \rho(\mathcal{F})$. □

Since I-maximal extension-sets $\mathbb{S}$ satisfy $S = \mathrm{up}_{\mathbb{S}}(S)$ for each $S \in \mathbb{S}$, we obtain that each semantics that satisfies I-maximality satisfies cautious closure as well.

Proposition 6.36. *I-maximality implies cautious closure.*

We obtain that i-grounded, h-preferred and h-naive semantics satisfy cautious closure even in the general case.

Proposition 6.37. *I-grounded, h-preferred and h-naive semantics satisfy cautious closure in $\mathfrak{C}_u$.*

For the remaining semantics, we consider the following counter-example:

Example 6.38. We consider the extension-set $\mathbb{S} = \{\{a, b\}, \{b, c\}, \{a, c\}, \{a\}\}$. The set $\mathbb{S}$ is not cautiously closed: indeed, the upper union of $\{a\}$ is given by $\{a, b, c\}$ and thus contains $\{b, c\}$. Nevertheless, $\{a, b, c\}$ is not contained in $\mathbb{S}$.

We generate the following CAF $\mathcal{F}$ by introducing an argument c_i for each claim c, for each claim-set $S_i \in \mathbb{S}$. Moreover, $cl(c_i) = c$. The attack-relation is defined as follows: two arguments c_i, d_j attack each other iff $i \neq j$.

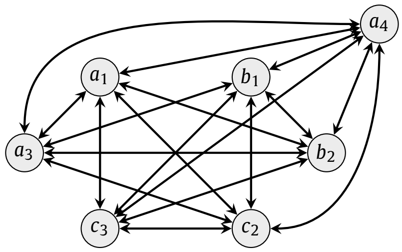

The construction ensures that each claim-set has its unique realization that attacks all remaining arguments. In $\mathcal{F}$, all attacks are symmetric and thus admissible-based semantics and semantics based on conflict-freeness coincide. We obtain that all considered semantics ρ apart from i-grounded, h-naive, and h-preferred semantics satisfy $\mathbb{S} \subseteq \rho(\mathcal{F})$. Moreover, the set $\{a, b, c\}$ is not accepted with respect to any of the considered semantics. It follows that cautious closure is violated by all semantics under consideration (apart from i-grounded, h-naive, and h-preferred semantics).

Cautious closure is satisfied by several semantics if restricted to well-formed CAFs. First, by Proposition 6.36, we obtain that h-naive and all variants of stable, preferred, semi-stable, and stage semantics satisfy cautious closure.

Proposition 6.39. *I-grounded, h-naive, and all variants of stable, preferred, semi-stable, and stage semantics satisfy cautious closure in $\mathfrak{C}_{wf}$.*

Next we show that i-admissible semantics satisfy this principle in $\mathfrak{C}_{wf}$.

Proposition 6.40. *I-admissible semantics satisfy cautious closure in $\mathfrak{C}_{wf}$.*

Proof. Given a well-formed CAF $\mathcal{F}$ and let $S, T \in ad_i(\mathcal{F})$ with $S \subseteq \mathrm{up}_{\rho(\mathcal{F})}(T)$. We show that $S \cup T \in ad_i(\mathcal{F})$.

Consider ad-realizations $E, D \subseteq A$ of S and T, respectively. By Dung's fundamental lemma, the union $E \cup D$ defends itself in F. Now assume there is a conflict in $E \cup D$, i.e., there are arguments $x, y \in E \cup D$ such that $(x, y) \in R$. w.l.o.g., let $x \in E$ and $y \in D$ (as both E, D are admissible it is not the case that both arguments x, y are contained in either E or D). Since $S \subseteq \mathrm{up}_{\rho(\mathcal{F})}(T)$ there is some i-admissible superset $T' \supseteq T$ such that $T \cup \{cl(x)\} \subseteq T'$. Let D' denote an ad-realization of T' and let $x' \in D'$ denote the occurrence of $cl(x)$ in D', that is, $cl(x') = cl(x)$. Then $(x', y) \in R$ by well-formedness. Since D defends itself, there is an argument $z \in D$ that attacks x'. Let $z' \in D'$ denote the occurrence of claim $cl(z')$ in D', that is, $cl(z') = cl(z)$. By well-formedness, we have that $(z', x') \in R$, which contradicts $D' \in ad(F)$. □

Complete, i-conflict-free and i-naive semantics do not satisfy cautious closure. Example 6.33 serves as a counter-example for i-conflict-free and i-naive semantics; for i-complete semantics, we consider the following counter-example.

Example 6.41. Consider the following CAF where each argument is assigned its unique argument name (i.e., $cl = id$):

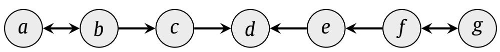

Both $\{b\}$ and $\{f\}$ are complete, but their union $\{b, f\}$ is not complete as it defends the argument d.

We consider a relaxation of cautious closure.

Principle 6.42 *(Weak cautious closure). A semantics ρ is* weakly cautiously closed *iff for every CAF $\mathcal{F}$, for every $S, T \in \rho(\mathcal{F})$, if $\mathrm{up}_{\rho(\mathcal{F})}(T)$ then there is $U \in \rho(\mathcal{F})$ with $S \cup T \subseteq U$.*

First, we observe that each semantics that satisfies cautious closure also satisfies weak cautious closure.

Proposition 6.43. *Cautious closure implies weak cautious closure.*

We thus obtain the following result.

Proposition 6.44. *I-grounded, h-preferred and h-naive semantics satisfy weak cautious closure in $\mathfrak{C}_u$.*

Example 6.38 serves as counter-example for the remaining semantics in the general case.

For well-formed CAFs, we obtain that i-complete semantics satisfy this weaker version of cautious closure.

Proposition 6.45. *Complete semantics satisfy weak cautious closure in $\mathfrak{C}_{wf}$.*

Proof. To show that $co_i(\mathcal{F})$ is weakly cautiously closed for each well-formed CAF $\mathcal{F}$, consider two claim-sets $S, T \in co_i(\mathcal{F})$ with $\mathrm{up}_{\rho(\mathcal{F})}(T)$. Clearly, S and T are i-admissible in $\mathcal{F}$. By Proposition 6.40, we obtain $S \cup T \in ad_i(\mathcal{F})$. Thus, there is some i-complete claim-set $U \in co_i(\mathcal{F})$ with $S \cup T \subseteq U$. □

By Proposition 6.43, we additionally obtain the following result.

Proposition 6.46. *I-grounded, i-admissible, h-naive, and all variants of stable, preferred, semi-stable, and stage semantics satisfy weak cautious closure in $\mathfrak{C}_{wf}$.*

Example 6.33 shows that weak cautious closure is violated by i-naive and i-conflict-free semantics for well-formed CAFs.

Let us next consider a principle that characterizes a crucial property of i-complete semantics. If two extensions S, T are contained in some other extension U, i.e., $S \cup T \subseteq U$, then there is a unique $\subseteq$-minimal extension that contains $S \cup T$. For this, it will be useful to define so-called *completion-sets* of a given set of claims.

Definition 6.47. Given a CAF $\mathcal{F}$, a semantics ρ and a set of claims $S \subseteq cl(A)$, we let $\mathrm{c}_{\rho(\mathcal{F})}(S) = \{T \in \rho(\mathcal{F}) \mid S \subseteq T, \nexists T' \in \rho(\mathcal{F}) : S \subseteq T' \subset T\}$ denote the minimal completion-sets of S in $\mathcal{F}$.

If $|\mathrm{c}_{\rho(\mathcal{F})}(S)| = 1$ we slightly abuse notation and write $\mathrm{c}_{\rho(\mathcal{F})}(S)$ to denote the unique minimal completion-set of S.

Principle 6.48 *(Unique completion). A semantics ρ satisfies* unique completion *in $\mathfrak{C}$ iff for every CAF $\mathcal{F} \in \mathfrak{C}$, for every $S, T \in \rho(\mathcal{F})$, $|\mathrm{c}_{\rho(\mathcal{F})}(S \cup T)| \leq 1$.*

Proposition 6.49. *Cautious closure implies unique completion.*

Proof. The unique completion of two extensions $S, T \in \rho(\mathcal{F})$ in question is given by the union $T \cup S$. In case there are several completions of $T \cup S$, we have that $S \subseteq \mathrm{up}_{\rho(\mathcal{F})}(T)$ and thus $S \cup T \in \rho(\mathcal{F})$. □

We thus obtain that unique completion is satisfied by i-grounded, h-naive, and h-preferred semantics in the general case and additionally by i-admissible and all versions of stable, semi-stable, and stage semantics in $\mathfrak{C}_{wf}$.

Proposition 6.50. *I-grounded, h-naive, and h-preferred semantics satisfy unique completion in $\mathfrak{C}_u$. Moreover, i-admissible, h-naive, and all variants of stable, preferred, semi-stable, and stage semantics satisfy unique completion in $\mathfrak{C}_{wf}$.*

For general CAFs, the principle is not satisfied by any of the remaining semantics: a counter-example is given by Example 6.38, here, $\{a\}$ has two minimal completions $\{a, b\}$ and $\{a, c\}$.

Likewise, neither i-naive nor i-conflict-free semantics satisfy unique completion in $\mathfrak{C}_{wf}$: in Example 6.33, the sets $\{a, b\}, \{a, c\}, \{b, c\}$ as well as the singletons $\{a\}, \{b\}, \{c\}$ are i-conflict-free and i-naive claim-sets. Thus, each singleton has two minimal completions.

We end this section by showing that for well-formed CAFs, unique completion is satisfied by i-complete semantics.

Proposition 6.51. *Complete semantics satisfy unique completion in $\mathfrak{C}_{wf}$.*

Proof. Recall that in well-formed CAFs, each realization of a claim-set attacks the same arguments. Thus, every realization of $T \cup S$ for two extensions $S, T \in co_i(\mathcal{F})$ in a well-formed CAF $\mathcal{F}$ defends the same arguments. It follows that $S \cup T$ admits a unique completion in case $T \cup S$ is ad-realizable in $\mathcal{F}$. □

Table 5
Extension-set principles w.r.t. general CAFs.

	Downw. Closure	Tight	Conflict-sensitive	Cautious Closure	w-Cautious Closure	Unique Compl.
cf_i	✓	✗	✗	✗	✗	✗
ad_i	✗	✗	✗	✗	✗	✗
gr_i	✗	✓	✓	✓	✓	✓
co_i	✗	✗	✗	✗	✗	✗
pr_i	✗	✗	✗	✗	✗	✗
pr_h	✗	✗	✗	✓	✓	✓
stb_i	✗	✗	✗	✗	✗	✗
$cf\text{-}stb_h$	✗	✗	✗	✗	✗	✗
$ad\text{-}stb_h$	✗	✗	✗	✗	✗	✗
ss_i	✗	✗	✗	✗	✗	✗
ss_h	✗	✗	✗	✗	✗	✗
na_i	✗	✗	✗	✗	✗	✗
na_h	✗	✗	✗	✓	✓	✓
stg_i	✗	✗	✗	✗	✗	✗
stg_h	✗	✗	✗	✗	✗	✗

Table 6
Extension-set principles w.r.t. well-formed CAFs.

	Downw. Closure	Tight	Conflict-sensitive	Cautious Closure	w-Cautious Closure	Unique Compl.
cf_i	✓	✗	✗	✗	✗	✗
ad_i	✗	✗	✗	✓	✓	✓
gr_i	✗	✓	✓	✓	✓	✓
co_i	✗	✗	✗	✗	✓	✓
pr_i	✗	✗	✗	✓	✓	✓
pr_h	✗	✗	✗	✓	✓	✓
stb_i	✗	✗	✗	✓	✓	✓
$cf\text{-}stb_h$	✗	✗	✗	✓	✓	✓
$ad\text{-}stb_h$	✗	✗	✗	✓	✓	✓
ss_i	✗	✗	✗	✓	✓	✓
ss_h	✗	✗	✗	✓	✓	✓
na_i	✗	✗	✗	✗	✗	✗
na_h	✗	✗	✗	✓	✓	✓
stg_i	✗	✗	✗	✓	✓	✓
stg_h	✗	✗	✗	✓	✓	✓

Summary We summarize our results in Table 5 and Table 6 for general and well-formed CAFs, respectively. Apart from i-grounded semantics which satisfies almost all set-theoretical principles under consideration by definition, only h-naive and h-preferred semantics satisfy (weak) cautious closure and unique completion in the general case. Tightness and conflict-sensitivity are also not satisfied in the well-formed case. Cautious closure, on the other hand, is satisfied by all but i-complete admissible-based semantics; weak cautious closure and unique completion are satisfied by i-complete semantics as well. Overall, our extension-set principles indicate a higher expressiveness of claim semantics, compared to AF semantics, even for the well-formed case.

7. Expressiveness

In this section, we investigate the expressive power of the considered semantics. As already observed in the previous section, claim-based semantics are in general more expressive than their AF counterparts. Indeed, several semantics violate I-maximality in the general case. Moreover, it is possible to construct (well-formed) CAFs that violate tightness and conflict-sensitivity which is impossible for e.g., preferred and admissible semantics, respectively, as shown by Dunne et al. [23].

In order to study the expressive power of the considered semantics, we provide characterizations of the *signatures* of the semantics [23]. The signature captures all possible outcomes which can be obtained by argumentation frameworks when evaluated under a semantics and thus characterizes the expressiveness of a semantics.

Formally, the signature Σ_σ^{AF} of an AF-semantics σ is defined as $\Sigma_\sigma^{AF} = \{\sigma(F) \mid F \text{ is an AF}\}$. We adapt the concept to CAFs and well-formed CAFs as follows.

Definition 7.1. Given a semantics τ, the *signature* of τ with respect to general and well-formed CAFs, respectively, is given by

$$\Sigma_\tau^{CAF} = \{\tau(\mathcal{F}) \mid \mathcal{F} \text{ is a CAF}\}$$

$$\Sigma_\tau^{wf} = \{\tau(\mathcal{F}) \mid \mathcal{F} \text{ is a well-formed CAF}\}.$$

Note that Σ_σ^{AF} yields a collection of sets of arguments while Σ_τ^{CAF} and Σ_τ^{wf} yield a collection of sets of claims. In order to compare argument-based signatures with their claim-based variants, we assume that the domains for arguments and claims coincide, i.e. $U = C$. We identify AFs with CAFs where each argument is assigned its unique argument name (i.e., $cl = id$) as done in Section 6. For any AF-semantics σ, it holds that

$$\Sigma_\sigma^{AF} \subseteq \Sigma_{\sigma_i}^{wf} \subseteq \Sigma_{\sigma_i}^{CAF} \quad \text{and} \quad \Sigma_\sigma^{AF} \subseteq \Sigma_{\sigma_h}^{wf} \subseteq \Sigma_{\sigma_h}^{CAF}$$

since each AF corresponds to a (well-formed) CAF with an unique claim per argument. Moreover, each well-formed CAF is indeed a CAF.

7.1. Expressiveness of CAF semantics

We begin our investigations with the class of general CAFs. As we will see, almost every extension-set can be expressed with only very soft restrictions, i.e., CAF semantics are in general very expressive, as the following theorem shows:

Theorem 7.2. *The following characterizations hold:*

$$\Sigma_{gr_i}^{CAF} = \{\mathbb{S} \subseteq 2^C \mid |\mathbb{S}| = 1\}$$
$$\Sigma_{cf_i}^{CAF} = \{\mathbb{S} \subseteq 2^C \mid \mathbb{S} \neq \emptyset, \mathbb{S} \textit{ is downwards closed}\}$$
$$\Sigma_{ad_i}^{CAF} = \{\mathbb{S} \subseteq 2^C \mid \emptyset \in \mathbb{S}\}$$
$$\Sigma_{co_c}^{CAF} = \{\mathbb{S} \subseteq 2^C \mid \mathbb{S} \neq \emptyset, \textstyle\bigcap_{S \in \mathbb{S}} S \in \mathbb{S}\}$$
$$\Sigma_{\rho}^{CAF} = \{\mathbb{S} \subseteq 2^C \mid \mathbb{S} \neq \emptyset, \mathbb{S} \textit{ is I-maximal}\},\ \rho \in \{pr_h, na_h\}$$
$$\Sigma_{\rho}^{CAF} = \{\mathbb{S} \subseteq 2^C \mid \mathbb{S} = \{\emptyset\} \textit{ or } \emptyset \notin \mathbb{S}\},\ \rho \in \{stb_i, cf\text{-}stb_h, ad\text{-}stb_h\}$$
$$\Sigma_{\rho}^{CAF} = \Sigma_{stb_c}^{CAF} \setminus \{\emptyset\},\ \rho \in \{pr_i, na_i, ss_i, ss_h, stg_i, stg_h\}$$

From Section 6, we know that i-conflict-free semantics are downwards closed and that h-preferred and h-naive semantics satisfy I-maximality. This confirms that it is impossible to construct CAFs where conflict-free extension-sets are not downwards-closed or, e.g., h-naive semantics violate I-maximality, as postulated in the theorem. Moreover, the i-grounded extension is always unique, the empty set is always i-admissible, the intersection of all i-complete sets is i-complete, and all variants of stable semantics might return empty extension-sets.

In the remaining part of this section, we show that for each extension-set $\mathbb{S}$ which obeys the ρ-specific requirements, we can construct a CAF $\mathcal{F}$ that returns exactly $\mathbb{S}$ as ρ-extensions, i.e., $\rho(\mathcal{F}) = \mathbb{S}$.

First, each extension-set $\mathbb{S}$ with $|\mathbb{S}| = 1$ is expressible under i-grounded semantics: it suffices to consider the CAF $\mathcal{F} = (\{c \in S \mid S \in \mathbb{S}\}, \emptyset, id)$ with no attacks. Second, in order to obtain $\mathbb{S} = \{\emptyset\}$ we consider the empty framework $\mathcal{F} = (\emptyset, \emptyset, cl)$ which satisfies $\rho(\mathcal{F}) = \mathbb{S}$ for all considered semantics. Third, stable semantics can express $\mathbb{S} = \emptyset$: as for AFs, it suffices to consider a single self-attacking argument; the CAF $\mathcal{F} = (\{a\}, \{(a, a)\}, id)$ thus yields an example for $stb_i(\mathcal{F}) = ad\text{-}stb_h(\mathcal{F}) = cf\text{-}stb_h(\mathcal{F}) = \emptyset$.

Next, we define a method which can be used to construct CAFs that return each non-empty extension-set $\mathbb{S}$ that obeys the semantics-specific requirements for all apart from i-admissible and i-complete semantics. Note that we have used the construction already in Section 6 in Example 6.38 to show that i-grounded, h-naive, and h-preferred semantics do not satisfy cautious closure in general. The basic idea is to add an argument c_i for each claim c from claim-set S_i in a given extension-set $\mathbb{S}$ that attacks all arguments not associated to claims in S_i. In this way, each claim-set realizes itself in the resulting CAF.

Construction 7.3. *Given a non-empty extension-set $\mathbb{S} = \{S_1, \ldots, S_n\} \subseteq 2^C$, we define $\mathcal{F}_{\mathbb{S}}^u = (A, R, cl)$ with*

$$A = \{c_i \mid S_i \in \mathbb{S}, c \in S_i\},$$
$$R = \{(c_i, d_j) \mid c_i, d_j \in A, i \neq j\},$$

and $cl(c_i) = c$ for all $c_i \in A$.

Proposition 7.4. *Given a non-empty extension-set $\mathbb{S} \subseteq 2^C$, $\emptyset \notin \mathbb{S}$, let $\mathcal{F}_{\mathbb{S}}^u$ be defined as in Construction 7.3, and let* `Sem` *denote the set of all semantics under consideration. It holds that*

1. *if $\emptyset \notin \mathbb{S}$, $\rho(\mathcal{F}_{\mathbb{S}}^u) = \mathbb{S}$ for $\rho \in \texttt{Sem} \setminus \{cf_i, ad_i, co_i, gr_i, pr_h, na_h\}$;*

2. *if* $\mathbb{S}$ *is I-maximal,* $\rho(\mathcal{F}^u_{\mathbb{S}}) = \mathbb{S}$ *for* $\mathtt{Sem} \setminus \{cf_i, ad_i, co_i, gr_i\}$*;*
3. *if* $\mathbb{S}$ *is downward closed,* $\rho(\mathcal{F}^u_{\mathbb{S}}) = \mathbb{S}$ *for* $\{cf_i, ad_i, co_i\}$*.*

Proof. Consider a non-empty extension-set $\mathbb{S} = \{S_1, \dots, S_n\}$ and let $\mathcal{F}^u_{\mathbb{S}}$ be constructed according to Construction 7.3. For each $S_i \in \mathbb{S}$, we consider the realization $E_i = \{c_i \mid c \in S_i\}$.

(1) To show that $\rho(\mathcal{F}) = \mathbb{S}$ for each of the considered semantics, we first observe that each attack is symmetric. We thus, obtain $pr_i(\mathcal{F}^u_{\mathbb{S}}) = na_i(\mathcal{F}^u_{\mathbb{S}})$ and $ss_i(\mathcal{F}^u_{\mathbb{S}}) = stg_i(\mathcal{F}^u_{\mathbb{S}})$; also, $ss_h(\mathcal{F}^u_{\mathbb{S}}) = stg_h(\mathcal{F}^u_{\mathbb{S}})$ and $cf\text{-}stb_h(\mathcal{F}^u_{\mathbb{S}}) = ad\text{-}stb_h(\mathcal{F}^u_{\mathbb{S}})$ (since $cf_i(\mathcal{F}^u_{\mathbb{S}}) = ad_i(\mathcal{F}^u_{\mathbb{S}})$).
Second, we observe that for each $S_i \in \mathbb{S}$, the realization E_i is stable in the underlying AF, therefore, $stb_i(\mathcal{F}^u_{\mathbb{S}}) \neq \emptyset$ and thus $stb_i(\mathcal{F}^u_{\mathbb{S}}) = ss_i(\mathcal{F}^u_{\mathbb{S}}) = stg_i(\mathcal{F}^u_{\mathbb{S}})$. We moreover obtain $\mathbb{S} \subseteq stb_i(\mathcal{F}^u_{\mathbb{S}})$. As the CAF possesses a stable extension, we furthermore conclude that $cf\text{-}stb_h(\mathcal{F}^u_{\mathbb{S}}) = stg_h(\mathcal{F}^u_{\mathbb{S}})$ (by Lemma 5.9) and thus $ss_h(\mathcal{F}^u_{\mathbb{S}}) = stg_h(\mathcal{F}^u_{\mathbb{S}}) = cf\text{-}stb_h(\mathcal{F}^u_{\mathbb{S}}) = ad\text{-}stb_h(\mathcal{F}^u_{\mathbb{S}})$.
Third, we observe that all stable variants coincide. It suffices to show that $cf\text{-}stb_h(\mathcal{F}^u_{\mathbb{S}}) \subseteq stb_i(\mathcal{F}^u_{\mathbb{S}})$. Consider a h-$cf$-stable set S and its $cf\text{-}stb_h$-realization E in $\mathcal{F}^u_{\mathbb{S}}$. We first observe that $S \subseteq S_i$ for some $S_i \in \mathbb{S}$ because all other claim-sets do not have a conflict-free realization in $\mathcal{F}^u_{\mathbb{S}}$. Moreover, $E \subseteq E_i$ because all other realizations of E are not conflict-free. E attacks all arguments with claims $c \notin S$. Now, assume there is an argument $a \in A \setminus E$ with $cl(a) \in S$ that is not attacked by E. This is the case only if $cl(a) \in S_i$. As each claim of the claim-set S_i has exactly one realization in E_i we have found a claim that is neither defeated nor contained in E, which contradicts our assumption E $cf\text{-}stb_h$-realizes S in $\mathcal{F}^u_{\mathbb{S}}$. Thus, the statement follows.
Finally, we observe that $pr_i(\mathcal{F}^u_{\mathbb{S}}) = stb_i(\mathcal{F}^u_{\mathbb{S}})$ since each $\subseteq$-maximal admissible set in F attacks all other arguments. As there are no other $\subseteq$-maximal admissible sets in the underlying AF we obtain $pr_i(\mathcal{F}^u_{\mathbb{S}}) \subseteq \mathbb{S}$. By $\mathbb{S} \subseteq stb_i(\mathcal{F}^u_{\mathbb{S}}) = pr_i(\mathcal{F}^u_{\mathbb{S}}) \subseteq \mathbb{S}$ we have shown that $\rho(\mathcal{F}^u_{\mathbb{S}}) = \mathbb{S}$ for all considered semantics as required.

(2) Let us now assume that $\mathbb{S}$ is I-maximal. By (1), we obtain the statement for all semantics in $\mathtt{Sem} \setminus \{cf_i, ad_i, co_i, gr_i, pr_h, na_h\}$. Since h-preferred and h-naive semantics can be equivalently defined based on preferred and naive argument-extensions, respectively (cf. Proposition 4.4), it holds that $\rho(\mathcal{F}^u_{\mathbb{S}}) = \mathbb{S}$ for $\rho \in \{pr_h, na_h\}$.

(3) Finally, let us assume that $\mathbb{S}$ is downward-closed. By (1), we obtain that $\mathbb{S} \setminus \{\emptyset\} = \rho(\mathcal{F}^u_{\mathbb{S}})$ for all semantics in $\mathtt{Sem} \setminus \{cf_i, ad_i, co_i, gr_i, pr_h, na_h\}$. As each subset of i-naive claim-sets is i-conflict-free, we obtain $cf_i(\mathcal{F}^u_{\mathbb{S}}) = \mathbb{S}$ as required. As observed in (1), i-conflict-free and i-admissible semantics coincide in $\mathcal{F}^u_{\mathbb{S}}$. Moreover, $\emptyset = \bigcap_{S \in \mathbb{S}} S$ is contained in $\mathbb{S}$, furthermore, each realization E_i of S_i contains all arguments it defends, consequently, we obtain $co_i(\mathcal{F}^u_{\mathbb{S}}) = \mathbb{S}$. □

It remains to give constructions for i-admissible and i-complete semantics. Note that evaluating $\mathcal{F}^u_{\mathbb{S}}$ under i-admissible and i-complete semantics might yield additional claim-sets. As observed in the proof of Proposition 7.4, $ad_i(\mathcal{F}^u_{\mathbb{S}})$ is downwards-closed for each extension-set $\mathbb{S}$. Moreover, the i-grounded extension is always empty in $\mathcal{F}^u_{\mathbb{S}}$ since there are no arguments that are unattacked. Consequently, $\mathbb{S} \cup \{\emptyset\} \subseteq co_i(\mathcal{F}^u_{\mathbb{S}})$ for each extension-set $\mathbb{S}$.

We observe however that in both cases, the construction produces a CAF that accepts *at least* all claim-sets in $\mathbb{S}$ with respect to i-admissible and i-complete semantics.

Proposition 7.5. *Consider an extension-set* $\mathbb{S}$ *and let* $\mathcal{F}^u_{\mathbb{S}}$ *be defined as in Construction 7.3. It holds that* $\mathbb{S} \subseteq \rho(\mathcal{F}^u_{\mathbb{S}})$ *for* $\rho \in \{ad_i, co_i\}$*.*

For i-complete semantics, we adapt the construction appropriately. Let $\mathbb{S}$ be an extension-set with $\mathbb{S} \neq \emptyset$, and the intersection $\bigcap_{S \in \mathbb{S}} S \in \mathbb{S}$. First, we apply Construction 7.3 to $\mathbb{S} \setminus \{\bigcap_{S \in \mathbb{S}} S\}$. In a second step, we add isolated arguments for all claims in $\bigcap_{S \in \mathbb{S}} S$.

The resulting CAF realizes $\mathbb{S}$ under i-complete semantics, as desired.

Proposition 7.6. *Given a non-empty extension-set* $\mathbb{S} \subseteq 2^C$ *with* $\bigcap_{S \in \mathbb{S}} S \in \mathbb{S}$*. Let* $\mathbb{T} = \mathbb{S} \setminus \{\bigcap_{S \in \mathbb{S}} S\}$ *and let* $\mathcal{F}^u_{\mathbb{T}} = (A, R, cl)$ *be defined as in Construction 7.3. We define* $\mathcal{F} = (A \cup A', R, cl')$ *with* $A' = \{a_c \mid c \in \bigcap_{S \in \mathbb{S}} S\}$ *and* $cl'(a_c) = c$ *for* $a_c \in A'$ *and* $cl'(a) = cl(a)$ *otherwise. It holds that* $co_i(\mathcal{F}) = \mathbb{S}$*.*

Proof. Consider an extension-set $\mathbb{S} = \{S_1, \dots, S_n\}$. We first observe that all arguments in A' are not attacked and thus contained in each complete set in F.

Second, we show that each claim-set $S_i \in \mathbb{S}$ is co_i-realized in $\mathcal{F}$: For $S_i = \bigcap_{S \in \mathbb{S}} S$, we observe that $\mathcal{F}$ contains precisely one argument a_c with claim c for all claims $c \in \bigcap_{S \in \mathbb{S}} S$. The set that contains all these arguments—the set A'—defends itself as it is unattacked. Moreover, it does not defend any other arguments as it has no outgoing attacks. Consequently, $\bigcap_{S \in \mathbb{S}} S \in co_i(\mathcal{F})$. We furthermore note that no subset of $\bigcap_{S \in \mathbb{S}} S$ is i-complete.

In case $S_i \neq \bigcap_{S \in \mathbb{S}} S$, we consider the realization $E_i = \{c_i \mid c \in S_i\} \cup A'$ of S_i. Observe that E_i is conflict-free and attacks all remaining arguments by construction. Thus, E_i is stable and, in particular, complete in F. Moreover, no subset of S_i is i-complete since each argument in E_i attacks all arguments in $A \setminus E_i$ and thus defends all arguments in E_i. Finally, we note that no superset of E_i is complete in F. Consequently, $co(F) = \{E_i \mid i \leq n\}$. We thus obtain $co_i(\mathcal{F}) = \mathbb{S}$, as desired. □

It remains to give a construction for i-admissible semantics. We let $[\mathbb{S}] = \bigcup_{S \in \mathbb{S}} S$ denote the set of all claims that appear in $\mathbb{S}$.

Construction 7.7. *Given a set $\mathbb{S} \subseteq 2^C$, we define $\mathcal{F}^{u_{ad}}_{\mathbb{S}} = (A, R, cl)$ with*

$$A = \{x_S \mid S \in \mathbb{S}, S \neq \emptyset\} \cup \{x_c, d_c \mid c \in [\mathbb{S}]\},$$
$$R = \{(x_S, x_T) \mid S, T \in \mathbb{S}, S \neq T\} \cup \{(x_S, x_c) \mid S \in \mathbb{S}, c \in [\mathbb{S}] \setminus S\} \cup$$
$$\{(x_c, d_c), (d_c, d_c) \mid c \in [\mathbb{S}]\} \cup \{(d_c, x_S) \mid S \in \mathbb{S}, c \in S\},$$

$cl(x_c) = cl(d_c) = c$ and $cl(x_S) \in S$, i.e., for x_S we pick an arbitrary claim from the set S.

Example 7.8. Consider a claim-set $\mathbb{S} = \{\emptyset, \{a\}, \{a, b\}, \{a, c\}\}$. Following Construction 7.7, we introduce an argument x_S for each claim-set in $S \in \mathbb{S}$. Moreover, we add attacks between all such arguments x_S and x_T, $T \neq S$. Each such argument belongs to the admissible extension that realizes S in the resulting CAF. We moreover introduce two arguments for each of the claims a, b, c that appear in $[\mathbb{S}]$: we add an argument x_c with claim c and a self-attacking argument d_c. The resulting CAF $\mathcal{F}^{u_{ad}}_{\mathbb{S}}$ looks as follows (claims are omitted, arguments that represent claims are filled white):

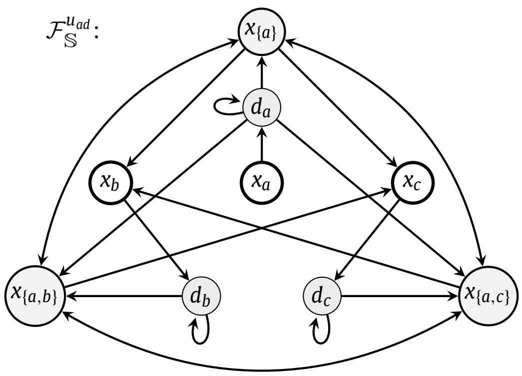

The set $\{x_{\{a,b\}}, x_a, x_b\}$ is admissible in $\mathcal{F}^{u_{ad}}_{\mathbb{S}}$: the argument $x_{\{a,b\}}$ defends the argument x_b against the attacks from the set-arguments $x_{\{a\}}$ and $x_{\{a,c\}}$. Moreover, the arguments x_a and x_b attack d_a and d_b, respectively, and thus defend the argument $x_{\{a,b\}}$. It follows that $\{a, b\}$ is admissible realizable in $\mathcal{F}^{u_{ad}}_{\mathbb{S}}$. It can be checked that $ad_i(\mathcal{F}^{u_{ad}}_{\mathbb{S}}) = \mathbb{S}$.

Proposition 7.9. *Given a set $\mathbb{S} \subseteq 2^C$ such that $\emptyset \in \mathbb{S}$, let $\mathcal{F}^{u_{ad}}_{\mathbb{S}}$ be defined as in Construction 7.7. It holds that $ad_i(\mathcal{F}^{u_{ad}}_{\mathbb{S}}) = \mathbb{S}$.*

Proof. We denote the underlying AF of $\mathcal{F}^{u_{ad}}_{\mathbb{S}}$ by F. First, let us show that each $S \in \mathbb{S}$ is admissible realizable in F. Indeed, the set $E = \{x_S\} \cup \{x_c \mid c \in S\}$ is admissible in F and satisfies $cl(E) = S$: E is conflict-free by construction; moreover, each argument x_c defends x_S against the attack from d_c. Furthermore, x_S attacks all remaining set-arguments. Thus E is admissible in F.

Next, we show that no proper superset of E is admissible in F: as each other set-argument is attacked, it holds that $E \cup \{x_T\}$ is conflicting for each x_T, $T \neq S$. Moreover, each dummy argument d_c is self-attacking. Thus, $E \cup \{d_c\}$ is conflicting for each $c \in [\mathbb{S}]$. Finally, since each claim-argument x_c with $c \notin S$ is attacked by $x_S \in E$, we obtain that no proper superset of E is conflict-free.

It remains to show that no proper subset of E is admissible. First, we observe that $E \setminus \{x_S\}$ is not admissible as it does not defend itself. In case we remove some argument x_c for some $c \in S$, we have that x_S is no longer defended against the attack from d_c. Consequently, we obtain $ad_i(\mathcal{F}^{u_{ad}}_{\mathbb{S}}) = \mathbb{S}$. □

7.2. Expressiveness of well-formed CAFs

Turning now to well-formed CAFs, we have already seen in Sections 4, 5, and 6 that the semantics under considerations admit a different behavior compared to the general case when restricted to this CAF-class. I-maximality is satisfied by h-naive and all variants of stable, preferred, semi-stable, and stage semantics; moreover, i-admissible and i-complete semantics satisfy cautious respectively weak cautious closure, indicating that not all extension-sets are expressible with respect to well-formed CAFs.

Our characterization results for well-formed CAFs can be summarized as follows:

Theorem 7.10. *The following characterizations hold:*

$$\Sigma^{wf}_{gr_i} = \{\mathbb{S} \subseteq 2^C \mid |\mathbb{S}| = 1\}$$
$$\Sigma^{wf}_{cf_i} = \{\mathbb{S} \subseteq 2^C \mid \mathbb{S} \neq \emptyset, \mathbb{S} \text{ is downwards-closed}\}$$
$$\Sigma^{wf}_{ad_i} = \{\mathbb{S} \subseteq 2^C \mid \emptyset \in \mathbb{S}, \mathbb{S} \text{ is cautiously closed}\}$$
$$\Sigma^{wf}_{co_c} = \{\mathbb{S} \subseteq 2^C \mid \mathbb{S} \neq \emptyset, \bigcap_{S\in\mathbb{S}} S \in \mathbb{S}, \mathbb{S} \text{ is weak-cautiously closed and satisfies unique completion}\}$$
$$\Sigma^{wf}_{\rho} = \{\mathbb{S} \subseteq 2^C \mid \mathbb{S} \text{ is I-maximal}\},\ \ \rho \in \{stb_i, cf\text{-}stb_h, ad\text{-}stb_h\}$$
$$\Sigma^{wf}_{\rho} = \Sigma^{wf}_{stb_c} \setminus \{\emptyset\},\ \ \rho \in \{pr_i, pr_h, na_h, ss_i, ss_h, stg_i, stg_h\}$$

Remark 7.11. We remark that signature characterizations for well-formed CAFs for some of the semantics, i.e., for i-conflict-free, h-naive, i-grounded, i-admissible, i-complete, h-semi-stable, h-stage, all variants of stable and preferred semantics, can also be obtained through recent expressiveness results for *AFs with collective attacks (SETAFs)* [51] and their relation to well-formed CAFs: SETAF signature characterizations provided in [52] translate to well-formed CAFs via the semantics-preserving transformation presented in [53,29]. It follows that the signatures for the aforementioned semantics coincide with their SETAF counter-part. However, in order to obtain a well-formed CAF having specific extensions, it is necessary to first construct a SETAF, determine its normal form, and apply the procedure in [53]. However, the so obtained CAF might have exponentially many arguments (in the number of claims). In order to avoid this detour over SETAFs, we will present genuine signature constructions for well-formed CAFs from Theorem 7.10 in the subsequent part of this section. We moreover note that for i-admissible and i-complete semantics, the formulations of the signature characterizations slightly differ: in [52], the distinctive characteristics of admissible and complete semantics are *set-conflict-sensitivity* and *set-com-closure*, respectively. The constructions furthermore show that our formulation in terms of (weak) cautious closure and unique completion are indeed equivalent to the SETAF formulation, thus offering an alternative view on admissible and complete semantics in SETAFs.

As the attentive reader might have noticed, Theorem 7.10 does not speak about i-naive semantics. Indeed, the characterization of the signature for well-formed CAFs for i-naive semantics remains an open problem. We discuss several observations and known (im)possibility-results at the end of this section.

Signatures for i-grounded and i-conflict-free semantics coincide with those for general CAFs using $\Sigma^{AF}_{\sigma} \subseteq \Sigma^{wf}_{\sigma_i} \subseteq \Sigma^{CAF}_{\sigma_i}$ and the coincidence of $\Sigma^{AF}_{\sigma} = \Sigma^{CAF}_{\sigma_i}$ for $\sigma \in \{cf, gr\}$.

I-maximality characterizes h-naive and all variants of stable, preferred, semi-stable, and stage semantics, as we show next. To do so, we consider a construction that has been used already in Section 6 in Example 6.33 to show that tightness and conflict-sensitivity is not satisfied by any of the (non-single-status) semantics under consideration. Now, let us formally introduce the construction:

Construction 7.12. *Given a set* $\mathbb{S} = \{S_1, \dots, S_n\} \subseteq 2^C$, *we define* $\mathcal{F}^{\text{I-max}}_{\mathbb{S}} = (A, R, cl)$ *with*

$$A = \{c_i \mid c \in S_i, 1 \le i \le n\},$$
$$R = \{(c_i, d_j) \mid 1 \le i, j \le n, c \notin S_j\},$$

and $cl(c_i) = c$ *for all* $c_i \in A$.

The construction yield well-formed CAFs as arguments with the same claim attack the same arguments.

Example 7.13. Consider the extension-set $\mathbb{S} = \{\{a, c\}, \{b, c\}, \{a, b, d\}\}$. Applying Construction 7.12 yields the following CAF:

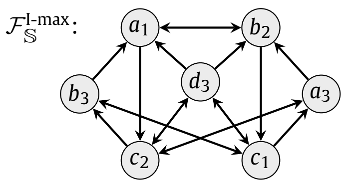

Next we show that each I-maximal non-empty extension-set can be obtained under h-naive and all variants of stable, preferred, semi-stable, and stage semantics when applying Construction 7.12. For the case $\mathbb{S} = \emptyset$, we consider again the CAF that contains a single self-attacking argument only. The following proposition thus proves signature characterizations from Theorem 7.10 for all of the aforementioned semantics.

Proposition 7.14. *Given an I-maximal non-empty extension-set* $\mathbb{S} \subseteq 2^C$*, let* $\mathcal{F}_{\mathbb{S}}^{\text{I-max}}$ *be defined as in Construction 7.12. It holds that* $\rho(\mathcal{F}_{\mathbb{S}}^{\text{I-max}}) = \mathbb{S}$ *for* $\rho \in \{stb_i, cf\text{-}stb_h, ad\text{-}stb_h, pr_i, pr_h, na_h, ss_i, ss_h, stg_i, stg_h\}$.

Proof. Let $\mathbb{S} = \{S_1, \ldots, S_n\}$. For each claim-set $S_i \in \mathbb{S}$ we denote its canonical realization in $\mathcal{F}_{\mathbb{S}}^{\text{I-max}}$ by $E_i = \{c_i \mid c \in S_i\}$. Moreover, we write F to denote the underlying AF of $\mathcal{F}_{\mathbb{S}}^{\text{I-max}}$.

First, we show the statement for the admissible-based semantics. Since $stb(F) \subseteq pr(F)$ holds, it remains to show (1) $\{\{c_i \mid c \in S_i\} \mid S_i \in \mathbb{S}\} \subseteq stb(F)$ and (2) $pr(F) \subseteq \{\{c_i \mid c \in S_i\} \mid S_i \in \mathbb{S}\}$.

(1) By construction, $E_i = \{c_i \mid c \in S_i\}$ is conflict-free in F for each $S_i \in \mathbb{S}$. Moreover, E_i attacks all d_j with $j \neq i$ since S_i and S_j are incomparable, hence there is an $c \in S_i$ which does not occur in S_j. Thus E_i is a stable extension of F.

(2) Consider a preferred set $E \in pr(F)$. We show that E is a subset of $\{c_i \mid c \in S_i\}$ for some $i \leq n$. First, we observe that $cl(E) \subseteq S_i$ for some $S_i \in \mathbb{S}$, otherwise, E is conflicting: if E realizes a claim d that does not occur in S_i then each argument $c_i \in S_i$ is attacked by arguments with claim d by construction. Thus $cl(E) \subseteq S_i$ for some $i \leq n$.
Now, towards a contradiction, assume that there is an argument $c_j \in E$ with $i \neq j$. As S_i and S_j are incomparable there is a claim $d \in S_i \setminus S_j$ that attacks c_j (i.e., each argument with claim d attacks c_j), in particular, the argument d_i attacks c_j. Since $cl(E) \subseteq S_i$, there is no argument in E that attacks d_i, otherwise S_i would be conflicting. Consequently, $E \subseteq E_i$. From (1), we already know that $E_i \in pr(F)$ for each $S_i \in \mathbb{S}$ (since each stable extension is preferred). Hence, by the $\subseteq$-maximality of preferred extensions, it holds that $E = E_i$.

By (1) & (2) we obtain $\mathbb{S} \subseteq stb_i(\mathcal{F}_{\mathbb{S}}^{\text{I-max}}) \subseteq pr_i(\mathcal{F}_{\mathbb{S}}^{\text{I-max}}) \subseteq \mathbb{S}$, thus

$$stb_i(\mathcal{F}_{\mathbb{S}}^{\text{I-max}}) = ss_i(\mathcal{F}_{\mathbb{S}}^{\text{I-max}}) = ss_h(\mathcal{F}_{\mathbb{S}}^{\text{I-max}}) = pr_i(\mathcal{F}_{\mathbb{S}}^{\text{I-max}}) = \mathbb{S}.$$

Recall that in well-formed CAFs, all variants of stable semantics coincide. Likewise, all variants of preferred semantics yield the same outcome.

Next, we show that (3) $na_h(\mathcal{F}_{\mathbb{S}}^{\text{I-max}}) \subseteq \mathbb{S}$. First, we observe that each $S_i \in \mathbb{S}$ is cf_i-realizable via E_i. Second, there is no $E \subseteq A$ with $cl(E) \supset S_i$: as already observed in (2), there is no set of arguments $E \subseteq A$ with $cl(E) \supset S_i$ that is conflict-free in F.

By (1) & (3) we obtain $\mathbb{S} \subseteq stb_i(\mathcal{F}_{\mathbb{S}}^{\text{I-max}}) \subseteq na_h(\mathcal{F}_{\mathbb{S}}^{\text{I-max}}) \subseteq \mathbb{S}$. Thus,

$$stb_i(\mathcal{F}_{\mathbb{S}}^{\text{I-max}}) = stg_i(\mathcal{F}_{\mathbb{S}}^{\text{I-max}}) = stg_h(\mathcal{F}_{\mathbb{S}}^{\text{I-max}}) = na_h(\mathcal{F}_{\mathbb{S}}^{\text{I-max}}) = \mathbb{S}.$$

This concludes the proof of the proposition. □

It remains to provide proofs for the signature characterizations for i-admissible and i-complete semantics for well-formed CAFs. We show that the signature for i-admissible semantics is characterized by cautious closure and empty-set-acceptance. Moreover, we show that i-complete semantics can express each extension-set $\mathbb{S}$ that is weakly cautiously closed, satisfies unique completion and contains $\bigcap_{S\in\mathbb{S}} \mathbb{S}$.

We start by introducing a construction that will serve as basis to express extension-sets under i-admissible and i-complete semantics. For this, it will be convenient to introduce a function $\mathtt{min}_{\mathbb{S}}(c)$ that returns, for a given extension-set $\mathbb{S}$ and a claim $c \in [\mathbb{S}]$, the $\subseteq$-minimal sets in $\mathbb{S}$ that contain c.

Definition 7.15. For an extension-set $\mathbb{S} \subseteq 2^C$ and claim $c \in [\mathbb{S}]$, we define $\mathtt{min}_{\mathbb{S}}(c) = \{M \in \mathbb{S} \mid c \in M, \nexists S \in \mathbb{S}(S \subset M \wedge c \in S)\}$.

For I-maximal extension-sets, the function $\mathtt{min}_{\mathbb{S}}(c)$ will return all sets in extension-set $\mathbb{S}$ that contain the claim $c \in [\mathbb{S}]$. Indeed, if $\mathbb{S} \setminus \{\emptyset\}$ is incomparable, then $\mathtt{min}_{\mathbb{S}}(c) = \{M \in \mathbb{S} | c \in M\}$ for each $M \in \mathbb{S}$.

Example 7.16. Consider the extension-set $\mathbb{S} = \{\emptyset, \{a, c\}, \{b, c\}, \{c\}, \{a, b, d\}\}$. The $\subseteq$-minimal sets relative to claims in $[\mathbb{S}]$ are given by

$$\mathtt{min}_{\mathbb{S}}(a) = \{\{a, c\}, \{a, b, d\}\} \qquad \mathtt{min}_{\mathbb{S}}(b) = \{\{b, c\}, \{a, b, d\}\}$$
$$\mathtt{min}_{\mathbb{S}}(c) = \{\{c\}\} \qquad \mathtt{min}_{\mathbb{S}}(d) = \{\{a, b, d\}\}$$

Now, consider the I-maximal extension-set $\mathbb{S}' = \mathbb{S} \setminus \{\emptyset, \{c\}\}$. We obtain

$$\mathtt{min}_{\mathbb{S}'}(a) = \{\{a, c\}, \{a, b, d\}\} \qquad \mathtt{min}_{\mathbb{S}'}(b) = \{\{b, c\}, \{a, b, d\}\}$$
$$\mathtt{min}_{\mathbb{S}'}(c) = \{\{a, c\}, \{b, c\}\} \qquad \mathtt{min}_{\mathbb{S}'}(d) = \{\{a, b, d\}\}$$

We are ready to present our construction that will serve as basis to characterize i-admissible and i-complete semantics.

Construction 7.17. *Given an extension-set $\mathbb{S} \subseteq 2^C$, we define $\mathcal{F}_\mathbb{S} = (A, R, cl)$ with*

$$A = \{c_M \mid c \in [\mathbb{S}], M \in \mathtt{min}_\mathbb{S}(c)\},$$
$$R = \{(c_M, c'_{M'}) \mid c_M, c'_{M'} \in A, c \notin \mathtt{up}_\mathbb{S}(M')\},$$

and $cl(c_M) = c$ for all $c_M \in A$.

$\mathcal{F}_\mathbb{S}$ is well-formed since each attack depends on the claim of the attacking argument. Moreover, in case $\mathbb{S} \setminus \{\emptyset\}$ is incomparable, we have $\mathtt{min}_\mathbb{S}(c) = \{M \in \mathbb{S} | c \in M\}$ and $\mathtt{up}_\mathbb{S}(M) = M$ for each $M \in \mathbb{S}$. Hence, $\mathcal{F}_\mathbb{S}$ can be written as

$$A = \{c_S \mid S \in \mathbb{S}, c \in S\},$$
$$R = \{(c_S, c'_{S'}) \mid c_S, c'_{S'} \in A, c \notin S'\},$$

with cl as defined above. Note that this construction corresponds to the CAF $\mathcal{F}^{\text{I-max}}_\mathbb{S}$ from Construction 7.3. In this way, $\mathcal{F}_\mathbb{S}$ generalizes $\mathcal{F}^{\text{I-max}}_\mathbb{S}$ which extends to extension-sets that are not I-maximal.

Example 7.18. Consider the extension-sets $\mathbb{S} = \{\emptyset, \{a, c\}, \{b, c\}, \{c\}, \{a, b, d\}\}$ and $\mathbb{S}' = \mathbb{S} \setminus \{\emptyset, \{c\}\}$ from Example 7.16. We note that both $\mathbb{S}$ and $\mathbb{S}'$ are cautiously closed. Construction 7.17 yields the following CAFs:

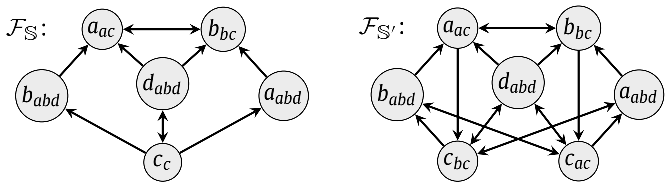

Note that $\mathcal{F}_{\mathbb{S}'}$ corresponds to the CAF from Example 7.13. We observe that there is only one single argument c_c in $\mathcal{F}_\mathbb{S}$ with claim c while $\mathcal{F}_{\mathbb{S}'}$ yields two arguments c_{bc} and c_{ac} with claim c.

Attacks of $\mathcal{F}_\mathbb{S}$ and $\mathcal{F}_{\mathbb{S}'}$ are constructed as follows: For each minimal set M that induces an argument c_M, c_M is attacked by all claims that are *not* contained in $\mathtt{up}_\mathbb{S}(M)$. For $M = \{a, c\}$, we have $\mathtt{up}_\mathbb{S}(\{a, c\}) = \{a, c\}$ as there are no proper supersets of $\{a, c\}$. It follows that the argument a_{ac} is attacked by all arguments having claim b or d. The set $\{c\}$ on the other hand, is contained in all non-empty sets of $\mathbb{S}$ except $\{a, b, d\}$, yielding $\mathtt{up}_\mathbb{S}(\{c\}) = \{a, b, c\}$; consequently, c_c is attacked only by the unique argument d_{abd} having claim d.

We show that each set $S \in \mathbb{S}$ is i-admissible in $\mathcal{F}_\mathbb{S}$ in case $\mathbb{S}$ is weakly cautiously closed and contains $\emptyset$.

Proposition 7.19. *Given a set $\mathbb{S} \subseteq 2^C$ that is weakly cautiously closed and contain $\emptyset$, and let $\mathcal{F}_\mathbb{S}$ be defined as in Construction 7.17. Then $\mathbb{S} \subseteq ad_i(\mathcal{F}_\mathbb{S})$.*

Proof. Let $S \in \mathbb{S}$, and let $E = \{c_M \in A | M \subseteq S\}$. Clearly, $cl(E) = S$; moreover, E is conflict-free since $c \in \mathtt{up}_\mathbb{S}(M')$ for each $c_M, c'_{M'} \in E$ using $M' \subseteq S \subseteq \mathtt{up}_\mathbb{S}(M')$. It remains to show that S defends itself. Let c_N denote an argument with claim c that attacks E. We proceed by case distinction: (i) $S \subseteq \mathtt{up}_\mathbb{S}(N)$ and (ii) $S \nsubseteq \mathtt{up}_\mathbb{S}(N)$.

(i) In case $S \subseteq \mathtt{up}_\mathbb{S}(N)$, there is $T \in \mathbb{S}$ such that $N \cup S \subseteq T$ since $\mathbb{S}$ is weakly cautiously closed. Thus we obtain a contradiction to c_N attacks E by construction of $\mathcal{F}_\mathbb{S}$.

(ii) In case $S \nsubseteq \mathtt{up}_\mathbb{S}(N)$, there is some $d \in S$ such that $d \notin T$ for all upper sets $T \supseteq N$ of N in $\mathbb{S}$, i.e., $d \notin \mathtt{up}_\mathbb{S}(N)$. Thus, by construction of $\mathcal{F}_\mathbb{S}$, all arguments with claim d attack c_N. It remains to show that E contains an argument with claim d. Again, by construction of $\mathcal{F}_\mathbb{S}$, each claim in S appears as claim of some subset S' of S. Thus, there is an argument $d_{S'}$, $d \in S'$ for some $S' \subseteq S$, with claim d that attacks c_N. □

Since the cautious closure principle is stricter and therefore a special case of the weak cautious closure principle, the statement also holds true if $\mathbb{S}$ is cautiously closed. The other direction does not hold as the CAF $\mathcal{F}_\mathbb{S}$ in Example 7.18 demonstrates: Here, the argument d_{abd} defends itself. Therefore, $\{d\}$ is i-admissible in $\mathcal{F}_\mathbb{S}$ although $\{d\} \notin \mathbb{S}$.

Next we show a property of $\mathcal{F}_\mathbb{S}$ that is crucial towards expressing suitable extension-sets under i-complete semantics: if $\mathbb{S}$ is weakly cautiously closed, then each admissible set E in $F_\mathbb{S}$ satisfies $\bigcup_{c_M \in E} M \subseteq S$ for some $S \in \mathbb{S}$.

Proposition 7.20. *Given a weakly cautiously closed extension-set $\mathbb{S} \subseteq 2^C$, then for all $E \in ad(F_\mathbb{S})$, there is $S \in \mathbb{S}$ such that $\bigcup_{c_M \in E} M \subseteq S$.*

Proof. Consider some $E \in ad(F_{\mathbb{S}})$. Then $cl(E) \subseteq \mathrm{up}_{\mathbb{S}}(M)$ for each $M \in \mathbb{S}$ with $c_M \in E$, otherwise there is $d \in cl(E)$ that attacks c_M, which contradicts conflict-freeness of E.

We show that for all arguments $c_M \in E$, for each claim $d \in M$, it holds that d does not attack E. Consider an argument $c_M \in E$. We proceed by case distinction: (i) $M \subseteq cl(E)$ and (ii) $M \not\subseteq cl(E)$.

(i) First assume $M \subseteq cl(E)$. As observed above, $cl(E) \subseteq \mathrm{up}_{\mathbb{S}}(M')$ for each argument $c'_{M'} \in E$, thus $d \in \mathrm{up}_{\mathbb{S}}(M')$ for each $d \in M$ and each argument $c'_{M'} \in E$. By construction of $\mathcal{F}_{\mathbb{S}}$, no $d \in M$ attacks E.
(ii) Now assume $M \not\subseteq cl(E)$. Towards a contradiction, let us assume that there is a claim $d \in M \setminus cl(E)$ that attacks E. That is, there is some argument d_N with claim d that attacks E and $N \subseteq M$ (since $d \in M$, there is $N \subseteq M$ such that N is a $\subseteq$-minimal set containing d in $\mathbb{S}$). Since E defends itself, there is some argument having claim $e \in cl(E)$ satisfying $e \notin \mathrm{up}_{\mathbb{S}}(N)$ (then e attacks d_N by construction of $\mathcal{F}_{\mathbb{S}}$). But then we obtain $e \notin \mathrm{up}_{\mathbb{S}}(N) \subseteq \mathrm{up}_{\mathbb{S}}(M)$, which contradicts $cl(E) \subseteq \mathrm{up}_{\mathbb{S}}(M)$.

We have shown that for all arguments $c_M \in E$, for each claim $d \in M$, it holds that d does not attack E. This means that for every two arguments $c_M, c'_{M'} \in E$, it holds that $M \subseteq \mathrm{up}_{\mathbb{S}}(M')$. By successive application of the weak cautious closure criteria, we obtain that there is $S \in \mathbb{S}$ with $\bigcup_{c_M \in E} M \subseteq S$. □

Moreover, in case $\mathbb{S}$ furthermore satisfies unique completion, then each union of two sets in $\mathbb{S}$ defends all 'missing elements' of its completion-set in $\mathcal{F}_{\mathbb{S}}$.

Proposition 7.21. *Given a weakly cautiously closed extension-set $\mathbb{S} \subseteq 2^C$ that satisfies unique completion, let $S, T \in \mathbb{S}$ and let $\mathcal{F}_{\mathbb{S}}$ be defined as in Construction 7.17. Then $S \cup T$ defends all arguments c_M that satisfy (1) $c \in \mathrm{C}_{\mathbb{S}}(S \cup T) \setminus (S \cup T)$ and (2) $M \subseteq \mathrm{C}_{\mathbb{S}}(S \cup T)$.*

Proof. Given $S, T \in \mathbb{S}$ and consider an argument c_M with $c \in \mathrm{C}_{\mathbb{S}}(S \cup T) \setminus (S \cup T)$ and $M \subseteq \mathrm{C}_{\mathbb{S}}(S \cup T)$, and let $c'_{M'}$ be and attacker of c_M in $\mathcal{F}_{\mathbb{S}}$. Consequently, $c' \notin \mathrm{up}_{\mathbb{S}}(M)$. Now assume c_M is not defended against the attack from $c'_{M'}$ by $S \cup T$. This is the case only if $S \cup T$ is contained in the union of all upper sets of M', i.e., $S \cup T \subseteq \mathrm{up}_{\mathbb{S}}(M')$. Since $\mathbb{S}$ is weakly closed, there is some set $U \in \mathbb{S}$ that contains $S \cup T \cup M'$; by unique completion we may furthermore assume that $\mathrm{C}_{\mathbb{S}}(S \cup T) \subseteq U$. But then we have $c' \in U \subseteq \mathrm{up}_{\mathbb{S}}(M)$, which contradicts our initial assumption $c'_{M'}$ attacks c_M. □

Next we show that each weakly closed extension-set $\mathbb{S}$ that satisfies unique completion and contains $\bigcap \mathbb{S}$ is a superset of $co_i(\mathcal{F})$. A crucial property is that arguments that correspond to the same minimal set (i.e., they possess the same subscript) are attacked by the same arguments.

Proposition 7.22. *Given a set $\mathbb{S} \subseteq 2^C$ that is weakly cautiously closed, satisfies unique completion and contains $\bigcap \mathbb{S}$, and let $\mathcal{F}_{\mathbb{S}}$ be defined as in Construction 7.17. Then $\mathbb{S} \supseteq co_i(\mathcal{F}_{\mathbb{S}})$.*

Proof. Assume there is $S \in co_i(\mathcal{F}_{\mathbb{S}})$ such that $S \notin \mathbb{S}$. Let E be a co-realization of S in $\mathcal{F}_{\mathbb{S}}$, then by Proposition 7.20, there is $T \in \mathbb{S}$ such that $\bigcup_{c_M \in E} M \subseteq T$.

Since E is complete, we have $S = \bigcup_{c_M \in E} M$: Consider some argument $c_M \in E$. By design of $\mathcal{F}_{\mathbb{S}}$, each argument d_M, $d \in M$, possesses the same attacker as c_M. Thus, d_M is defended by E because c_M is defended by E. It is evident that d_M is not attacked by any argument $a \in E$ (otherwise, a attacks c_M); moreover, d_M does not attack any argument $c'_{M'} \in E$ since in this case, E attacks d_M and thus also c_M, which contradicts conflict-freeness of E. By Proposition 7.21, we have that $S = \bigcup_{c_M \in E} M$ contains all arguments $c'_{M'}$ with $c' \in \mathrm{C}_{\mathbb{S}}(S) \setminus S$ and $M' \subseteq \mathrm{C}_{\mathbb{S}}(S)$. Thus, we obtain $S = \mathrm{C}_{\mathbb{S}}(S)$ and thus $S \in \mathbb{S}$. □

Although $\mathcal{F}_{\mathbb{S}}$ possesses characteristics that are necessary for realizing extension-sets under i-admissible and i-complete semantics, we observe that the construction is not sufficient to express all suitable extension-sets under i-admissible or i-complete semantics, respectively:

- $\mathcal{F}_{\mathbb{S}}$ does not realize extension-sets under i-admissible semantics (assuming $\mathbb{S}$ is cautiously closed and contains $\emptyset$): As already mentioned, constructing $\mathcal{F}_{\mathbb{S}}$ might yield additional i-admissible claim-sets that are not contained in $\mathbb{S}$ (cf. Example 7.18, here, $S = \{d_{a,b,d}\}$ is admissible in $\mathcal{F}_{\mathbb{S}}$ but $S \notin \mathbb{S}$).
- $\mathcal{F}_{\mathbb{S}}$ does not realize i-complete extension-sets (assuming $\mathbb{S}$ is weakly cautiously closed, satisfies unique completion, and contains $\bigcap \mathbb{S}$): While $\mathcal{F}_{\mathbb{S}}$ might produce too many extensions for i-admissible semantics, the opposite is the case for i-complete semantics: Let $\mathbb{S} = \{\emptyset, \{a\}, \{b\}, \{a, b, c\}\}$, then $\mathcal{F}_{\mathbb{S}} = (\{a_a, b_b, c_{abc}\}, \emptyset, cl)$ which yields $co_i(\mathcal{F}_{\mathbb{S}}) = \{\{a, b, c\}\}$. Thus for i-complete extensions, the challenge lies in differentiating all i-complete subsets.

First, we extend $\mathcal{F}_{\mathbb{S}}$ to capture i-admissible claim-sets.

Construction 7.23. *For set $\mathbb{S} \subseteq 2^C$, let $\mathcal{F}_\mathbb{S} = (A, R, cl)$ be defined as in Construction 7.17. We define $\mathcal{F}^{ad}_\mathbb{S} = (A^{ad}, R^{ad}, cl^{ad})$ with*

$$A^{ad} = A \cup \{x^d_{c_M} | c_M \in A, d \in M\},$$

$$R^{ad} = R \cup \{(d_{M'}, x^d_{c_M}), (x^d_{c_M}, x^d_{c_M}), (x^d_{c_M}, c_M) | c_M \in A, d \in M\},$$

and $cl^{ad}(c_M) = cl(c_M) = c$ for all $c \in [\mathbb{S}]$ and $cl^{ad}(x^d_{c_M}) = x^d_{c_M}$ otherwise.

Example 7.24. Let $\mathbb{S} = \{\emptyset, \{a\}, \{b, c\}, \{a, b, c\}\}$. First, we construct the corresponding CAF $\mathcal{F}_\mathbb{S}$ that contains no attacks; additionally, we get $|M|$ new (self-attacking) arguments for each $c_M \in A$ that attack c_M and are attacked by each argument having claim $d \in M$. The resulting framework is thus given as follows:

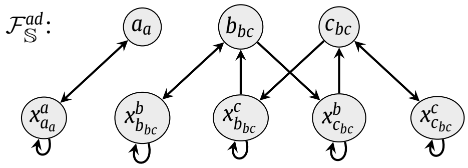

The following lemma will be useful.

Lemma 7.25. *For an extension-set $\mathbb{S}$, let $\mathcal{F}^{ad}_\mathbb{S}$ be defined as in Construction 7.23, and let $E \subseteq A$. Then*

1. *if an argument $c_M \in A$ is defended by E then it holds that $M \subseteq cl(E)$;*
2. *$ad_i(\mathcal{F}^{ad}_\mathbb{S}) \subseteq ad_i(\mathcal{F}_\mathbb{S})$;*
3. *$E \in ad(F^{ad}_\mathbb{S})$ implies $cl(E) = \bigcup_{c_M \in E} M$.*

Proof. (1) follows since only arguments with claim d defend c_M against the attack from $x^d_{c_M}$ for all $d \in M$. To show (2), consider a set $S \in ad_i(\mathcal{F}^{ad}_\mathbb{S})$ and an ad-realization E of S in F. Then E defends itself against all attackers in A, thus $S \in ad_i(\mathcal{F}_\mathbb{S})$.

For (3), let us first observe that each admissible set $E \in ad(\mathcal{F}_\mathbb{S})$ is contained in the union of all minimal sets M that are associated to arguments in E, i.e., $cl(E) \subseteq \bigcup_{c_M \in E} M$. This follows from the fact that $c \in M$ for every argument $c_M \in A$. Moreover, $E \in ad(F^{ad}_\mathbb{S})$ implies $E \in ad(F_\mathbb{S})$ implies that $cl(E) \subseteq \bigcup_{c_M \in E} M$. By (1) we obtain equality since each argument c_M requires $d \in cl(E)$ for all $d \in M$. □

Proposition 7.26. *Let $\mathbb{S}$ be cautiously closed and contain $\emptyset$. Then $\mathbb{S} = ad_i(\mathcal{F}^{ad}_\mathbb{S})$.*

Proof. We first prove that each set $S \in \mathbb{S}$ is indeed i-admissible: First, in case $S = \emptyset$ we are done since the empty set is always admissible. Now, let $S \in \mathbb{S}$ be non-empty. We show that $E = \{c_M \in A | M \subseteq S, c \in S\}$ is an admissible realization of $\mathbb{S}$ in $\mathcal{F}^{ad}_\mathbb{S}$. It is easy to see that $cl(E) = S$. Moreover, E is conflict-free since for every two arguments $c_M, c'_{M'} \in E$, it holds that $c \in \mathrm{up}_\mathbb{S}(M')$ since $M' \subseteq S \subseteq \mathrm{up}_\mathbb{S}(M')$. Moreover, E defends itself: Consider some argument $x \in A$ that attacks an argument $c_M \in E$. In case x is of the form $x^d_{c_M}$, it holds that E defends itself since $M \subseteq S$. In case x is of the form $c'_{M'}$ for some claim c', we proceed analogous as in the proof of Proposition 7.19 and obtain that E defends itself against each attack.

The other direction is by Proposition 7.20 and by Lemma 7.25: Given an admissible set $E \in ad(F^{ad}_\mathbb{S})$ we have $cl(E) = \bigcup_{c_M \in E} M$. By Proposition 7.20, there is some $S \in \mathbb{S}$ that contains $cl(E)$; since $\mathbb{S}$ is cautiously closed, we obtain that $cl(E) \in \mathbb{S}$ since S serves as witness for $M \in \mathrm{up}_\mathbb{S}(M')$ for every sets $M, M' \in \mathbb{S}$ that are associated to arguments in $c_M, c'_{M'} \in E$. □

Next we show that $co_i(\mathcal{F}^{ad}_\mathbb{S}) = ad_i(\mathcal{F}^{ad}_\mathbb{S})$ iff $\mathbb{S}$ is cautiously closed and contains $\emptyset$.

Proposition 7.27. *Let $\mathbb{S}$ be cautiously closed and contain $\emptyset$. Then $\mathbb{S} = co_i(\mathcal{F}^{ad}_\mathbb{S})$.*

Proof. We have shown in Lemma 7.25 that each i-admissible set $S \in \mathcal{F}^{ad}_\mathbb{S}$ is realized by $E = \{c_M \in A \mid M \subseteq S\}$. In case E defends some argument $c_M \notin E$, we have $M \not\subseteq S$, that is, there is some argument $x^d_{c_M}$ that attacks c_M and is defended by $d \in M \setminus S$ but not by S. Thus the statement follows. □

In case $\mathbb{S}$ is *weakly* cautiously closed we observe that $\mathbb{S} \neq co_i(\mathcal{F}^{ad}_\mathbb{S})$. Note that the empty set is complete in $\mathcal{F}^{ad}_\mathbb{S}$ since each argument has an attacker. Moreover, in case the minimal completion set of $S \cup T$ contains additional arguments for two sets $S, T \in \mathbb{S}$, i.e., in case $\mathrm{C}_\mathbb{S}(S \cup T) \notin \{\emptyset, S \cup T\}$, we have that $S \cup T$ is also i-complete in $\mathcal{F}^{ad}_\mathbb{S}$.

In order to deal with this issue, we adapt a concept from [23]. We use defense formulas to determine which arguments are needed to defend a given claim c.

Definition 7.28. Given an extension-set $\mathbb{S} \subseteq 2^C$ and a claim $c \in [\mathbb{S}]$, we let $def_{\mathbb{S}}(c) = \{S \cup T \mid S, T \in \mathbb{S}, c \in \mathbb{C}_{\mathbb{S}}(S \cup T) \setminus (S \cup T)\}$. The *DNF defense formula* of c is defined as $\mathcal{D}^c_{\mathbb{S}} = \bigvee_{S \in def_{\mathbb{S}}(c)} \bigwedge_{d \in S} d$.

Example 7.29. We consider a set $\mathbb{S} = \{\{a\}, \{a, c\}, \{a, b\}, \{a, b, c, d\}\}$. $\mathbb{S}$ is weakly cautiously closed, moreover, $\bigcap_{S \in \mathbb{S}} S = \{a\}$ is contained in $\mathbb{S}$. We obtain $def_{\mathbb{S}}(a) = def_{\mathbb{S}}(b) = def_{\mathbb{S}}(c) = \emptyset$ and $def_{\mathbb{S}}(d) = \{\{a, b, c\}\}$. For a, b, and c, the corresponding DNF formula corresponds to $\perp$; for d, we have $\mathcal{D}^d_{\mathbb{S}} = (a \wedge b \wedge c)$.

We are ready to present the construction for i-complete semantics.

Construction 7.30. *Given a set $\mathbb{S} \subseteq 2^C$ and let $\mathcal{F}_{\mathbb{S}} = (A, R, cl)$ be defined as in Construction 7.17. For every argument $c_M \in A$, we consider the extended DNF defense formula $\mathcal{D}^c_{\mathbb{S}} \vee \bigwedge_{d \in M} d$ and denote by $\mathcal{CD}^{c_M}_{\mathbb{S}}$ the corresponding CNF formula. We define $\mathcal{F}^{co}_{\mathbb{S}} = (A^{co}, R^{co}, cl^{co})$ as follows*

$$A^{co} = A \cup \{x^{\gamma}_{c_M} | c_M \in A, M \neq \bigcap_{S \in \mathbb{S}} \mathbb{S}, \gamma \in \mathcal{CD}^{c_M}_{\mathbb{S}}\},$$

$$R^{co} = R \cup \{(d_{M'}, x^{\gamma}_{c_M}), (x^{\gamma}_{c_M}, x^{\gamma}_{c_M}), (x^{\gamma}_{c_M}, c_M) | c_M \in A, d \in \gamma\},$$

and $cl^{co}(c_M) = cl(c_M) = c$ for all $c \in [\mathbb{S}]$ and $cl^{co}(x^{\gamma}_{c_M}) = x^{\gamma}_{c_M}$ otherwise.

Observe that the i-grounded extension is realized by arguments that are unattacked in case it is non-empty: auxiliary arguments for an argument C_M are only constructed in case $M \neq \bigcap_{S \in \mathbb{S}} \mathbb{S}$. In every other case, c_M is attacked by argument(s) $x^{\gamma}_{c_M}$ determined by the extended attack formula. Let us consider an example.

Example 7.31. Let us consider the set $\mathbb{S} = \{\{a\}, \{a, c\}, \{a, b\}, \{a, b, c, d\}\}$ from Example 7.29. First, when constructing $\mathcal{F}_{\mathbb{S}}$, we generate four arguments, one for each claim: a_a, b_{ab}, c_{ac}, and d_{abcd}. Observe that none of this arguments are attacking each other.

We proceed by generating the auxiliary arguments: For the claims a, b, and c, the DNF defense formula is empty. The extended DNF defense formula for the arguments a_a, b_{ab}, and c_{ac} thus corresponds to the conjunction of the respective sets in the subscript: $\mathcal{D}^a_{\mathbb{S}} = (a)$, $\mathcal{D}^b_{\mathbb{S}} = (a \wedge b)$, and $\mathcal{D}^c_{\mathbb{S}} = (a \wedge c)$. The corresponding CNF formulae are thus $\{\{a\}\}$, $\{\{a\}, \{b\}\}$, and $\{\{a\}, \{c\}\}$, respectively. For claim d, the DNF defense formula is given by $def_{\mathbb{S}}(d) = \{\{a, b, c\}\}$, thus the extended DNF defense formula corresponding to the argument d_{abcd} is given by $\mathcal{D}^d_{\mathbb{S}} \vee \bigwedge_{x \in M} x = (a \wedge b \wedge c) \vee (a \wedge b \wedge c \wedge d)$. Clearly, this formula can be simplified to the single clause $(a \wedge b \wedge c)$. The corresponding CNF is $\mathcal{CD}^{d_{abcd}}_{\mathbb{S}} = \{\{a\}, \{b\}, \{c\}\}$.

We are ready to give the construction. Note that no auxiliary arguments are generated for a_a since $\{a\} = \bigcap_{S \in \mathbb{S}} S$. The resulting CAF is depicted below:

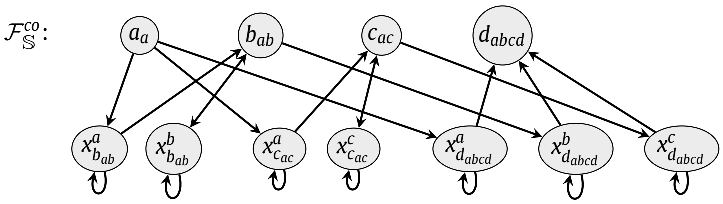

The argument a_a is unattacked and does not defend any other argument. Thus, $\{a\}$ is the i-grounded extension, as desired. It can be checked that the i-complete claim-sets coincide with $\mathbb{S}$ (e.g., a_a and b_{ab} jointly defend the argument b_{ab}).

Observe that the only difference between $\mathcal{F}^{ad}_{\mathbb{S}}$ and $\mathcal{F}^{co}_{\mathbb{S}}$ for the extension-set $\mathbb{S}$ is that $\mathcal{F}^{ad}_{\mathbb{S}}$ would contain an additional self-attacking node $x^d_{d_{abcd}}$ that attacks and is counter-attacked d_{abcd}. In $\mathcal{F}^{ad}_{\mathbb{S}}$, the set $\{a_a, b_{ab}, c_{ac}\}$ does therefore not defend d_{abcd}, consequently, $\{a, b, c\}$ is i-complete in $\mathcal{F}^{ad}_{\mathbb{S}}$. In $\mathcal{F}^{co}_{\mathbb{S}}$, on the other hand, d_{abcd} is defended by $\{a_a, b_{ab}, c_{ac}\}$ in $\mathcal{F}^{co}_{\mathbb{S}}$ and we obtain $co_i(\mathcal{F}^{co}_{\mathbb{S}}) = \mathbb{S}$.

In case $\mathbb{S}$ is cautiously closed and $\bigcap \mathbb{S} = \emptyset$, the construction yields a CAF identical to $\mathcal{F}^{ad}_{\mathbb{S}}$. In this sense the construction refines Construction 7.23. We note that we lose a useful property of $\mathcal{F}^{ad}_{\mathbb{S}}$: While in $\mathcal{F}^{ad}_{\mathbb{S}}$, each i-complete set S is realized by $\{c_M \mid M \subseteq S\}$, the extended construction might cause the defense of additional arguments c_M such that $M \not\subseteq S$. By Lemma 7.20, this affects only arguments c_M such that $c \in S$ and $M \cup S$ possesses a completion-set in $\mathbb{S}$ (all other arguments c_M with claim $c \in S$ are attacked by some arguments in $\{c_M \mid M \subseteq S\}$).

We are ready to show our last characterization result.

Proposition 7.32. *Let $\mathbb{S}$ be weakly cautiously closed, satisfy unique completion and contain $\bigcap \mathbb{S}$. Then $\mathbb{S} = co_i(\mathcal{F}^{co}_{\mathbb{S}})$.*

Proof. Consider a set $S \subseteq \mathbb{S}$ and let $E' = \{c_M \in A \mid M \subseteq S\}$; moreover, let $E = E' \cup \{c_M \in A \mid c \in S, \mathrm{c}_{\mathbb{S}}(S \cup M) = 1, \exists T, U \subseteq S : c \in \mathrm{c}_{\mathbb{S}}(T \cup U) \setminus (T \cup U)\}$. Observe that E is conflict-free since for every two arguments $c_M, c'_{M'} \in E$ we have $c \in S \subseteq \mathrm{up}_{\mathbb{S}}(M')$ (in case $M' \notin S$ we have $\mathrm{c}_{\mathbb{S}}(S \cup M') = 1$ thus the statement holds also in this case).

Next we show that E defends itself: Consider some argument $x \in A$ that attacks an argument $c_M \in E$. The case x is of the form $c'_{M'}$ for some $c' \in [\mathbb{S}]$ is analogous to the case distinction in the proof of Proposition 7.19. In case x is of the form $x^{\gamma}_{c_M}$ and $M \subseteq S$, E defends itself since $\gamma \cap M \neq \emptyset$. In case $M \not\subseteq S$, there are $T, U \subseteq S$ with $c \in \mathrm{c}_{\mathbb{S}}(T \cup U) \setminus (T \cup U)$; by construction of $\mathcal{F}^{co}_{\mathbb{S}}$, $T \cup U \in \mathrm{def}^c_{\mathbb{S}}$, we thus obtain $\gamma \cap (T \cup U) \neq \emptyset$. We obtain that E defends itself against all attacker.

Moreover, E contains all arguments it defends: Assume there is an argument $c_M \in A$ that is not contained in E but defended by E. We show that there is $\gamma \in \mathcal{CD}^{c_M}_{\mathbb{S}}$ such that $\gamma \cap S = \emptyset$. It suffices to show that for all $T \in \mathrm{def}_{\mathbb{S}}(c)$, there is $d \in T$ such that $d \notin S$ (we note that by definition of E, we have $M \not\subseteq S$; thus, there is a claim $d \in M \setminus S$).

First note that in case $c \in S$ and there is $T \in \mathrm{def}_{\mathbb{S}}(c)$ with $T \subseteq S$ we have $c_M \in E$: By assumption c_M is defended by E we have (1) E does not attack c_M. Thus, we obtain $S \subseteq \mathrm{up}_{\mathbb{S}}(M)$ and therefore $\mathrm{c}_{\mathbb{S}}(S \cup M) = 1$ is satisfied; and (2) there are sets $A, B \subseteq S$ with $T = A \cup B$ that defend c.

In case $c \in S$ and there is no $T \in \mathrm{def}_{\mathbb{S}}(c)$ with $T \subseteq S$ we are done: In this case, there is $\gamma \in \mathcal{CD}^{c_M}_{\mathbb{S}}$ such that $\gamma \cap S = \emptyset$ and thus c_M is not defended against the attack $x^{\gamma}_{c_M}$.

Let us now consider the case $c \notin S$. In case there is no $T \in \mathrm{def}_{\mathbb{S}}(c)$ with $T \subseteq S$ we are done: In this case, there is $\gamma \in \mathcal{CD}^{c_M}_{\mathbb{S}}$ such that $\gamma \cap S = \emptyset$ and thus c_M is not defended against the attack $x^{\gamma}_{c_M}$.

In case $c \notin S$ and there is $T \in \mathrm{def}_{\mathbb{S}}(c)$ with $T \subseteq S$. Thus there are sets $A, B \subseteq S$ with $T = A \cup B$ that defend c. Consequently, $\mathrm{c}_{\mathbb{S}}(A \cup B) \not\subseteq S$, which contradicts unique completion.

For the other direction, consider a set $E \in co(F^{co}_{\mathbb{S}})$. We show that $cl(E) \in \mathbb{S}$. In case $E = \emptyset$, there is no argument in E that is unattacked. By construction of $\mathcal{F}^{co}_{\mathbb{S}}$, this is the case only if $\bigcap_{S \in \mathbb{S}} S = \emptyset$, i.e., if $\emptyset \in \mathbb{S}$.

Now assume $E \neq \emptyset$. It holds that E contains all arguments c_M with $M \subseteq cl(E)$ since each such argument is defended by M. Thus there is some $S \in \mathbb{S}$ such that $cl(E) \subseteq S$ by Lemma 7.20. Now assume $cl(E) \notin \mathbb{S}$. In this case, $T = \mathrm{c}_{\mathbb{S}}(\bigcup_{c_M \in E} M)$ is a proper superset of $cl(E)$. Observe that E is not constructed from a single $\subseteq$-minimal set M, i.e., E contains arguments c_M, $c'_{M'}$ with $M \neq M'$ (since no proper subset of such a set M is i-complete). Now, by design of $\mathcal{F}^{co}_{\mathbb{S}}$, there are sets $U, V \in \mathbb{S}$ with $U, V \subseteq cl(E)$ and there is $c \in T \setminus cl(E)$ such that $U \cup V$ defend all arguments with claim c against the attacks of arguments of the form $x^{\gamma}_{c_M}$ for an arbitrary $\subseteq$-minimal set $M \subseteq T$ containing c. Now, let $M \subseteq \mathrm{c}_{\mathbb{S}}(U \cup V)$ be a $\subseteq$-minimal set in $\mathbb{S}$ that contains c. Then c_M is defended by $U \cup V$ against attacks from arguments in A by Proposition 7.21 (since $c \in \mathrm{c}_{\mathbb{S}}(U \cup V) \setminus (U \cup V)$ and $M \subseteq \mathrm{c}_{\mathbb{S}}(U \cup V)$ is satisfied).

Consequently, E defends c_M against all attacks, moreover, $E \cup \{c_M\}$ is conflict-free since $M \subseteq \mathrm{c}_{\mathbb{S}}(U \cup V)$. Thus, E is not complete in $F^{co}_{\mathbb{S}}$, which contradicts our assumption. □

Inherited naive semantics Naive semantics are often perceived as the conflict-free counter-part of preferred semantics as they have many common characteristics. It is thus surprising that the semantics admit several differences when considered with respect to the claims of the arguments. The variants of naive semantics differ even on well-formed CAFs while preferred semantics suggest that maximization on argument-level and maximization on claim-level coincide in this case (recall that both variants of preferred semantics coincide on well-formed CAFs).

We recall that i-naive semantics does not satisfy I-maximality, not even on well-formed CAFs (cf. Example 4.2). On the other hand, it is not possible to express all i-maximal extension-sets, as we show next. Let us first observe that, for each well-formed CAF $\mathcal{F}$, the set of all (non-self-attacking) occurrences of a claim c is contained in some naive extension in the underlying AF F.

Proposition 7.33. *Let $\mathcal{F}$ be a well-formed CAF. Then, for each $c \in \bigcup_{S \in na_i(\mathcal{F})} S$ there is an extension $E \in na(F)$ such that all (non-self-attacking) $a \in A$ with $cl(a) = c$ are contained in E.*

Proof. As $c \in \bigcup_{S \in na_i(\mathcal{F})} S$, there is an argument with claim c that is not self-attacking in F. As $\mathcal{F}$ is well-formed, the set $\{a \in A \mid cl(a) = c, (a, a) \notin R\}$ is conflict-free in F and thus contained in some $E \in na(F)$. □

Lemma 7.34. *For well-formed CAFs, the set $\mathbb{S} = \{\{a, b\}, \{a, c\}, \{b, c\}\}$ cannot be realized with inherited na semantics, i.e. $\mathbb{S} \notin \Sigma^{wf}_{na_i}$.*

Proof. Towards a contradiction assume there is a CAF $\mathcal{F}$ with $na_i(\mathcal{F}) = \mathbb{S}$. By Proposition 7.33 there are sets $E_a, E_b, E_c \in na(\mathcal{F})$ containing all arguments with claim a, b, and c respectively. Let us first assume that all three sets E_a, E_b, E_c are different and have different claim sets, i.e. $cl(E_a)$, $cl(E_b)$, $cl(E_c)$ are mutually distinct. w.l.o.g. we can assume that $cl(E_a) = \{a, b\}$, $cl(E_b) = \{b, c\}$ and $cl(E_c) = \{a, c\}$. That is, (a) there is an argument $b_i \in E_a$ that is not in conflict with any argument with claim a; (b) there is $c_j \in E_b$ that is not in conflict with any argument with claim b; and (c) there is $a_k \in E_c$ that is not in conflict with any argument with claim c. Now consider the set $\{a_k, b_i\}$ which is conflict-free by (a). As $\{a, b, c\} \notin \mathbb{S}$ the set $\{a_k, b_i\}$ has a conflict with c_j. By (c) the conflict has to be between b_i and c_j. However, from (b) we have that c_j is not in conflict with b_i. That is, $\{a_k, b_i, c_j\} \in cf(\mathcal{F})$ and thus $\{a, b, c\} \in na_i(\mathcal{F})$, a contradiction to $na_i(\mathcal{F}) = \mathbb{S}$.

The remaining cases, i.e. (i) E_a, E_b, E_c are different but two of the sets have the same claim-set, and (ii) at least two of the sets E_a, E_b, E_c coincide, can be shown to lead to a contradiction by similar arguments. □

Although i-naive semantics are not I-maximal, it is not possible to express all extension-sets under naive semantics, in particular, it is not possible to express each I-maximal extension-set. This shows that the signatures of i-naive and h-naive semantics are incomparable. As summarized in Table 6, i-naive semantics satisfy none of the known principles for AF or CAF semantics. The precise characterization of naive semantics remains an open problem.

8. Discussion and conclusion

We conclude by summarizing our findings as well as by discussing related and future work directions.

8.1. Summary

In this work, we thoroughly investigated argumentation semantics in the realm of claim-based reasoning. Our study includes the adaption of classical concepts of abstract argumentation to claim-based semantics on the one hand, and a principle-based analysis—complemented by expressiveness results—on the other hand.

We proposed novel semantics for CAFs by shifting classical concepts of abstract argumentation semantics to claim-level. We focused on *claim-set maximization* of conflict-free, admissible, and range-based semantics, yielding novel variants of *naive*, *preferred*, *stage*, and *semi-stable semantics*. Range-based semantics in the realm of claim-based reasoning naturally require a concept of *claim-defeat* that furthermore gave rise to two different versions of *stable semantics*. We settled the relation between the semantics in Sections 4 and 5. We showed that for well-formed CAFs, stable and preferred variants coincide, while naive, stage, and semi-stable variants differ. The latter highlights the fundamental difference between maximization on claim- and on argument-level, in particular, for range-based semantics. Thus, hybrid semantics give an alternative view on claim justification in the spirit of abstract argumentation semantics. They furthermore constitute an argumentation-based formalization of conclusion-focused knowledge representation formalisms such as logic programs (cf. Section 2 and 4.4). By doing so, we deepen the close connection of logic programming semantics and argumentation semantics; in particular, we succeed to capture L-stable semantics with h-semi-stable semantics which is—under standard instantiation methods—impossible for AFs without claims. Thus, our hybrid semantics incorporate evaluation methods which are common to conclusion-based knowledge representation formalisms and add a novel perspective to argumentation semantics by putting the focus on claim acceptance (via claim-set maximization and claim-defeat). With this, we broaden the argumentation semantics landscape and increase the flexibility of the abstract model. Our hybrid semantics make the abstract model accessible to a broader range of conclusion-oriented knowledge representation formalisms. Instantiations of default logic, assumption-based argumentation, logic programs, and other conclusion-based formalisms can now benefit from a semantics-preserving graph-based representation which is user-friendly and easy to understand. Moreover, our hybrid semantics give rise to an alternative evaluation method for ASPIC$^+$.

Our principle-based analysis includes a wide range of genuine and fundamental principles for claim-based reasoning on the one hand as well as the adaption of many well-investigated principles lifted to claim-level on the other hand. Our results show that well-formed CAFs retain many of the desired properties like (CF-)reinstatement and I-maximality on claim-level. Set-theoretical principles like conflict-sensitivity and tightness are however violated, which already indicates the higher expressiveness of (well-formed) CAFs when compared to AFs. Our findings moreover reveal that the behavior of claim-based semantics with respect to general CAFs is more difficult to capture by means of existing principles; in particular, inherited semantics successfully withstand traditional analysis methods. Exceptions are those principles that require the existence of a set of arguments with specific properties (e.g., the defense principle which requires that a set of claims has a realization that defends itself); notable is also the justified rejection principle which is satisfied by stable and conflict-free-based semantics also in the general case. The difficulty indicates that the 'right' principles that characterize the behavior of some of the inherited semantics when considered with respect to general CAFs have yet to be found; we consider this as an important point on our future agenda.

Finally, our expressiveness study (in terms of signature characterization) confirms that claim-based semantics are more expressive than their AF counterpart, already when restricted to the class of well-formed CAFs. In general CAFs, the restrictions are marginal. Indeed, almost each extension-set can be expressed by most of the semantics apart from h-preferred and h-naive which are constrained by I-maximality. This property also characterizes many semantics in well-formed CAFs. We have furthermore identified generalizing properties (i.e., (weak) cautious closure and unique completion) that are characteristic for i-admissible and i-complete semantics, respectively, and presented constructions to realize extension-sets that satisfy to these properties. By doing so, we provide explicit algorithms to construct a (well-formed) CAF that models a desired situation. Moreover, our signature results can prove useful when considering changes in argumentation frameworks or their underlying knowledge bases following certain constraints since expressiveness characterizations are the basis for certain (im-)possibility results regarding changes of the extensions (cf. [24]).

8.2. Related work

Claims in knowledge representation formalisms The evaluation of argumentative scenarios in terms of claims has been discussed in several different contexts. The conclusion-based view appears frequently in structured argumentation formalisms and instantiations of knowledge bases. Apart from our running example of instantiations based on logic programs in which conclusions are key to connect LP semantics with argumentation semantics [7], we mention the conclusion-based view of assumption-based argumentation [11, Definition 2.10] and ASPIC$^+$ [10, Definition 2.18] as two prominent examples. In that aspect, let us furthermore note that in Dung's seminal paper [5], the extraction of the claims of the arguments is a crucial step to embed the argumentative view in the context of logic programs and instances of default logic. Moreover, postulates of logic-based argumentation and deductive argumentation consider the conclusion-based outcome of an argumentation framework in terms of consistency and closure postulates [8,54,15,55,56].

Claims in (semi-)abstract argumentation In recent years, the focus on claims has furthermore reached the realm of abstract argumentation. Several abstract (or semi-abstract) models of argumentation feature claims as part of the abstract level. Apart from CAFs, we mention the labeling-based model of Baroni, Governatori, and Riveret [57,18] that allows for fine-grained evaluation of statement justification. Fermüller and Corsi considered in [58] the connection of logic and abstract argumentation via so-called attack principles. They considered a framework that is similar to CAFs but with a logical formula as claim instead of an atomic entity. We furthermore mention recent work on dynamics in argumentation [59,60] in which arguments are identified with tuples (V, p) where V is a set of vulnerabilities, i.e., the *weak points* or *attack points* of the argument while p is the argument's claim.

All these models demonstrate the increased interest into the claim-based view of argumentation. Indeed, claims shape the structure of an instantiated framework as they often determine the outgoing attacks (assuming well-formedness of the attack relation, but also when incorporating preferences as shown in [27]). The advantage of these semi-abstract models is that they make this influence of the internal structure of the arguments explicit. They bridge the gap between structured and abstract argumentation. We note that, in contrast to the present work, all of the aforementioned semi-abstract models rely on classical (abstract) argumentation semantics in their analysis. A close connection to the present work can be found in the multiple stages labeling formalism by Baroni, Governatori, and Riveret. In [57], they identify two different ways to statement justification (under the credulous or skeptical reasoning mode): having settled the acceptance of arguments, one can either

1. identify justified arguments and extract justified claims; or
2. identify accepted claims and extract justified claims.

They show that their system is able to capture the different treatment of floating conclusions (cf. Remark 6.10) in several knowledge representation formalisms. They compare their model to different formalisms and show that ASPIC$^+$ employs approach 1 while assumption-based argumentation follows approach 2 for statement justification. They show, in particular, that the two approaches are incomparable. In [18], they generalize their model to arbitrary stages of labellings. In contrast to the present work, they put the main emphasis on statement justification but keep argument acceptance and claim acceptance as two separate stages. However, we believe that multi-labeling systems are a powerful tool to investigate argument-focused and conclusion-focused evaluation methods on a very general level and deserves further research. Here, one could, for instance, consider the evaluation of a semantics itself as a multi-stage process (e.g., treating maximization of admissible sets as separate step which can be performed on argument- or on claim-level).

By now, CAFs and generalizations thereof have been studied in several different aspects; e.g., in terms of strong equivalence [45], equipped with preferences [27], and to explore the connection between logic and argumentation [58]. Particularly relevant to the present work is the computational complexity analysis of hybrid and inherited claim-based argumentation conducted in [19,20,61]. Moreover, we refer to [62] for ASP-based instantiations of all claim-based semantics considered in this work.[3] We furthermore note that CAFs have been successfully used to transfer L-stable model semantics to probabilistic argumentation: Rocha and Cozman [63] adapt h-semi-stable semantics to probabilistic argumentation and study the connection to probabilistic logic programs. In [64], they generalize CAFs by incorporating a support relation and show that this model one-to-one corresponds to normal LPs.

An axiomatic view on argumentation Principles, postulates and properties of argumentation semantics have been considered in different facets for different (structured and abstract) argumentation formalisms, e.g., [21,22,55,15,56,31,65]. Likewise, expressiveness of argumentation semantics is an important topic that has been considered for different abstract formalisms [23,52]. In contrast to most of the aforementioned works which investigate principles and expressiveness in terms of arguments, our studies focus on claim-based semantics. While there is naturally a close correspondence if not dependence between these two viewpoints the differences are considerable as shown in the present work. We also want to highlight in this regard in particular the work by Amgoud, Caminada, Gorogiannis, and Hunter [15,56,55] which study rationality postulates for logic-based argumentation systems also in terms of the conclusion-based outcome. In contrast to our analysis they focus on consistency and closure properties. In our work, claims are considered abstract in order to investigate structural properties of the claim-based outcome.

[3] Encodings can be found at https://www.dbai.tuwien.ac.at/proj/argumentation/systempage/caf.html.

On the relation between well-formed CAFs and SETAFs We furthermore mention results on expressiveness and principle-based investigations for AFs with collective attacks (SETAFs): As noted in Remark 7.11, well-formed CAFs and SETAFs are closely related [53]. We thus obtain an alternative characterization of the signatures for well-formed CAFs from signature results presented in [52]. In particular, we obtain that the respective properties coincide, i.e., set-conflict-sensitivity coincides with cautious closure and set-com-closure is equivalent to weak cautious closure and unique completion. While set-conflict-sensitivity and set-com-closure are formalized in terms of potential conflicts, our formulations are conflict-independent and yield an alternative view on the SETAF characterizations.

The close relation between well-formed CAFs and SETAFs furthermore reveals interesting parallels between our principle-based analysis for well-formed CAFs and the principle-based analysis of SETAF semantics recently conducted in [65]. Indeed, we obtain similar results regarding the common principles we investigated, i.e., for conflict-freeness, defense, admissibility, (CF-)reinstatement, h-naivety, and I-maximality. Apart from these principles, they put their focus on the investigation of modularization, non-interference principles, and SCC-recursiveness utilizing the so-called reduct [66], while we conducted set-theoretical investigations and considered genuine principles for claim-based reasoning.

To conclude, the close connection between SETAFs and well-formed CAFs can be of mutual interest and should be investigated further. However, we want to point out that so far, the only known translation between these two formalisms might cause an exponential blowup. Therefore, the computational advantage is still under question. Moreover, we want to emphasize that the relation between SETAFs and claim-augmented frameworks holds exclusively for well-formed CAFs. There is no known semantics-preserving translation between SETAFs and unrestricted CAFs, and according to our signature-results, such translations cannot exist for most of the semantics considered in this work.

8.3. Future work

As already expressed above, the principles and properties formulated in this work capture the behavior of the considered claim-based semantics to a different extent; in particular, inherited semantics in unrestricted CAFs lack principles that characterize their distinct behavior. One point on our future agenda is thus to deepen the principle-based analysis on inherited semantics. Moreover, we plan to adapt more classical AF principles to the realm of claim-based reasoning. Although the principle-based investigation we conducted in the present work already collects many of the classical principles that have been considered in the literature, there are a lot of other principles left that are worth studying in the context of claims (we refer to, e.g., directionality and non-interference principles [21,31]). Thereby, we consider the development of general schemes to transfer principles between AFs and CAFs as a crucial point on our agenda.

We furthermore want to extend our principle-based analysis to related formalisms that can be captured by CAFs. A principled approach to semantics for logic programs has been conducted in the context of dynamics [67]. It would be interesting to deepen these investigations and to study principles for logic programs the static case. Our work can serve as a promising starting point for such an agenda.

Another future work direction would be to extend our analysis to other CAF classes. While general CAFs capture all models that express claim-based argumentation and include, in particular, all possible ways to deviate from well-formedness, there are certain restrictions that are imposed by preference incorporation when considering concrete formalisms and methods. As shown in [27] typical methods for preference incorporation (cf. [26]) give rise to different CAF classes that lie between well-formed and unrestricted CAFs. It would be interesting to investigate the behavior of the semantics with respect to this classes. In this regard, we mention studies on the relation of structured argumentation formalisms and CAFs [30], which show that the instantiation of the rule-based argumentation formalism ASPIC$^+$ [10] that handles preferences on the level of rules is not contained by any of the aforementioned classes. However, precise characterizations of CAF classes arising from several structured argumentation formalisms such as assumption-based argumentation with preferences, ASPIC$^+$, or logic-based approaches have still to be found; tackling these open problems would be a challenging avenue for future work.

Extending our investigations to further claim-based semantics would be also a promising endeavor. In this regard, we consider studies on other inherited semantics based on e.g., strong or weak admissibility [21,68] worth investigating. It would be furthermore interesting to explore new notions of claim-acceptance and -defeat, e.g., by relaxing the defeat notion to allow for partial defeat or by strengthening claim-acceptance. Here, in-depth studies of conclusion-based evaluation methods in related formalisms, extending previous work [30,38], can be a promising starting point which is definitely an item on our future agenda.

From a computational point of view, we want to point out that the constructions that we have established in the scope of our expressiveness results can serve as interesting starting point to develop tools to model different viewpoints within argumentation. Having in mind that similar constructions play a role in the context of dynamics [24,69], it would be certainly of interest to pursue this direction of future work.

It would be furthermore interesting to adapt a ranking-based approach [70] in the context of claims. Another potential future research direction would be to generalize the claim function to allow for arguments that support several claims.

Apart from these endeavors, we identify the investigation of claim-based evaluation methods in the context of applied argumentation techniques (e.g., for case studies or argument mining [13,14,71]) as an interesting avenue for future work.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Data availability

No data was used for the research described in the article.

Acknowledgements

We thank the anonymous reviewers of the present and the conference version [29] for their thorough and considerate reviews and their helpful comments, which contributed significantly to improving this paper. This research has been supported by WWTF through project ICT19-065, and FWF through projects P30168 and W1255-N23.

Appendix A. On the relation between CAFs and logic programs

In this section, we prove our main result from Section 4.4. We furthermore give the results for the remaining semantics that we have not captured in this section. For this, let us formally define the LP semantics under consideration (we recall p-stable and L-stable model semantics).

Definition A.1. A 3-valued interpretation $I = (T, F)$ of P is

- partially stable (p-stable) iff I is a 3-valued model of P/I;
- well-founded iff I is p-stable with $\subseteq$-minimal T;
- regular iff I is p-stable with $\subseteq$-maximal T;
- stable iff I is p-stable and $T \cup F = \mathcal{L}(P)$;
- L-stable iff I is p-stable and $T \cup F$ is $\subseteq$-maximal among all p-stable models of P.

Utilizing results from [7], we obtain the following semantics correspondence:

Proposition A.2. *Let P be a logic program and $I = (T, F)$ be a 3-valued interpretation. I is P-stable iff $T \in co_i(\mathcal{F}_P)$; well-founded iff $T \in gr_c(\mathcal{F}_P)$; regular iff $T \in pr_c(\mathcal{F}_P)$; stable iff $T \in stb_i(\mathcal{F}_P)$.*

The remaining part of this appendix is devoted to the proof of Theorem 4.27.

Given two rules r and s with $head(s) \in body(r)$, we apply *rule-chaining* to obtain the rule r' by replacing the atom $head(s)$ with $body(s)$, i.e., r' is a rule with $head(r') = head(r)$ and $body(r') = (body(r) \setminus head(s)) \cup body(s)$. A rule r is called *atomic* if $pos(r) = \emptyset$. A program P is called *atomic* iff each rule in P is atomic.

Definition A.3. Let P be a logic program. An atom a in P is called *reachable* in P iff it is possible to construct an atomic rule r from rules in P by successive rule-chaining with $head(r) = a$. Atom a is called *unreachable* in P iff a is not reachable in P.

Each acceptable (w.r.t. 3-valued model semantics) atom is reachable.

Proposition A.4. *Let P be a logic program. It holds that all atoms in T of a 3-valued model $I = (T, F)$ of P are reachable in P.*

Proof. Let $I = (T, F)$ denote a 3-valued model of P and let U denote the set of unreachable atoms in P. We show that there is a 3-valued model $I' = (T', F)$ of P with $T' \subseteq T$ and $T' \cap U = \emptyset$. Since $I' \leq I$, it follows that $I' = I$ and thus T contains no unreachable atoms. We construct I via fixed point iteration:

$$I^0 = (T^0, F) = (T \setminus U, F)$$
$$I^{n+1} = (T^{n+1}, F) = (\{a \in T^n \mid \nexists r \in P/I : (a = head(r) \wedge pos(r) \subseteq T^n)\}$$

Starting with the set of unreachable atoms in P, we remove in each step atoms from T which require atoms outside of T to satisfy condition (a); one could say, we shrink T until we reach a state in which all atoms in T are reachable within T. The procedure has a fixed point (worst case we remove all atoms from T) and is thus guaranteed to terminate. We denote this fixed point by $I' = (T', F)$.

We show that I' is a 3-valued model of P. First observe that I' satisfies condition (b) since (b) is satisfied by I and since the fixed point iteration did not change atoms that are set to false in I. Moreover, I' satisfies condition (a):

($\Rightarrow$): Consider an atom $a \in T'$. That is, a is reachable in P with atoms from T. By construction, there is a rule r in the reduct P/I with $head(r) = a$ and $pos(r) \subseteq T'$, consequently the condition is satisfied.

($\Leftarrow$): Consider an atom $a \in \mathcal{L}(P)$ such that there is a rule $r \in P/I$ with $a = head(r)$ and $pos(r) \subseteq T'$. Since $pos(r) \subseteq T' \subseteq T$ it holds that $a \in T$ (by assumption I is a 3-valued model of P); consequently, $a \in T'$ as required.

Thus I' satisfies (a) and (b), moreover, we have $I' \leq I$ by construction. It follows that $I' = I$ and thus T contains no unreachable atoms. $\square$

Being reachable is a necessary but not a sufficient criteria for an atom a to appear in a p-stable model of a given program P (e.g., consider the program $P = \{a \leftarrow \text{not}\, a.\}$, then the atom a is reachable but not contained in a p-stable model of P).

Next we show that unreachable atoms are always false.

Proposition A.5. *Let P be a logic program and let a denote an atom which is unreachable in P. For all 3-valued models $I = (T, F)$ of P, it holds that $a \in F$.*

Proof. Consider an unreachable atom $a \in \mathcal{L}(P)$ and a 3-valued Herbrand interpretation $I = (T, F)$ with $a \notin F$. By Proposition A.4, it holds that $a \notin T$. Then $I' = (T, F \cup \{a\})$ is a Herbrand interpretation satisfying conditions (a) and (b) in the reduct P/I for all atoms $a \in \mathcal{L}(P)$, moreover, it holds that $I' < I$. Thus I is not a 3-valued model of P. $\square$

Definition A.6. Let P be a logic program. Set $P^0 = P$ and let

$$P^{i+1} = \{head(s) \leftarrow (body(s) \setminus \{head(r)\}) \cup body(r) \mid r, s \in P^i, head(r) \in body(s)\}$$
$$\cup \{r \in P^i \mid r \text{ is atomic in } P^i\}.$$

By P^∞ we denote the fixed point of this procedure, i.e., $P^\infty = P^i = P^{i+1}$ for some large enough $i \in \mathbb{N}$.

Next, we show that rule-chaining is a syntactic operation that does not change the semantics of a program.

Proposition A.7. *Let P be a logic program. $I = (T, F)$ denote a 3-valued model of P iff I is a 3-valued model of P^∞.*

Proof. First, we note that the addition of a rule s' which is obtained by replacing the atom $head(r) \in body(s)$ with $body(r)$ for given rules $r, s \in P$ does not affect the semantics. That is,

(1) $I = (T, F)$ is a 3-valued model of P iff I is a 3-valued model of $P' = P \cup \{head(s) \leftarrow (body(s) \setminus \{head(r)\}) \cup body(r)\}$ for rules $r, s \in P$.

Proof of (1). Consider rules $r, s \in P$ with $p = head(r)$ and $p \in body(s)$. Let s' denote the rule $head(s) \leftarrow (body(s) \setminus \{p\}) \cup body(r)$ and let $head(s) = head(s') = a$. First, we observe that $P/I \subseteq P'/I$ (since P' properly extends P by rule s') for any model I of P and P'.

First, consider a 3-valued model $I = (T, F)$ of P. Note that conditions (a) and (b) are satisfied in P'/I for each atom $b \neq a$. It thus suffices to check the conditions for atom a. In case $a \in T$, there is a rule $t \in P/I$ with $head(t) = a$ and $pos(t) \subseteq T$. Since P'/I is a superset of P/I, it holds that $t \in P'/I$. Now assume $a \in F$ and let us assume that (a modified version of) s' is contained in P'/I (otherwise, we are done as $P'/I = P/I$ in this case). Let s'' denote the modified version. It holds that $s'' \in P/I$. Since $a \in F$ we have $pos(s'') \cap F \neq \emptyset$. In case there is some $b \in pos(s'') \cap F$ different from p (i.e., $b \neq p = head(r)$), we are done: in this case, $b \in pos(s')$. Now assume that $p \in pos(s'') \cap F$ is the unique atom contained in the intersection. But then $pos(r) \cap F \neq \emptyset$ since $p \in F$. Consequently, we obtain that $pos(s') \cap F \neq \emptyset$.

For the other direction, let us assume that I is a 3-valued model of P'/I. Again, conditions (a) and (b) are satisfied in P/I for each atom $b \neq a$. Let us now consider the atom a. In case $a \in T$, there is a rule $t \in P'/I$ with $head(t) = a$ and $pos(t) \subseteq T$. In case $t \neq s''$ for s'' being the modified version of s' in the reduct P'/I we are done because then it holds that $t \in P/I$ as well. In case $t = s''$ for s'' being the modified version of s' in the reduct P'/I, it holds that (the modified version of) s serves as witness for $a \in T$ in P/I: indeed, $head(s) = a$ and $pos(s) \subseteq pos(s') \subseteq T$. Now, suppose $a \in F$ Hence, for each rule $t \in P'/I$ with $head(t) = a$ we have $pos(t) \cap F \neq \emptyset$. From $P/I \subseteq P'/I$ we obtain that condition (b) is satisfied in P/I as well. $\diamond$

Next, we show that replacing an atom $p \in body(s)$ with the body of *each* rule r_i with $head(r_i) = p$ (thus generating a new rule s_i for each such rule r_i) also allows for deletion of the rule s.

(2) Given $s \in P$ with $p \in body(s)$, and let $R = \{r_1, \dots, r_m\} \subseteq P$ denote the set of rules with rule head p. For each $i \leq m$, we let s_i denote the rule obtained from replacing p in $body(s)$ with $body(r_i)$, i.e., s_i is of the form $head(s) \leftarrow (body(s) \setminus \{p\}) \cup body(r_i)$. It holds that $I = (T, F)$ is a 3-valued model of P iff I is a 3-valued model of $P' = (P \setminus \{s\}) \cup \{s_1, \dots, s_m\}$.

Proof of (2). From (1) we know that the addition of rules $s_1, \dots, s_m$ to P does not affect the semantics. Let $P^* = P \cup \{s_1, \dots, s_m\}$. Then I is a 3-valued model of P iff I is a 3-valued model of P^*. The programs P' and P^* differ in exactly one rule, namely rule s. Let $head(s) = a$. We show that the deletion of s preserves 3-valued models. Similar as in (1), it suffices to discuss conditions (a) and (b) for atom a.

First, assume $I = (T, F)$ is a 3-valued model of P (and thus of P^*). Observe that $P'/I \subseteq P^*/I$ (in case $T \cap neg(s) = \emptyset$ we have $P'/I = P^*/I$). Let $a \in T$. Then there is a rule $t \in P^*/I$ with $head(t) = a$ and $pos(t) \subseteq T$. Again, we are done in case $t \neq s$ because then $t \in P'/I$ holds. Now assume $t = s$. Then $pos(s) \subseteq T$ and (a modified version of) s is contained in the reduct P^*. That is, $neg(s) \cap T \neq \emptyset$. From $pos(s) \subseteq T$ we obtain $p \in T$. Thus there is a rule $r_i' \in P^*/I$ with $head(r_i') = p$ and $pos(r_i') \subseteq T$ where r_i' is a modified version of rule $r_i \in P^*$ with head p. Thus there is a rule $s_i' \in P'/I$ with $head(s_i') = a$ and $pos(s_i') \subseteq T$ which corresponds to the rule $s_i \in P'$ obtained by replacing $p \in body(s)$ by $body(r_i)$. Consequently, condition (a) is satisfied. In case $a \in F$ it holds that condition (b) is satisfied in P'/I because $P'/I \subseteq P^*/I$.

For the other direction, assume $I = (T, F)$ is a model of P'. Similar as above, in case $a \in T$ we obtain that condition (a) is satisfied in P^*/I because $P'/I \subseteq P^*/I$. Now assume $a \in F$. That is, each rule t with $head(t) = a$ satisfies $pos(t) \cap F \neq \emptyset$. We show that the modified version s' of s in P^*/I satisfies the condition as well. Each s_i' (where s_i' being the modified version of s_i in the reduct P'/I) satisfies condition (b). In case there is $b \in pos(s_i')$ with $b \notin pos(r_i)$ for some $i \leq m$ we are done. In this case, $b \in pos(s')$. Otherwise, it holds that for all rules $r_i' \in P'/I$ with $head(r_i') = p$ there is some $c \in pos(r_i') \cap F$. As $r_i' \in P'/I$ iff $r_i' \in P^*/I$ we obtain $p \in F$. Consequently, $pos(s') \cap F \neq \emptyset$ and we obtain that condition (b) is satisfied. ◇

Given P^i we obtain P^{i+1} as follows: for each rule $s \in P^i$, for each $p \in pos(s)$, we replace s with the set of rules obtained by replacing p with the body of all rules in P^i with head p. In case s is atomic we add it to P^{i+1}. As shown in (2), replacing rules does not change the 3-valued models of a program. □

Reachability can be alternatively defined via P^∞: An atom a is reachable if there exists an atomic rule $r \in P^\infty$ with $head(r) = a$. We note that the rules in P^∞ which are not atomic can be deleted without changing the semantics in case each atom in P^∞ is reachable. Intuitively, such rules do not carry any additional information which has not been incorporated yet. Recall that unreachable atoms are set to false. We thus obtain the following result.

Proposition A.8. *For each logic program P with unreachable atoms $U \subseteq \mathcal{L}(P)$, there exists an atomic program P' such that $I' = (T, F)$ is a 3-valued model of P' iff $I = (T, F \cup U)$ is a 3-valued model of P.*

We are ready to prove our main result of this section.

Theorem 4.27. *Let P be a logic program, $\mathcal{F}_P$ the associated CAF, and $I = (T, F)$ be a 3-valued interpretation. Then I is L-stable in P iff $T \in ss_h(\mathcal{F}_P)$.*

Proof. By Proposition A.5, it suffices to consider logic programs without unreachable atoms: indeed, if atom a is unreachable, then we have that $a \in F$ for each model $I = (T, F)$. Removing unreachable atoms therefore does not change $\subseteq$-maximality of $T \cup F$.

Consider a logic program P without unreachable atoms. Notice that the corresponding CAF $\mathcal{F}_P$ contains (at least) one argument for each atom in P. By Proposition A.2, we have $T \in co(\mathcal{F}_P)$ iff $I = (T, F)$ is p-stable in P. We obtain the correspondence of L-stable semantics with h-semi-stable semantics by observing that defeated claims (in $\mathcal{F}_P$) correspond to (reachable) atoms that are set so false (in P).

By Proposition A.7, we obtain that moving from P to P^∞ does not change the semantics of P, i.e., I is a 3-valued model of P iff I is a 3-valued model of P^∞. It thus suffices to show $F = T^*_{\mathcal{F}_P}$ for all p-stable models $I = (T, F)$ of P^∞. By assumption each atom is reachable we observe that each rule in P^∞ is atomic. As each atomic rule induces exactly one argument, there is a one-to-one correspondence between the arguments constructed from P and the rules in P^∞.

Let $I = (T, F)$ denote a 3-valued model of P.

First, we show that all arguments in the corresponding CAF $\mathcal{F}_P$ with claims in F are attacked by T. Consider some $p \in F$ and let r denote a rule of P^∞ with $head(r) = p$. The rule r is of the form $p \leftarrow \text{not } b_1, \ldots, \text{not } b_m$. Since $p \in F$ and since $pos(r) = \emptyset$ it holds that $T \cap neg(r) \neq \emptyset$. By definition of an argument in $\mathcal{F}_P$, each $b \in neg(r)$ is a vulnerability of A, i.e., $b \in \textsc{Vul}(A)$. By definition of the attack relation, it holds that each argument with claim b attacks A.

For the other direction, consider some claim p that is attacked by T in $\mathcal{F}_P$. That is, for each argument A with claim p, it holds that $\textsc{Vul}(A) \cap T \neq \emptyset$. Thus for each rule r with $head(r) = p$, it holds that $T \cap neg(r) \neq \emptyset$. Consequently, P^∞ does not contain rules with head p. It follows that $p \in F$. □

References

[1] T.J.M. Bench-Capon, P.E. Dunne, Argumentation in artificial intelligence, Artif. Intell. 171 (10–15) (2007) 619–641, https://doi.org/10.1016/j.artint.2007.05.001.

[2] K. Atkinson, P. Baroni, M. Giacomin, A. Hunter, H. Prakken, C. Reed, G.R. Simari, M. Thimm, S. Villata, Towards artificial argumentation, AI Mag. 38 (3) (2017) 25–36.

[3] P. Baroni, D.M. Gabbay, M. Giacomin, L. van der Torre, Handbook of Formal Argumentation, College Publications, 2018, https://books.google.at/books?id=_OnTswEACAAJ.

[4] F.H. van Eemeren, B. Verheij, Argumentation theory in formal and computational perspective, in: Handbook of Formal Argumentation, College Publications, 2018, pp. 2–73, Ch. 1.

[5] P.M. Dung, On the acceptability of arguments and its fundamental role in nonmonotonic reasoning, logic programming and n-person games, Artif. Intell. 77 (2) (1995) 321–358.

[6] Y. Wu, M. Caminada, D.M. Gabbay, Complete extensions in argumentation coincide with 3-valued stable models in logic programming, Stud. Log. 93 (2–3) (2009) 383–403, https://doi.org/10.1007/s11225-009-9210-5.

[7] M. Caminada, S. Sá, J. Alcântara, W. Dvořák, On the equivalence between logic programming semantics and argumentation semantics, Int. J. Approx. Reason. 58 (2015) 87–111, https://doi.org/10.1016/j.ijar.2014.12.004.

[8] P. Besnard, A. Hunter, A review of argumentation based on deductive arguments, in: Handbook of Formal Argumentation, College Publications, 2018, pp. 436–484, Ch. 9.

[9] L. Amgoud, P. Besnard, A formal analysis of logic-based argumentation systems, in: Proceedings of the 4th International Conference on Scalable Uncertainty Management, SUM 2010, in: Lecture Notes in Computer Science, vol. 6379, Springer, 2010, pp. 42–55, https://doi.org/10.1007/978-3-642-15951-0_10.

[10] S. Modgil, H. Prakken, Abstract rule-based argumentation, in: Handbook of Formal Argumentation, College Publications, 2018, pp. 287–364, Ch. 6.

[11] K. Cyras, X. Fan, C. Schulz, F. Toni, Assumption-based argumentation: disputes, explanations, preferences, in: Handbook of Formal Argumentation, College Publications, 2018, pp. 365–408, Ch. 7, also appears in IfCoLog J. Log. Appl. 4 (8) (2017) 2407–2456.

[12] T.J.M. Bench-Capon, K. Atkinson, A. Chorley, Persuasion and value in legal argument, J. Log. Comput. 15 (6) (2005) 1075–1097, https://doi.org/10.1093/logcom/exi058.

[13] T.J.M. Bench-Capon, Representation of case law as an argumentation framework, in: Proceedings of the 15th Annual Conference on Legal Knowledge and Information Systems, JURIX 2002, IOS Press, Amsterdam, 2002, pp. 103–112.

[14] A. Hunter, M. Williams, Aggregating evidence about the positive and negative effects of treatments, Artif. Intell. Med. 56 (3) (2012) 173–190, https://doi.org/10.1016/j.artmed.2012.09.004.

[15] L. Amgoud, Postulates for logic-based argumentation systems, Int. J. Approx. Reason. 55 (9) (2014) 2028–2048, https://doi.org/10.1016/j.ijar.2013.10.004.

[16] J.F. Horty, Skepticism and floating conclusions, Artif. Intell. 135 (1–2) (2002) 55–72, https://doi.org/10.1016/S0004-3702(01)00160-6.

[17] M. Beirlaen, J. Heyninck, C. Straßer, Reasoning by cases in structured argumentation, in: Proceedings of the Symposium on Applied Computing, SAC 2017, ACM, 2017, pp. 989–994, https://doi.org/10.1145/3019612.3019716.

[18] P. Baroni, R. Riveret, Enhancing statement evaluation in argumentation via multi-labelling systems, J. Artif. Intell. Res. 66 (2019) 793–860, https://doi.org/10.1613/jair.1.11428.

[19] W. Dvořák, S. Woltran, Complexity of abstract argumentation under a claim-centric view, Artif. Intell. 285 (2020) 103290, https://doi.org/10.1016/j.artint.2020.103290.

[20] W. Dvořák, A. Greßler, A. Rapberger, S. Woltran, The complexity landscape of claim-augmented argumentation frameworks, Artif. Intell. 317 (2023) 103873, https://doi.org/10.1016/j.artint.2023.103873.

[21] P. Baroni, M. Giacomin, On principle-based evaluation of extension-based argumentation semantics, Artif. Intell. 171 (10–15) (2007) 675–700.

[22] L. van der Torre, S. Vesic, The principle-based approach to abstract argumentation semantics, IfCoLog J. Log. Appl. 4 (8) (2017), http://www.collegepublications.co.uk/downloads/ifcolog00017.pdf.

[23] P.E. Dunne, W. Dvořák, T. Linsbichler, S. Woltran, Characteristics of multiple viewpoints in abstract argumentation, Artif. Intell. 228 (2015) 153–178, https://doi.org/10.1016/j.artint.2015.07.006.

[24] R. Baumann, D.M. Gabbay, O. Rodrigues, Forgetting an argument, in: Proceedings of the 34th AAAI Conference on Artificial Intelligence, AAAI 2020, AAAI Press, 2020, pp. 2750–2757, https://ojs.aaai.org/index.php/AAAI/article/view/5662.

[25] M. Berthold, A. Rapberger, M. Ulbricht, Forgetting aspects in assumption-based argumentation, in: Proceedings of the 20th International Conference on Principles of Knowledge Representation and Reasoning, KR 2023, 2023, pp. 86–96, https://doi.org/10.24963/kr.2023/9.

[26] S. Kaci, L.W.N. van der Torre, S. Vesic, S. Villata, Preference in Abstract Argumentation, Handbook of Formal Argumentation, vol. 2, College Publications, 2021, pp. 211–248.

[27] M. Bernreiter, W. Dvořák, A. Rapberger, S. Woltran, The effect of preferences in abstract argumentation under a claim-centric view, in: Proceedings of the 37th AAAI Conference on Artificial Intelligence, AAAI 2023, AAAI Press, 2023, pp. 6253–6261, https://ojs.aaai.org/index.php/AAAI/article/view/25770.

[28] A. Rapberger, Defining argumentation semantics under a claim-centric view, in: Proceedings of the 9th European Starting AI Researchers' Symposium, STAIRS 2020, in: CEUR Workshop Proceedings, CEUR-WS.org, vol. 2655, 2020, http://ceur-ws.org/Vol-2655/paper2.pdf.

[29] W. Dvořák, A. Rapberger, S. Woltran, Argumentation semantics under a claim-centric view: properties, expressiveness and relation to SETAFs, in: Proceedings of the 17th International Conference on Principles of Knowledge Representation and Reasoning, KR 2020, 2020, pp. 341–350, https://doi.org/10.24963/kr.2020/35.

[30] A. Rapberger, Unpacking the argument – a claim-centric view on abstract argumentation, Ph.D. thesis, TU Wien, 2023.

[31] M. Caminada, W.A. Carnielli, P.E. Dunne, Semi-stable semantics, J. Log. Comput. 22 (2012) 1207–1254, https://doi.org/10.1093/logcom/exr033.

[32] B. Verheij, Two approaches to dialectical argumentation: admissible sets and argumentation stages, in: Proceedings of the 8th Dutch Conference on Artificial Intelligence, NAIC'96, 1996, pp. 357–368.

[33] D. Saccà, C. Zaniolo, Deterministic and non-deterministic stable models, J. Log. Comput. 7 (5) (1997) 555–579, https://doi.org/10.1093/logcom/7.5.555.
[34] T. Eiter, N. Leone, D. Saccà, On the partial semantics for disjunctive deductive databases, Ann. Math. Artif. Intell. 19 (1–2) (1997) 59–96, https://doi.org/10.1023/A:1018947420290.
[35] R. Reiter, A logic for default reasoning, Artif. Intell. 13 (1–2) (1980) 81–132, https://doi.org/10.1016/0004-3702(80)90014-4.
[36] P. Baroni, M. Caminada, M. Giacomin, An introduction to argumentation semantics, Knowl. Eng. Rev. 26 (4) (2011) 365–410, https://doi.org/10.1017/S0269888911000166.
[37] M. Caminada, S. Sá, J. Alcântara, W. Dvořák, On the difference between assumption-based argumentation and abstract argumentation, IfCoLog J. Log. Appl. 2 (1) (2015) 15–34.
[38] M. König, A. Rapberger, M. Ulbricht, Just a matter of perspective: intertranslating expressive argumentation formalisms, in: Proceedings of the 9th International Conference on Computational Models of Argument, COMMA 2022, in: Frontiers in Artificial Intelligence and Applications, vol. 353, IOS Press, 2022, pp. 212–223.
[39] G. Brewka, T. Eiter, M. Truszczynski, Answer set programming at a glance, Commun. ACM 54 (12) (2011) 92–103, https://doi.org/10.1145/2043174.2043195.
[40] T.C. Przymusinski, The well-founded semantics coincides with the three-valued stable semantics, Fundam. Inform. 13 (4) (1990) 445–463.
[41] M. Gelfond, V. Lifschitz, The stable model semantics for logic programming, in: Proceedings of the 5th International Conference and Symposium on Logic Programming, ICLP/SLP 1988, MIT Press, 1988, pp. 1070–1080.
[42] P. Baroni, M. Caminada, M. Giacomin, Abstract argumentation frameworks and their semantics, in: Handbook of Formal Argumentation, College Publications, 2018, pp. 159–236, Ch. 4.
[43] L. Amgoud, P. Besnard, Logical limits of abstract argumentation frameworks, J. Appl. Non-Class. Log. 23 (3) (2013) 229–267, https://doi.org/10.1080/11663081.2013.830381.
[44] E. Bonzon, J. Delobelle, S. Konieczny, N. Maudet, A comparative study of ranking-based semantics for abstract argumentation, in: Proceedings of the 30th AAAI Conference on Artificial Intelligence, AAAI 2016, AAAI Press, 2016, pp. 914–920, http://www.aaai.org/ocs/index.php/AAAI/AAAI16/paper/view/12465.
[45] R. Baumann, A. Rapberger, M. Ulbricht, Equivalence in argumentation frameworks with a claim-centric view: classical results with novel ingredients, J. Artif. Intell. Res. 77 (2023) 891–948, https://doi.org/10.1613/jair.1.14625.
[46] D. Makinson, K. Schlechta, Floating conclusions and zombie paths: two deep difficulties in the "directly skeptical" approach to defeasible inheritance nets, Artif. Intell. 48 (2) (1991) 199–209, https://doi.org/10.1016/0004-3702(91)90061-N.
[47] H. Prakken, Intuitions and the modelling of defeasible reasoning: some case studies, in: Proceedings of the 9th International Workshop on Non-Monotonic Reasoning, NMR 2002, 2002, pp. 91–102.
[48] J.F. Horty, Skepticism and floating conclusions, Artif. Intell. 135 (1–2) (2002) 55–72, https://doi.org/10.1016/S0004-3702(01)00160-6.
[49] P.E. Dunne, P. Marquis, M. Wooldridge, Argument aggregation: basic axioms and complexity results, in: Proceedings of the 4th International Conference on Computational Models of Argument, COMMA 2012, in: Frontiers in Artificial Intelligence and Applications, vol. 245, IOS Press, 2012, pp. 129–140.
[50] R. Baumann, On the nature of argumentation semantics: existence and uniqueness, expressibility, and replaceability, in: Handbook of Formal Argumentation, College Publications, 2018, pp. 839–936, Ch. 17.
[51] S.H. Nielsen, S. Parsons, A generalization of Dung's abstract framework for argumentation: arguing with sets of attacking arguments, in: Proceedings of the 3rd International Workshop on Argumentation in Multi-Agent Systems, ArgMAS 2006, Revised Selected and Invited Papers, in: Lecture Notes in Computer Science, vol. 4766, Springer, 2006, pp. 54–73, https://doi.org/10.1007/978-3-540-75526-5_4.
[52] W. Dvořák, J. Fandinno, S. Woltran, On the expressive power of collective attacks, in: Proceedings of the 7th International Conference on Computational Models of Argument, COMMA 2018, 2018, pp. 49–60, https://doi.org/10.3233/978-1-61499-906-5-49.
[53] W. Dvořák, A. Rapberger, S. Woltran, On the relation between claim-augmented argumentation frameworks and collective attacks, in: Proceedings of the 24th European Conference on Artificial Intelligence, ECAI 2020, in: Frontiers in Artificial Intelligence and Applications, vol. 325, IOS Press, 2020, pp. 721–728, https://doi.org/10.3233/FAIA200159.
[54] L. Amgoud, P. Besnard, Bridging the gap between abstract argumentation systems and logic, in: 3rd International Conference on Scalable Uncertainty Management, SUM 2009, in: Lecture Notes in Computer Science, vol. 5785, Springer, 2009, pp. 12–27, https://doi.org/10.1007/978-3-642-04388-8_3.
[55] N. Gorogiannis, A. Hunter, Instantiating abstract argumentation with classical logic arguments: postulates and properties, Artif. Intell. 175 (9–10) (2011) 1479–1497, https://doi.org/10.1016/j.artint.2010.12.003.
[56] M. Caminada, L. Amgoud, On the evaluation of argumentation formalisms, Artif. Intell. 171 (5–6) (2007) 286–310, https://doi.org/10.1016/j.artint.2007.02.003.
[57] P. Baroni, G. Governatori, R. Riveret, On labelling statements in multi-labelling argumentation, in: Proceedings of the 22nd European Conference on Artificial Intelligence, ECAI 2016, in: Frontiers in Artificial Intelligence and Applications, vol. 285, IOS Press, 2016, pp. 489–497, https://doi.org/10.3233/978-1-61499-672-9-489.
[58] E.A. Corsi, C.G. Fermüller, Connecting fuzzy logic and argumentation frames via logical attack principles, Soft Comput. 23 (7) (2019) 2255–2270, https://doi.org/10.1007/s00500-018-3513-2.
[59] A. Rapberger, M. Ulbricht, On dynamics in structured argumentation formalisms, J. Artif. Intell. Res. 77 (2023) 563–643, https://doi.org/10.1613/jair.1.14481.
[60] H. Prakken, Relating abstract and structured accounts of argumentation dynamics: the case of expansions, in: Proceedings of the 20th International Conference on Principles of Knowledge Representation and Reasoning, KR 2023, 2023, pp. 562–571, https://doi.org/10.24963/kr.2023/55.
[61] J.K. Fichte, M. Hecher, Y. Mahmood, A. Meier, Quantitative reasoning and structural complexity for claim-centric argumentation, in: Proceedings of the Thirty-Second International Joint Conference on Artificial Intelligence, IJCAI 2023, ijcai.org, 2023, pp. 3212–3220, https://doi.org/10.24963/ijcai.2023/358.
[62] A. Greßler, Argumentation frameworks with claims and collective attacks: Complexity results and answer-set programming encodings, Master's thesis, TU Wien, 2019.
[63] V.H.N. Rocha, F.G. Cozman, A credal least undefined stable semantics for probabilistic logic programs and probabilistic argumentation, in: Proceedings of the 19th International Conference on Principles of Knowledge Representation and Reasoning, KR 2022, 2022, https://proceedings.kr.org/2022/31/.
[64] V.H.N. Rocha, F.G. Cozman, Bipolar argumentation frameworks with explicit conclusions: connecting argumentation and logic programming, in: Proceedings of the 20th International Workshop on Non-Monotonic Reasoning, NMR 2022, CEUR-WS.org, in: CEUR Workshop Proceedings, vol. 3197, 2022, pp. 49–60, http://ceur-ws.org/Vol-3197/paper5.pdf.
[65] W. Dvorák, M. König, M. Ulbricht, S. Woltran, Rediscovering argumentation principles utilizing collective attacks, in: Proceedings of the 19th International Conference on Principles of Knowledge Representation and Reasoning, KR 2022, 2022, pp. 122–131, https://doi.org/10.24963/kr.2022/13.
[66] R. Baumann, G. Brewka, M. Ulbricht, Comparing weak admissibility semantics to their dung-style counterparts - reduct, modularization, and strong equivalence in abstract argumentation, in: Proceedings of the 17th International Conference on Principles of Knowledge Representation and Reasoning, KR 2020, 2020, pp. 79–88, https://doi.org/10.24963/kr.2020/9.

[67] J.J. Alferes, F. Banti, A. Brogi, J.A. Leite, The refined extension principle for semantics of dynamic logic programming, Stud. Log. 79 (1) (2005) 7–32, https://doi.org/10.1007/s11225-005-0492-y.

[68] R. Baumann, G. Brewka, M. Ulbricht, Revisiting the foundations of abstract argumentation - semantics based on weak admissibility and weak defense, in: The 34th AAAI Conference on Artificial Intelligence, AAAI 2020, AAAI Press, 2020, pp. 2742–2749, https://ojs.aaai.org/index.php/AAAI/article/view/5661.

[69] C. Cayrol, F.D. de Saint-Cyr, M. Lagasquie-Schiex, Change in abstract argumentation frameworks: adding an argument, J. Artif. Intell. Res. 38 (2010) 49–84, https://doi.org/10.1613/jair.2965.

[70] K. Skiba, T. Rienstra, M. Thimm, J. Heyninck, G. Kern-Isberner, Ranking extensions in abstract argumentation, in: Proceedings of the 30th International Joint Conference on Artificial Intelligence, IJCAI 2021, ijcai.org, 2021, pp. 2047–2053, https://doi.org/10.24963/ijcai.2021/282.

[71] J. Lawrence, C. Reed, Argument mining: a survey, Comput. Linguist. 45 (4) (2019) 765–818, https://doi.org/10.1162/coli_a_00364.