Journal of Artificial Intelligence Research 77 (2023) 891-948 Submitted 01/2023; published 07/2023

Equivalence in Argumentation Frameworks with a Claim-centric View: Classical Results with Novel Ingredients

Ringo Baumann BAUMANN@INFORMATIK.UNI-LEIPZIG.DE
Department of Computer Science,
Leipzig University, Germany

Anna Rapberger ANNA.RAPBERGER@TUWIEN.AC.AT
Institute of Logic and Computation,
TU Wien, Austria

Markus Ulbricht MULBRICHT@INFORMATIK.UNI-LEIPZIG.DE
Department of Computer Science,
Leipzig University, Germany

Abstract

A common feature of non-monotonic logics is that the classical notion of equivalence does not preserve the intended meaning in light of additional information. Consequently, the term strong equivalence was coined in the literature and thoroughly investigated. In the present paper, the knowledge representation formalism under consideration is claim-augmented argumentation frameworks (CAFs) which provide a formal basis to analyze conclusion-oriented problems in argumentation by adapting a claim-focused perspective. CAFs extend Dung AFs by associating a claim to each argument representing its conclusion. In this paper, we investigate both ordinary and strong equivalence in CAFs. Thereby, we take the fact into account that one might either be interested in the actual arguments or their claims only. The former point of view naturally yields an extension of strong equivalence for AFs to the claim-based setting while the latter gives rise to a novel equivalence notion which is genuine for CAFs. We tailor, examine and compare these notions and obtain a comprehensive study of this matter for CAFs. We conclude by investigating the computational complexity of naturally arising decision problems.

1. Introduction

Equivalence is an important subject of research in knowledge representation and reasoning. Given a knowledge base $\mathcal{K}$, finding an equivalent one, say $\mathcal{K}'$, helps to obtain a better understanding or more concise representation of $\mathcal{K}$. From a computational point of view, equivalence is particularly interesting whenever a certain subset of a collection of information can be replaced without changing the intended meaning. In propositional logic, for example, replacing a subformula ϕ of ψ with an equivalent one, say ϕ', yields a formula $\psi[\phi/\phi']$ equivalent to ψ. That is, we may view ϕ as an independent module of ψ. Within the KR community it is well known that this is usually not the case for non-monotonic logics (Truszczyński, 2006; Baumann & Strass, 2022).

Motivated by this observation, the notion of strong equivalence was introduced in the literature. In a nutshell, strong equivalence requires the aforementioned property by design: $\mathcal{K}$ and $\mathcal{K}'$ are strongly equivalent if for any $\mathcal{H}$, the knowledge bases $\mathcal{K} \cup \mathcal{H}$ and $\mathcal{K}' \cup \mathcal{H}$ are equivalent. Although a naive implementation would require to iterate over an infinite number of possible $\mathcal{H}$, researchers discovered techniques to decide strong equivalence of two knowledge bases efficiently, most notably for logic programming (Lifschitz, Pearce, & Valverde, 2001) and argumentation frameworks (AFs) (Oikarinen & Woltran, 2011). The possibility to replace parts of a framework in a semantically neutral way is particularly important whenever dynamics in argumentation are considered. The latter topic is rightly one of the most active research areas within the community at the moment (Gabbay, Giacomin, Simari, & Thimm, 2021). In this paper, we extend this line of research to a recent extension of AFs, called *claim-augmented argumentation frameworks (CAFs)* (Dvořák & Woltran, 2020).

Dung boosted the research in abstract argumentation (Dung, 1995) which by now can be considered a classical area in knowledge representation and reasoning. AFs have been thoroughly investigated since then and various extensions have been proposed; e.g., the addition of supports (Cayrol & Lagasquie-Schiex, 2005), recursive (Baroni, Cerutti, Giacomin, & Guida, 2011) and collective (Nielsen & Parsons, 2006) attacks, or probabilities (Thimm, 2012) to mention a few. A popular line of research that emerged in recent years is centered around conclusion-oriented reasoning in argumentation (Baroni & Riveret, 2019; Dvořák & Woltran, 2020). While traditional argumentation formalisms focus on the identification of acceptable arguments, the emphasis in claim-focused argumentation lies on the argument's conclusions (*claims*).

Building on the observation that a claim can be supported by different arguments, it becomes evident that the traditional argument-focused perspective is often insufficient to capture claim-based reasoning. Claim-augmented argumentation frameworks (CAFs) address this issue by extending AFs with a function that assigns a claim to each argument. They are ideally suited to analyze instantiation-based approaches, e.g., instantiations of logic programs (Caminada, Sá, Alcântara, & Dvořák, 2015), rule-based formalisms, e.g., assumption-based argumentation (Bondarenko, Toni, & Kowalski, 1993; Cyras, Fan, Schulz, & Toni, 2018), or ASPIC$^+$ (Modgil & Prakken, 2018), as well as logic-based instantiations (Besnard & Hunter, 2001; Gorogiannis & Hunter, 2011), where the focus lies on the claims of the constructed arguments. In a nutshell, such an instantiation procedure starts from a knowledge base $\mathcal{K}$ by constructing arguments and identifying conflicts between them. In the next step, one abstracts away from the internal structure of the arguments and analyzes the resulting abstract framework. It then becomes possible to analyze the initial problem in terms of the claims directly on the abstract level.

Our main motivation to investigate the behavior of CAFs is that they can help streamlining such instantiation procedures. To illustrate this, we consider an example within *assumption-based argumentation* (ABA).[1]

Example 1.1. We consider an instantiation of an ABA framework $D = (\mathcal{L}, \mathcal{R}, \mathcal{A}, ^-)$ with atoms $\mathcal{L} = \{a, b, c, d, \overline{a}, \overline{b}, \overline{c}, \overline{d}, p\}$, assumptions $\mathcal{A} = \{a, b, c, d\}$, and their contraries $\overline{a}, \overline{b}, \overline{c}, \overline{d}$. We furthermore assume we are given five rules

$$r_1 : \bar{a} \leftarrow a, b. \qquad r_2 : p \leftarrow a. \qquad r_3 : \bar{b} \leftarrow c. \qquad r_4 : \bar{c} \leftarrow b. \qquad r_5 : \bar{c} \leftarrow d.$$

1. We do not discuss ABA in detail in this paper. Confer the *tutorial on assumption-based argumentation* by Toni (2014) for a comprehensive introduction. The example is given in a way that no ABA knowledge is required in order to follow our reasoning.

We obtain the associated AF F_D as follows: each assumption in $\mathcal{A}$ yields a corresponding argument and each rule r_i yields an argument x_i. Attacks depend on the conclusion of the attacking argument, e.g., x_3 attacks x_4 because $\overline{b}$ is the contrary of b.

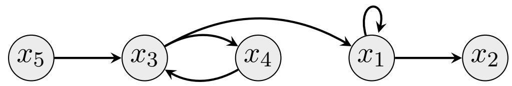

Thereby, both arguments x_4 and x_5 are associated with the conclusion $\bar{c}$. Since Dung-style AFs are not tailored to capture such a relationship between two arguments, some more technical machinery is required where we need to make use of information encoded in the underlying knowledge base; the AF itself does not contain sufficient information to assess whether e.g. a certain conclusion is credulously accepted.

CAFs provide a natural solution to this problem by extending AFs in a way that to each argument an associated claim is assigned as well. In recent years, CAFs have been studied under various aspects, e.g., in terms of computational complexity (Dvořák & Woltran, 2020; Dvořák, Greßler, Rapberger, & Woltran, 2023), their expressiveness (Dvořák, Rapberger, & Woltran, 2020a), and their relation to other formalisms (Dvořák et al., 2020a; König, Rapberger, & Ulbricht, 2022; Rocha & Cozman, 2022); in particular, it has been shown that with CAFs it is possible to capture semantics which cannot be modeled by AFs (Rapberger, 2020). As pointed out by Dvořák and Woltran (2020), using CAFs to instantiate knowledge representation formalisms streamlines the evaluation.

Example 1.2 (Example 1.1 ctd.)**.** Let us consider the same instantiation procedure, but this time augment the arguments with their respective conclusions.

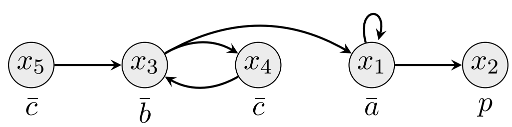

We now see that the relationship between x_4 and x_5 is encoded in the graph and we can assess the behavior of the knowledge base by investigating the CAF only.

CAFs thus provide a more robust translation and better preserve properties of the instantiated knowledge base, making CAFs a promising subject to investigation. As a first step, we will therefore examine ordinary and strong equivalence for CAFs. It will turn out that, similar to related research for AFs, we can characterize strong equivalence for CAFs in terms of semantics-preserving normal forms, called *kernels*.

We will then re-assess these results from the point of view of claim-based reasoning: While our characterizations are solid whenever we are interested in both the arguments as well as their claims, one could also interpret CAFs as a tool to reason *solely* with claims. That is, we view a CAF as a representation of a multi-set of claims and their interactions.

Example 1.3 (Example 1.1 ctd.)**.** Suppose we are only interested in the claims of each argument, and abstract away their names. This would yield the following situation:

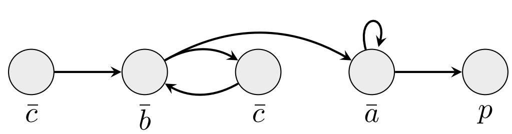

This is captured by our notion of *renamings*, which allow us to characterize isomorphisms between CAFs. We will tailor a strong equivalence notion for this interpretation, called *strong equivalence up to renaming* and see that characterizing this notion builds upon our aforementioned kernel characterizations and the well-known graph isomorphism problem.

For many instantiation procedures in the literature, out-going attacks of arguments are characterized by their conclusions. Arguments typically consist of a *claim* and a *support*, and two arguments attack each other if the claim of an argument contradicts the support of another argument. Prominent examples are argumentation formalisms based on logic; here, arguments are pairs $\langle \Phi, \alpha \rangle$ consisting of a set of formulas Φ and claim α which determines the outgoing attacks (Besnard & Hunter, 2018, Definition 5.5); ABA where arguments are tree-derivations that attack each other if the contrary of an assumption is derived; or instantiations of logic programs and default logic (Caminada et al., 2015; Dung, 1995; Wu, Caminada, & Gabbay, 2009). This motivated the notion of so-called *well-formed CAFs* (Dvořák & Woltran, 2020). In such frameworks, arguments with the same claim have the same outgoing attacks. Driven by this observation, we will investigate well-formed CAFs as a special case throughout this work.

Our main contributions can be summarized as follows:

- We discuss ordinary equivalence for general and well-formed CAFs and present dependencies between semantics.
- We develop the notion of strong equivalence between CAFs and provide characterization results via semantics-dependent kernels. We achieve this for each semantics which has been considered in the literature so far.
- We introduce novel equivalence concepts based on argument renaming which are genuine for CAFs. We show that ordinary equivalence up to renaming coincides with ordinary equivalence while strong equivalence up to renaming can be characterized via kernel isomorphism.
- We present a rigorous complexity analysis of deciding equivalence between two CAFs for all of the aforementioned equivalence notions. We show that deciding ordinary equivalence can be computationally hard, up to the third level of the polynomial hierarchy. Moreover, we show that strong equivalence up to renaming is as hard as the graph isomorphism problem.
- We show that deciding whether two well-formed CAFs are isomorphic is tractable. Building upon this observation, we infer tractability of strong equivalence up to renaming for this class of CAFs.

The present paper is an extended version of the conference version (Baumann, Rapberger, & Ulbricht, 2022). Besides providing full proofs and a stronger intuition about our technical details, the present version extends the previous conference publication by a comprehensive analysis of the equivalence behavior of well-formed CAFs: we present novel results regarding ordinary and strong equivalence, and strong equivalence up to argument renaming for this important sub-class. We furthermore extend our computational complexity analysis to well-formed CAFs. Moreover, the absence of space limit gives us the chance to better put our results in context and discuss related work in more detail.

2. Background

We start by giving necessary background on abstract and claim-based argumentation.

Abstract Argumentation. We fix a non-finite background set $\mathcal{U}$. An argumentation framework (AF) (Dung, 1995) is a directed graph $F = (A, R)$ where $A \subseteq \mathcal{U}$ is a finite[2] set of arguments and $R \subseteq A \times A$ models *attacks* between them. We use $\mathcal{AF}$ to denote the set of all AFs.

For two arguments $a, b \in A$, if $(a, b) \in R$ we say that a *attacks* b as well as the set $E \subseteq A$ *attacks* b if $a \in E$. Analogously, a attacks E if $(a, b) \in R$ for some $b \in E$. The *range* of a set $E \subseteq A$ is defined as $E^{\oplus}_F = E \cup E^+_F$ where $E^+_F = \{a \in A \mid E \text{ attacks } a\}$. E is conflict-free in F (for short, $E \in cf(F)$) iff for no $a, b \in E$, $(a, b) \in R$. E *defends* an argument a if any attacker of a is attacked by E. A *semantics* is a function $\sigma : \mathcal{AF} \to 2^{2^{\mathcal{U}}}$ with $F \mapsto \sigma(F) \subseteq 2^A$. This means, given an AF $F = (A, R)$ a semantics returns a set of subsets of A. These subsets are called σ-*extensions*.

In this paper we consider so-called *naive*, *admissible*, *complete*, *grounded*, *preferred*, *stable*, *semi-stable* and *stage* semantics (abbr. *na*, *ad*, *co*, *gr*, *pr*, *stb*, *ss*, *stg*). Apart from naive, semi-stable and stage semantics (Verheij, 1996; Caminada, 2006), all mentioned semantics were already introduced by Dung (1995).

Definition 2.1. Let $F = (A, R)$ be an AF and $E \in cf(F)$.

1. $E \in na(F)$ iff E is $\subseteq$-maximal in $cf(F)$,
2. $E \in ad(F)$ iff E defends all its elements,
3. $E \in co(F)$ iff $E \in ad(F)$ and for any a defended by E we have, $a \in E$,
4. $E \in gr(F)$ iff E is $\subseteq$-minimal in $co(F)$, and
5. $E \in pr(F)$ iff E is $\subseteq$-maximal in $ad(F)$,
6. $E \in stb(F)$ iff $E^{\oplus}_F = A$,
7. $E \in ss(F)$ iff $E \in ad(F)$ and $E^{\oplus}_F$ is $\subseteq$-maximal in $\{D^{\oplus}_F \mid D \in ad(F)\}$,
8. $E \in stg(F)$ iff $E \in cf(F)$ and $E^{\oplus}_F$ is $\subseteq$-maximal in $\{D^{\oplus}_F \mid D \in cf(F)\}$.

Claim-based Argumentation. A *claim-augmented argumentation framework (CAF)* (Dvořák & Woltran, 2020) is a triple $\mathcal{F} = (A, R, cl)$ where $F = (A, R)$ is an AF. We call F the *underlying framework*. Similarly, we use G as underlying AF for a given CAF $\mathcal{G}$. The function $cl : A \to \mathcal{C}$ assigns a so-called *claim* to each argument; we assume that $\mathcal{C}$ is a countable infinite set of claims. The claim-function is extended to sets in the natural way. This means, for $E \subseteq A$ we set $cl(E) = \{cl(a) \mid a \in E\}$. A CAF $\mathcal{F}$ is called *well-formed* if $a^+_F = b^+_F$ for all $a, b \in A$ with $cl(a) = cl(b)$. For a claim c, we call an argument $x \in A$ with claim $cl(x) = c$ an *occurrence of c in $\mathcal{F}$*.

There are several ways in which semantics for AFs extend to CAFs. The most basic one is to choose an appropriate AF semantics and consider the claims of the induced extensions.

2. Note that Dung does not restrict AFs to the realm of finiteness. Properties in the unrestricted case have been studied by, e.g., Baumann and Spanring (2017).

In this way, CAFs inherit all semantics introduced for Dung AFs; we thus call this variant *inherited semantics.*

Definition 2.2. For a CAF $\mathcal{F} = (A, R, cl)$, $F = (A, R)$, and a semantics σ, we define the inherited variant of σ *(i-σ)* as $\sigma_i(\mathcal{F}) = \{cl(E) \mid E \in \sigma(F)\}$. We call $E \in \sigma(F)$ with $cl(E) = S$ a σ_i-realization of S in $\mathcal{F}$.

Example 2.3. Consider the following CAF $\mathcal{F}$:

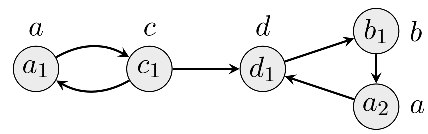

First, let us focus on stable semantics. For the underlying AF F we have the unique stable extension $E = \{c_1, b_1\}$ since c_1 attacks a_1 and d_1 and b_1 attacks a_2. Hence we obtain $stb_i(\mathcal{F}) = \{\{c, b\}\}$. Moreover, E is a stb_i-realization of $\{c, b\}$.

Coming to naive semantics, we obtain the argument-sets $\{a_1, d_1\}$, $\{a_1, b_1\}$, $\{a_1, a_2\}$, $\{c_1, b_1\}$, and $\{c_1, a_2\}$, yielding the i-naive claim-sets $\{a, d\}$, $\{a, b\}$, $\{a\}$, $\{c, b\}$, and $\{a, c\}$. Here, we experience an important difference to classical AF semantics: the claim-extensions form not necessarily an anti-chain, i.e., $\{a\}$ is a proper subset of $\{a, d\}$, $\{a, b\}$, and $\{a, c\}$.

Let us now turn to a family of semantics operating on the level of claims instead of focusing on the underlying arguments. In order to do so we need to generalize the notion of defeat to claims. A set of arguments $E \subseteq A$ *defeats* a claim $c \in cl(A)$ in $\mathcal{F}$ if E attacks every $a \in A$ with $cl(a) = c$ (in F). We use

$$E^{*}_{\mathcal{F}} = \{c \in cl(A) \mid E \text{ defeats } c \text{ in } \mathcal{F}\}$$

to denote the set of all claims which are defeated by E in $\mathcal{F}$. The claim-range of a set E of arguments is denoted by $E^{\circledast}_{\mathcal{F}} = cl(E) \cup E^{*}_{\mathcal{F}}$. If a singleton $\{x\}$ defeats a claim c, we simply write x defeats c. We say that E has *full claim-range* iff $E^{\circledast}_{\mathcal{F}} = cl(A)$.

Example 2.4 (Example 2.3 ctd.)**.** Consider again the CAF $\mathcal{F}$. Although c_1 attacks a_1, it does not defeat the claim a since it does not attack the argument a_2. However, $E = \{c_1, b_1\}$ defeats a, i.e. $a \in E^{*}_{\mathcal{F}}$. The claim-range of E is thus $E^{\circledast}_{\mathcal{F}} = \{a, b, c, d\}$.

Let us now turn to the CAF semantics which make direct use of claims. *Hybrid semantics* as introduced by Dvořák et al. (2020a) provide an alternative evaluation method for semantics based on maximization, e.g., for naive or preferred semantics, and for range-based semantics such as stable or semi-stable semantics.

Let us discuss the case for stable semantics. Following Dung's definition, a set is considered stable if each argument in the framework is either contained in or attacked by the set. Adapting this concept to claim-level we say a set is stable if each claim is either accepted or defeated by a given set of arguments. Now, consider again the CAF $\mathcal{F}$ from our running example. As observed above, the set $E = \{c_1, b_1\}$ has full claim-range, i.e., $E^{\circledast}_{\mathcal{F}} = \{a, b, c, d\}$. We say that $\{c, b\}$ is *hybrid stable* because it has a realization with full claim-range. We observe that there is another set of arguments with full claim-range: The set $D = \{a_1, d_1\}$

contains claims a and d and defeats claims c and b, that is, $D^{\circledast}_{\mathcal{F}} = \{a, b, c, d\}$ as well. Therefore, the claim-set $\{a, d\}$ is stable on the level of claims although it has no realization that is stable in the underlying AF. We note that the realization D with full claim-range is conflict-free but not admissible. Following this observation, we identify two different ways to lift stability to claim-level: the first variant requires that the underlying set of arguments is conflict-free; the second variant additionally requires admissibility of the realization.

Below, we define *h-preferred*, *h-naive*, *h-cf-stable*, *h-ad-stable*, *h-semi-stable* and *h-stage* semantics (abbr. pr_h, na_h, cf-stb_h, ad-stb_h, ss_h, stg_h).

Definition 2.5. Let $\mathcal{F} = (A, R, cl)$ be a CAF with underlying AF $F = (A, R)$. For a set of claims $S \subseteq cl(A)$,

- $S \in pr_h(\mathcal{F})$ iff S is $\subseteq$-maximal in $ad_i(\mathcal{F})$;
- $S \in na_h(\mathcal{F})$ iff S is $\subseteq$-maximal in $cf_i(\mathcal{F})$;
- $S \in \tau$-$stb_h(\mathcal{F})$, $\tau \in \{cf, ad\}$, iff there is a τ_i-realization E of S in $\mathcal{F}$ and $E^{\circledast}_{\mathcal{F}} = cl(A)$;
- $S \in ss_h(\mathcal{F})$ iff there is an ad_i-realization E of S in $\mathcal{F}$ which is $\subseteq$-maximal in $\{D^{\circledast}_{\mathcal{F}} \mid D \in ad(F)\}$;
- $S \in stg_h(\mathcal{F})$ iff there is an cf_i-realization E of S in $\mathcal{F}$ which is $\subseteq$-maximal in $\{D^{\circledast}_{\mathcal{F}} \mid D \in cf(F)\}$.

A set $E \subseteq A$ *σ_h-realizes* the claim-set S in $\mathcal{F}$ if $cl(E) = S$ and E satisfies the respective requirements; e.g., $E \in cf(F)$ and $E^{\circledast}_{\mathcal{F}} = cl(A)$ for h-cf-stable semantics. We call E a σ_h-realization of S in $\mathcal{F}$.

Remark 2.6. Originally, hybrid semantics have been introduced under the name *claim-level semantics* (Rapberger, 2020; Dvořák et al., 2020a). Inherited σ-claim-sets of a CAF $\mathcal{F}$ have been denoted by $\sigma_c(\mathcal{F})$ and the hybrid (former: claim-level) σ-claim-sets by cl-$\sigma(\mathcal{F})$. Following more recent publications (Dvořák et al., 2023; Rapberger, 2023), we adapted the wording and the notation: we use σ_i instead of σ_c to denote the inherited variant of the semantics σ, and σ_h instead of cl-σ to denote the hybrid variant, respectively.

Example 2.7. Consider the semantics cf-stb_h. We have that $S = \{c, b\} \in cf$-$stb_h(\mathcal{F})$ since the realization $E = \{c_1, b_1\}$ for S has full claim-range as we already observed before. Moreover, $S' = \{a, d\} \in cf$-$stb_h(\mathcal{F})$ as well: As observed above, the realization $D = \{a_1, d_1\}$ defeats the claims c and b and hence, $D^{\circledast}_{\mathcal{F}} = \{a, b, c, d\}$. Note that D is not a stable extension of the underlying AF.

Comparing i-naive and h-naive semantics, we observe that the claim-set $\{a\}$ is not h-naive since it is not $\subseteq$-maximal in $cf_i(\mathcal{F})$.

Basic relations between i-semantics carry over from AF semantics. For each CAF $\mathcal{F}$,

$$stb_i(\mathcal{F}) \subseteq ss_i(\mathcal{F}) \subseteq pr_i(\mathcal{F}) \subseteq co_i(\mathcal{F}) \subseteq ad_i(\mathcal{F}) \subseteq cf_i(\mathcal{F}),$$
$$stb_i(\mathcal{F}) \subseteq stg_i(\mathcal{F}) \subseteq na_i(\mathcal{F}) \subseteq cf_i(\mathcal{F});$$

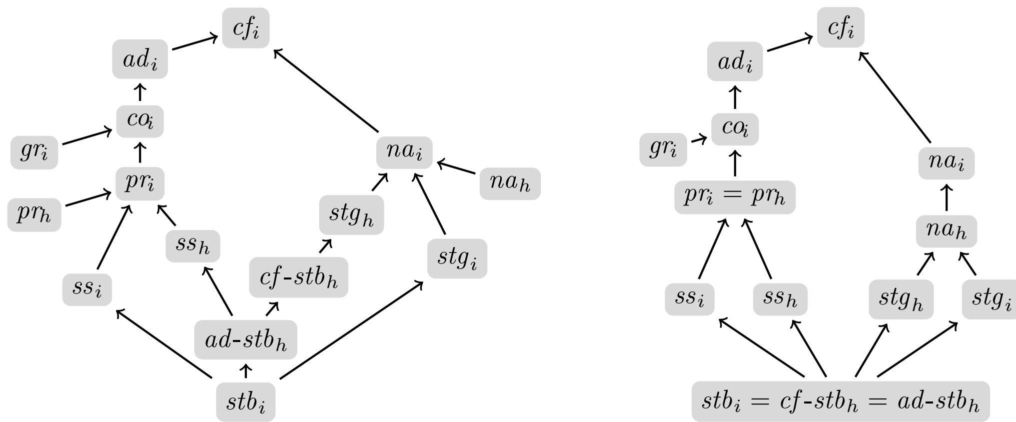

(a) Relations between CAF semantics. (b) Relations for well-formed CAFs.

Figure 1: Relations between semantics for general (a) and well-formed CAFs (b)(Dvořák et al., 2020a). An arrow from σ to τ indicates $\sigma(\mathcal{F}) \subseteq \tau(\mathcal{F})$ for each (well-formed) CAF $\mathcal{F}$.

moreover, $gr_i(\mathcal{F})$ is unique and contained in (the intersection of) $co_i(\mathcal{F})$. Furthermore, $co_i(\mathcal{F})$ forms a meet semilattice with respect to the subset relation. As shown by Dvořák et al. (2020a), it holds that

$$stb_i(\mathcal{F}) \subseteq ad\text{-}stb_h(\mathcal{F}) \subseteq cf\text{-}stb_h(\mathcal{F}) \subseteq stg_h(\mathcal{F}) \subseteq na_i(\mathcal{F}),$$
$$ad\text{-}stb_h(\mathcal{F}) \subseteq ss_h(\mathcal{F}) \subseteq pr_i(\mathcal{F})$$

for each $\mathcal{F}$. Moreover, for $\sigma \in \{pr, na\}$, each σ_h-claim-set of $\mathcal{F}$ is $\subseteq$-maximal in $\sigma_i(\mathcal{F})$, i.e.

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

For well-formed CAFs, the variants of preferred as well as the variants of stable semantics collapse (Dvořák et al., 2020a). That is, for every well-formed CAF $\mathcal{F}$, it holds that

$$stb_i(\mathcal{F}) = ad\text{-}stb_h(\mathcal{F}) = cf\text{-}stb_h(\mathcal{F}) \text{ and } pr_i(\mathcal{F}) = pr_h(\mathcal{F}).$$

Figure 1 gives an overview over the relations between the semantics.

Notation. We write $\mathcal{F} = (F, cl)$ as an abbreviation for $\mathcal{F} = (A, R, cl)$ with AF $F = (A, R)$ (similar for CAFs $\mathcal{G}$ or $\mathcal{H}$ for which we denote the corresponding AFs by G and H, respectively). Also, we use the subscript-notation $A_{\mathcal{F}}$, $R_{\mathcal{F}}$, and $cl_{\mathcal{F}}$ to refer to the arguments, attack relations, and claims of a given CAF $\mathcal{F}$.

3. Ordinary Equivalence

The distinction between explicit and implicit information is essential in knowledge representation. The former is interpreted according to the underlying semantics of the considered formalism, i.e. the set of models in case of classical propositional logic or the set of extensions in case of classical AFs. In contrast, the implicit information of an knowledge base comes to light if it undergoes dynamic changes. Both concepts come along with an induced notion of equivalence, namely *ordinary* or *strong equivalence*, respectively. We start our analysis by investigating ordinary equivalence for CAFs.

Definition 3.1. Two CAFs $\mathcal{F}$ and $\mathcal{G}$ are *ordinarily equivalent* w.r.t. semantics ρ, in symbols $\mathcal{F} \equiv_o^\rho \mathcal{G}$, if we have $\rho(\mathcal{F}) = \rho(\mathcal{G})$.

Example 3.2. Consider the following CAFs $\mathcal{F}$ and $\mathcal{G}$. Note that they disagree on the attack relation between a_1 and a_2 only.

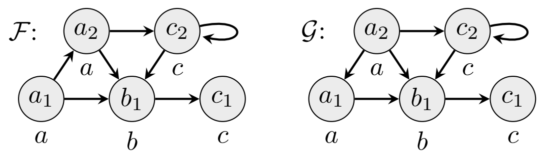

We have $stb(F) = \emptyset$ and $stb(G) = \{a_2, c_1\}$. Consequently, the inherited variants are $stb_i(\mathcal{F}) = \emptyset$ and $stb_i(\mathcal{G}) = \{a, c\}$ justifying $\mathcal{F} \not\equiv_o^{stb_i} \mathcal{G}$. If we consider instead the claim-based versions, we observe that the two CAFs agree on their outcome: More precisely, due to $stb_i(\mathcal{G}) \subseteq ad\text{-}stb_h(\mathcal{G}) \subseteq cf\text{-}stb_h(\mathcal{G})$ we obtain $\{a, c\} \in ad\text{-}stb_h(\mathcal{G}), cf\text{-}stb_h(\mathcal{G})$. Moreover, we have that $\{a, c\} \in ad\text{-}stb_h(\mathcal{F}), cf\text{-}stb_h(\mathcal{F})$ since the set $\{a_1, c_1\}$ is admissible (thus, conflict-free) and defeats every remaining claim. As a side remark, we mention that the claim-set $\{a, c\}$ has two cf_i-realizations in $\mathcal{F}$ and $\mathcal{G}$ since both of the sets $\{a_1, c_1\}$, $\{a_2, c_1\}$ are conflict-free and have full claim-range. It can be checked that no other claim-set than $\{a, c\}$ satisfies the requirements of the claim-based stable versions. Consequently, $\mathcal{F}$ and $\mathcal{G}$ are ordinarily equivalent with respect to $ad\text{-}stb_h$ and $cf\text{-}stb_h$ semantics, in symbols: $\mathcal{F} \equiv_o^{ad\text{-}stb_h} \mathcal{G}$ and $\mathcal{F} \equiv_o^{cf\text{-}stb_h} \mathcal{G}$.

In the following we consider (non-)relations between ordinary equivalences w.r.t. different semantics. We will see that the inherited variants behave differently in comparison to claim-based versions. Let us recap the case of Dung-style AFs. It was shown that sharing the same admissible/conflict-free sets guarantees no difference regarding preferred/naive extensions. Moreover, equivalence with respect to naive sets implies that the conflict-free sets coincide. Also, possessing the same complete extensions implies coinciding grounded and preferred extensions (Oikarinen & Woltran, 2011, Proposition 1).

Let us start with the relations between inherited semantics.

Proposition 3.3. *Consider two CAFs $\mathcal{F}$ and $\mathcal{G}$. It holds that*

1. $\mathcal{F} \equiv_o^{co_i} \mathcal{G} \Rightarrow \mathcal{F} \equiv_o^{gr_i} \mathcal{G}$,
2. $\mathcal{F} \equiv_o^{na_i} \mathcal{G} \Rightarrow \mathcal{F} \equiv_o^{cf_i} \mathcal{G}$.

Proof. The relations can be transferred from the case for the respective AF semantics. The first item is due to the fact that the i-grounded claim-set is, per definition, the $\subseteq$-minimal i-complete claim-set. For the second item, assume that $\mathcal{F}$ and $\mathcal{G}$ agree on their i-naive extensions, and let $S \in cf_i(\mathcal{F})$ be a i-conflict-free set in $\mathcal{F}$. Since it holds that each subset of every i-naive extension can be realized with a conflict-free set in both $\mathcal{F}$ and $\mathcal{G}$, we obtain that S is i-conflict-free in $\mathcal{G}$ as well. Hence we obtain $cf_i(\mathcal{F}) = cf_i(\mathcal{G})$. □

Interestingly, we observe that not all relations for AF semantics carry over to inherited semantics. This is due to the fact that i-preferred (i-naive) semantics are not necessarily $\subseteq$-maximal i-admissible (i-conflict-free) claim-sets. Let us consider the following example.

Example 3.4. Assume we are given two CAFs as follows:

We have $ad_i(\mathcal{F}) = ad_i(\mathcal{G}) = \{\emptyset, \{a\}, \{b\}, \{a,b\}\}$. Hence, $\{a,b\}$ is the unique i-preferred claim-set of $\mathcal{F}$. However, $pr_i(\mathcal{G}) = \{\{a\}, \{a,b\}\}$ witnessed by the extensions $\{a_1, a_2\}$ and $\{a_1, b_1\}$. Thus $\mathcal{F} \equiv_o^{ad_i} \mathcal{G} \not\Rightarrow \mathcal{F} \equiv_o^{pr_i} \mathcal{G}$. This also shows $\mathcal{F} \equiv_o^{cf_i} \mathcal{G} \not\Rightarrow \mathcal{F} \equiv_o^{na_i} \mathcal{G}$ since cf_i and ad_i as well as the respective variants of naive and preferred semantics coincide in $\mathcal{F}$ and $\mathcal{G}$.

Let us next consider relations between inherited and claim-based semantics. Overall, we observe that equivalence with respect to h-preferred semantics can be decided by looking either at i-admissible, i-complete, or i-preferred semantics. Moreover, coincidence of i-naive extension implies equivalence with respect to h-naive semantics. Also, inherited conflict-free sets coincide if and only if h-naive semantics yield the same claim-sets.

Proposition 3.5. *Consider two CAFs $\mathcal{F}$ and $\mathcal{G}$. It holds that*

1. $\mathcal{F} \equiv_o^{ad_i} \mathcal{G} \Rightarrow \mathcal{F} \equiv_o^{pr_h} \mathcal{G}$,
2. $\mathcal{F} \equiv_o^{pr_i} \mathcal{G} \Rightarrow \mathcal{F} \equiv_o^{pr_h} \mathcal{G}$,
3. $\mathcal{F} \equiv_o^{co_i} \mathcal{G} \Rightarrow \mathcal{F} \equiv_o^{pr_h} \mathcal{G}$,
4. $\mathcal{F} \equiv_o^{cf_i} \mathcal{G} \Leftrightarrow \mathcal{F} \equiv_o^{na_h} \mathcal{G}$,
5. $\mathcal{F} \equiv_o^{na_i} \mathcal{G} \Rightarrow \mathcal{F} \equiv_o^{na_h} \mathcal{G}$.

Proof. First, assume $\mathcal{F} \equiv_o^{ad_i} \mathcal{G}$. By definition, h-preferred extensions are the $\subseteq$-maximal i-admissible extensions, hence $\mathcal{F} \equiv_o^{pr_h} \mathcal{G}$ follows. Since h-preferred extensions coincide with the $\subseteq$-maximal i-preferred and i-complete claim-sets for each CAF, we obtain that $\mathcal{F} \equiv_o^{pr_i} \mathcal{G}$ and $\mathcal{F} \equiv_o^{co_i} \mathcal{G}$ imply $\mathcal{F} \equiv_o^{pr_h} \mathcal{G}$.

Let us next consider the relation between i-conflict-free and h-naive semantics. By definition, h-naive extensions are the $\subseteq$-maximal i-conflict-free extensions, hence we obtain $\mathcal{F} \equiv_o^{cf_i} \mathcal{G}$ implies $\mathcal{F} \equiv_o^{na_h} \mathcal{G}$. For the other direction, note that each subset of a h-naive extension has a conflict-free realization, hence the statement follows.

Finally, we note that h-naive extensions are precisely the $\subseteq$-maximal i-naive extensions, which implies the equivalence in the last item. $\square$

For well-formed CAFs, we obtain the following relations between ordinary equivalences.

Proposition 3.6. *For any two well-formed CAFs $\mathcal{F}$ and $\mathcal{G}$, it holds that*

- $\mathcal{F} \equiv_o^{\rho} \mathcal{G} \Rightarrow \mathcal{F} \equiv_o^{pr_i} \mathcal{G}$, $\rho \in \{ad_i, co_i, pr_h\}$;
- $\mathcal{F} \equiv_o^{stb_i} \mathcal{G} \Leftrightarrow \mathcal{F} \equiv_o^{ad\text{-}stb_h} \mathcal{G} \Leftrightarrow \mathcal{F} \equiv_o^{cf\text{-}stb_h} \mathcal{G}$.

Proof. The relations follow since the variants of preferred as well as the variants of stable semantics collapse for well-formed CAFs. □

Let us now turn to the non-relations between the semantics. Negative results (i.e., counter-examples) generalize to CAFs from the corresponding AF semantics.

Lemma 3.7. *For two AF semantics σ and τ, if $\sigma(F) = \sigma(G) \not\Rightarrow \tau(F) = \tau(G)$ for some AFs F, G, then $\sigma_c(\mathcal{F}) = \sigma_c(\mathcal{G}) \not\Rightarrow \tau_c(\mathcal{F}) = \tau_c(\mathcal{G})$ for some CAFs $\mathcal{F}$, $\mathcal{G}$.*

Indeed, when identifying AFs with CAFs where each claim is unique (i.e., taking $cl = id$), we obtain counter-examples from known results for AFs (Oikarinen & Woltran, 2011). We furthermore recall that in this case, hybrid semantics coincide with their inherited counterparts. Thus remains to provide counter-examples for naive, semi-stable, and stage semantics as well as for stable semantics in the general case.

For naive semantics, we observe that in both CAFs from Example 3.4, preferred and naive semantics coincide in both variants. To separate the stable variants, we consider the following examples.

Example 3.8. Consider the following CAFs $\mathcal{F}_1$, $\mathcal{F}_2$, and $\mathcal{G}$:

It holds that $stb_{hcf}(\mathcal{F}_1) = stb_{hcf}(\mathcal{G}) = \{\{a\}, \{b\}\}$ but $\rho(\mathcal{F}_1) \neq \rho(\mathcal{G})$ for $\rho \in \{stb_c, stb_{had}\}$; moreover, $\rho(\mathcal{F}_2) = \rho(\mathcal{G}) = \{\{a\}\}$ and $stb_{hcf}(\mathcal{F}_2) \neq stb_{hcf}(\mathcal{G})$.

Example 3.9. Consider the following CAFs $\mathcal{F}_1$, $\mathcal{F}_2$, and $\mathcal{G}$:

It holds that $stb_{had}(\mathcal{F}_1) = stb_{had}(\mathcal{G}) = \{\{a\}, \{b\}\}$ but $stb_c(\mathcal{F}_1) \neq stb_c(\mathcal{G})$; moreover, $stb_c(\mathcal{F}_2) = stb_c(\mathcal{G}) = \{\{a\}\}$ but $stb_{had}(\mathcal{F}_2) \neq stb_{had}(\mathcal{G})$.

It remains to consider semi-stable and stage semantics.

Example 3.10. Consider the following (well-formed) CAF $\mathcal{F}$:

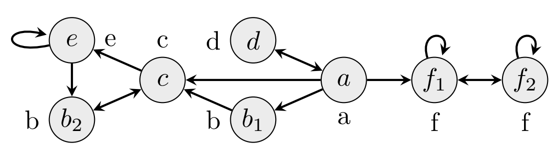

In $\mathcal{F}$, it holds that $ss_c(\mathcal{F}) = \{\{a\}\}$, $ss_h(\mathcal{F}) = \{\{b, d\}\}$, $stg_c(\mathcal{F}) = \{\{c\}, \{a\}\}$, and $stg_h(\mathcal{F}) = \{\{b, d\}, \{c\}\}$. To obtain counter-examples for the involved semantics, it suffices to construct a (well-formed) CAF $\mathcal{G}$ in which both variants agree on one of the aforementioned claim-sets of $\mathcal{F}$. First, let $\mathcal{G}_1 = (\{a\}, \emptyset, id)$, then all considered semantics return claim-set $\{a\}$. Thus $ss_c(\mathcal{F}) = ss_c(CG)$ but $ss_h(\mathcal{F}) \neq ss_h(\mathcal{G})$. Likewise, we let $\mathcal{G}_2 = (\{b, d\}, \emptyset, id)$ to obtain a counter-example for the other direction. For stage semantics, we consider the CAFs $\mathcal{G}_3 = (\{a, c\}, \{(a, c), (c, a)\}, id)$ and $\mathcal{G}_4 = (\{b, c, d\}, \{(b, c), (c, d), (d, c), (c, d)\}, id)$ instead.

This concludes our study of relations between semantics with respect to ordinary equivalence. We considered both general and well-formed CAFs. Similar as for AFs, we observe that ordinary equivalence for CAF semantics are largely independent of each other.

4. Strong Equivalence

A crucial observation is that ordinary equivalence is not robust when it comes to expansions of the frameworks, e.g., if an update in the knowledge base induces new arguments or attacks. Let us illustrate this by the following example:

Example 4.1. Assume we are given an updated version of $\mathcal{F}$ and $\mathcal{G}$ from Example 3.2 where an additional argument d_1 has been introduced. Let $\mathcal{F}'$ and $\mathcal{G}'$ be given as follows:

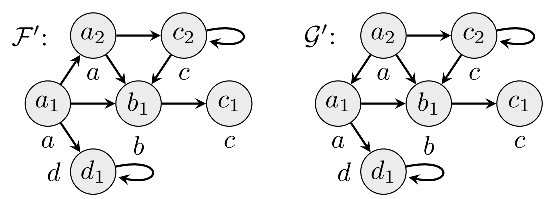

$\mathcal{F}'$ and $\mathcal{G}'$ no longer agree on their admissible-based h-stable claim-sets: In $\mathcal{G}'$, the set $\{a_2, c_1\}$ does not defeat claim d, thus $ad\text{-}stb_h(\mathcal{G}') = \emptyset$ while $\{a, c\}$ remains ad-based h-stable in $\mathcal{F}'$ due to $\{a_1, c_1\}$.

This is not due to CAF-specific properties, but a rather common behavior for many non-monotonic logics considered in the literature. Driven by this observation, *strong equivalence* has been introduced and investigated in other non-monotonic formalisms (Lifschitz et al., 2001; Oikarinen & Woltran, 2011; Baumann & Strass, 2022). Strong equivalence is a more restrictive notion which is tailored to handle equivalence between knowledge bases in a dynamical setting, i.e. the behavior remains the same even if we update our information.

Two AFs F and G are said to be *strongly equivalent* w.r.t. a semantics σ if and only if it holds that $\sigma(F \cup H) = \sigma(G \cup H)$ for each AF H (denoted by $F \equiv_s^\sigma G$). Judging this property merely by its definition, it appears to be computational hard at first glance: Naively iterating over each conceivable H only yields semi-decidability of checking that F and G are *not* strongly equivalent; if they are, the algorithm will not terminate.

Remarkably, this problem is tractable for all main AF semantics (Oikarinen & Woltran, 2011; Baumann, Linsbichler, & Woltran, 2016). It turned out that strong equivalence can be characterized via syntactical identity of so-called (semantics-dependent) *kernels*. These kernels are obtained by attack modifications of the given frameworks and are therefore straightforward to compute. Let us recall the definitions of the stable, admissible, complete, grounded, and naive kernel (Oikarinen & Woltran, 2011; Baumann et al., 2016).

Definition 4.2. For an AF $F = (A, R)$, we define the *stable kernel* $F^{sk} = (A, R^{sk})$; *admissible kernel* $F^{ak} = (A, R^{ak})$; the *complete kernel* $F^{ck} = (A, R^{ck})$; *grounded kernel*

$F^{gk} = (A, R^{gk})$; and the *naive kernel* $F^{nk} = (A, R^{nk})$ with

$$\begin{aligned}
R^{sk} &= R \setminus \{(a,b) \mid a \neq b, (a,a) \in R\} \\
R^{ak} &= R \setminus \{(a,b) \mid a \neq b, (a,a) \in R, \{(b,a),(b,b)\} \cap R \neq \emptyset\}; \\
R^{ck} &= R \setminus \{(a,b) \mid a \neq b, (a,a), (b,b) \in R\}; \\
R^{gk} &= R \setminus \{(a,b) \mid a \neq b, (b,b) \in R, \{(b,a),(a,a)\} \cap R \neq \emptyset\}; \\
R^{nk} &= R \cup \{(a,b) \mid a \neq b, \{(a,a),(b,b),(b,a)\} \cap R \neq \emptyset\}.
\end{aligned}$$

For a CAF $\mathcal{F} = (F, cl)$, we write $\mathcal{F}^k$ to denote (F^k, cl) for $k \in \{sk, ak, ck, gk, nk\}$.

We recall the characterization results of strong equivalence for AF semantics.

Theorem 4.3 ((Oikarinen & Woltran, 2011; Baumann et al., 2016)). *For any two AFs F and G,*

$$\begin{aligned}
F \equiv_s^\sigma G &\text{ iff } F^{sk} = G^{sk} \text{ for } \sigma \in \{stb, stg\}, \\
F \equiv_s^\sigma G &\text{ iff } F^{ak} = G^{ak} \text{ for } \sigma \in \{ad, pr, ss\} \\
F \equiv_s^{co} G &\text{ iff } F^{ck} = G^{ck} \\
F \equiv_s^{gr} G &\text{ iff } F^{gk} = G^{gk} \\
F \equiv_s^\sigma G &\text{ iff } F^{nk} = G^{nk} \text{ for } \sigma \in \{cf, na\}
\end{aligned}$$

For an AF F, we write $F^{k(\sigma)}$ to denote the kernel which characterizes strong equivalence for the semantics σ.

Example 4.4. Consider the following AFs F and G:

Since a_2 is a self-attacking argument, out-going attacks are removed when constructing the *stb*-kernel F^{sk}. Therefore F^{sk} and G^{sk} coincide.

Due to Theorem 4.3, we infer $F \equiv_s^{stb} G$ without checking a single candidate H by hand.

In the following, our goal is to establish analogous results for CAFs. We characterize strong equivalence for all considered CAF semantics by identifying appropriate kernels. In brief, our findings reveal that all semantics apart from *cf*-based h-stable semantics can be characterized with the kernels of their AF semantics counterpart. We identify a novel kernel for *cf*-based h-stable semantics, which exhibits interesting overlaps with the stable and the naive kernel for AF semantics.

4.1 Strong Equivalence in CAFs

Before formally introducing strong equivalence for CAFs we require an additional concept which ensures that expansions are well-defined.

Definition 4.5. Two CAFs $\mathcal{F}$ and $\mathcal{G}$ are *compatible* if $cl_{\mathcal{F}}(a) = cl_{\mathcal{G}}(a)$ for each $a \in A_{\mathcal{F}} \cap A_{\mathcal{G}}$.

Given two compatible CAFs we define the union $\mathcal{F} \cup \mathcal{G}$ as usual, namely componentwise as $\mathcal{F} \cup \mathcal{G} = (A_{\mathcal{F}} \cup A_{\mathcal{G}}, R_{\mathcal{F}} \cup R_{\mathcal{G}}, cl_{\mathcal{F}} \cup cl_{\mathcal{G}})$. We are now ready to define strong equivalence for CAFs.

Definition 4.6. Two CAFs $\mathcal{F}$ and $\mathcal{G}$ are *strongly equivalent* w.r.t. a semantics ρ, in symbols $\mathcal{F} \equiv_s^\rho \mathcal{G}$, iff

1. $\mathcal{F}$ and $\mathcal{G}$ are compatible; and
2. For each CAF $\mathcal{H}$ compatible with $\mathcal{F}$ and $\mathcal{G}$ we have, $\rho(\mathcal{F} \cup \mathcal{H}) = \rho(\mathcal{G} \cup \mathcal{H})$.

Our first step is to discuss some general observations that turn out to be useful when providing our characterization results. We will show that (i) two CAFs are strongly equivalent to each other only if they agree on their arguments; and (ii) strongly equivalent CAFs have the same self-attacking arguments.

We will first show that two CAFs with different arguments are not strongly equivalent.

Lemma 4.7. *For any two compatible CAFs $\mathcal{F}$ and $\mathcal{G}$, $A_{\mathcal{F}} \neq A_{\mathcal{G}}$ implies $\mathcal{F} \not\equiv_s^\rho \mathcal{G}$ for any considered semantics ρ.*

Proof. W.l.o.g., we may assume that there is $a \in A_{\mathcal{F}}$ with $a \notin A_{\mathcal{G}}$. Let $cl_{\mathcal{F}}(a) = c$. We distinguish the following cases: (a) $(a, a) \notin R_{\mathcal{F}}$ and (b) $(a, a) \in R_{\mathcal{F}}$.

- In case $(a, a) \notin R_{\mathcal{F}}$, we consider the following construction: For a fresh argument x and a fresh claim d, let $\mathcal{H} = (A_{\mathcal{H}}, R_{\mathcal{H}}, cl_{\mathcal{H}})$ with

$$A_{\mathcal{H}} = (A_{\mathcal{F}} \cup A_{\mathcal{G}} \cup \{x\}) \setminus \{a\};$$
$$R_{\mathcal{H}} = \{(x, b) \mid b \in (A_{\mathcal{F}} \cup A_{\mathcal{G}}) \setminus \{a\}\};$$

  and $cl_{\mathcal{H}}(b) = cl_{\mathcal{F}}(b)$ for $b \in A_{\mathcal{F}} \cup A_{\mathcal{G}}$ and $cl_{\mathcal{H}}(x) = d$; that is, we introduce a new argument having a fresh claim d which attacks every argument except a. Observe that $\{c, d\} \in gr_i(\mathcal{F} \cup \mathcal{H})$ and $\{c, d\} \in stb_i(\mathcal{F} \cup \mathcal{H})$ since $\{a, x\}$ is conflict-free, and x is unattacked and attacks all remaining arguments except a in $F \cup H$; thus there is $S \in \rho(\mathcal{F} \cup \mathcal{H})$ with $\{c, d\} \subseteq S$ for every semantics ρ under consideration by Lemma 4.9. On the other hand, $\{c, d\} \notin cf(\mathcal{G} \cup \mathcal{H})$ since x attacks every occurrence of $cl_{\mathcal{H}}(a)$ in $\mathcal{G}$; therefore, $\{c, d\} \notin \rho(\mathcal{G} \cup \mathcal{H})$.

- Now, let $(a, a) \in R_{\mathcal{F}}$. We construct our counter-example as follows: For a fresh argument x and a fresh claim d, let $\mathcal{H} = (A_{\mathcal{H}}, R_{\mathcal{H}}, cl_{\mathcal{H}})$ with

$$A_{\mathcal{H}} = A_{\mathcal{F}} \cup A_{\mathcal{G}} \cup \{x\};$$
$$R_{\mathcal{H}} = \{(x, b) \mid b \in (A_{\mathcal{F}} \cup A_{\mathcal{G}}) \setminus \{a\}\};$$

and $cl_{\mathcal{H}}(b) = cl_{\mathcal{F}}(b)$ for $b \in A_{\mathcal{F}}$, $cl_{\mathcal{H}}(b) = cl_{\mathcal{G}}(b)$ for $b \in A_{\mathcal{G}}$; and $cl_{\mathcal{H}}(x) = d$; that is, the new argument x attacks every argument in $A_{\mathcal{F}} \cup A_{\mathcal{G}}$ except a. Observe that a is unattacked in $\mathcal{G} \cup \mathcal{H}$ since a is a newly introduced argument in $\mathcal{G} \cup \mathcal{H}_1$ by assumption $a \notin A_{\mathcal{G}}$. Therefore $\{c, d\} \in gr_i(\mathcal{G} \cup \mathcal{H})$ since $\{a, x\}$ is conflict-free and unattacked; moreover, $\{c, d\} \in stb_i(\mathcal{G} \cup \mathcal{H})$ since $\{a, x\}$ is conflict-free and attacks all remaining arguments in $G \cup H$. By Lemma 4.9, $\{c, d\}$ is thus contained in some ρ-claim-set for every semantics ρ under consideration. On the other hand, $\{c, d\} \notin cf(\mathcal{F} \cup \mathcal{H})$ since every realisation of $\{c, d\}$ is conflicting: a is self-attacking and x attacks every other occurrence of c. Thus $\{c, d\} \notin \rho(\mathcal{F} \cup \mathcal{H}_{\mathcal{H}})$ for each considered semantics ρ.

In both cases, we found a witness $\mathcal{H}$ showing that $\rho(\mathcal{F} \cup \mathcal{H}) \neq \rho(\mathcal{G} \cup \mathcal{H})$. □

Next we show that two strongly equivalent CAFs $\mathcal{F}$ and $\mathcal{G}$ possess the same self-attackers.

Lemma 4.8. *For any two compatible CAFs $\mathcal{F}$ and $\mathcal{G}$, $(a, a) \in R_{\mathcal{F}} \Delta R_{\mathcal{G}}$ implies $\mathcal{F} \not\equiv^{\rho}_{s} \mathcal{G}$ for any semantics ρ under consideration.*

Proof. By Lemma 4.7, we may assume that $A_{\mathcal{F}} = A_{\mathcal{G}} (= A)$, i.e., a is contained in both CAFs $\mathcal{F}$ and $\mathcal{G}$. W.l.o.g., let $(a, a) \in R_{\mathcal{F}}$ and $(a, a) \notin R_{\mathcal{G}}$. Let $cl_{\mathcal{F}}(a) = cl_C G(a) = c$. Now, for a fresh argument x and fresh claim d, consider the CAF $\mathcal{H} = (A, R_{\mathcal{H}}, cl_{\mathcal{H}})$ with

$$R_{\mathcal{H}} = \{(x, b) \mid b \in A \setminus \{a\}\}$$

and $cl_{\mathcal{H}}(b) = cl_{\mathcal{F}}(b)$ for $b \in A$ and $cl_{\mathcal{H}}(x) = d$. Then $\{c, d\}$ has no cf-realisation in $\mathcal{F} \cup \mathcal{H}$ since a is self-attacking and x attacks every remaining occurrence of c in $\mathcal{F} \cup \mathcal{H}$. On the other hand, $\{c, d\} \in gr_i(\mathcal{G} \cup \mathcal{H})$ and $\{c, d\} \in stb_i(\mathcal{G} \cup \mathcal{H})$ since $\{a, x\}$ is conflict-free and attacks every other argument, moreover, x is unattacked. By Lemma 4.9, for all semantics ρ, there is $S \in \rho(\mathcal{G} \cup \mathcal{H})$ which contains $\{c, d\}$. Thus $\mathcal{F} \not\equiv^{\rho}_{s} \mathcal{G}$. □

Furthermore, the next auxiliary lemma will be convenient throughout the next subsections.

Lemma 4.9. *For a CAF $\mathcal{F}$ and a set of claims $S \subseteq cl(A)$, it holds that $S \subseteq S'$ for some $S' \in stb_i(\mathcal{F})$ implies that there is some $S'' \in \rho(\mathcal{F})$ with $S \subseteq S''$ for all semantics $\rho \neq gr_i$ under consideration.*

Proof. For all except h-preferred and h-naive semantics, the statement follows directly from known relations between semantics (since $stb_i(\mathcal{F}) \subseteq \rho(\mathcal{F})$ in this case). Let $\rho \in \{pr_h, na_h\}$ and consider some claim-set $S \subseteq cl(A)$ such that $S \subseteq S'$ for some $S' \in stb_i(\mathcal{F})$ $(\subseteq \tau(\mathcal{F})$ for $\tau \in \{pr_i, na_i\})$. Since h-preferred and h-naive claim-sets are precisely the $\subseteq$-maximal i-preferred resp. i-naive claim-sets, there is $T \in pr_h(\mathcal{F})$ $(T \in na_h(\mathcal{F}))$ with $T \subseteq S'$. □

4.2 Inherited Semantics

We start with discussing strong equivalence w.r.t. inherited semantics. The main result is that inherited semantics can be characterized via already known AF kernels. More precisely, a characterizing kernel for a specific AF semantics also serve for its inherited variant in the realm of CAFs. The following theorem expresses this intuition in a formal way. The remainder of this section collects propositions witnessing the truth of this claim.

Theorem 4.10. *For any two compatible CAFs $\mathcal{F}$ and $\mathcal{G}$, and each considered AF semantics σ, the following statements are equivalent:*

- $\mathcal{F} \equiv_s^{\sigma_i} \mathcal{G}$,
- $F \equiv_s^{\sigma} G$,
- $F^{k(\sigma)} = G^{k(\sigma)}$.

Recall that $F \equiv_s^{\sigma} G$ iff $F^{k(\sigma)} = G^{k(\sigma)}$ holds by former charaterization results (Oikarinen & Woltran, 2011; Baumann et al., 2016). Moreover, for two compatible CAFs $\mathcal{F}$ and $\mathcal{G}$ with strongly equivalent underlying AFs, i.e. $F \equiv_s^{\sigma} G$, we immediately deduce $\mathcal{F} \equiv_s^{\sigma_i} \mathcal{G}$ as $\sigma(F \cup H) = \sigma(G \cup H)$ implies $\sigma_i(\mathcal{F} \cup \mathcal{H}) = \sigma_i(\mathcal{G} \cup \mathcal{H})$ for any CAF $\mathcal{H}$ compatible with $\mathcal{F}$ and $\mathcal{G}$.

It remains to show that being strongly equivalent on CAF-level implies the syntactical identity of the associated kernels of the underlying AFs, i.e.$\mathcal{F} \equiv_s^{\sigma_i} \mathcal{G}$ implies $F^{k(\sigma)} = G^{k(\sigma)}$. We will show the contrapositive. Assume $F^{k(\sigma)} \neq G^{k(\sigma)}$. Applying Lemma 4.7 we may further assume that $A_{\mathcal{F}} = A_{\mathcal{G}} (= A)$. This means, $\mathcal{F}$ and $\mathcal{G}$ agree on their arguments. Consequently, there must be some attack $(a, b) \in R_{\mathcal{F}}^{k(\sigma)} \Delta R_{\mathcal{G}}^{k(\sigma)}$. Without loss of generality we assume $(a, b) \in R_{\mathcal{F}}^{k(\sigma)} \setminus R_{\mathcal{G}}^{k(\sigma)}$.

Let us discuss each kernel separately. We start with the most prominent one, the stable kernel. This kernel deletes outgoing attacks from self-defeating arguments.

Proposition 4.11. *Given two CAFs $\mathcal{F}$ and $\mathcal{G}$ satisfying $(a, a) \in R_{\mathcal{F}}$ iff $(a, a) \in R_{\mathcal{G}}$ and $A_{\mathcal{F}} = A_{\mathcal{G}} (= A)$. Then $(a, b) \in R_{\mathcal{F}}^{sk} \setminus R_{\mathcal{G}}^{sk}$ implies $\mathcal{F} \not\equiv_s^{\sigma_i} \mathcal{G}$ for $\sigma \in \{stb, stg\}$.*

Proof. Since $(a, b) \in R_{\mathcal{F}}^{sk}$, we conclude that a is not self-attacking in F (which implies $(a, a) \notin R_{\mathcal{G}}$ by Lemma 4.8). We construct our counter-example as follows: for fresh arguments x, y, z and fresh claims c, d, e, let $\mathcal{H} = (A \cup \{x, y, z\}, R, cl)$ with

$$R = \{(b, z)\} \cup \{(x, h) \mid h \in (A \cup \{y\}) \setminus \{a, b\}\} \cup \{(y, h) \mid h \in A \cup \{x, z\}\}$$

and $cl(h) = cl_{\mathcal{F}}(h)$ for $h \in A$, $cl(x) = c$, $cl(y) = d$, and $cl(z) = e$. First observe that $\{y\}$ is stable in both $stb(F^{sk} \cup H)$ and $stb(G^{sk} \cup H)$, thus $stb(F^{sk} \cup H) = stg(F^{sk} \cup H)$ and $stb(G^{sk} \cup H) = stg(G^{sk} \cup H)$. Moreover, $\{a, x, z\} \in stb(F^{sk} \cup H)$ since x attacks each remaining argument; thus $\{cl(a), c, e\} \in stb_i(\mathcal{F}^{sk} \cup \mathcal{H})$. On the other hand, $\{cl(a), c, e\}$ has no stb_i-realisation in $\mathcal{G}^{sk} \cup \mathcal{H}$ since $\{a, x, z\}$ does not attack b; every other realisation of $\{cl(a), c, e\}$ in $\mathcal{G}^{sk} \cup \mathcal{H}$ is conflicting. □

We proceed with the admissible kernel which serves as characterizing kernel for admissible, preferred as well as semi-stable semantics.

Proposition 4.12. *Given two CAFs $\mathcal{F}$ and $\mathcal{G}$ satisfying $(a, a) \in R_{\mathcal{F}}$ iff $(a, a) \in R_{\mathcal{G}}$ and $A_{\mathcal{F}} = A_{\mathcal{G}} (= A)$. Then $(a, b) \in R_{\mathcal{F}}^{ak} \setminus R_{\mathcal{G}}^{ak}$ implies $\mathcal{F} \not\equiv_s^{\sigma_i} \mathcal{G}$ for $\sigma \in \{ad, pr, ss\}$.*

Proof. Since $(a, b) \in R_{\mathcal{F}}^{ak}$, it holds that either (a) $(a, a) \notin R_{\mathcal{F}}^{ak}$; or (b) $(a, a) \in R_{\mathcal{F}}^{ak}$ and $\{(b, a), (b, b)\} \notin R_{\mathcal{F}}^{ak}$.

(a) In case $(a,a) \notin R_{\mathcal{F}}$, consider construction $\mathcal{H}$ from the proof of Proposition 4.11. Then $\{cl(a), c, e\} \in \sigma_i(\mathcal{F}^{ak} \cup \mathcal{H})$ since $\{cl(a), c, e\} \in stb_i(\mathcal{F}^{ak} \cup \mathcal{H})$; on the other hand, $\{cl(a), c, e\}$ has no ad-realisation in $\mathcal{G}^{ak} \cup \mathcal{H}_1$ since z is not defended against b; every other realisation of $\{cl(a), c, e\}$ in $\mathcal{G}^{ak} \cup \mathcal{H}_1$ is conflicting since z is attacked by b and x attacks every remaining argument.

(b) For a fresh argument x and a fresh claim c, let

$$\mathcal{H}_2 = (A \cup \{x\}, \{(x,h) \mid h \in A \setminus \{a,b\}\}, cl_2)$$

with $cl_2(h) = cl_{\mathcal{F}}(h)$ for $h \in A$ and $cl_2(x) = c$. Then $\{b, x\} \in ad(G^{ak} \cup H_2)$ since b is not attacked by a in G^{ak} and defended against any other potential attack by x; moreover, $\{b, x\}$ semi-stable in $G^{ak} \cup H_2$ since there is no other set $D \subseteq A \cup \{x\}$ with $x \in D^{\oplus}_{\mathcal{G}^{ak} \cup \mathcal{H}_2}$ (besides $\{x\}$ which is a proper subset of $\{b, x\}$). Thus $\{cl_2(b), c\} \in \sigma_i(\mathcal{G}^{ak} \cup \mathcal{H}_1)$. On the other hand, $\{b, x\} \notin ad(F^{ak} \cup H_2)$ since b is not defended against a in $F^{ak} \cup H_2$. Thus $\{cl_2(b), c\} \notin \sigma_i(\mathcal{F}^{ak} \cup \mathcal{H}_2)$. □

We now turn to complete semantics. In this case, only attacks between two self-defeating arguments are rendered redundant.

Proposition 4.13. *Given two CAFs $\mathcal{F}$ and $\mathcal{G}$ satisfying $(a,a) \in R_{\mathcal{F}}$ iff $(a,a) \in R_{\mathcal{G}}$ and $A_{\mathcal{F}} = A_{\mathcal{G}} (= A)$. Then $(a,b) \in R^{ck}_{\mathcal{F}} \setminus R^{ck}_{\mathcal{G}}$ implies $\mathcal{F} \not\equiv^{co_i}_{s} \mathcal{G}$.*

Proof. We have either $(a,a) \notin R^{ck}_{\mathcal{F}}$ or $(b,b) \notin R^{ck}_{\mathcal{F}}$. The case $(a,a) \notin R^{ak}_{\mathcal{F}}$ is analogous to the case (a) in the proof of Proposition 4.12. It remains to discuss the case $(b,b) \notin R^{ck}_{\mathcal{F}}$. For fresh arguments x, y and fresh claims c, d, let

$$\mathcal{H}_3 = (A \cup \{x,y\}, \{(y,a),(y,y)\} \cup \{(x,h) \mid h \in A \setminus \{a,b\}\}, cl_3)$$

with $cl_3(h) = cl_{\mathcal{F}}(h)$ for $h \in A$, $cl_3(x) = c$, $cl_3(y) = d$. Then $\{cl_3(b), c\} \in co_i(\mathcal{G}^{ck} \cup \mathcal{H}_3)$ since $\{b, x\}$ is conflict-free and x defends b against each attack; moreover, a is not defended by $\{b, x\}$ against y. On the other hand, $\{cl_3(b), c\} \notin co_i(\mathcal{F}^{ck} \cup \mathcal{H}_3)$ since the only conflict-free sets containing x are $\{b, x\}$, which is not defended against a; $\{x\}$, which does not realise $cl_3(b)$; and $\{a, x\}$, which is not defended against y (and a has potentially a different claim than b). □

The uniquely defined grounded extension is the most sceptical one among all complete extensions. It is therefore not surprising that grounded semantics offers more potential for redundancy than complete semantics. In fact, an attack (a,b) can also be deleted if both, a is self-defeating and b counterattacks a.

Proposition 4.14. *Given two CAFs $\mathcal{F}$ and $\mathcal{G}$ satisfying $(a,a) \in R_{\mathcal{F}}$ iff $(a,a) \in R_{\mathcal{G}}$ and $A_{\mathcal{F}} = A_{\mathcal{G}} (= A)$. Then $(a,b) \in R^{gk}_{\mathcal{F}} \setminus R^{gk}_{\mathcal{G}}$ implies $\mathcal{F} \not\equiv^{gr_i}_{s} \mathcal{G}$.*

Proof. Either (a) $(b,b) \in R^{gk}_{\mathcal{F}}$ and $\{(b,a),(a,a)\} \notin R^{gk}_{\mathcal{F}}$; or (b) $(b,b) \notin R^{gk}_{\mathcal{F}}$. Since (b) is analogous to the case considered in the proof in Proposition 4.13 where we constructed an expansion $\mathcal{H}_3$ yielding different i-grounded claim-sets in $\mathcal{F}^{gk} \cup \mathcal{H}_3$ and $\mathcal{G}^{gk} \cup \mathcal{H}_3$, it remains to discuss case $(b,b) \in R^{gk}_{\mathcal{F}}$. For fresh arguments x, y and fresh claims c, d, let

$$\mathcal{H}_4 = (A \cup \{x,y\}, \{(b,y)\} \cup \{(x,h) \mid h \in A \setminus \{a,b\}\}, cl_4)$$

with $cl_4(h) = cl_{\mathcal{F}}(h)$ for $h \in A$, $cl_4(x) = c$, $cl_4(y) = d$. Then x is unattacked and defends a in $\mathcal{F}^{gk} \cup \mathcal{H}_4$, which in turn defends y. Thus $\{cl_4(a), c, d\} \in gr_i(\mathcal{F}^{gk} \cup \mathcal{H}_4)$. On the other hand, we have $\{cl_4(a), c, e\} \notin gr_i(\mathcal{G}^{gk} \cup \mathcal{H}_4)$ since y is not defended against b. □

This concludes the proof for the semantics $\sigma \in \{stb, stg, ad, pr, ss, gr, co\}$: in every case, we found a witness $\mathcal{H}$ showing $\sigma_i\left(\mathcal{F}^{k(\sigma)} \cup \mathcal{H}\right) \neq \sigma_i\left(\mathcal{G}^{k(\sigma)} \cup \mathcal{H}\right)$. By Lemma 4.16, we get

$$\sigma_i(\mathcal{F} \cup \mathcal{H}) = \sigma_i((\mathcal{F} \cup \mathcal{H})^{k(\sigma)}) = \sigma_i(\mathcal{F}^{k(\sigma)} \cup \mathcal{H}) \neq \sigma_i(\mathcal{G}^{k(\sigma)} \cup \mathcal{H}) = \sigma_i((\mathcal{G} \cup \mathcal{H})^{k(\sigma)}) = \sigma_i(\mathcal{G} \cup \mathcal{H}).$$

Hence, $\mathcal{F} \not\equiv^s_{\sigma_i} \mathcal{G}$ is indeed shown.

Finally, let us consider conflict-free and naive semantics both characterized by the naive kernel. This kernel behaves differently than any other considered kernel as it is the only one that adds attacks instead of deleting them (Baumann, 2018, Section 4.2).

Proposition 4.15. *Given two CAFs $\mathcal{F}$ and $\mathcal{G}$ satisfying $(a, a) \in R_{\mathcal{F}}$ iff $(a, a) \in R_{\mathcal{G}}$ and $A_{\mathcal{F}} = A_{\mathcal{G}} (= A)$. Then $F^{nk} \neq G^{nk}$ implies $\mathcal{F} \not\equiv^{\sigma_i}_s \mathcal{G}$ for $\sigma \in \{cf, na\}$.*

Proof. For $\sigma \in \{cf, na\}$, first notice that we can assume $\sigma_i(\mathcal{F}) = \sigma_i(\mathcal{G})$ otherwise let $\mathcal{H} = (\emptyset, \emptyset, \emptyset)$; furthermore, we can assume $\sigma(F) \neq \sigma(G)$; otherwise consider instead $\mathcal{F} \cup \mathcal{H}$ and $\mathcal{G} \cup \mathcal{H}$ for a compatible CAF $\mathcal{H}$ with $\sigma_i(\mathcal{F} \cup \mathcal{H}) \neq \sigma_i(\mathcal{G} \cup \mathcal{H})$.

First consider the case that there is some $E \in \sigma(F)\Delta\sigma(G)$ such that E is not conflict-free in F (or G, respectively). W.l.o.g., let $E \in \sigma(F)$ such that E is subset-minimal among $\sigma(F)\Delta\sigma(G)$, i.e., there is no $E' \in \sigma(F)\Delta\sigma(G)$ with $E' \subsetneq E$; otherwise, exchange the roles of F and G. For a fresh argument x and a fresh claim c, let $\mathcal{H}_5 = (A \cup \{x\}, \{(x, b) \mid b \in A \setminus E, cl_5)$ with $cl_5(b) = cl_{\mathcal{F}}(b)$ for $b \in A$ and $cl_5(x) = c$. Then $cl_5(E) \cup \{c\} \in na(\mathcal{F} \cup \mathcal{H}_5)$ but $\{cl_5(E) \cup \{c\}$ has no cf-realisation in $\mathcal{G} \cup \mathcal{H}_5$ since every subset of E is conflicting and x attacks all remaining arguments, thus $cl_5(E) \cup \{c\} \notin \sigma_i(\mathcal{G} \cup \mathcal{H}_5)$. Observe that this suffices to conclude the proof for conflict-free semantics.

For naive semantics, assume that for all $E \in \sigma(F)\Delta\sigma(G)$, $E \in cf(F) \cap cf(G)$. We derive a contradiction: W.l.o.g., let $E \in \sigma(F)$ such that E is subset-minimal among $\sigma(F)\Delta\sigma(G)$. Since E is conflict-free in G, there is some $E' \in na(G)$ with $E \subseteq E'$. But then $E' \in cf(G)$ and thus $E \in cf(F)$ by assumption, contradiction to E being a subset-maximal conflict-free extension in F. We have shown $\mathcal{F} \not\equiv^{\sigma_i}_s \mathcal{G}$ for $\sigma \in \{cf, na\}$ concluding the proof. □

4.3 Hybrid Semantics and AF Kernels

In this section we consider hybrid semantics. In particular, we are interested in h-ad-stable, h-semi-stable, h-preferred, and h-naive semantics. We will show that strong equivalence with respect to the first three semantics can be characterized via the classical admissible kernel whereas h-naive semantics is characterized by the naive kernel for AFs.

We first consider h-ad-stable and h-semi-stable semantics. The following result show that that claim-extensions remain unchanged if turning to the kernelized versions of the initial CAFs.

Lemma 4.16. *For any CAF $\mathcal{F}$ and any semantics $\rho \in \{ad\text{-}stb_h, ss_h\}$, $\rho(\mathcal{F}) = \rho(\mathcal{F}^{ak})$.*

Proof. Let $\mathcal{F} = (A, R, cl)$. As shown by Oikarinen and Woltran (2011), it holds that $\sigma(F) = \sigma(F^a k)$ for $\sigma \in \{ad, pr, semi\}$. Hence, it follows that $\sigma_i(\mathcal{F}) = \sigma_i(\mathcal{F}^{ak})$; in particular, i-admissible semantics are preserved. It remains to prove that the range of every admissible set $E \subseteq A$ remains unchanged in F^{ak}. By definition, we delete only attacks (a, b) with $(a, a) \in R$ when constructing the kernel. Hence, it follows that $E^+_F = E^+_{F^{ak}}$ for every conflict-free (thus also for every admissible set) E. We obtain $E^*_{\mathcal{F}} = E^*_{\mathcal{F}^{ak}}$ and $E^{\circledast}_{\mathcal{F}} = E^{\circledast}_{\mathcal{F}^{ak}}$ for every admissible set. Therefore, h-ad-stable and h-semi-stable semantics are preserved. □

Next we will prove the central characterization theorem. Two compatible CAFs are strongly equivalent under h-ad-stable as well as h-semi-stable semantics iff their admissible kernels coincide.

Theorem 4.17. *For any two compatible CAFs $\mathcal{F}$ and $\mathcal{G}$,*

$$\mathcal{F} \equiv^{\rho}_{s} \mathcal{G} \text{ iff } F^{ak} = G^{ak} \text{ for } \rho \in \{ad\text{-}stb_h, ss_h\}.$$

Proof. First suppose $F^{ak} = G^{ak}$ and let $\mathcal{H}$ be a CAF compatible with $\mathcal{F}$, $\mathcal{G}$. By Lemma 4.16, and since $F \cup H = (F \cup H)^{ak}$ by known results for AF (Oikarinen & Woltran, 2011, Lemma 5), we obtain $\rho(\mathcal{F} \cup \mathcal{H}) = \rho((\mathcal{F} \cup \mathcal{H})^{ak}) = \rho((\mathcal{G} \cup \mathcal{H})^{ak}) = \rho(\mathcal{G} \cup \mathcal{H})$. Therefore, $\mathcal{F} \equiv^{\rho}_{s} \mathcal{G}$.

Now assume $\mathcal{F}^{ak} \neq \mathcal{G}^{ak}$. We may assume $\rho(\mathcal{F}^{ak}) = \rho(\mathcal{G}^{ak})$ by Lemma 4.16; also, $A_{\mathcal{F}} = A_{\mathcal{G}} (= A)$ by Lemma 4.7 and $\mathcal{F}$ and $\mathcal{G}$ contain the same self-attacks by Lemma 4.8. Thus there is $(a, b) \in R^{ak}_{\mathcal{F}} \Delta R^{ak}_{\mathcal{G}}$; w.l.o.g., let $(a, b) \in R^{ak}_{\mathcal{F}}$. We distinguish three cases: (a) $(a, a) \notin R_{\mathcal{F}^{ak}}$; (b) $(a, a) \in R_{\mathcal{F}^{ak}}$ and $cl(a) \neq cl(b)$; and (c) $(a, a) \in R_{\mathcal{F}^{ak}}$ and $cl(a) = cl(b)$.

(a) In case $(a, a) \notin R_{\mathcal{F}^{ak}}$, let $\mathcal{H}1 = (A \cup \{x, y\}, R_1, cl_1)$ with

$$R_1 = \{(b, y)\} \cup \{(x, h) \mid h \in A \setminus \{a, b\}\}$$

and $cl_1(h) = cl_{\mathcal{F}}(h)$ if $h \in A$ and $cl_1(x) = c$, $cl_1(y) = d$ for newly introduced arguments x, y and fresh claims c, d. Note that $\{a, x, y\} \in stb(F^{ak} \cup H_1)$ since a defends y against b and x attacks every remaining argument. Consequently, $\{cl_1(a), c, d\} \in stb_i(\mathcal{F}^{ak} \cup \mathcal{H}_1) \subseteq \rho(\mathcal{F}^{ak} \cup \mathcal{H}_1)$.

On the other hand, we have that $\{cl_1(a), c, d\}$ is not admissible in $\mathcal{G}^{ak} \cup \mathcal{H}_1$ since it has no ad-realisation in $G^{ak} \cup H_1$: Clearly, every candidate set must contain x, y, which are the only arguments having claims c, d. The only cf-realisation of $\{cl_1(a), c, d\}$ is $\{a, x, y\}$ since every other argument is attacked by x. Observe that y is not defended against b by $\{a, x, y\}$ in $G^{ak} \cup H_1$, thus $\{cl_1(a), c, d\} \notin \rho(\mathcal{G}^{ak} \cup \mathcal{H}_1)$.

(b) In case $(a, a) \in R_{\mathcal{F}^{ak}}$, $cl(a) \neq cl(b)$, let $\mathcal{H}_2 = (A \cup \{x\}, R_2, cl_2)$ with

$$R_2 = \{(x, h) \mid h \in A \setminus \{a, b\}\}$$

for a fresh argument x with $cl_2(h) = cl_{\mathcal{F}}(h)$ for $h \in A$ and $cl_2(x) = cl_{\mathcal{F}}(a)$. First observe that $(b, b) \notin R^{ak}_{\mathcal{F}}$ (and thus also not in $R^{ak}_{\mathcal{G}}$), otherwise $(a, b) \notin R^{ak}_{\mathcal{F}}$ by definition. Hence $E = \{b, x\}$ is admissible in $G^{ak} \cup H_2$ since a does not attack b and x attacks each remaining argument. Let $S = cl_2(E)$ and observe that $S \cup E^*_{\mathcal{G}^{ak} \cup \mathcal{H}_2} = S \cup cl_2(A \setminus \{a\}) = cl_2(A)$ since $cl_2(a) \in S$. Thus $S \in \rho(\mathcal{G}^{ak} \cup \mathcal{H}_2)$.

On the other hand, $S \notin ad_i(\mathcal{F}^{ak} \cup \mathcal{H}_2)$: Consider a cf-realisation D of S. In case $x \notin D$, we have that D is not defended against x in $F^{ak} \cup H_2$ since x attacks any potential realization of $cl_2(a)$ in F which is not self-attacking. Now assume $x \in D$, then also $b \in D$, since x attacks any other possible choice of $cl_2(b)$ in F. In this case we have that D is not defended against a in $G^{ak} \cup H_2$ and thus $S \notin ad_i(\mathcal{F}^{ak} \cup \mathcal{H}_2)$. It follows that $\rho(\mathcal{F}^{ak} \cup \mathcal{H}_2) \neq \rho(\mathcal{G}^{ak} \cup \mathcal{H}_2)$.

(c) Now assume $(a, a) \in R_{\mathcal{F}^{ak}}$ and $cl(a) = cl(b)$. Let $\mathcal{H}_3 = (A \cup \{x, y\}, R_3, cl_3)$ with

$$R_3 = \{(x, y), (y, x)\} \cup \{(y, h \mid h \in A \cup \{x\}\} \cup \{(x, h) \mid h \in A \setminus \{a, b\}\}$$

and $cl_3(h) = cl_{\mathcal{F}}(h)$ if $h \in A$ and $cl_3(x) = c$, $cl_3(y) = d$ for newly introduced arguments x, y and fresh claims c, d, that is, $\mathcal{H}_3$ coincides with the construction $\mathcal{H}_1$ from case (a) in the Proof of Theorem 4.24. The argument y guarantees that $ad\text{-}stb_h(\mathcal{F}^{ak} \cup \mathcal{H}_3) \neq \emptyset$ and $ad\text{-}stb_h(\mathcal{G}^{ak} \cup \mathcal{H}_3) \neq \emptyset$ since in both $\mathcal{F}^{ak} \cup \mathcal{H}_3$ and $\mathcal{G}^{ak} \cup \mathcal{H}_3$, the claim-set $\{d\}$ is i-stable. Moreover, we have that $\{cl_3(b), c\} \in ad\text{-}stb_h(G^{ak} \cup H_3)$ (and thus $\{cl_3(b), c\} \in ss_h(G^{ak} \cup H_3)$) since $\{b, x\}$ is conflict-free and defends itself in $G^{ak} \cup H_3$—recall that $(b, b), (a, b) \notin R^{ak}_{\mathcal{G}}$ and x attacks every remaining argument except a. Since $cl_3(a) = cl_3(b)$ it follows that $\{b, x\}$ has full claim-range. On the other hand, we have that $\{cl_3(b), c\}$ has no ad-realisation in $F^{ak} \cup H_3$: Clearly, each candidate must contain x which is the only argument having claim c. Thus $\{b, x\}$ is the only cf-realisation of $\{cl_3(b), c\}$ in $F^{ak} \cup H_3$. Observe that $\{b, x\}$ is not admissible since b is not defended against the attack from a. We obtain $\rho(\mathcal{F}^{ak} \cup \mathcal{H}_3) \neq \rho(\mathcal{G}^{ak} \cup \mathcal{H}_3)$.

In every case, we have found a witness $\mathcal{H}$ showing $\rho(\mathcal{F}^{ak} \cup \mathcal{H}) \neq \rho(\mathcal{G}^{ak} \cup \mathcal{H})$. By Lemma 4.16, we get $\rho(\mathcal{F} \cup \mathcal{H}) = \rho((\mathcal{F} \cup \mathcal{H})^{ak}) = \rho(\mathcal{F}^{ak} \cup \mathcal{H}) \neq \rho(\mathcal{G}^{ak} \cup \mathcal{H}) = \rho((\mathcal{G} \cup \mathcal{H})^{ak}) = \rho(\mathcal{G} \cup \mathcal{H})$. It follows that $\mathcal{F} \not\equiv^s_\rho \mathcal{G}$. □

We now show that deciding strong equivalence w.r.t. h-naive and h-preferred semantics coincides with deciding strong equivalence w.r.t. their inherited counterparts.

Theorem 4.18. *For any two compatible CAFs $\mathcal{F}$ and $\mathcal{G}$,*

$$\mathcal{F} \equiv^{\sigma_h}_s \mathcal{G} \textit{ iff } \mathcal{F} \equiv^{\sigma_i}_s \mathcal{G} \textit{ for } \sigma \in \{na, pr\}.$$

Proof. If $\mathcal{F} \equiv^{\sigma_i}_s \mathcal{G}$, then $\sigma_i(\mathcal{F} \cup \mathcal{H}) = \sigma_i(\mathcal{G} \cup \mathcal{H})$ for every compatible CAF $\mathcal{H}$. $\mathcal{F} \equiv^{\sigma_h}_s \mathcal{G}$ follows since $\sigma_h(\mathcal{F} \cup \mathcal{H})$ are the subset-maximal i-naive claim-sets of $\mathcal{F} \cup \mathcal{H}$ and, analogously, $\sigma_h(\mathcal{G} \cup \mathcal{H})$ are the subset-maximal i-naive claim-sets of $\mathcal{G} \cup \mathcal{H}$.

Now assume $\mathcal{F} \not\equiv^{\sigma_i}_s \mathcal{G}$ and let $\sigma = pr$ (the proof for $\sigma = na$ is analogous). We may assume $A_{\mathcal{F}} = A_{\mathcal{G}} (= A)$ (by Lemma 4.7); also, $pr_i(\mathcal{F}) \neq pr_i(\mathcal{G})$ (otherwise consider instead $\mathcal{F} \cup \mathcal{H}$ and $\mathcal{G} \cup \mathcal{H}$ for a compatible CAF $\mathcal{H}$ with $pr_i(\mathcal{F} \cup \mathcal{H}) \neq pr_i(\mathcal{G} \cup \mathcal{H})$). Hence $ad(F) \neq ad(G)$. Consider a $\subseteq$-minimal set $E \in ad(F) \Delta ad(G)$. W.l.o.g., let $E \in ad(F)$.

In case there is no $D \in ad(F) \cap ad(G)$ with $D \subsetneq E$, we consider the following construction: For a fresh argument x and a fresh claim c, let

$$\mathcal{H}_1 = ((A \cup \{x\}, \{(x, b) \mid b \in (A \setminus E\}, cl_1)$$

with $cl_1(b) = cl_{\mathcal{F}}(b)$ for $b \in A$ and $cl_1(x) = c$. Then $E \cup \{x\} \in ad(F \cup H)$ since $E \cup \{x\}$ is conflict-free and defends itself, thus $cl(E) \cup \{c\} \in ad_i(\mathcal{F} \cup \mathcal{H}_1)$. Also observe that there is no other admissible set D with $D \nsubseteq E \cup \{x\}$ which contains x, thus $cl(E) \cup \{x\}$ is a subset-maximal i-admissible set in $\mathcal{F} \cup \mathcal{H}_1$. On the other hand, $cl(E) \cup \{x\}$ has no ad-realisation in $\mathcal{G} \cup \mathcal{H}_1$ since no subset of E is admissible in G by minimality of E and x attacks every remaining argument. Thus $cl(E) \cup \{c\} \notin pr_h(\mathcal{G} \cup \mathcal{H}_1)$.

Observe that for naive semantics, this concludes the proof since by minimality of E, we can always find a conflict-free set E such that there is no $D \in cf(F) \cap cf(G)$ with $D \subsetneq E$.

In case of preferred semantics, we now assume that the assumption is not satisfied, i.e., there is $D \in ad(F) \cap ad(G)$ with $D \subsetneq E$. There is some $a \in E$ such that $a \notin D$ for any $D \in ad(F) \cap ad(G)$ with $D \subsetneq E$: Otherwise every argument $a \in E$ is contained in some admissible set $D \subsetneq E$, and thus $\bigcup\{D \in ad(G) \cap ad(F) \mid D \subsetneq E\} = E$, i.e., the union of all admissible sets contained in E coincides with E, which implies E is admissible in G, contradiction to the assumption. We consider the following construction: For fresh arguments x, y and fresh claims c, d, let

$$\mathcal{H}_2 = (A \cup \{x, y\}, \{(a, y)\} \cup \{(y, b) \mid b \in E\} \cup \{(x, b) \mid b \in (A \setminus E)\}, cl_2)$$

with $cl_1(b) = cl_{\mathcal{F}}(b)$ for $b \in A$, $cl_2(y) = d$ and $cl_2(x) = c$. First observe that there is no $D \subsetneq E$ such that $D \in ad(F \cup H_2)$ (or $D \in ad(G \cup H_2)$ by the choice of a: y attacks every argument $b \in E$ and a is the only argument which defends E against y. Similar as above, we conclude that $cl(E) \cup \{c\} \in pr_h(\mathcal{F} \cup \mathcal{H}_2)$ since E is admissible in $F \cup H_2$ and x attacks every remaining argument; on the other hand, $cl(E) \cup \{c\} \notin pr_h(\mathcal{G} \cup \mathcal{H}_2)$ since no subset D of E is admissible in G.

In every case, we have found a witness $\mathcal{H}$ such that $\sigma_h(\mathcal{F} \cup \mathcal{H}) \neq \sigma_h(\mathcal{G} \cup \mathcal{H})$. □

4.4 *cf*-based h-stable Semantics

In this section we will introduce a novel characterizing kernel. Let us consider the so-called *cf*-based h-stable semantics. It is not hard to see that outgoing attacks from self-attacking arguments can be semantically neutral removed (apart from the self-attack itself) as such an argument cannot be part of a $cf\text{-}stb_h$-realization E, and moreover, it is not necessary that E defends itself against such attacks.

While the removal of outgoing attacks from self-attacking arguments has been already observed in the context of Dung AFs as integral part of many kernels (and defines the stable kernel, cf. Definition 4.2), we observe a specific behavior regarding arguments with the same claims: Coming back to our CAFs $\mathcal{F}'$ and $\mathcal{G}'$ from Example 4.1, we recall that they yield the same h-*cf*-stable claim-sets even after the argument d_1 has been added. The reason is that the direction of the attack between the arguments a_1 and a_2 is irrelevant since both arguments possess the same claim a. Thus it suffices to include one of them in a h-*cf*-stable claim-set in case not both of them are attacked.

Inspired by these observations, we introduce the novel *cf-stable kernel* for CAFs which takes into account:

- *remove* all outgoing attacks (a, b) with $a \neq b$, if a is a self-attacking argument, and
- *add* attacks between different arguments a and b, if both carry the same claim.

Definition 4.19. For a CAF $\mathcal{F} = (A, R, cl)$, we define the *cf-stable kernel* $\mathcal{F}^{csk} = (A, R^{csk}, cl)$ with

$$R^{csk} = R \cup \{(a,b) \mid a \neq b, (a,a) \in R \vee (cl(a) = cl(b) \wedge \{(b,a),(b,b)\} \cap R \neq \emptyset)\}.$$

We denote the underlying AF (A, R^{csk}) by F^{csk}.

Remark 4.20. The *cf*-stable kernel can be seen as a combination of the stable and naive kernel for AFs. The claim-independent part stems from the stable kernel and the case where two arguments have the same claim relates to the naive kernel. In a nutshell, it is save to introduce an attack between different arguments a and b if a is self-attacking without changing stable semantics. This is because attacks of this form neither interfere with the conflict-free extensions of an AF nor change the range of a conflict-free set. In case two arguments have the same claim, it is irrelevant which of these arguments is included in an extension. It is thus save to introduce attacks between two arguments in case their union is conflicting.

Example 4.21. Consider again our previous CAF $\mathcal{F}$. Below we depict the stable kernel $\mathcal{F}^{sk}$, the naive kernel $\mathcal{F}^{nk}$, and the newly introduced *cf*-stable kernel $\mathcal{F}^{csk}$:

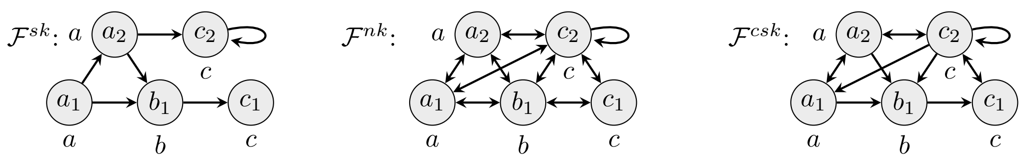

In what follows, we will prove that the *cf*-kernel characterizes strong equivalence for hybrid *cf*-stable and stage semantics. For this, we will first show that (i) a CAF admits the same h-*cf*-stable (h-stage) claim-sets as its *cf*-stable kernel and (ii) syntactical identity of the kernels implies that the kernels coincide under any possible expansion.

Lemma 4.22. *For any CAF $\mathcal{F}$ and any semantics $\rho \in \{cf\text{-}stb_h, stg_h\}$, $\rho(\mathcal{F}) = \rho(\mathcal{F}^{csk})$.*

Proof. We show (a) $cf(F) = cf(F^{csk})$ and (b) for all $E \in cf(F)$, $E^*_{\mathcal{F}} = E^*_{\mathcal{F}^{csk}}$.

To show (a), first observe that $cf(F^{csk}) \subseteq cf(F)$ since no new attacks between two unconflicting arguments are introduced. Moreover, we remove only attacks (a,b) where either a or b is self-attacking, thus we obtain $cf(F) \subseteq cf(F^{csk})$.

To show (b), let $E \in cf(F)$. It holds that $E^*_{\mathcal{F}} \subseteq E^*_{\mathcal{F}^{csk}}$. Now, let $c \in E^*_{\mathcal{F}^{csk}}$ and assume $c \notin E^*_{\mathcal{F}}$, i.e., there is $b \in A$ with $cl(b) = c$ which is not attacked by E in $\mathcal{F}$ but there is $a \in E$ such that $(a,b) \in R^{csk}$. Hence either $(a,a) \in R$ or $cl(a) = cl(b)$ and $(b,a) \in R$ or $(b,b) \in R$, contradiction to E being conflict-free in F^{csk}. □

Next, we show that syntactical identity of *cf*-stable kernels of two CAFs $\mathcal{F}$ and $\mathcal{G}$ implies that the kernels of $\mathcal{F} \cup \mathcal{H}$ and $\mathcal{G} \cup \mathcal{H}$ coincide for any compatible $\mathcal{H}$.

Lemma 4.23. *For any two compatible CAFs $\mathcal{F}$ and $\mathcal{G}$, $\mathcal{F}^{csk} = \mathcal{G}^{csk}$ implies $(\mathcal{F} \cup \mathcal{H})^{csk} = (\mathcal{G} \cup \mathcal{H})^{csk}$ for any CAF $\mathcal{H}$ compatible with $\mathcal{F}$ and $\mathcal{G}$.*

Proof. First observe that (i) $\mathcal{F} \cup \mathcal{H} \subseteq \mathcal{F}^{csk} \cup \mathcal{H}^{csk} \subseteq (\mathcal{F} \cup \mathcal{H})^{csk}$ holds for every two CAFs $\mathcal{F}$ and $\mathcal{H}$. Moreover, (ii) $\mathcal{F}^{csk} = \mathcal{G}^{csk}$ implies that $\mathcal{F}$, $\mathcal{G}$ contain the same self-attacks by definition of the *cf*-stable kernel.

Now, suppose $\mathcal{F}^{csk} = \mathcal{G}^{csk}$ and let $(a,b) \in (\mathcal{F} \cup \mathcal{H})^{csk}$. We show that $(a,b) \in (\mathcal{G} \cup \mathcal{H})^{csk}$ (the other direction is analogous): In case $(a,b) \in \mathcal{F} \cup \mathcal{H}$, we have $(a,b) \in \mathcal{F}^{csk} \cup \mathcal{H}^{csk}$ by (i). Since $\mathcal{F}^{csk} \cup \mathcal{H}^{csk} = \mathcal{G}^{csk} \cup \mathcal{H}^{csk}$ we conclude $(a,b) \in (\mathcal{G} \cup \mathcal{H})^{csk}$. In case $(a,b) \notin \mathcal{F} \cup \mathcal{H}$, either $(a,a) \in \mathcal{F} \cup \mathcal{H}$ or $cl(a) = cl(b)$ and $\{(b,b),(b,a)\} \cap (\mathcal{F} \cup \mathcal{H}) \neq \emptyset$. In case $(a,a) \in \mathcal{F} \cup \mathcal{H}$ ($(b,b) \in \mathcal{F} \cup \mathcal{H}$), we are done since $(a,a) \in \mathcal{G} \cup \mathcal{H}$ ($(b,b) \in \mathcal{G} \cup \mathcal{H}$) by (ii). Now, suppose $cl(a) = cl(b)$ and $(b,a) \in \mathcal{F} \cup \mathcal{H}$, then $(b,a) \in \mathcal{F}^{csk} \cup \mathcal{H}^{csk}$ by (i), thus also $(b,a) \in \mathcal{G}^{csk} \cup \mathcal{H}^{csk}$ by assumption $\mathcal{F}^{csk} = \mathcal{G}^{csk}$. In case $(b,a) \in \mathcal{G} \cup \mathcal{H}$, we get $(a,b) \in (\mathcal{G} \cup \mathcal{H})^{csk}$; else we have $cl(a) = cl(b)$ and $\{(a,a),(b,b),(a,b)\} \cap (\mathcal{G} \cup \mathcal{H}) \neq \emptyset$. By definition of the *cf*-stable kernel we obtain $(a,b) \in (\mathcal{G} \cup \mathcal{H})^{csk}$. □

We are now ready to prove our first main result stating that two CAFs $\mathcal{F}$ and $\mathcal{G}$ are strongly equivalent to each other w.r.t. h-*cf*-stable and h-stage semantics if and only if their h-stable kernels coincide. Let us sketch the idea.

First note that we obtain the 'if'-direction from Lemma 4.22 and 4.23: indeed, in case $\mathcal{F}^{csk} = \mathcal{G}^{csk}$ holds for two compatible CAFs $\mathcal{F}$ and $\mathcal{G}$, it holds that $(\mathcal{F} \cup \mathcal{H})^{csk} = (\mathcal{G} \cup \mathcal{H})^{csk}$ for any compatible CAF $\mathcal{H}$ by Lemma 4.23. From Lemma 4.22, we infer $\rho(\mathcal{F} \cup \mathcal{H}) = \rho((\mathcal{F} \cup \mathcal{H})^{csk})$ as well as $\rho((\mathcal{G} \cup \mathcal{H})^{csk}) = \rho(\mathcal{G} \cup \mathcal{H})$, hence we obtain $\rho(\mathcal{F} \cup \mathcal{H}) = \rho(\mathcal{G} \cup \mathcal{H})$.

For the 'only if'-direction, we will assume that the kernels disagree. By Lemma 4.7 and 4.8, it holds that $\mathcal{F}$ and $\mathcal{G}$ have the same arguments and in particular the same self-attackers. It thus remains to provide counter-examples for the case that the kernels of $\mathcal{F}$ and $\mathcal{G}$ disagree on an attack (a,b) for $a \neq b$. Figure 2 illustrates the counter-example for the case $cl(a) = cl(b)$ (case (b) in the proof of Theorem 4.24).

Theorem 4.24. *For any two compatible CAFs $\mathcal{F}$ and $\mathcal{G}$,*

$$\mathcal{F}^{csk} = \mathcal{G}^{csk} \text{ iff } \mathcal{F} \equiv_s^{\rho} \mathcal{G} \text{ for } \rho \in \{cf\text{-}stb_h, stg_h\}.$$

Proof. We obtain $\mathcal{F}^{csk} = \mathcal{G}^{csk}$ implies $\mathcal{F} \equiv_s^{\rho} \mathcal{G}$ from Lemma 4.22 and 4.23 as outlined above. It remains to prove the other direction. To do so, we suppose $\mathcal{F}^{csk} \neq \mathcal{G}^{csk}$. By Lemma 4.22 we may assume $\rho(\mathcal{F}^{csk}) = \rho(\mathcal{G}^{csk})$; and $A_{\mathcal{F}} = A_{\mathcal{G}} (= A)$ by Lemma 4.7. Thus it holds that $R_{\mathcal{F}^{csk}} \neq R_{\mathcal{G}^{csk}}$. W.l.o.g., let $(a,b) \in R_{\mathcal{F}^{csk}} \setminus R_{\mathcal{G}^{csk}}$; we have $a \neq b$ by Lemma 4.8. Moreover, $(a,a) \notin R_{\mathcal{G}}^{csk}$ (and thus, $(a,a) \notin R_{\mathcal{F}}^{csk}$), otherwise, $(a,b) \in R_{\mathcal{G}^{csk}}$ by definition. We distinguish the following cases: (a) $cl(a) \neq cl(b)$, and (b) $cl(a) = cl(b)$.

(a) In case $cl(a) \neq cl(b)$, consider two newly introduced arguments x, y and fresh claims c, d. We consider the AF $\mathcal{H}_1 = (A \cup \{x,y\}, R_1, cl_1)$ where

$$R_1 = \{(x,y)\} \cup \{(y,h) \mid h \in A \cup \{x\}\} \cup \{(x,h) \mid h \in A \setminus \{a,b\}\},$$

and the function cl_1 is given as follows: $cl_1(x) = c$, $cl_1(y) = d$, and the other claims coincide with the given ones, i.e. $cl_1(h) = cl_{\mathcal{F}}(h)$ if $h \in A$. First observe that $\{d\}$ is i-stable in both $\mathcal{F}^{csk} \cup \mathcal{H}_1$ and $\mathcal{G}^{csk} \cup \mathcal{H}_1$ and thus guarantees that $\rho(\mathcal{F}^{csk} \cup \mathcal{H}_1)$ and $\rho(\mathcal{G}^{csk} \cup \mathcal{H}_1)$ are non-empty. It can be checked that $S = \{cl(a), c\}$ is h-*cf*-stable and h-stage in $\mathcal{F}^{csk} \cup \mathcal{H}_1$ (since $\{a,x\}$ is stable); on the other hand, $S \notin \rho(\mathcal{G}^{csk} \cup \mathcal{H}_1)$ since b is not defeated by $\{a,x\}$. However, this is our only candidate since S has no other *cf*-realization in $\mathcal{G}^{csk} \cup \mathcal{H}_1$.

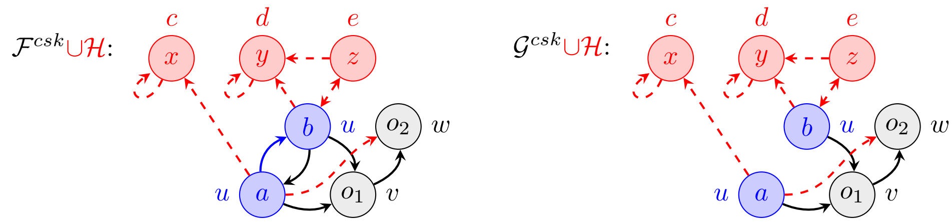

Figure 2: Counter-example for the case $(a, b) \in R_{\mathcal{F}^{csk}} \setminus R_{\mathcal{G}^{csk}}$ (case (b) in the proof of Theorem 4.24). New arguments introduced by $\mathcal{H}$ are $\{x, y, z\}$ (in red), new attacks are dashed (and red). The claim-set $\{u\}$ is h-*cf*-stable (h-stage) in $\mathcal{G}^{csk} \cup \mathcal{H}$ (since the set $\{a, b\}$ is stable in the underlying AF) but not in $\mathcal{F}^{csk} \cup \mathcal{H}$.

(b) Now consider the case $cl(a) = cl(b)$ and observe that $(a, a), (b, b), (b, a) \notin R^{csk}_{\mathcal{G}}$ (otherwise $(a, b) \in R_{\mathcal{G}^{csk}}$). Since $\mathcal{F}$ and $\mathcal{G}$ contain the same self-attacks, we furthermore have $(a, a), (b, b) \notin R_{\mathcal{F}^{csk}}$. Having established this situation let us construct $\mathcal{H}_2$ as follows: For fresh arguments x, y, z and fresh claims c, d, e, we consider $\mathcal{H}_2 = (A \cup \{x, y, z\}, R_2, cl_2)$ where

$$R_2 = \{(a, h) \mid h \in (A \cup \{x\}) \setminus \{a, b\}\} \cup \{(x, x), (b, y), (y, y), (z, b), (b, z), (z, y)\}$$

and as before we let $cl_2(h) = cl_{\mathcal{F}}(h)$ for $h \in A$; for the fresh arguments let $cl_2(x) = c$, $cl_2(y) = d$, as well as $cl_2(z) = e$. It can be checked that each CAF admits a stable extension; thus it suffices to show that the h-*cf*-stable claim-sets disagree. First observe that we now have $\{cl_2(a)\} \in \rho(\mathcal{G}^{csk} \cup \mathcal{H}_2)$ since $\{a, b\}$ is a stable extension in $G^{csk} \cup H_2$. On the other hand, we have that $\{cl_2(a)\}$ is neither cf-stb_h-realizable nor stg_h-realizable in $\mathcal{F}^{csk} \cup \mathcal{H}_2$. Figure 2 illustrates the construction.

In every case, we have found some $\mathcal{H}$ showing $\rho(\mathcal{F}^{csk} \cup \mathcal{H}) \neq \rho(\mathcal{G}^{csk} \cup \mathcal{H})$. By Lemma 4.22, we get $\rho(\mathcal{F} \cup \mathcal{H}) = \rho((\mathcal{F} \cup \mathcal{H})^{csk}) = \rho(\mathcal{F}^{csk} \cup \mathcal{H}) \neq \rho(\mathcal{G}^{csk} \cup \mathcal{H}) = \rho((\mathcal{G} \cup \mathcal{H})^{csk}) = \rho(\mathcal{G} \cup \mathcal{H})$. Thus it holds that $\mathcal{F} \not\equiv^s_\rho \mathcal{G}$. □

4.5 Strong Equivalence for Well-formed CAFs

We end this section with a brief discussion on strong equivalence for well-formed CAFs, i.e. the two input CAFs $\mathcal{F}$ and $\mathcal{G}$ are well-formed but no further restriction is imposed on the expansion $\mathcal{H}$. Recall that $\mathcal{F}$ is well-formed if $a^+_F = b^+_F$ for all $a, b \in A$ with $cl(a) = cl(b)$.

Although the variants of stable and preferred semantics coincide for well-formed CAFs, we observe that this is in general not the case when investigating strong equivalence.

Example 4.25. Consider the following two well-formed CAFs $\mathcal{F}$ and $\mathcal{G}$ depicted below:

The set $\{a, b\}$ is stable in both CAFs; also, $\mathcal{F}^{sk} = \mathcal{G} = \mathcal{G}^{sk}$ hence $\mathcal{F} \equiv^{stb_i}_s \mathcal{G}$ by Theorem 4.10. However, if we add a novel argument x with claim a that attacks a_1, we have that $\{a, b\}$ is h-ad-stable in the expansion of $\mathcal{F}$ (witnessed by $\{x, b_1\}$) but $\{a, b\}$ is not even admissible in $\mathcal{G} \cup \{x\}$ (we already used this construction in the proof of Theorem 4.17).

Interestingly, we observe a close correspondence of i-stable and h-cf-stable semantics.

Proposition 4.26. *$F^{csk} = G^{csk}$ iff $F^{sk} = G^{sk}$ for every two well-formed CAFs $\mathcal{F}$, $\mathcal{G}$.*

Proof. First note that, for each well-formed CAF, the set $\{a, b\}$ with $cl(a) = cl(b)$ is conflicting iff $(a, a) \in R$ or $(b, b) \in R$. By Lemma 4.7 and 4.8, $\mathcal{F}$ and $\mathcal{G}$ have the same (self-attacking) arguments. Hence $(a, b) \in R_{\mathcal{F}}$ iff $(b, b) \in R_{\mathcal{F}}$ iff $(b, b) \in R_{\mathcal{G}}$ iff $(a, b) \in R_{\mathcal{G}}$ for all $a, b \in A_C F$ with $cl(a) = cl(b)$. Hence if $\mathcal{F}$ and $\mathcal{G}$ agree on their cf-stable kernels then the restriction to arguments with the same claims yields identical graphs. Since the stable and the cf-stable kernel both delete the same attacks between arguments not having the same claim, the statement follows. □

It follows that two well-formed CAFs are strongly equivalent w.r.t. inherited stable semantics iff they are strongly equivalent w.r.t. h-cf-stable semantics.

Corollary 4.27. *$\mathcal{F} \equiv_s^{stb_i} \mathcal{G}$ iff $\mathcal{F} \equiv_s^{cf\text{-}stb_h} \mathcal{G}$ for every two well-formed CAFs $\mathcal{F}$, $\mathcal{G}$.*

4.6 Summary

In this section, we have established strong equivalence characterizations for all CAF semantics that have been considered in the literature so far. We have shown that most of the semantics can be characterized by known kernels. We present a novel kernel to characterize strong equivalence for h-cf-stable and h-stage semantics.

Our results can be summarized as follows:

Theorem 4.28. *For any two CAFs $\mathcal{F}$ and $\mathcal{G}$,*

$$\begin{aligned}
\mathcal{F} \equiv_s^{\rho} \mathcal{G} &\text{ iff } \mathcal{F}^{sk} = \mathcal{G}^{sk} \text{ for } \rho \in \{stb_i, stg_i\} \text{ (Thm 4.10)}\\
\mathcal{F} \equiv_s^{\rho} \mathcal{G} &\text{ iff } \mathcal{F}^{csk} = \mathcal{G}^{csk} \text{ for } \sigma \in \{cf\text{-}stb_h, stg_h\} \text{ (Thm 4.24)}\\
\mathcal{F} \equiv_s^{\rho} \mathcal{G} &\text{ iff } \mathcal{F}^{ak} = \mathcal{G}^{ak} \text{ for } \sigma \in \{ad_i, pr_i, ss_i, pr_h, ad\text{-}stb_h, ss_h\} \text{ (Thms 4.10, 4.17, 4.18)}\\
\mathcal{F} \equiv_s^{co_i} \mathcal{G} &\text{ iff } \mathcal{F}^{ck} = \mathcal{G}^{ck} \text{ (Thm 4.10)}\\
\mathcal{F} \equiv_s^{gr_i} \mathcal{G} &\text{ iff } \mathcal{F}^{gk} = \mathcal{G}^{gk} \text{ (Thm 4.10)}\\
\mathcal{F} \equiv_s^{\rho} \mathcal{G} &\text{ iff } \mathcal{F}^{nk} = \mathcal{G}^{nk} \text{ for } \sigma \in \{cf_i, na_i, na_h\} \text{ (Thms 4.10, 4.18)}
\end{aligned}$$

For well-formed CAFs, we additionally have shown that the stable and the cf-stable kernel coincide (cf. Proposition 4.26). We obtain the following results for well-formed CAFs:

Theorem 4.29. *For any two CAFs $\mathcal{F}$ and $\mathcal{G}$,*

$$\begin{aligned}
\mathcal{F} \equiv_s^{\rho} \mathcal{G} &\text{ iff } \mathcal{F}^{sk} = \mathcal{G}^{sk} \text{ for } \rho \in \{stb_i, stg_i, cf\text{-}stb_h, stg_h\}\\
\mathcal{F} \equiv_s^{\rho} \mathcal{G} &\text{ iff } \mathcal{F}^{ak} = \mathcal{G}^{ak} \text{ for } \sigma \in \{ad_i, pr_i, ss_i, pr_h, ad\text{-}stb_h, ss_h\}\\
\mathcal{F} \equiv_s^{co_i} \mathcal{G} &\text{ iff } \mathcal{F}^{ck} = \mathcal{G}^{ck}\\
\mathcal{F} \equiv_s^{gr_i} \mathcal{G} &\text{ iff } \mathcal{F}^{gk} = \mathcal{G}^{gk}\\
\mathcal{F} \equiv_s^{\rho} \mathcal{G} &\text{ iff } \mathcal{F}^{nk} = \mathcal{G}^{nk} \text{ for } \sigma \in \{cf_i, na_i, na_h\}
\end{aligned}$$

5. Renaming and Equivalence

The equivalence notions we investigated so far were operating on the given arguments together with their claims. However, as we already mentioned in the introduction, a key motivation behind CAFs is the investigation of claim-based reasoning. It therefore makes sense to consider an equivalence notion which abstracts from the underlying arguments and thus focuses on the claims and their relationships only. Let us consider the following illustrative example.

Example 5.1. Assume we are given two CAFs $\mathcal{F}$ and $\mathcal{G}$ (cf. Example 3.2) which both stem from instantiating the same knowledge base using different argument naming schemes – the CAF $\mathcal{F}$ relates argument names with the corresponding claim (e.g., arguments with claim a are named a_i) while $\mathcal{G}$ uses a consecutive numbering for all arguments:

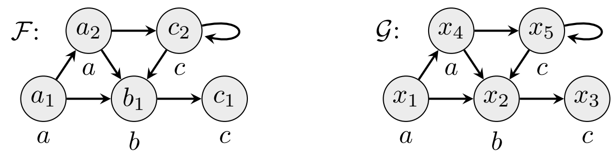

It is evident that $\mathcal{F}$ and $\mathcal{G}$ are ordinary equivalent w.r.t. all considered semantics despite the mismatch in argument names. However, when we consider equivalence in a dynamic setting, we observe that different argument naming patterns can cause undesired effects. To illustrate this let us suppose we are given $\mathcal{H}$ in a way that a novel argument e_1 with claim e is given which attacks x_1 and x_4:

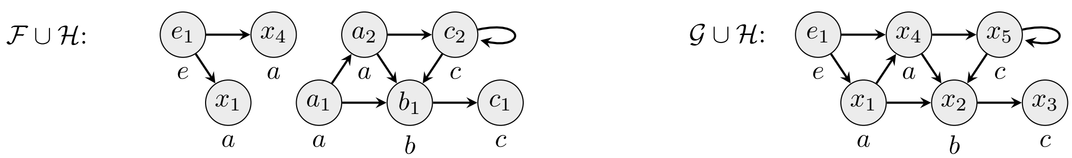

Evidently, the structure of $\mathcal{F} \cup \mathcal{H}$ and $\mathcal{G} \cup \mathcal{H}$ do not match (although this was the case for $\mathcal{F}$ and $\mathcal{G}$). This has several issues: First, from an intuitive point of view, it does not make much sense to disrupt the similarity between $\mathcal{F}$ and $\mathcal{G}$ that harshly; second, the modified $\mathcal{F} \cup \mathcal{H}$ and $\mathcal{G} \cup \mathcal{H}$ do not correspond to the same modification of the underlying knowledge base anymore (hence the disruption); third, CAFs are designed to reason with the claims and not the underlying arguments and thus, if $\mathcal{F}$ and $\mathcal{G}$ agree on the interaction of the claims (as they do), our expansion notion should preserve this property.

The technical problem revealed by this example is that $\mathcal{F}$ and $\mathcal{G}$ are not considered strongly equivalent: Indeed, a is accepted in $\mathcal{F} \cup \mathcal{H}$, but not in $\mathcal{G} \cup \mathcal{H}$. The goal of this section is to overcome this issue and tailor our strong equivalence notion suitably for claim-based reasoning.

5.1 Basic Notions

As the above example suggests, the usual notion of strong equivalence does not handle situations where we are interested in claims only very well. Preferably, we would like to have a situation where argument names are ignored.

Example 5.2 (Example 5.1 ctd.)**.** Let us consider our running example in which $\mathcal{F}$ and CG both stem from instantiating the same knowledge base. We want to abstract away from the names of the arguments:

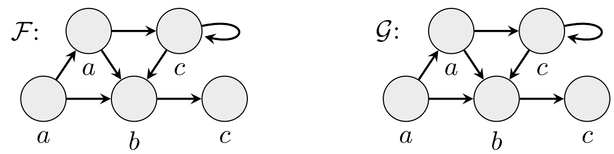

Since we formally work with CAFs, we require some technical machinery in order to take arguments names out of the equation. Instead of removing them, we allow for changing argument names arbitrarily. Formally, this idea yields the notion of a renaming.

Definition 5.3. For a CAF $\mathcal{F}$ and a set A' of arguments we call a bijective mapping $f : A_\mathcal{F} \to A'$ a *renaming for* $\mathcal{F}$. By $f(\mathcal{F})$ we denote the induced CAF (A_f, R_f, cl_f) where

- $A_f = A'$,
- $R_f = \{(a', b') \in A' \times A' \mid \left(f^{-1}(a'), f^{-1}(b')\right) \in R_\mathcal{F}\}$
- $cl_f(a') = cl_\mathcal{F}\left(f^{-1}(a')\right)$

Due to the required bijection the latter both conditions can be reformulated in a more eye-catching way as

- $(a, b) \in R_\mathcal{F}$ iff $(f(a), f(b)) \in R_f$ and
- $cl_\mathcal{F}(a) = cl_f(f(a))$.

Example 5.4. Consider again our previous CAF $\mathcal{F}$. Let us assume we are given $A' = \{x_1, x_2, y_1, z_1, z_2\}$. The renaming f with $a_i \mapsto x_i$, $b_1 \mapsto y_1$ and $c_i \mapsto z_i$ induces the following CAF $f(\mathcal{F})$:

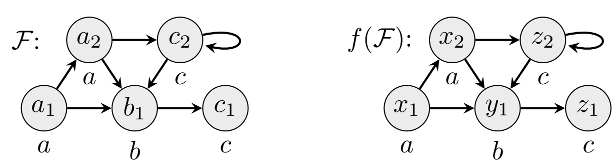

We observe that f does not change the structure of $\mathcal{F}$ on the claim-level. In particular, we observe that $\rho(\mathcal{F}) = \rho(f(\mathcal{F}))$ for all considered semantics ρ.

The last observation we made was no coincidence in the specific situation. More precisely, for the semantics considered in this paper, renaming does not change the meaning of our CAF.

Proposition 5.5. *For any CAF* $\mathcal{F}$ *and any renaming* f *of it, we have:* $\rho(\mathcal{F}) = \rho(f(\mathcal{F}))$ *for any semantics* ρ *considered in this paper.*

For convenience, we call this property *renaming robust*, i.e. ρ is renaming robust if for each CAF $\mathcal{F}$ and any renaming f of $\mathcal{F}$ we have that $\rho(\mathcal{F}) = \rho(f(\mathcal{F}))$. With these observations in mind, we can now proceed as follows when comparing two CAFs: We synchronize the argument names in order to make the CAFs better comparable. Due to Proposition 5.5 we can be certain that this procedure does not alter the meaning of the CAFs. The resulting (renamed) CAFs can then be checked for (strong) equivalence.

Example 5.6. Recall the CAFs $\mathcal{F}$ and $\mathcal{G}$:

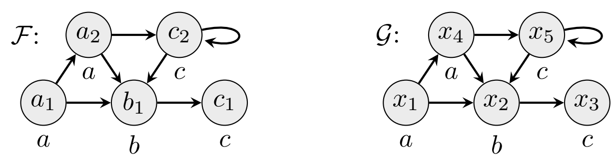

Indeed, when applying the renaming f with $a_1 \mapsto x_1$, $a_2 \mapsto x_4$, $b_1 \mapsto x_2$, $c_1 \mapsto x_3$, $c_2 \mapsto x_5$ we get two identical CAFs. It is therefore easy to see that $f(\mathcal{F})$ and $\mathcal{G}$ are (strongly) equivalent.

This more desirable notion is formalized in the following definition of equivalence up to renaming.

Definition 5.7. Two CAFs $\mathcal{F}$ and $\mathcal{G}$ are *ordinary equivalent up to renaming* for semantics ρ, in symbols $\mathcal{F} \equiv^{\rho}_{or} \mathcal{G}$, iff there are renamings f and g for $\mathcal{F}$ and $\mathcal{G}$, s.t. $f(\mathcal{F}) \equiv^{\rho}_{o} g(\mathcal{G})$.

So, informally speaking, Definition 5.7 requires that $\mathcal{F}$ and $\mathcal{G}$ are equivalent, at least after the underlying arguments are relabeled in a suitable way. However, in Proposition 5.5 we have actually already established that this adjustment is superfluous for our semantics. Consequently, we infer the following result.

Proposition 5.8. *For two CAFs $\mathcal{F}$ and $\mathcal{G}$, and any considered semantics ρ we have: $\mathcal{F} \equiv^{\rho}_{or} \mathcal{G}$ iff $\mathcal{F} \equiv^{\rho}_{o} \mathcal{G}$.*

5.2 Strong Equivalence up to Renaming

Now we utilize the notion of a renaming in order to define an appropriated notion of strong equivalence. We thereby proceed as in Definition 5.7. This means, we first allow for renaming the given CAFs and then, secondly, using the modified versions, we check for the desired property.

Definition 5.9. Two CAFs $\mathcal{F}$ and $\mathcal{G}$ are *strongly equivalent up to renaming* for semantics ρ, in symbols $\mathcal{F} \equiv^{\rho}_{sr} \mathcal{G}$, iff there are renamings f and g for $\mathcal{F}$ and $\mathcal{G}$, s.t. $f(\mathcal{F}) \equiv^{\rho}_{s} g(\mathcal{G})$.

Replacing the strong equivalence requirement with its definition yields the following two conditions:

1. $f(\mathcal{F})$ and $g(\mathcal{G})$ are compatible with each other; and
2. $\rho(f(\mathcal{F}) \cup \mathcal{H}) = \rho(g(\mathcal{G}) \cup \mathcal{H})$ for each CAF $\mathcal{H}$ which is compatible with $f(\mathcal{F})$ and $g(\mathcal{G})$.

Let us reconsider our motivating Example 5.1.

Example 5.10. Recall the CAFs $\mathcal{F}$ and $\mathcal{G}$, the renaming f with $f(a_1) = x_1$, $f(b_1) = x_2$, $f(c_1) = x_3$, $f(a_2) = x_4$, and $f(c_2) = x_5$ and let $g = id$.

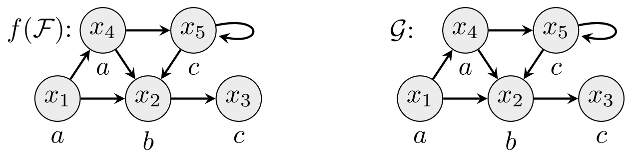

Augmenting both $f(\mathcal{F})$ and $\mathcal{G}$ with the CAF $\mathcal{H}$ considered above, we obtain the following desired situation:

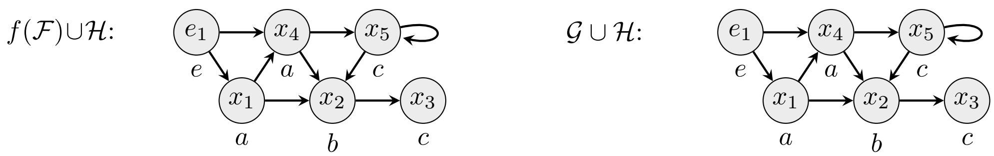

Notice that Proposition 5.5 ensures that our renaming for $\mathcal{F}$ only prevents $\mathcal{H}$ from introducing novel arguments, while preserving the semantics of $\mathcal{F}$.

Please note that strong equivalence up to renaming faithfully generalizes strong equivalence as $f = g = id$ can be chosen. This inside is expressed in the following proposition.

Proposition 5.11. *For any two CAFs $\mathcal{F}$ and $\mathcal{G}$ and any possible semantics ρ we have: If $\mathcal{F} \equiv^{\rho}_{s} \mathcal{G}$, then $\mathcal{F} \equiv^{\rho}_{sr} \mathcal{G}$.*

Moreover, as expected, strong equivalence up to renaming is a stricter notion than ordinary equivalence up to renaming.

Proposition 5.12. *For any two CAFs $\mathcal{F}$ and $\mathcal{G}$ and any possible semantics ρ we have: If $\mathcal{F} \equiv^{\rho}_{s} \mathcal{G}$, then $\mathcal{F} \equiv^{\rho}_{or} \mathcal{G}$.*

Let us now show how to characterize strong equivalence up to renaming using our kernels. Since this equivalence notion allows for changing the names of the arguments, we expect our kernels to behave similarly. More specifically, we also need to consider renamed versions of the CAFs before evaluating the kernels. However, checking strong equivalence up to renaming will certainly require to take the structure of the CAFs into consideration. We thus define what we mean by a CAF isomorphism.

Definition 5.13. Two CAFs $\mathcal{F}$ and $\mathcal{G}$ are *isomorphic* to each other iff there is a bijection $f : A_{\mathcal{F}} \to A_{\mathcal{G}}$ such that for all $a, b \in A_{\mathcal{F}}$, $cl_{\mathcal{F}}(a) = cl_{\mathcal{G}}(f(a))$ and $(f(a), f(b)) \in R_{\mathcal{G}}$ iff $(a, b) \in R_{\mathcal{F}}$; f is called *isomorphism* between $\mathcal{F}$ and $\mathcal{G}$.

CAFs $\mathcal{F}$ and $f(\mathcal{F})$ from Example 5.4 are isomorphic. The given renaming f naturally is a CAF-isomorphism between $\mathcal{F}$ and $f(\mathcal{F})$. More generally, it is easy to see that renamings characterize *all* conceivable isomorphism.

Remark 5.14. Let $\mathcal{F}$ and $\mathcal{G}$ be two CAFs. Then $\mathcal{F}$ and $\mathcal{G}$ are isomorphic to each other iff there is a renaming f for $\mathcal{F}$ s.t. $f(\mathcal{F}) = \mathcal{G}$.

As it turns out, we can decide $\mathcal{F} \equiv_{sr}^{\rho} \mathcal{G}$ in the following way: We choose the appropriate kernel for ρ, compute the kernels of $\mathcal{F}$ as well as $\mathcal{G}$, and then check whether those are isomorphic to each other. To this end we require the following auxiliary result stating that computing the kernel and applying an isomorphism commutes.

Proposition 5.15. *Let $\mathcal{F}$ be a CAF and f be a renaming for $\mathcal{F}$. The for any kernel $k(\rho)$ considered in this paper we have $f(\mathcal{F})^{k(\rho)} = f(\mathcal{F}^{k(\rho)})$.*

Proof. The following statements hold trivially:

- a is a self-attacker in $\mathcal{F}$ iff a is a self-attacker in $f(\mathcal{F})$;
- a attacks b in $\mathcal{F}$ iff $f(a)$ attacks $f(b)$ in $f(\mathcal{F})$;
- $cl(a) = c$ in $\mathcal{F}$ iff $cl(f(a)) = c$ in $f(\mathcal{F})$.

However, these properties characterize all kernels considered in this paper. Thus, (a, b) is an attack in $\mathcal{F}^{\rho(k)}$ iff $(f(a), f(b))$ is an attack in $f(\mathcal{F})^{\rho(k)}$. By definition of a renaming, (a, b) is an attack in $\mathcal{F}^{\rho(k)}$ iff $(f(a), f(b))$ is an attack in $f(\mathcal{F}^{\rho(k)})$. Hence the following statements are equivalent:

- $(f(a), f(b))$ is an attack in $f(\mathcal{F}^{\rho(k)})$,
- (a, b) is an attack in $\mathcal{F}^{\rho(k)}$,
- $(f(a), f(b))$ is an attack in $f(\mathcal{F})^{\rho(k)}$. □

Theorem 5.16. *For any two CAFs $\mathcal{F}$ and $\mathcal{G}$, and any semantics ρ under consideration, $\mathcal{F} \equiv_{sr}^{\rho} \mathcal{G}$ iff $\mathcal{F}^{k(\rho)}$ and $\mathcal{G}^{k(\rho)}$ are isomorphic.*

Proof. (⇐) Let $\mathcal{F}^{k(\rho)}$ and $\mathcal{G}^{k(\rho)}$ be isomorphic, witnessed by the isomorphism f. We have $f(\mathcal{F}^{k(\rho)}) = \mathcal{G}^{k(\rho)}$; and hence by Proposition 5.15 we infer $f(\mathcal{F})^{k(\rho)} = \mathcal{G}^{k(\rho)}$. By the results from Section 4, $f(\mathcal{F})$ and $\mathcal{G}$ are strongly equivalent. Hence $\mathcal{F} \equiv_{sr}^{\rho} \mathcal{G}$.

(⇒) Now assume the kernels $\mathcal{F}^{k(\rho)}$ and $\mathcal{G}^{k(\rho)}$ are not isomorphic, i.e. for any two renamings f and g, $f(\mathcal{F}^{k(\rho)}) \neq g(\mathcal{G}^{k(\rho)})$. Hence (again by Proposition 5.15) we find $f(\mathcal{F})^{k(\rho)} \neq g(\mathcal{G})^{k(\rho)}$ for any such f and g. Again by the results from Section 4 we find a suitable counter-example $\mathcal{H}$ for each conceivable pair f and g of renamings. Thus $\mathcal{F} \not\equiv_{sr}^{\rho} \mathcal{G}$. □

Example 5.17. For our CAFs $\mathcal{F}$ and $\mathcal{G}$ from Example 5.1 we see that their kernels are isomorphic. Hence $\mathcal{F}$ and $\mathcal{G}$ are strongly equivalent up to renaming w.r.t. all semantics considered in this paper.

5.3 Alternative Definitions

Our Definition 5.9 allowed for renaming both $\mathcal{F}$ and $\mathcal{G}$ and did impose any restriction on the co-domain of the respective renamings f and g. In this section we discuss useful variations.

First, we show that renaming either of the two CAFs is already sufficient.

Proposition 5.18. *For two CAFs $\mathcal{F}$ and $\mathcal{G}$ and a semantics ρ, it holds that $\mathcal{F} \equiv_{sr}^{\rho} \mathcal{G}$ iff $f(\mathcal{F}) \equiv_{s}^{\rho} \mathcal{G}$ for some renaming $f : A_{\mathcal{F}} \to A$.*

Proof. The ($\Leftarrow$) direction is obtained by setting $g = id$.

($\Rightarrow$) Take suitable renamings f and g. By definition (and the characterizations of Section 4), the kernels of $f(\mathcal{F})$ and $g(\mathcal{G})$ coincide. Therefore, the kernels of $g^{-1}(f(\mathcal{F}))$ and $\mathcal{G}$ coincide due to Proposition 5.15. □

We can further refine the strong renaming equivalence notion. As the following proposition formalizes, the renaming of f must be chosen in a way that the arguments of $f(\mathcal{F})$ and $\mathcal{G}$ coincide.

Proposition 5.19. *Given two CAFs $\mathcal{F}$ and $\mathcal{G}$ and a renaming $f : A_{\mathcal{F}} \to A$ for $\mathcal{F}$. Let ρ be any semantics under consideration. If $\rho(f(\mathcal{F}) \cup \mathcal{H}) = \rho(\mathcal{G} \cup \mathcal{H})$ for each CAF $\mathcal{H}$ which is compatible with $f(\mathcal{F})$ and $\mathcal{G}$, then $A = A_{\mathcal{G}}$.*

Proof. We assume that $f(\mathcal{F})$ and $\mathcal{G}$ are strongly equivalent w.r.t. the semantics ρ. By Lemma 4.7, we obtain $f(A_{\mathcal{F}}) = A = A_{\mathcal{G}}$. □

Consequently, the choice of f can be further restricted.

Corollary 5.20. *For two CAFs $\mathcal{F}$ and $\mathcal{G}$ and a semantics ρ, it holds that $\mathcal{F} \equiv^{\rho}_{sr} \mathcal{G}$ iff there is a renaming $f : A_{\mathcal{F}} \to A_{\mathcal{G}}$ for $\mathcal{F}$ s.t. $f(\mathcal{F}) \equiv^{\rho}_{s} \mathcal{G}$.*

6. Computational Complexity

In this section we examine the computational complexity of deciding equivalence between two CAFs $\mathcal{F}$ and $\mathcal{G}$ for every equivalence notion which has been established in this paper. We assume the reader to be familiar with the polynomial hierarchy (Papadimitriou, 1994). Moreover, by $\text{QSAT}^{\exists}_{n}$ ($\text{QSAT}^{\forall}_{n}$) we denote the generic Σ^{P}_{n}-complete (Π^{P}_{n}-complete) problem, i.e. checking validity of a corresponding QBF.

Our results reveal that ordinary equivalence can be computationally hard, up to the third level of the polynomial hierarchy for both variants of semi-stable and stage semantics as well as for i-preferred semantics. For the remaining semantics under consideration, the problem is Π^{P}_{2}-complete; the only exception is i-grounded semantics for which deciding ordinary equivalence is in P. Moreover, we show that deciding strong equivalence up to renaming extends the list of problems which lie in NP but are not known to be NP-complete.

6.1 Ordinary Equivalence

In this section, we discuss our complexity results for deciding ordinary equivalence for CAFs. As we will see, deciding whether two CAFs admit the same claim-extensions can be computationally hard, ranging up to the third level of the polynomial hierarchy. For well-formed CAFs, the problem turns out to be easier: deciding ordinary equivalence drops one level in the polynomial hierarchy for all (except grounded) semantics.

General CAFs. First we present our complexity results for ordinary equivalence regarding general CAFs. We formulate the following decision problem:

Ver-OE$_{\rho}$	
Input:	Two CAFs $\mathcal{F}$ and $\mathcal{G}$.
Output:	TRUE iff $\mathcal{F}$ and $\mathcal{G}$ are ordinary equivalent w.r.t. ρ.

The goal of this subsection is to formally prove the following complexity bounds for the various semantics.

Theorem 6.1. Ver-OE$_\rho$ *is*

- *in* P *for* $\rho = gr_i$;
- Π_2^{P}*-complete for* $\rho \in \{cf_i, ad_i, co_i, na_i, pr_h, na_h, stb_i, cf\text{-}stb_h, ad\text{-}stb_h,\}$*; and*
- Π_3^{P}*-complete for* $\rho \in \{pr_i, ss_i, stg_i, stg_h, ss_h\}$.

Let us note that deciding Ver-OE$_{gr_i}$ is in P since computing the unique grounded extensions of F and G and comparing the claims can be done in polynomial time (Dvořák & Woltran, 2020). In the following, we will discuss the remaining results from Theorem 6.1 in more detail. To begin with, we present membership proofs.

Proposition 6.2. *The problem* Ver-OE$_\rho$ *is*

- *in* Π_2^{P} *for* $\rho \in \{cf_i, ad_i, co_i, na_i, pr_h, na_h, stb_i, cf\text{-}stb_h, ad\text{-}stb_h,\}$*; and*
- *in* Π_3^{P} *for* $\rho \in \{pr_i, ss_i, stg_i, stg_h, ss_h\}$.

Proof. Membership proofs for Ver-OE$_\rho$, $\rho \neq gr_i$ are by standard guess-and-check procedures for the complementary problems: Guess a set of claims S and (w.l.o.g.) check whether it holds that $S \in \rho(\mathcal{F})$ as well as $S \notin \rho(\mathcal{G})$.

For the semantics $\sigma_i \in \{cf_i, ad_i, co_i, na_i, stb_i\}$, we proceed as follows:

- guess a set S of claims and a set E of arguments in $\mathcal{F}$ ($\exists$-quantifier);
- verify that $S = cl(E)$ and $E \in \sigma(\mathcal{F})$ (polynomial);
- check for each set E' of arguments in $\mathcal{G}$ that E' it is *not* the case that $S = cl(E')$ and $E' \in \sigma(\mathcal{G})$ ($\forall$-quantifier);

this yields a Σ_2^{P}-algorithm for the complementary problem.

For the semantics $\rho \in \{pr_h, na_h\}$ we first recall that verification for h-preferred and h-naive semantics is in $\mathsf{D}_1^{\mathsf{P}}$ (Dvořák, Greßler, Rapberger, & Woltran, 2021). Thus we

- guess a set S of claims ($\exists$-quantifier);
- require two $\mathsf{D}_1^{\mathsf{P}}$-oracle calls to check $S \in \rho(\mathcal{F})$ and $S \notin \rho(\mathcal{G})$

yielding Π_3^{P}-procedures for the decision problem Ver-OE$_\rho$.

For the semantics $\rho \in \{cf\text{-}stb_h, ad\text{-}stb_h\}$, we proceed as follows:

- guess a set S of claims and a set E of arguments in $\mathcal{F}$ ($\exists$-quantifier);
- verify that $S = cl(E)$ and $E \in ad(\mathcal{F})$ resp. $E \in cf(\mathcal{F})$ attacks all claims outside S (polynomial);
- check for each set E' of arguments in $\mathcal{G}$ that E' does *not* satisfy the above conditions ($\forall$-quantifier).

Again, this is a Σ_2^{P}-algorithm for the complementary problem.

For the semantics $\rho \in \{pr_i, ss_i, stg_i, ss_h, stg_h\}$, we

- guess a set S of claims ($\exists$-quantifier);
- require two Σ_2^{P}-oracle calls to check $S \in \rho(\mathcal{F})$ and $S \notin \rho(\mathcal{G})$

yielding Π_3^{P}-procedures for the decision problem Ver-OE$_\rho$. □

To show hardness of Ver-OE$_\rho$ for $\rho \neq gr_i$, we present reductions from QSAT$_2^\forall$ or QSAT$_3^\exists$, respectively. The overall idea is to construct two CAFs $\mathcal{F}$ and $\mathcal{G}$ where $\rho(\mathcal{F})$ depends on the particular instance of the source problem while $\mathcal{G}$ serves as controlling entity. Let us outline the idea for our Π_2^{P}-hardness proofs.

For a given instance $\Psi = \forall Y \exists Z \varphi$ of QSAT$_2^\forall$, we construct two CAFs $\mathcal{F}$ and $\mathcal{G}$ as follows:

- First, the claim-extensions (under some given semantics ρ) of both CAFs $\mathcal{F}$ and $\mathcal{G}$ should be of the form $Y' \cup \bar{Y}' \cup Z$ for some subset $Y' \subseteq Y$ and its complement $\bar{Y}' = \{\bar{y} \mid y \notin Y'\}$ (note that $\bar{y}$ represents $\neg y$, as usual).
- Second, we construct $\mathcal{F}$ such that the models of φ determine the claim-extensions of $\mathcal{F}$. That is, given an arbitrary subset $Y' \subseteq Y$ and its complement $\bar{Y}'$, we want that $Y' \cup \bar{Y}' \cup Z$ is a claim-extension of $\mathcal{F}$ if and only if there exists a subset $Z' \subseteq Z$ such that $Y' \cup Z'$ is a model of φ.

  Then it holds that $Y' \cup \bar{Y}' \cup Z$ is a claim-extension of $\mathcal{F}$ for all $Y' \subseteq Y$ if and only if the formula Ψ is valid.
- Finally, we construct our controlling CAF $\mathcal{G}$. This CAF is independent of the validity of Ψ. It realizes *all* claim-extensions $Y' \cup \bar{Y}' \cup Z$ for each subset $Y' \subseteq Y$ by default.

  Thus it holds that $\mathcal{F}$ and $\mathcal{G}$ yield the same claim-extensions if and only if Ψ is valid.

Below, we present the Π_2^{P}-hardness proof for inherited and both variants of hybrid stable semantics. The hardness proofs for the remaining semantics for which verifying ordinary equivalence is Π_2^{P}-complete proceed in a similar way; they can be found in the appendix.

Proposition 6.3. *Deciding* Ver-OE$_\rho$ *is* Π_2^{P}*-hard for* $\rho \in \{stb_i, cf\text{-}stb_h, ad\text{-}stb_h\}$.

Proof. Let $\rho = stb_i$ and let $\Psi = \forall Y \exists Z \varphi(Y, Z)$ be an instance of QSAT$_2^\forall$ where φ is given by a set of clauses C over atoms in $V = Y \cup Z$. We define two CAFs $\mathcal{F} = (A_\mathcal{F}, R_\mathcal{F}, cl_\mathcal{F})$, $\mathcal{G} = (A_\mathcal{G}, R_\mathcal{G}, id)$ as follows: For $\mathcal{F}$, we let

$$\begin{aligned} A_\mathcal{F} &= V \cup \bar{V} \cup C \text{ with } \bar{V} = \{\bar{v} \mid v \in V\}; \\ R_\mathcal{F} &= \{(v, cl) \mid cl \in C, v \in cl\} \cup \{(cl, cl) \mid cl \in C\} \cup \\ &\quad \{(\bar{v}, cl) \mid cl \in C, \neg v \in cl\} \cup \{(v, \bar{v}), (\bar{v}, v) \mid v \in V\}, \end{aligned}$$

with claim-function $cl_\mathcal{F}(z) = cl_\mathcal{F}(\bar{z}) = z$ for $z \in Z$ and $cl_\mathcal{F}(a) = a$ otherwise. We note that this reduction has been introduced byDvořák et al. (2023); it is a variant of the *standard translation* for AFs (Dvořák & Dunne, 2018, Reduction 3.6).

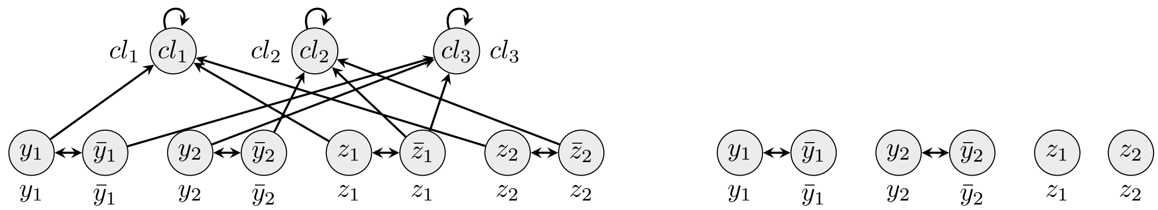

Figure 3: CAFs $\mathcal{F}$ (left) and $\mathcal{G}$ (right) illustrating the reduction from the Proof of Proposition 6.3 for the formula $\Psi = \forall Y \exists Z \varphi(Y, Z)$ where $\varphi(Y, Z)$ is given by the clauses $\{\{y_1, z_1, z_2\}, \{\bar{z}_1, \bar{z}_2, \bar{y}_2)\}, \{\bar{y}_1, \bar{z}_1, y_2\}\}$.

The CAF $\mathcal{G}$ is given by

$$A_{\mathcal{G}} = Y \cup \bar{Y} \cup Z;$$
$$R_{\mathcal{G}} = \{(y, \bar{y}), (\bar{y}, y) \mid y \in Y\}$$

and $cl_{\mathcal{G}} = id$, i.e., $cl_{\mathcal{G}}(x) = x$ for all $x \in A_{\mathcal{G}}$. An example of the two CAFs is given in Figure 3. We observe that $stb_i(\mathcal{G}) = \{Y' \cup \{\bar{y} \mid y \notin Y'\} \cup Z \mid Y' \subseteq Y\}$.

We show that Ψ is valid iff $stb_i(\mathcal{F}) = stb_i(\mathcal{G})$.

First assume Ψ is valid, let $Y' \subseteq Y$ and consider a model $M = Y' \cup Z'$ of φ. Then the set of arguments $E = M \cup \{\bar{v} \mid v \notin M\}$ is stable in $\mathcal{F}$: We observe that E is conflict-free; moreover, E attacks every $cl \in C$ since every clause cl is satisfied by M: In case there is $v \in M$ with $v \in cl$ we have $v \in E$ with $(v, cl) \in R_{\mathcal{F}}$; in case there is $\neg v \in M$ we have $\bar{v} \in E$ with $(\bar{v}, cl) \in R_{\mathcal{F}}$. Since $cl_{\mathcal{F}}(E) = Y' \cup \{\bar{y} \mid y \notin Y'\} \cup Z$ we have shown that every such claim-set is contained in $stb_i(\mathcal{F})$; $stb_i(\mathcal{F}) = stb_i(\mathcal{G})$ thus follows.

Now assume $stb_i(\mathcal{F}) = stb_i(\mathcal{G})$. Let $Y' \subseteq Y$, let E be a stb_i-realization of $Y' \cup \{\bar{y} \mid y \notin Y'\} \cup Z$ and let $Z' = E \cap Z$. We show that $M = Y' \cup Z'$ is a model of φ: Consider an arbitrary clause $cl \in C$. By assumption that E is stable in F there is some $a \in E$ such that $(a, cl) \in R_{\mathcal{F}}$. In case $a = v$ for some atom $v \in V$ we have $v \in cl$; in this case $v \in M$ and thus cl is satisfied. In case $a = \bar{v}$ for some atom v we have $\neg v \in cl$; in this case $v \notin M$ since $\bar{v} \in E$ and thus cl is satisfied. We obtain that M is a model of φ. We have shown that for any $Y' \subseteq Y$ there is $Z' \subseteq Z$ such that $Y' \cup Z'$ is a model of φ; i.e., Ψ is valid.

Π_2^{P}-hardness of Ver-OE$_\rho$ for $\rho \in \{cf\text{-}stb_h, ad\text{-}stb_h\}$ follows since $stb_i(\mathcal{F}) = cf\text{-}stb_h(\mathcal{F}) = ad\text{-}stb_h(\mathcal{F})$ and $stb_i(\mathcal{G}) = cf\text{-}stb_h(\mathcal{G}) = ad\text{-}stb_h(\mathcal{G})$. □

Turning now to the Π_3^{P}-hardness results, we slightly adapt our strategy. Similarly as for our Π_2^{P}-hardness proofs, we construct two CAFs $\mathcal{F}$ and $\mathcal{G}$ such that the claim-extensions of $\mathcal{F}$ depend on the validity of an instance $\Psi = \exists X \forall Y \exists Z \varphi(X, Y, Z)$ of QSAT$_3^{\exists}$ while the claim-extensions of $\mathcal{G}$ are independent of Ψ. Now, our target is to construct $\mathcal{F}$ in a way such that $\rho(\mathcal{F}) \neq \rho(\mathcal{G})$ iff Ψ is valid.

To show Π_3^{P}-hardness of Ver-OE$_{ss_i}$ and Ver-OE$_{stg_i}$, we will make use of the following reduction (Dvořák et al., 2021).

Reduction 6.4. Let $\Psi = \exists X \forall Y \exists Z \varphi(X, Y, Z)$ be an instance of $QSAT_3^{\exists}$, where φ is given by a set of clauses C over atoms in $V = X \cup Y \cup Z$. Let $V' = X \cup Y$ and let $\bar{x}$ denote $\neg x$.

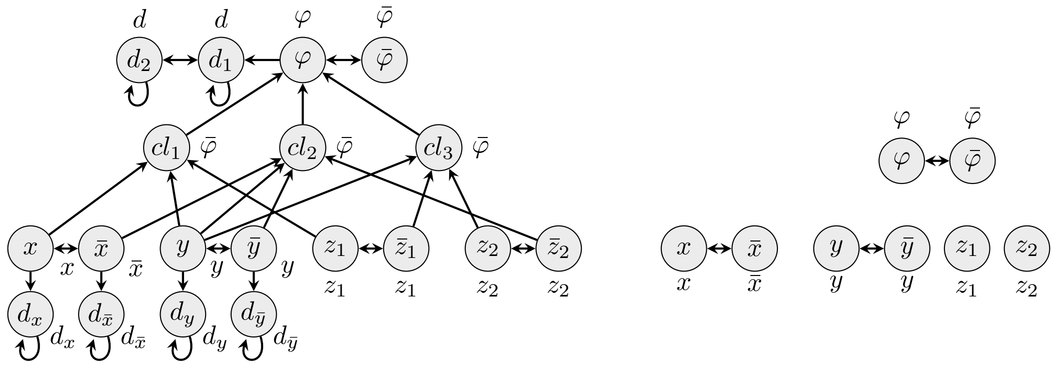

Figure 4: CAF $\mathcal{F}$ ((left) and CAF $\mathcal{G}$ (right) from the proof of Proposition 6.5 for the formula $\exists X \forall Y \exists Z \varphi(X, Y, Z)$ with clauses $\{\{z_1, x, y\}, \{\neg x, \neg y, \neg z_2, y\}, \{\neg z_1, z_2, y\}\}$.

We can assume that there is $y_0 \in Y$ with $y_0 \in cl$ for all $cl \in C$ (otherwise we can add such a y_0 without changing the validity of Ψ). Let $\mathcal{F} = (A, R, cl)$ be given by

$$\begin{aligned} A &= V \cup \bar{V} \cup \mathcal{C} \cup \{d_1, d_2, \varphi, \bar{\varphi}\} \cup \{d_v, d_{\bar{v}} \mid v \in V'\}; \\ R &= \{(a, cl) \mid cl \in C, a \in cl, a \in V \cup \bar{V}\} \cup \{(cl, \varphi) \mid cl \in C\} \cup \\ &\quad \{(a, d_a), (d_a, d_a) \mid a \in V' \cup \bar{V}'\} \cup \{(d_i, d_j) \mid i = 1, 2\} \\ &\quad \cup \{(v, \bar{v}), (\bar{v}, v) \mid v \in V\} \cup \{(\varphi, \bar{\varphi}), (\bar{\varphi}, \varphi), (\varphi, d_1)\}; \end{aligned}$$

and $cl(v) = cl(\bar{v}) = v$ for $v \in Y \cup Z$; $cl(cl_i) = \bar{\varphi}$ for $i \leq n$; $cl(d_1) = cl(d_2) = d$; and $cl(a) = a$ else.

Proposition 6.5. *Deciding* Ver-OE$_\rho$ *is* Π_3^{P}*-hard,* $\rho \in \{ss_i, stg_i\}$.

Proof. Consider an instance $\Psi = \exists X \forall Y \exists Z \varphi(X, Y, Z)$ of $\text{QSAT}_3^{\exists}$, where φ is given by a set of clauses C over atoms in $V = X \cup Y \cup Z$. We will first discuss the case for i-semi-stable semantics.

For the CAF $\mathcal{F}$, we apply Reduction 6.4. The reduction has been used to show Π_3^{P}-hardness of the *concurrence problem* for semi-stable (and stage) semantics (Dvořák et al., 2021): Given a CAF $\mathcal{F}$, does it hold that $ss_i(\mathcal{F}) = ss_h(\mathcal{F})$ (resp. $stg_i(\mathcal{F}) = stg_h(\mathcal{F})$)?

Dvořák et al. (2021) showed that Ψ is not valid iff

$$ss_i(\mathcal{F}) = ss_h(\mathcal{F}) = \{X' \cup \{\bar{x} \mid x \notin X'\} \cup Y \cup Z \cup \{e\} \mid X' \subseteq X, e \in \{\varphi, \bar{\varphi}\}\}.$$

Hence it suffices to construct the CAF $\mathcal{G}$ in such a way that $ss_i(\mathcal{G}) = ss_h(\mathcal{F})$. Then Ψ is not valid iff $ss_i(\mathcal{G}) = ss_h(\mathcal{F}) = ss_i(\mathcal{F})$.

We construct such a CAF $\mathcal{G} = (A_\mathcal{G}, R_\mathcal{G}, id)$ by setting

$$\begin{aligned} A_\mathcal{G} &= X \cup \bar{X} \cup Y \cup Z \cup \{\varphi, \bar{\varphi}\}, \text{ and} \\ R_\mathcal{G} &= \{(x, \bar{x}), (\bar{x}, x) \mid x \in X\} \cup \{(\varphi, \bar{\varphi}), (\bar{\varphi}, \varphi)\}. \end{aligned}$$

Figure 4 provides an illustrative example of $\mathcal{F}$ and $\mathcal{G}$. It is easy to see that $\mathcal{G}$ possesses exactly the desired i-semi-stable claim-sets.

This concludes the proof for i-semi-stable semantics. For i-stage semantics, we note that $ss_i(\mathcal{F}) = stg_i(\mathcal{F})$ and $ss_h(\mathcal{F}) = stg_h(\mathcal{F})$ (Dvořák et al., 2021). Hence, Π_3^{P}-hardness of $\textsc{Ver-OE}_{stg_i}$ follows from the additional observation that $ss_i(\mathcal{G}) = stg_i(\mathcal{G})$. □

The Π_3^{P}-hardness proofs for i-preferred, h-stage, and h-semi-stable semantics can be found in the appendix.

Observe that the computational complexity results from Theorem 6.1 extend to ordinary equivalence up to renaming by Proposition 5.8 for any semantics under consideration.

Well-formed CAFs. Let us now discuss ordinary equivalence for well-formed CAFs. By $\textsc{Ver-OE}_\rho^{wf}$ we denote the problem restricted to well-formed CAFs, as follows:

$\textsc{Ver-OE}_\rho^{wf}$	
Input:	Two well-formed CAFs $\mathcal{F}$, $\mathcal{G}$.
Output:	TRUE iff $\mathcal{F}$, $\mathcal{G}$ are ordinary equivalent w.r.t. ρ.

In general, we observe that the computational complexity of deciding ordinary equivalence drops one level in the polynomial hierarchy for all considered semantics (except for grounded semantics) when considering well-formed CAFs only. Our results can be summarized as follows.

Theorem 6.6. $\textsc{Ver-OE}_\rho^{wf}$ *is*

- *in* P *for* $\rho = gr_i$;
- coNP*-complete for* $\rho \in \{cf_i, ad_i, co_i, na_i, na_h, stb_i, cf\text{-}stb_h, ad\text{-}stb_h\}$; *and*
- Π_2^{P}*-complete for* $\rho \in \{pr_i, ss_i, stg_i, pr_h, stg_h, ss_h\}$.

First, we recall that all variants of stable semantics coincide for well-formed CAFs, i.e., $stb_i(\mathcal{F}) = cf\text{-}stb_h(\mathcal{F}) = ad\text{-}stb_h(\mathcal{F})$ for each well-formed CAF $\mathcal{F}$, likewise, both preferred variants coincide. Hence it suffices to establish the complexity of one of the variants.

Membership results are obtained in the same way as for general CAFs. To this end we have to make two additional observations: First, we have fewer cases here because for well-formed CAFs some of our semantics coincide. Second, for all semantics except grounded, we go down one level in the polynomial hierarchy. This is due to the lower computational complexity of the verification problem (Dvořák & Woltran, 2020; Dvořák & Dunne, 2018), for example in case of stb_i we get the following procedure:

- iterate over each set S of claims ($\forall$-quantifier);
- verify that either $S \in stb_i(\mathcal{F})$ and $S \in stb_i(\mathcal{G})$ or $S \notin stb_i(\mathcal{F})$ and $S \notin stb_i(\mathcal{G})$ (polynomial).

The other membership results follow analogously.

The idea for the hardness proof is similar to the general case: for an instance φ of SAT (an instance $\Psi = \forall Y \exists Z \varphi(Y, Z)$ of $\text{QSAT}_2^\forall$), we construct two CAFs $\mathcal{F}$ and $\mathcal{G}$ such that

- $\rho(\mathcal{F}) = \rho(\mathcal{G})$ iff φ is unsatisfiable for $\rho \in \{cf_i, ad_i, co_i, na_i, na_h, stb_i\}$; or

- $\rho(\mathcal{F}) = \rho(\mathcal{G})$ iff Ψ is valid for $\rho \in \{pr_i, ss_i, stg_i, stg_h, ss_h\}$, respectively.

For admissible, complete, and stable semantics, we can utilize coNP-hardness of a a related decision problem which is a special case of deciding ordinary equivalence. The *non-emptiness problem* (Dvořák & Dunne, 2018) is formulated as follows: given a CAF $\mathcal{F}$, does there exist an extension $S \in \rho(\mathcal{F})$ such that $S \neq \emptyset$?

Since deciding non-emptiness of claim-extensions is a special case of deciding ordinary equivalence we obtain coNP-hardness for admissible, complete, and stable semantics.

Proposition 6.7. Ver-OE$_\rho^{wf}$ *is* coNP*-hard for* $\rho \in \{ad_i, co_i, stb_i\}$.

Proof. Let $\mathcal{G}$ be a CAF for which $\rho(\mathcal{G}) = \{\emptyset\}$ holds (e.g., by setting $\mathcal{G} = (\emptyset, \emptyset, \emptyset)$). Then it holds that $\mathcal{F}$ and $\mathcal{G}$ are ordinary equivalent to each other iff $\mathcal{F}$ has no non-empty ρ-extension. Since deciding non-emptiness of an extension is coNP-complete for admissible, complete, and stable semantics (Dvořák & Dunne, 2018), we obtain coNP-hardness of Ver-OE$_\rho^{wf}$. □

The remaining hardness proofs make use of (modified) constructions by Dvořák and Woltran (2020), Dvořák et al. (2023), Kiesel and Rapberger (2021), and Dvořák and Dunne (2018), they can be found in the appendix.

6.2 Strong Equivalence

Having established complexity results for ordinary equivalence it remains to discuss the computational complexity of strong equivalence and its renaming version.

Ver-SE$_\rho$	
Input:	Two CAFs $\mathcal{F}$, $\mathcal{G}$.
Output:	true iff $\mathcal{F}$, $\mathcal{G}$ are strongly equivalent w.r.t. ρ.

Recall that in Section 4, we have shown that strong equivalence of two CAFs $\mathcal{F}$ and $\mathcal{G}$ can be characterized via syntactic equivalence of their kernels. Since the computation and comparison of the kernels of $\mathcal{F}$ and $\mathcal{G}$ can be done in polynomial time, we obtain tractability of strong equivalence for every semantics under consideration.

Theorem 6.8. *The problem* Ver-SE$_\rho$ *can be solved in polynomial time for any semantics* ρ *considered in this paper.*

Finally, we consider strong equivalence up to renaming. An analogous decision problem can be formulated as follows:

Ver-SER$_\rho$	
Input:	Two CAFs $\mathcal{F}$, $\mathcal{G}$.
Output:	true iff $\mathcal{F}$, $\mathcal{G}$ are strongly equivalent up to renaming w.r.t. ρ.

As outlined above, the computation of the kernels lies in P and is therefore negligible; the complexity of verifying strong equivalence up to renaming thus stems entirely from deciding whether two labelled graphs (i.e., the kernels of the given CAFs) are isomorphic. As a consequence we obtain that the complexity of Ver-SER$_\rho$ coincides with the complexity of the well-known graph isomorphism problem.

Theorem 6.9. *The problem* Ver-SER$_\rho$ *is exactly as hard as the graph isomorphism problem for any semantics* ρ *considered in this paper.*

Proof. For a reduction of the graph isomorphism problem to Ver-SER$_\rho$, consider two undirected, unlabelled graphs $F = (V, E)$ and $G = (V', E')$. We define the CAFs $\mathcal{F}$ and $\mathcal{G}$ by replacing each undirected edge by a symmetric one, moreover, each argument is labelled with the same claim. Formally, $\mathcal{F} = (V, \{(v, v'), (v', v) \mid \{v, v'\} \in E\}, cl)$ and $\mathcal{G} = (V', \{(v, v'), (v', v) \mid \{v, v'\} \in E'\}, cl)$ with $cl(v) = c$ for a fixed claim c. Observe that for any semantics ρ considered in this paper, the ρ-kernel of $\mathcal{F}$ ($\mathcal{G}$) coincides with $\mathcal{F}$ ($\mathcal{G}$, respectively): the CAFs do not contain self-attacking arguments; moreover, each conflict between two arguments with the same claim is already symmetric (i.e., $(a, b) \in R$ iff $(b, a) \in R$), thus no new attacks are introduced by computing the stb_h-kernel. We obtain that F is isomorphic to G iff $\mathcal{F}$ and $\mathcal{G}$ are isomorphic iff $\mathcal{F}$ and $\mathcal{G}$ are strongly equivalent up to renaming w.r.t. ρ.

For the other direction, observe that CAF isomorphism corresponds to the labelled variant of the graph isomorphism problem that is both edge- and label-preserving. We reduce Ver-SER$_\rho$ by setting $F := \mathcal{F}^{k(\rho)}$ and $G := \mathcal{G}^{k(\rho)}$ for two given CAFs $\mathcal{F}$ and $\mathcal{G}$. □

6.3 Well-Formed CAFs and Isomorphisms

In this subsection we will see that we can decide whether two well-formed CAFs are isomorphic in polynomial time. However, in order to lift this result to deciding strong equivalence up to renaming, we have to be careful: For this, we need to decide whether the kernels are isomorphic, but the kernel of a well-formed CAF is not necessarily well-formed itself. We will first discuss the case of well-formed CAFs and then see how to transition the underlying observations to comparing the kernels of well-formed CAFs.

In the following, we let

$$x^- = \{y \in A \mid (y, x) \in R\},$$

i.e. x^- is the set of arguments attacking x.

The high level idea for deciding whether two well-formed CAFs $\mathcal{F}$ and $\mathcal{G}$ are is isomorphic is that we can utilize the claims of the arguments for our guidance. We map arguments x in $A_\mathcal{F}$ with some claim c to arguments y in $A_\mathcal{G}$ with the same claim s.t. $cl(x^-) = cl(y^-)$, i.e. x and y have the same claim and are attacked by the same claims. Due to well-formedness, this information suffices to render the two arguments equivalent.

Let us apply this approach to our running renaming example, with some adjustments to obtain two well-formed CAFs.

Example 6.10. Let $\mathcal{F}$ and $\mathcal{G}$ be the following CAFs.

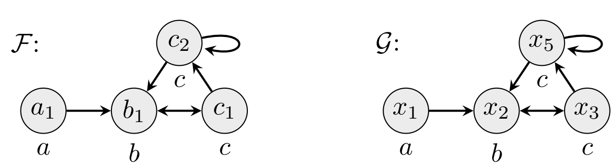

Let us start with $a_1 \in A_\mathcal{F}$; this argument is unattacked and has claim a: it can therefore be mapped to $x_1 \in A_\mathcal{G}$ which has the same properties. We continue with b_1, characterized by $cl(b_1) = b$ and $cl(b^-) = \{a, c\}$. The counterpart to this argument is $x_2 \in A_\mathcal{G}$. Regarding the two arguments with claim c, we have $cl(c_1^-) = \{b\}$ and $cl(c_2^-) = \{c\}$ corresponding to x_3 and x_5, respectively. Hence our isomorphism is given by

$$a_1 \mapsto x_1 \qquad b_1 \mapsto x_2 \qquad c_1 \mapsto x_3 \qquad c_2 \mapsto x_5,$$

corresponding to the way we depicted the two graphs.

The following theorem states that this procedure works for any two well-formed CAFs $\mathcal{F}$ and $\mathcal{G}$.

Theorem 6.11. *Deciding whether two well formed CAFs $\mathcal{F}$ and $\mathcal{G}$, are isomorphic is tractable.*

Proof. Consider the following algorithm: Check whether $cl(A(\mathcal{F})) = cl(A(\mathcal{G}))$ and then for each claim c occurring in $\mathcal{F}$, check if $|\{x \in A(\mathcal{F}) \mid cl(x) = c\}| = |\{y \in A(\mathcal{G}) \mid cl(y) = c\}|$; if not, stop. Otherwise for each claim c, proceed as follows:

1. chose some unmarked $x \in A(\mathcal{F})$ with $cl(x) = c$, compute $cl(x^-)$;
2. find an unmarked $y \in A(\mathcal{G})$ with $cl(y) = c$, and $cl(y^-) = cl(x^-)$;
   - if such y does not exist, stop;
   - else mark x and y as mapped to each other and go to 1.

If the algorithm successfully maps each x with $cl(x) = c$ to some y with $cl(y) = c$ for each claim c occurring in both CAFs, the mapping suggested by the algorithm is an isomorphism. If not, then there is some claim c and some set C of claims s.t.

$$|\{x \in A(\mathcal{F}) \mid cl(x) = c, cl(x^-) = C\}| \neq |\{y \in A(\mathcal{G}) \mid cl(y) = c, cl(y^-) = C\}|,$$

i.e. no isomorphism exists. □

The problem with Theorem 6.11 is however that in order to decide strong equivalence of two well-formed CAFs, it does not suffice to construct an isomorphism between the two (well-formed) CAFs. Rather, we need to decide whether the *kernels* are isomorphic to each other. However, these frameworks are not necessarily well-formed anymore, hence we cannot directly apply our algorithm.

For the kernels which remove attacks (i.e., for the stable, admissible, complete, and grounded kernel), we can circumvent this issue simply by considering the well-formed completion of the kernels instead.

Definition 6.12. For a CAF $\mathcal{F} = (A, R, cl)$ we define the *well-forming operator* $wf(\mathcal{F}) = (A, R \cup R', cl)$ where $R' \subseteq A \times A$ is $\subseteq$-minimal in $\{R'' \subseteq A \times A \mid (A, R \cup R'', cl) \text{ is well-formed}\}$. We call $wf(\mathcal{F})$ the well-formed completion of $\mathcal{F}$.

We observe that wf is a function.

Lemma 6.13. *$wf(\mathcal{F})$ is unique for every CAF $\mathcal{F} = (A, R, cl)$.*

Proof. Towards a contradiction, assume there are two sets of attacks $R_1, R_2 \subseteq A \times A$, $R_1 \neq R_2$, such that $(A, R \cup R_1, cl)$ and $(A, R \cup R_2, cl)$ are both well-formed and $\subseteq$-minimal in $\{R'' \subseteq A \times A \mid (A, R \cup R'', cl) \text{ is well-formed}\}$. Wlog, let $(a, b) \in R_1 \setminus R_2$. Then there is $(c, b) \in R$ with $cl(a) = cl(c)$ (otherwise, $(A, R \cup (R_1 \setminus \{(a, b)\}), cl)$ is well-formed, contradiction to $\subseteq$-minimality of R_1). Therefore, $(a, b) \in R_2$ as well (otherwise, $(A, R \cup R_2, cl)$ is not well-formed). □

Moreover, when applying the well-forming operator to the kernel of a well-formed CAF, we obtain a subgraph of the original CAF (except for the naive kernel for obvious reasons). In general, we obtain the following subset-relation between the well-forming operator and most of the kernels:

Lemma 6.14. *Let $\mathcal{F}$ be a well-formed CAF and let $k \in \{sk, ak, ck, gk\}$. It holds that $\mathcal{F}^k \subseteq wf(\mathcal{F}^k) \subseteq \mathcal{F}$.*

We furthermore observe that $\mathcal{F}^k = wf(\mathcal{F}^k)^k$ for $k \in \{sk, ak, ck, gk\}$. This follows from the observation that we can remove the attacks iteratively for the respective kernels.

Lemma 6.15. *Let $\mathcal{F} = (A, R, cl)$ be a CAF and let $k \in \{sk, ak, ck, gk\}$. Let $(x, y) \in \mathcal{F} \setminus \mathcal{F}^k$ and let $\mathcal{F}' = (A, R \setminus \{(x, y)\}, cl)$. It holds that $(\mathcal{F}')^k = \mathcal{F}^k$.*

Proof. Let $k = sk$. We have

$$\begin{aligned} R_{\mathcal{F}^{sk}} &= R \setminus \{(a, b) \mid a \neq b, (a, a) \in R\} \\ &= (R \setminus \{(x, y)\}) \setminus \{(a, b) \mid a \neq b, (a, a) \in (R \setminus \{(x, y)\}\} \\ &= R_{(\mathcal{F}')^{sk}} \end{aligned}$$

Note that we do not remove self-attackers, hence we can remove the attack (x, y) without affecting the deletion of other attacks. The proof is analogous for the remaining kernels. □

Corollary 6.16. *Let $\mathcal{F} = (A, R, cl)$ be a well-formed CAF and let $k \in \{sk, ak, ck, gk\}$. It holds that $\mathcal{F}^k = wf(\mathcal{F}^k)^k$.*

Proposition 6.17. *Let $\mathcal{F}$, $\mathcal{G}$ be two well-formed CAFs, and let $k \in \{sk, ak, ck, gk\}$. $\mathcal{F}^k$ and $\mathcal{G}^k$ are isomorphic to each other iff $wf(\mathcal{F}^k)$ and $wf(\mathcal{G}^k)$ are isomorphic to each other.*

Proof. ($\Rightarrow$): Let $h : A_\mathcal{F} \to A_\mathcal{G}$ denote an isomorphism between $\mathcal{F}^k$ and $\mathcal{G}^k$, i.e., $h(\mathcal{F}^k) = \mathcal{G}^k$. We show that h is an isomorphism for $wf(\mathcal{F}^k)$ and $wf(\mathcal{G}^k)$: (a) $cl_{wf(\mathcal{F}^k)}(a) = cl_{wf(\mathcal{G}^k)}(h(a))$ is satisfied since h is an isomorphism between the kernels and the well-forming operator does not add new arguments. (b) Let $(a, b) \in R_{wf(\mathcal{F}^k)}$. In case $(a, b) \in R_{\mathcal{F}^k}$, we are done; by assumption, h is an isomorphism between the sub-frameworks, hence $(h(a), h(b)) \in R_{\mathcal{G}^k} \subseteq R_{wf(\mathcal{G}^k)}$. Assume $(a, b) \in R_{wf(\mathcal{F}^k)} \setminus R_{\mathcal{F}^k}$. This implies that there is some attack $(c, b) \in R_{\mathcal{F}^k}$ with $cl(a) = cl(c)$ (by minimality of $wf(\mathcal{F}^k)$). Then $(h(c), h(b)) \in R_{\mathcal{G}^k}$. Since $h(a)$ and $h(c)$ have the same claims as a and c, we conclude that $(h(a), h(c)) \in R_{wf(\mathcal{G}^k)}$ as well (by well-formedness of $wf(\mathcal{G}^k)$). Hence $(a, b) \in R_{wf(\mathcal{F}^k)}$ implies $(h(a), h(b)) \in R_{wf(\mathcal{G}^k)}$. The other direction is analogous.

($\Leftarrow$): Let $h : A_\mathcal{F} \to A_\mathcal{G}$ denote an isomorphism between $wf(\mathcal{F}^k)$ and $wf(\mathcal{G}^k)$, i.e., $h(wf(\mathcal{F}^k)) = wf(\mathcal{G}^k)$. Let $\mathcal{H} := h(wf(\mathcal{F}^k)) = wf(\mathcal{G}^k)$. Constructing the kernel of $\mathcal{H}$ yields a unique framework, moreover, $\mathcal{H}^k = h(\mathcal{F}^k) = \mathcal{G}^k$ by Corollary 6.16, hence we obtain the desired result. □

Let us now turn to the naive kernel. Recall that we construct the naive kernel by making each attack symmetric; moreover, we add a symmetric attack (s, a) between each self-attacker s and each argument a.

In the following, we show that the functions $cl(x)$, $cl(x^-)$ still suffice to characterize the naive kernel uniquely.

Proposition 6.18. *Let $\mathfrak{C} = \{\mathcal{F}^{nk} \mid \mathcal{F} = (A, R, cl)$ is a well-formed CAF $\}$. The function $R : \mathfrak{C} \to \mathcal{U} \times \mathcal{C} \times 2^{\mathcal{C}}$ with $R(\mathcal{F}) = \{(x, cl(x), cl(x^-)) \mid x \in A_{\mathcal{F}}\}$ is injective.*

Proof. R is well-defined by definition. We show that R is injective. Let $R(\mathcal{F}) = R(\mathcal{G})$. Hence it holds that $A_{\mathcal{F}} = A_{\mathcal{G}} (=: A)$, moreover, the claim-function in $\mathcal{F}$ and $\mathcal{G}$ coincides, i.e., $cl_{\mathcal{F}}(x) = cl_{\mathcal{G}}(x)$ for all $x \in A$. It remains to prove that $R_{\mathcal{F}} = R_{\mathcal{G}}$. Wlog, let $(a, b) \in R_{\mathcal{F}}$. We show that $(a, b) \in R_{\mathcal{G}}$.

Recall that $\mathcal{F}$ is the naive kernel of some well-formed CAF $\mathcal{F}'$. Hence $(b, a) \in R_{\mathcal{F}}$ as well. Now, we proceed by case distinction: (a) a or b is self-attacking; or (b) a and b are not self-attacking.

Case (a): Wlog, assume $(a, a) \in R_{\mathcal{F}}$. Then a symmetrically attacks each other argument in $\mathcal{F}$ by construction of the naive kernel. By definition of R, we have $(a, cl(a), cl(A)) \in R(\mathcal{F})$. Hence $(a, cl(a), cl(A)) \in R(\mathcal{G})$ as well, therefore, $(a, a) \in R_{\mathcal{G}}$ and a attacks (and is attacked by) each other argument in A. Consequently, $(a, b) \in R_{\mathcal{G}}$.

Case (b): Now assume both a and b are not self-attacking. By construction of the naive kernel, either (a, b) or (b, a) is contained in the original CAF $\mathcal{F}'$. Since $\mathcal{F}'$ is well-formed, it holds that either $cl_{\mathcal{F}'}(a) \in cl_{\mathcal{F}'}(b^-_{\mathcal{F}'})$ (in case $(a, b) \in R_{\mathcal{F}'}$) or $cl(b) \in cl(a^-_{\mathcal{F}'})$ (in case $(b, a) \in R_{\mathcal{F}'}$). Wlog, let us assume that $cl(a) \in cl(b^-_{\mathcal{F}'})$.

Let us briefly clarify the relation between $cl(b^-_{\mathcal{F}}$ and $cl(b^-_{\mathcal{F}'}$: since we add attacks when constructing the naive kernel, we might get more claims, i.e., it holds that $cl(b^-_{\mathcal{F}'} \subseteq cl(b^-_{\mathcal{F}}$.

Hence we obtain $cl(a) \in cl(b^-_{\mathcal{F}})$ as well. Therefore, the set $R(\mathcal{F})$ contains a tuple $(b, cl(b), C)$ with $cl(a) \in C$. Since $R(\mathcal{F}) = R(\mathcal{G})$ we obtain that $R(\mathcal{G})$ contains the tuple as well. Consequently, each argument with claim $cl(a)$ attacks the argument b in $\mathcal{G}$. We obtain $(a, b) \in R_{\mathcal{G}}$, as desired. □

We have shown that well-formed CAFs and their naive kernels can be represented using the same information, namely the claim $cl(x)$ and all attacking claims $cl(x^-)$ for each argument x. Therefore, we can apply the algorithm from Theorem 6.11 to construct an isomorphism between two naive kernels.

We are ready to show that deciding strong equivalence up to renaming for well-formed CAFs is tractable for all considered semantics.

Theorem 6.19. *Let $\mathcal{F}$ and $\mathcal{G}$ be two well-formed CAFs and let ρ denote any semantics considered in this paper. Deciding whether $\mathcal{F}$ and $\mathcal{G}$ are strongly equivalent up to renaming w.r.t. ρ is tractable.*

Proof. By our results from the preceding sections, $\mathcal{F}$ and $\mathcal{G}$ are strongly equivalent up to renaming w.r.t. ρ iff their kernels corresponding to ρ are isomorphic.

First, let $k \in \{ak, sk, ck, gk\}$. We utilize Theorem 6.11 and Proposition 6.17: $\mathcal{F}^k$ and $\mathcal{G}^k$ are isomorphic iff $wf(\mathcal{F}^k)$ and $wf(\mathcal{G}^k)$ are isomorphic to each other. The latter can be decided in polynomial time.

Now, let $k = nk$. By Proposition 6.18, we can encode the naive kernel of well-formed CAFs analogous to well-formed CAFs. Hence, we can apply the algorithm from Theorem 6.11 to construct the desired isomorphism. □

7. Conclusion and Future Work

In this work, we examined several different equivalence notions for claim-augmented argumentation. We considered ordinary and strong equivalence as well as novel equivalence notions based on argument renaming for CAFs w.r.t. all semantics for CAFs which have been considered in the literature so far and provided a complexity analysis of all considered equivalence notions. We show that strong equivalence can be characterized via semantics-dependent kernels. Hence, we obtain tractability of strong equivalence w.r.t. all considered semantics. In contrast, ordinary equivalence can be computationally expensive, ranging up to the third level of the polynomial hierarchy. We furthermore show that strong equivalence up to renaming has the same complexity as the graph isomorphism problem and is thus presumably of higher complexity than classical strong equivalence. When restricting the problem to the class of well-formed CAFs, we can exploit the structure of the graphs sufficiently to compute an isomorphism in polynomial time. Hence we identified a tractable fragment of the graph isomorphism problem in the course of our complexity analysis of renaming strong equivalence.

7.1 Related Work

Ordinary equivalence for AFs has been studied by Baumann, Dvořák, Linsbichler, and Woltran (2019) and by Oikarinen and Woltran (2011). Similar to our setting, there are only few dependencies between the semantics. When comparing the computational complexity of deciding ordinary equivalence between CAFs and AFs, we observe that in general, the problem is one level harder for CAFs than for AFs, however, when moving to well-formed CAFs, the problem is of the same complexity as for AFs. This behavior can be observed for other well-known decision problems (e.g., for the verification problem) (Dvořák & Woltran, 2020; Dvořák et al., 2021).

Our characterization results for strong equivalence are in line with existing studies for related argumentation formalisms (Oikarinen & Woltran, 2011; Dvořák, Rapberger, & Woltran, 2020b): we show that strong equivalence can be characterized via syntactic equivalence of semantics-dependent kernels; by that, we obtain tractability of testing strong equivalence for CAFs. The notion of strong equivalence has been also tackled in several related formalisms, most famously in the context of logic programs (LP) (Lifschitz et al., 2001), but also in the context of structured argumentation (Amgoud, Besnard, & Vesic, 2014). In contrast to the abstract setting, deciding strong equivalence can be computationally hard: it is well-known that deciding strong equivalence is intractable for LPs (Pearce, Tompits, & Woltran, 2001; Lin, 2002); recently, it has been shown that deciding strong equivalence is intractable for flat ABA as well (Rapberger & Ulbricht, 2022). Amgoud et al. (2014) study equivalence of logic-based argumentation by adapting the classical equivalence notion of logical formulae to logical arguments. They show that under certain conditions on the underlying logic, unnecessary arguments can be removed while retaining strong equivalence. Similar studies have been carried out by Rapberger and Ulbricht (2022) in the context of assumption-based argumentation and instantiations of logic programming. They introduce the notion of *instantiated arguments* which are tuples consisting of the claim and the vulnerabilities of the arguments. They identify several semantics-dependent redundancy patterns for instantiated arguments and show that these arguments can be modified or even removed without affecting the respective semantics. Based on these redundancies, they identify fragments for which deciding strong equivalence becomes tractable. In contrast to the work by Amgoud et al. and Rapberger and Ulbricht, we do not assume any internal structure of the arguments which makes our approach applicable even beyond the scope of logic-based argumentation, ABA, and LP instantiations. In particular, both approaches consider only instances for which the attack relation is well-formed. However, there are several reasons why a framework might violate this condition, e.g., when taking preferences into account (Cyras & Toni, 2016; Modgil & Prakken, 2014; Bernreiter, Dvořák, Rapberger, & Woltran, 2022). The present work takes these observations into account. In our studies, we do not restrict ourselves to the class of well-formed CAFs. Moreover, our studies are independent of the underlying formalism of the instantiated argumentation system as we do not impose any further constraints on the arguments or their claims; in this way, it is even possible to test equivalence between argumentation systems stemming from entirely different base formalisms.

An interesting difference between the structured and the abstract approaches to strong equivalence lies in the different treatment of arguments and attacks: while in abstract formalisms, the attack structure is modified, we observe that in structured formalisms, arguments get modified. It would be interesting to study how these different approaches relate to each other.

With our studies on strong equivalence, we tackle the long-term behavior of CAFs, under the assumption that existing knowledge never gets lost. It is evident that the topic is closely related to several other aspects of dynamics in argumentation, e.g., argument revision (Coste-Marquis, Konieczny, Mailly, & Marquis, 2014; Baumann & Brewka, 2015; Snaith & Reed, 2017; Cayrol, de Saint-Cyr, & Lagasquie-Schiex, 2010; Alfano, Greco, & Parisi, 2021a), framework modifications (Liao, Jin, & Koons, 2011; Baumann, 2011), or enforcement (Baumann, 2012b; Wallner, Niskanen, & Järvisalo, 2017; Rapberger & Ulbricht, 2022). In this regard, we mention in particular the work by Cayrol et al. (2010) who study framework expansions in the context of AFs. They focus on the addition of a new argument to an AF which may interact with existing arguments. However, they do not consider claims in their studies. There are several approaches to dynamics in structured argumentation. Snaith and Reed (2017) consider revision operations in ASPIC+. Falappa, Kern-Isberner, and Simari (2002) study changes in logic-based argumentation systems and how the modification of strict to defeasible rules gives rise to the changing of arguments and their attack relation. Alfano, Greco, Parisi, Simari, and Simari (2021b) study the addition and removal of knowledge in defeasible logic programming frameworks. They develop methods to efficiently compute the warrant status of claims in evolving knowledge bases. Also, Rotstein et al. (2008, 2010) consider a framework specifically designed for handling dynamic changes in argumentation through the consideration of varying evidences. They develop dynamic argumentation frameworks where arguments have a richer structure, in particular, they keep track of the claims of the arguments. In their work, they consider the addition and the removal of arguments and study associated interactions. In contrast to our work, all the aforementioned approaches allow for the removal of arguments and attacks or for the modification of arguments. However, it would be interesting to explore revision operators in the context of CAFs. Since CAFs lie in-between structured and abstract approaches, they could turn out to the perfect playground for applying established techniques from the abstract universe in the structured setting; e.g., in the context of claim revision.

7.2 Future Work

For future work, we want to extend our strong equivalence studies by considering certain constraints of the framework modifications. What has been commonly investigated in the literature are *normal expansions* (Baumann, 2012a) where attacks can only be introduced if they involve newly added arguments (observe that in the proof of Theorem 4.24, the expansion in case (a) satisfy this criteria while $\mathcal{H}$ in case (b) introduces also new attacks between existing arguments).

We moreover plan to adapt our renaming strong equivalence notion to arbitrary CAFs, not only compatible ones, by relaxing the notion of framework expansions. By doing so, we expect to generalize the considered equivalence notions even further. In this aspect, it would be also interesting to investigate the connection between renaming equivalence and bisimulation as both notions constitute a relaxation of graph isomorphism and preserve semantics respectively truth values (when considering bisimulation in modal theories, for instance). Further studies in these directions could reveal interesting novel connections between different research areas.

On a more general note, we want to study other dynamic operators in the context of CAFs. We believe that studies on several aspects of belief revision in the context of claim-centric reasoning would be an interesting avenue for future research.

Acknowledgements

This work was supported by the German Federal Ministry of Education and Research (BMBF, 01/S18026A-F) by funding the competence center for Big Data and AI "ScaDS.AI" Dresden/Leipzig, by the Austrian Science Fund (FWF) through projects W1255-N23 and Y698 and by the Vienna Science and Technology Fund (WWTF) through project ICT19-065.

Appendix A. Computational Complexity of Ordinary Equivalence: Omitted Proofs from Section 6.1

Let us first provide the omitted proofs of the following theorem:

Theorem 6.1. Ver-OE$_\rho$ *is*

- *in* P *for* $\rho = gr_i$*;*
- Π_2^{P}*-complete for* $\rho \in \{cf_i, ad_i, co_i, na_i, pr_h, na_h,\ stb_i, cf\text{-}stb_h, ad\text{-}stb_h, \}$*; and*
- Π_3^{P}*-complete for* $\rho \in \{pr_i, ss_i, stg_i, stg_h,\ ss_h\}$.

We will make use (variants of) the following reduction (Dvořák & Dunne, 2018, Reduction 3.6).

Reduction A.1. Let φ be given by a set of clauses $C = \{cl_1, \dots, cl_n\}$ over atoms in X and let $\bar{X} = \{\bar{x} \mid x \in X\}$. We construct AF $F = (A, R)$ with

$$\begin{aligned} A &= X \cup \bar{X} \cup C \cup \{\varphi\} \\ R &= \{(x, cl) \mid cl \in C, x \in cl\} \cup \{(\bar{x}, cl) \mid cl \in C, \neg x \in cl\} \cup \\ &\quad \{(x, \bar{x}), (\bar{x}, x) \mid x \in X\} \cup \{(cl_i, \varphi) \mid i \leq n\} \end{aligned}$$

Let us furthermore recall the following reduction which we used in the proof of Proposition 6.3. This construction modifies the standard construction.

Reduction A.2. For a formula φ which is given by a set of clauses C over atoms in V we construct an AF $F = (A, R)$ with

$$\begin{aligned} A &= V \cup \bar{V} \cup C \text{ with } \bar{V} = \{\bar{v} \mid v \in V\}; \\ R &= \{(v, cl) \mid cl \in C, v \in cl\} \cup \{(cl, cl) \mid cl \in C\} \cup \\ &\quad \{(\bar{v}, cl) \mid cl \in C, \neg v \in cl\} \cup \{(v, \bar{v}), (\bar{v}, v) \mid v \in V\}. \end{aligned}$$

Now, we are ready to give the remaining proofs of Theorem 6.1. Let us start with i-naive semantics.

Proposition A.3. *Deciding* Ver-OE$_{na_i}$ *is* Π_2^{P}*-hard.*

Proof. Consider an instance $\Psi = \forall Y \exists Z \varphi(Y, Z)$ of $\text{QSAT}_2^{\forall}$, where φ is a 3-CNF, given by a set of clauses $C = \{cl_1, \dots, cl_n\}$ over atoms in $V = Y \cup Z$. We construct two CAFs $\mathcal{F} = (A_{\mathcal{F}}, R_{\mathcal{F}}, cl_{\mathcal{F}})$, $\mathcal{G} = (A_{\mathcal{G}}, R_{\mathcal{G}}, id)$, where $\mathcal{F}$ modifies the AF (A, R) obtained from Reduction A.2 as follows:

$$\begin{aligned} A_{\mathcal{F}} &= A \cup Y_2 \cup \bar{Y}_2 \cup Z_2; \\ R_{\mathcal{F}} &= (R \setminus \{(cl, cl) \mid cl \in C\}) \cup \\ &\quad \{(y_2, \bar{y}_2) \mid y_2 \in Y_2\} \cup (y, \bar{y}_2), (y_2, \bar{y}) \mid y \in Y\}; \end{aligned}$$

and $cl_{\mathcal{F}}(y_2) = y$, $cl_{\mathcal{F}}(\bar{y}_2) = \bar{y}$ for $y_2 \in Y_2$; $cl_{\mathcal{F}}(z_2) = cl_{\mathcal{F}}(\bar{z}_2) = z$ for $z_2 \in Z_2$; $cl_{\mathcal{F}}(cl) = \bar{\varphi}$ for $cl \in C$; $cl_{\mathcal{F}}(a) = a$ else.

Observe that $Y' \cup \{\bar{y} \mid y \notin Y'\} \cup Z \cup \{\bar{\varphi}\}$ is i-naive for every $Y' \subseteq Y$: Let $E = Y_2' \cup \{\bar{y}_2 \mid y_2 \notin Y_2'\} \cup Z_2 \cup C \cup E'$ with $Y_2' \subseteq Y_2$ and $E' \subseteq V \cup \bar{V}$ is a non-conflicting subset-maximal set of arguments which do not attack any $cl \in C$. E is conflict-free and subset-maximal by the choice of E'; moreover, $cl_{\mathcal{F}}(E) = Y' \cup \{\bar{y} \mid y \notin Y'\} \cup Z \cup \{\bar{\varphi}\}$.

We construct $\mathcal{G} = (A_{\mathcal{G}}, R_{\mathcal{G}}, cl_{\mathcal{G}})$ such that $na_i(\mathcal{G}) = \{Y' \cup \{\bar{y} \mid y \notin Y'\} \cup Z \cup \{\bar{\varphi}\} \mid Y' \subseteq Y\} \cup \{Y' \cup \{\bar{y} \mid y \notin Y'\} \cup Z \mid Y' \subseteq Y\}$. Let

$$\begin{aligned} A_{\mathcal{G}} &= Y_1 \cup \bar{Y}_1 \cup Y_2 \cup \bar{Y}_2 \cup Z \cup \{\bar{\varphi}\}; \\ R_{\mathcal{F}} &= \{(y_i, \bar{y}_i) \mid y_i \in Y_i, i \leq 2\} \cup \\ &\quad \{(a, b) \mid a \in Y_1 \cup \bar{Y}_1, b \in Y_2 \cup \bar{Y}_2 \cup \{\bar{\varphi}\}\}; \end{aligned}$$

and $cl_{\mathcal{G}}(y_i) = y$, $cl_{\mathcal{G}}(\bar{y}_i) = \bar{y}$ for $y_i \in Y_i$; $cl_{\mathcal{G}}(z) = z$, $cl_{\mathcal{G}}(\bar{z}) = \bar{z}$ for $z \in Z$; $cl_{\mathcal{G}}(\bar{\varphi}) = \bar{\varphi}$. See Figure 5 for an illustrative example of $\mathcal{F}$ and $\mathcal{G}$. It can be checked that $\mathcal{G}$ has precisely the desired i-naive extensions.

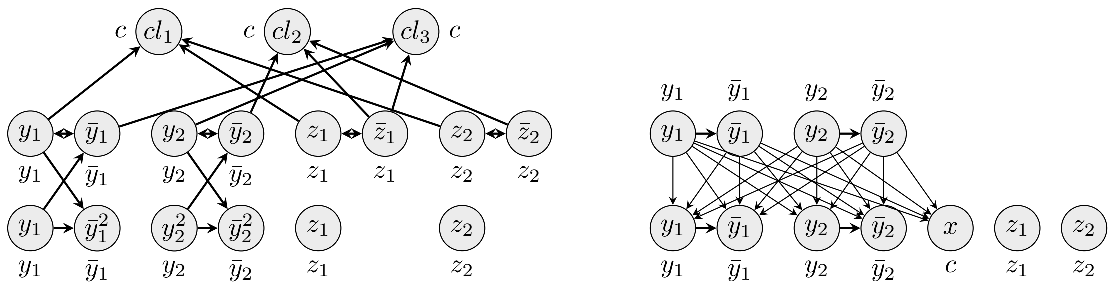

Figure 5: CAFs $\mathcal{F}$ (left) and $\mathcal{G}$ (right) illustrating the reduction from the Proof of Proposition A.3 for the formula $\Psi = \forall Y \exists Z \varphi(Y, Z)$ where $\varphi(Y, Z)$ is given by the clauses $\{\{y_1, z_1, z_2\}, \{\bar{z}_1, \bar{z}_2, \bar{y}_2)\}, \{\bar{y}_1, \bar{z}_1, y_2\}\}$.

We show that Ψ is valid iff $na_i(\mathcal{F}) = na_i(\mathcal{G})$. First, assume Ψ is valid and fix some $Y' \subseteq Y$. There is $Z' \subseteq Z$ such that $M = Y' \cup Z'$ is a model of φ. Let $E = M \cup \{\bar{v} \mid v \notin M\} \cup Y_2' \cup \{\bar{y}_2 \mid y_2 \notin Y_2'\} \cup Z_2$. E is conflict-free; moreover, E is subset-maximal among conflict-free sets since any other argument $a \in A_{\mathcal{F}} \setminus E$ is in conflict with E: Since M is a model of φ, we have that for each clause cl_i there is either a positive literal $v \in cl$ with $v \in M$ or a negative literal $\bar{v} \in cl$ with $v \notin M$; that is, each cl is attacked in $\mathcal{F}$. Also, E contains either v or $\bar{v}$ for any atom $v \in Y \cup Z \cup Y_2$, thus any argument representing a literal in $\mathcal{F}$ which is not a member of E is attacked by E. It follows that $Y' \cup \{\bar{y} \mid y \notin Y'\} \cup Z \in na_i(\mathcal{F})$ for every $Y' \subseteq Y$. Each i-naive claim-set is thus either of the form $Y' \cup \{\bar{y} \mid y \notin Y'\} \cup Z \cup \{\bar{\varphi}\}$ or $Y' \cup \{\bar{y} \mid y \notin Y'\} \cup Z$. Consequently, $na_i(\mathcal{F}) = na_i(\mathcal{G})$ in case Ψ is valid.

Now assume $na_i(\mathcal{F}) = na_i(\mathcal{G})$ and fix $Y' \subseteq Y$. Consider a na_i-realisation E of $Y' \cup \{\bar{y} \mid y \notin Y'\} \cup Z$ and let $Z' = E \cap Z$. We show $M = Y' \cup Z'$ is a model of φ: Consider an arbitrary clause $cl \in C$. Since E is a subset-maximal conflict-free set of arguments we have $E \cup \{cl\}$ is conflicting; that is, there is $a \in E$ such that a attacks cl. In case $a = v$ for some atom v we have $v \in cl$; in case $a = \bar{v}$ for some v we have $\bar{v} \in cl$. In the former case, $v \in M$ and thus cl is satisfied, in the latter case we have $v \notin M$ and thus cl is satisfied. We obtain that M is a model of φ. ☐

Proposition A.4. *Deciding* Ver-OE$_\rho$ *is* Π_2^{P}*-hard,* $\rho \in \{cf_i, ad_i na_h, pr_h\}$.

Proof. We will first show the statement for h-naive semantics: Consider an instance $\Psi = \forall Y \exists Z \varphi(Y, Z)$ of QSAT$_2^\forall$, where φ is a 3-CNF, given by a set of clauses $C = \{cl_1, \dots, cl_n\}$ over atoms in $V = Y \cup Z$. We construct two CAFs $\mathcal{F} = (A_{\mathcal{F}}, R_{\mathcal{F}}, cl_{\mathcal{F}})$, $\mathcal{G} = (A_{\mathcal{G}}, R_{\mathcal{G}}, id)$. The CAF $\mathcal{F}$ is given by

$$\begin{aligned} A_{\mathcal{F}} = {} & Y \cup \bar{Y} \cup \{v_i \mid v \in cl_i, cl_i \in C\} \cup \\ & \{\bar{v}_i \mid \neg v \in cl_i, cl_i \in C\}; \\ R_{\mathcal{F}} = {} & \{(v_i, \bar{v}_j), (\bar{v}_j, v_i), (v, \bar{v}_i), (\bar{v}_i, v), \\ & (v_i, \bar{v}), (\bar{v}, v_i) \mid v \in V; i, j \leq n\}; \end{aligned}$$

and $cl_{\mathcal{F}}(v_i) = cl_{\mathcal{F}}(\bar{v}_i) = i$, $cl_{\mathcal{F}}(y) = y$, and $cl_{\mathcal{F}}(\bar{y}) = \bar{y}$. We construct a CAF $\mathcal{G}$ having the h-naive claim-sets $Y' \cup \{\bar{y} \mid y \notin Y'\} \cup \{1, \dots, n\}$ for every $Y' \subseteq Y$ by setting $A_{\mathcal{G}} = Y \cup \bar{Y} \cup \{1, \dots, n\}$ and $R_{\mathcal{G}} = \{(y, \bar{y}), (\bar{y}, y) \mid y \in Y\}$.

First assume Ψ is valid. Fix some $Y' \subseteq Y$. Since Ψ is valid, there is $Z' \subseteq Z$ such that $M = Y' \cup Z'$ is a model of φ. Let $E = Y' \cup \{\bar{y} \mid y \notin Y'\} \cup \{v_i \mid v \in M\} \cup \{\bar{v}_i \mid v \notin M\}$. E is conflict-free since conflicts appear only between arguments representing negated literals; moreover, since M is a model of φ, we have that for each clause cl_i there is either a positive literal $v \in cl_i$ with $v \in M$ or a negative literal $\bar{v} \in cl_i$ with $v \notin M$. Thus $\{1, \ldots, n\} \subseteq cl_{\mathcal{F}}(E)$; moreover, $Y' \cup \{\bar{y} \mid y \notin Y'\} \subseteq cl_{\mathcal{F}}(E)$. $S = cl_{\mathcal{F}}(E)$ is a maximal h-conflict-free claim-set since $S \cup \{c\} \notin cf_i(\mathcal{F})$ for any $c \in (Y \cup \bar{Y}) \setminus S$ as each realization of $S \cup \{c\}$ contains y, $\bar{y}$ for some $y \in Y$. It follows that $Y' \cup \{\bar{y} \mid y \notin Y'\} \cup \{1, \ldots, n\} \in na_h(\mathcal{F})$ for every $Y' \subseteq Y$. Moreover, there is no other h-naive claim-set of $\mathcal{F}$ since every proper superset has no cf-realisation in $\mathcal{F}$ as outlined above. We have shown $na_h(\mathcal{F}) = na_h(\mathcal{G})$ in case Ψ is valid.

Now assume $na_h(\mathcal{F}) = na_h(\mathcal{G})$ and fix some $Y' \subseteq Y$. Let E be some cf-realisation of $S = Y' \cup \{\bar{y} \mid y \notin Y'\} \cup \{1, \ldots, n\}$, let $Z' = \{z \mid z_i \in E\}$ and let $M = Y' \cup Z'$. Now, consider an arbitrary clause cl_i. Since E cf-realises S, there is some argument with claim i in E, that is, either $v_i \in E$ or $\bar{v}_i \in E$ for some $v \in Y \cup Z$ (observe that $y_i \in E$ iff $y \in E$ and $\bar{y}_i \in E$ iff $\bar{y} \in E$, thus a further case distinction for $y \in Y$, $\bar{y} \in \bar{Y}$ is not required). In the former case, we have $v \in M$ and thus M satisfies cl_i, in the latter case $v \notin M$ and thus cl_i is satisfied. We obtain that M is a model of φ.

Since conflict-free semantics satisfy downward closure (each subset of a conflict-free set is conflict-free), we have $cf_i(\mathcal{F}) = cf_i(\mathcal{G})$ iff $na_h(\mathcal{F}) = na_h(\mathcal{G})$ and thus the statement follows for i-conflict-free semantics. By symmetry of $\mathcal{F}$ and $\mathcal{G}$ we furthermore have $ad(F) = cf(F)$ and $ad(G) = cf(G)$ which implies $ad_i(\mathcal{F}) = cf_i(\mathcal{F})$, $ad_i(\mathcal{G}) = cf_i(\mathcal{G})$, $pr_h(\mathcal{F}) = na_h(\mathcal{F})$, and $pr_h(\mathcal{G}) = na_h(\mathcal{G})$. Thus Π_2^{P}-hardness of Ver-OE$_\rho$ for i-admissible and h-preferred semantics follow. □

It remains to give the proof for complete semantics.

Proposition A.5. *Deciding* Ver-OE$_{co_i}$ *is* Π_2^{P}*-hard.*

Proof. Consider an instance $\Psi = \forall Y \exists Z \varphi(Y, Z)$ of QSAT$_2^\forall$, where φ is given by a set of clauses $C = \{cl_1, \ldots, cl_n\}$ over atoms in $V = Y \cup Z$. We may assume that $Z \neq \emptyset$; i.e., there is some $z_0 \in Z$. We construct two CAFs $\mathcal{F} = (A_{\mathcal{F}}, R_{\mathcal{F}}, cl_{\mathcal{F}})$, $\mathcal{G} = (A_{\mathcal{G}}, R_{\mathcal{G}}, cl_{\mathcal{G}})$, where $\mathcal{F}$ is a modification of the standard construction (A, R) (cf. Reduction A.1) with

$$\begin{aligned} A_{\mathcal{F}} &= A \cup \{\bar{\varphi}\} \cup \{d_v \mid v \in V\}; \\ R_{\mathcal{F}} &= R \cup \{(cl, cl) \mid cl \in C\} \cup \{(d_v, d_v), (v, d_v), (\bar{v}, d_v), \\ &\qquad (d_v, a) \mid v \in V, a \in V \cup \bar{V}\} \cup \{(\varphi, \bar{\varphi}), (\bar{\varphi}, \varphi)\}; \end{aligned}$$

$cl_{\mathcal{F}}(\bar{\varphi}) = z_0$, $cl_{\mathcal{F}}(z) = cl_{\mathcal{F}}(\bar{z}) = z$ and $cl_{\mathcal{F}}(a) = a$ else. We observe that $co_i(\mathcal{F})$ contains $Y' \cup \{\bar{y} \mid y \notin Y'\} \cup Z$ for each $Y' \subseteq Y$ as well as $\emptyset$. A witness is given by the complete extension $Y' \cup \{\bar{y} \mid y \notin Y'\} \cup Z \cup \{\bar{\varphi}\}$. Moreover, since at least one of $v, \bar{v}$ has to be contained in a complete extension E in order to be defended we observe that no subset of any $Y' \cup \{\bar{y} \mid y \notin Y'\} \cup Z$, $Y' \subseteq Y$, is i-complete in $\mathcal{F}$.

The CAF $\mathcal{G}$ is given by

$$\begin{aligned}A_{\mathcal{G}} =& Y \cup \bar{Y} \cup Z \cup \{\varphi, \bar{\varphi}, d_\varphi\} \cup \{d_y \mid y \in Y\};\\ R_{\mathcal{G}} =& \{(y, \bar{y}), (\bar{y}, y) \mid y \in Y\}) \cup \{(d_v, d_v), (v, d_v), (\bar{v}, d_v),\\ & (d_v, a) \mid v \in Y \cup \{\varphi\}, a \in A_{\mathcal{G}}\};\end{aligned}$$

and $cl_{\mathcal{G}}(\bar{\varphi}) = z_0$ and $cl_{\mathcal{G}}(a) = a$ else. Observe that $\mathcal{G}$ contains the i-complete claim-sets $Y' \cup \{\bar{y} \mid y \notin Y'\} \cup \{\varphi\} \cup Z$ and $Y' \cup \{\bar{y} \mid y \notin Y'\} \cup Z$ for $Y' \subseteq Y$ as well as the empty claim-set $\emptyset$. Given a complete extension $E \neq \emptyset$ of G, we observe that either y or $\bar{y}$ is contained in E for every $y \in Y$ since every $a \in Y \cup \bar{Y} \cup \{\varphi\} \cup Z$ must be defended against d_y; similarly, either φ or $\bar{\varphi}$ is contained in E. Thus there is some $Y' \subseteq Y$ such that $Y' \cup \{\bar{y} \mid y \notin Y'\} \subseteq E$. In case $\varphi \in E$ we have that E is of the form $Y' \cup \{\bar{y} \mid y \notin Y'\} \cup \{\varphi\} \cup Z$ for some $Y' \subseteq Y$ since each $Y' \cup \{\bar{y} \mid y \notin Y'\} \cup \{\varphi\}$ defends itself and Z in G; in case $\bar{\varphi} \in E$ we have that E is of the form $Y' \cup \{\bar{y} \mid y \notin Y'\} \cup \{\bar{\varphi}\} \cup Z$ for some $Y' \subseteq Y$ since each $Y' \cup \{\bar{y} \mid y \notin Y'\} \cup \{\bar{\varphi}\}$ defends itself and Z in G. In the former case, $cl_{\mathcal{G}}(E) = Y' \cup \{\bar{y} \mid y \notin Y'\} \cup \{\varphi\} \cup Z$, in the latter case, $cl_{\mathcal{G}}(E) = Y' \cup \{\bar{y} \mid y \notin Y'\} \cup Z$.

We show Ψ is valid iff $Y' \cup \{\bar{y} \mid y \notin Y'\} \cup Z \cup \{\varphi\} \in co_i(\mathcal{F})$ for each $Y' \subseteq Y$.

Assume Ψ is valid; fix some $Y' \subseteq Y$. Then there is $Z' \subseteq Z$ such that $M = Y' \cup Z'$ is a model of φ. We show that $E = M \cup \{\bar{v} \mid v \notin M\} \cup \{\varphi\}$ is complete in F: E is conflict-free; moreover, since M is a model of φ we have that each clause $cl \in C$ is attacked; consequently, E defends φ against each attack. E contains each defended argument since it attacks any remaining argument $a \notin E$ in F. Thus $cl_{\mathcal{F}}(E) = Y' \cup \{\bar{y} \mid y \notin Y'\} \cup \{\varphi\} \cup Z \in co_i(\mathcal{F})$. As Y' was arbitrary, we have shown $Y' \cup \{\bar{y} \mid y \notin Y'\} \cup Z \cup \{\varphi\} \in co_i(\mathcal{F})$ for each $Y' \subseteq Y$.

Now assume $Y' \cup \{\bar{y} \mid y \notin Y'\} \cup Z \cup \{\varphi\} \in co_i(\mathcal{F})$ for each $Y' \subseteq Y$. Fix some $Y' \subseteq Y$ and let E be the complete realization of $Y' \cup \{\bar{y} \mid y \notin Y'\} \cup Z \cup \{\varphi\}$ in $\mathcal{F}$. We show that $M = Y' \cup Z'$ with $Z' = E \cap Z$ is a model of φ: From $\varphi \in E$ we obtain that every clause $cl \in C$ is attacked; that is, for every $cl \in C$, there is $a \in E$ with $(a, cl) \in R_{\mathcal{G}}$. In case $a = v$ for some $v \in V$, we have $v \in M \cap cl$; in case $a = \bar{v}$ for some $v \in V$ we have $\neg v \in cl$ and $v \notin M$—in both cases, cl is satisfied, thus M is a model of φ. It follows that Ψ is valid.

As outlined above, $co_i(\mathcal{F})$ contains $Y' \cup \{\bar{y} \mid y \notin Y'\} \cup Z$ for each $Y' \subseteq Y$, moreover, $\emptyset \in co_i(\mathcal{F})$ and $Y' \cup \{\bar{y} \mid y \notin Y'\} \cup Z \cup \{\varphi\} \in co_i(\mathcal{F})$ for each $Y' \subseteq Y$ iff Ψ is valid. By design of $\mathcal{G}$ we obtain Ψ is valid iff $co_i(\mathcal{F}) = co_i(\mathcal{G})$. □

This ends our proofs for Π_2^{P}-hardness. It remains to prove Π_3^{P}-hardness for i-preferred, h-semi-stable, and h-stage semantics. We will make use of the following translations (Dvořák et al., 2021, Definition 8).

Definition A.6. For a CAF $\mathcal{F} = (A, R, cl)$, we define three translations:

- $Tr_1(\mathcal{F}) = (A', R', cl')$ with

$$\begin{aligned}A' =& A \cup \{a' \mid a \in A\}\\ R' =& R \cup \{(a, a'), (a', a') \mid a \in A\}\end{aligned}$$

  and $cl'(a) = cl(a)$ for $a \in A$, $cl(a') = c_a$ for $a' \in \{a' \mid a \in A\}$ with fresh claims $c_a \notin cl(A)$.

- $Tr_2(\mathcal{F}) = (A', R'_2, cl')$ with

$$\begin{aligned} A' =& A \cup \{a' \mid a \in A\} \\ R'_2 =& R' \cup \{(a, b') \mid (a, b) \in R\}; \end{aligned}$$

  and cl' as before.

- $Tr_3(CF) = (A', R'_3, cl')$ with

$$\begin{aligned} A' =& A \cup \{a' \mid a \in A\} \\ R'_3 =& R'_2 \cup \{(b, a) \mid (a, b) \in R\} \cup \{(a, b) \mid a \in A, (b, b) \in R\}; \end{aligned}$$

  and cl' as before.

As shown by Dvořák et al. (2021), it holds that (a) Tr_1 maps i-preferred semantics to h-semi-stable semantics, (b) Tr_2 maps inherited to hybrid stable semantics, and (c) Tr_3 maps inherited to hybrid stage semantics.

Lemma A.7. *For a CAF* $\mathcal{F} = (A, R, cl)$,

$$\begin{aligned} pr_i(\mathcal{F}) &= pr_i(Tr_1(\mathcal{F})) = ss_h(Tr_1(\mathcal{F})), \\ stb_i(\mathcal{F}) &= stb_i(Tr_2(\mathcal{F})) = \tau\text{-}stb_h(Tr_2(\mathcal{F})) \text{ for } \tau \in \{ad, cf\}, \\ stg_i(\mathcal{F}) &= stg_i(Tr_3(\mathcal{F})) = stg_h(Tr_3(\mathcal{F})). \end{aligned}$$

Since all translations can be computed in polynomial time, we obtain lower bounds for Ver-OE$_\rho$ $\rho \in \{ss_h, stg_h\}$: to decide $\sigma_h(\mathcal{F}) = \sigma_h(\mathcal{G})$ for two CAFs $\mathcal{F}$ and $\mathcal{G}$, we can apply the respective translation T_i and check $\sigma_i(T_i(\mathcal{F})) = \sigma_i(T_i(\mathcal{G}))$ instead.

Proposition A.8. *Deciding* Ver-OE$_\rho$ *is* Π_3^{P}*-hard for* $\rho \in \{ss_h, stg_h\}$.

Proposition A.9. *Deciding* Ver-OE$_{pr_i}$ *is* Π_3^{P}*-hard.*

Proof. We show hardness via a reduction from QSAT$_3^\exists$.

Consider an instance $\Psi = \exists X \forall Y \exists Z \varphi(X, Y, Z)$ of QSAT$_3^\exists$, where φ is given by a set of clauses C over atoms in $V = X \cup Y \cup Z$. W.l.o.g., we can assume there is $y_0 \in Y$ which is contained in each clause $cl \in C$. We construct two CAFs $\mathcal{F} = (A_\mathcal{F}, R_\mathcal{F}, cl_\mathcal{F})$, $\mathcal{G} = (A_\mathcal{G}, R_\mathcal{G}, cl_\mathcal{G})$. Let (A, R) be given as in Reduction A.2. We construct $\mathcal{F}$ with

$$\begin{aligned} A_\mathcal{F} &= A \cup \{\varphi, \bar{\varphi}\}; \\ R_\mathcal{F} &= R \cup \{(cl, \varphi) \mid cl \in C\} \cup \{(\varphi, \bar{\varphi}), (\bar{\varphi}, \bar{\varphi})\} \cup \\ &\quad\ \{(\bar{\varphi}, z), (\bar{\varphi}, \bar{z}) \mid z \in Z\}; \end{aligned}$$

and $cl_\mathcal{F}(y) = cl_\mathcal{F}(\bar{y}) = y$ for $y \in Y$ and $cl_\mathcal{F}(v) = v$ else; that is, $\mathcal{F}$ is the standard construction for preferred semantics on AF level.

We construct the CAF $\mathcal{G}$ such that $pr_i(\mathcal{G}) = \{V' \cup \{\bar{v} \mid v \notin V'\} \cup Y \cup \{\varphi\} \mid V' \subseteq X \cup Z\} \cup \{X' \cup \{\bar{x} \mid x \notin X'\} \cup Y \mid X' \subseteq X\}$. This can be achieved by setting

$$A_\mathcal{G} = X_i \cup \bar{X}_i \cup Y \cup Z \cup \bar{Z} \cup \{\varphi\}$$

for two copies X_i, $\bar{X}_i$, $i \leq 2$, of X and $\bar{X}$, respectively;

$$\begin{aligned} R_{\mathcal{G}} = \{&(v_i, \bar{v}_j), (\bar{v}_i, v_j) \mid v_i, v_j \in X_1 \cup X_2\} \cup \\ \{&(v, \bar{v}), (\bar{v}, v) \mid v \in Z\} \cup \\ \{&(a, b), (b, a) \mid a \in A' \cup \{\varphi\}, b \in X_2 \cup \bar{X}_2\} \end{aligned}$$

where $A' = X_1 \cup \bar{X}_1 \cup Z \cup \bar{Z}$; moreover, $cl_{\mathcal{G}}(x_i) = x$, $cl_{\mathcal{G}}(\bar{x}_i) = \bar{x}$, and $cl_{\mathcal{G}}(a) = a$ for all remaining $a \in A_{\mathcal{G}}$.

First observe that $\{V' \cup \{\bar{v} \mid v \notin V'\} \cup Y \cup \{\varphi\} \mid V' \subseteq X \cup Z\} \subseteq pr_i(\mathcal{F})$ since $y_0 \in cl$ for every clause cl, that is, for every atom $v \in V \setminus \{y_0\}$, we can choose either v or $\bar{v}$ as long as y_0 is contained in $E \subseteq A_{\mathcal{F}}$, we have that E defends φ against each attack.

In case Ψ is not valid, consider some $X' \subseteq X$. Since Ψ is not valid, there is some $Y' \subseteq Y$ such that for all $Z' \subseteq Z$, some clause $cl \in C$ is not satisfied. It follows that $E = X' \cup \{\bar{x} \mid x \notin X'\} \cup Y' \cup \{\bar{y} \mid y \notin Y'\}$ is preferred in F: Clearly, E is conflict-free and defends itself. Now assume there is $a \in A \setminus E$ such that $E \cup \{a\} \in ad(F)$. In case $a = \varphi$ we have that each $cl \in C$ is attacked, that is, for every clause $cl \in C$ there is $v \in X' \cup Y'$ such that either $v \in X' \cup Y'$ with $v \in cl$ or $v \notin X' \cup Y'$ with $\neq v \in cl$. Thus $X' \cup Y'$ is a model of φ, contradiction to Ψ being not valid. Observe that the case $a \in Z \cup \bar{Z}$ requires $\varphi \in E$, otherwise a is not defended against $\bar{\varphi}$. We have thus shown that $cl(E) = X' \cup \{\bar{x} \mid x \notin X'\} \cup Y \in pr_i(\mathcal{F})$ for every $X' \subseteq X$.

We show that every i-preferred set of $\mathcal{F}$ is either of the form (a) $V' \cup \{\bar{v} \mid v \notin V'\} \cup Y \cup \{\varphi\}$ for some $V' \subseteq X \cup Z$ or (b) $X' \cup \{\bar{x} \mid x \notin X'\} \cup Y$ for some $X' \subseteq X$. As outlined above, any such set is i-preferred in $\mathcal{F}$, thus it remains to show that there is no other i-preferred set in $\mathcal{F}$. First notice that each i-preferred claim-set of $\mathcal{F}$ contains $X' \cup \{\bar{x} \mid x \notin X'\} \cup Y$ for some $X' \subseteq X$ since every preferred set E of F contains $X' \cup \{\bar{x} \mid x \notin X'\} \cup Y' \cup \{\bar{y} \mid y \notin Y'\}$ for some $X' \subseteq X$, $Y' \subseteq Y$ by construction. Now assume there is $S \subseteq cl(A_{\mathcal{F}})$ such that $S \in pr_i(\mathcal{F})$ and S is not of the form (a) or (b). Let E be a pr_i-realisation of S. First assume $\varphi \notin E$. Then $z, \bar{z} \notin E$ for any $z \in Z$ since φ is the only argument which defends $z, \bar{z}$ against $\bar{\varphi}$. By the above consideration there are $X' \subseteq X$, $Y' \subseteq Y$ such that $X' \cup \{\bar{x} \mid x \notin X'\} \cup Y' \cup \{\bar{y} \mid y \notin Y'\} \subseteq E$. Observe that $a \notin E$ for any $a \in (X \setminus X') \cup \{\bar{x} \mid x \in X'\} \cup (Y \setminus Y') \cup \{\bar{y} \mid y \in Y'\}$ since $v, \bar{v}$ are mutually attacking for any $v \in X \cup Y$. Since every remaining argument is either attacked by E or self-attacking it follows that $S = X' \cup \{\bar{x} \mid x \notin X'\} \cup Y$. In case $\varphi \in E$, we have that every $z, \bar{z}$ is defended against $\bar{\varphi}$. Thus E contains either z or $\bar{z}$ for every $z \in Z$ by subset-maximality of E. Thus there is $Z' \subseteq Z$ such that $E = V' \cup \{\bar{v} \mid v \notin V'\} \cup \{\varphi\}$. Since every remaining argument is either attacked by E or self-attacking, we have $S = V' \cup \{\bar{v} \mid v \notin V'\} \cup Y \cup \{\varphi\}$ for some $V' \subseteq X \cup Z$. It follows that $pr_i(\mathcal{F}) = pr_i(\mathcal{G})$.

Now assume $pr_i(\mathcal{F}) = pr_i(\mathcal{G})$ and consider some $X' \subseteq X$. Let E be a pr_i-realisation of $X' \cup \{\bar{x} \mid x \notin X'\} \cup Y$ and let $Y' = E \cap Y$. We show that for all $Z' \subseteq Z$, $X' \cup Y' \cup Z'$ is not a model of φ. Fix some $Z' \subseteq Z$ and let $M = X' \cup Y' \cup Z'$. Since E is preferred in $\mathcal{F}$ we have that φ is not defended by $E \cup Z' \cup \{\bar{z} \mid z \notin Z'\}$; i.e., there is some $cl \in C$ such that $E \cup Z' \cup \{\bar{z} \mid z \notin Z'\}$ does not attack cl. Consequently, for all $v \in V$, in case $v \in cl$ we have $v \notin M$, and in case $\neq v \in cl$ we have $v \in M$. It follows that M is not a model of φ.

It follows that Ψ is not valid iff $pr_i(\mathcal{F}) = pr_i(\mathcal{G})$. □

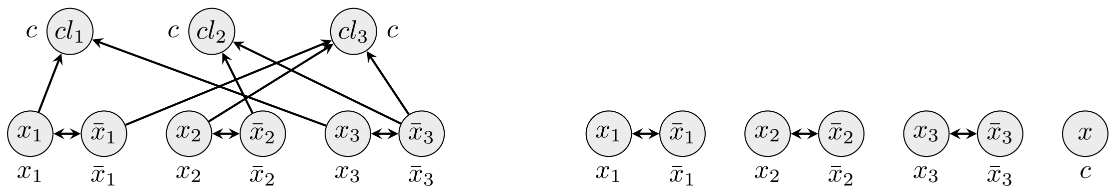

Figure 6: CAFs $\mathcal{F}$ (left) and $\mathcal{G}$ (right) illustrating the reduction from the Proof of Proposition A.10 for the formula φ given by the clauses $\{\{x_1, x_3\}, \{\bar{x}_2, \bar{x}_3)\}, \{\bar{x}_1, \bar{x}_3, x_2\}\}$.

Let us now turn to well-formed CAFs. We provide the omitted proofs of the following theorem:

Theorem 6.6. $\textsc{Ver-OE}^{wf}_{\rho}$ *is*

- *in* P *for* $\rho = gr_i$;
- coNP-*complete for* $\rho \in \{cf_i, ad_i, co_i, na_i, na_h, stb_i, cf\text{-}stb_h, ad\text{-}stb_h\}$; *and*
- Π_2^{P}-*complete for* $\rho \in \{pr_i, ss_i, stg_i, pr_h, stg_h, ss_h\}$.

For h-naive and conflict-free semantics, we utilize the standard construction once again.

Proposition A.10. $\textsc{Ver-OE}^{wf}_{\rho}$ *is* coNP-*hard for* $\rho \in \{cf_i, na_h\}$.

Proof. Consider a SAT instance φ given by a set of clauses C over atoms in X. We may assume that there is no clause $cl \in C$ such that $x, \bar{x} \in cl$ for any atom $x \in X$. We construct two CAFs. For $\mathcal{F}$, construct (A, R, cl) as follows:

$$A_{\mathcal{F}} = X \cup \bar{X} \cup C \text{ with } \bar{X} = \{\bar{x} \mid x \in X\};$$
$$R_{\mathcal{F}} = \{(x, cl) \mid cl \in C, v \in cl\} \cup \{(\bar{x}, cl) \mid cl \in C, \neg x \in cl\} \cup \{(x, \bar{x}), (\bar{x}, x) \mid x \in X\}$$

and $cl_{\mathcal{F}}(cl) = c$ for $cl \in C$ and $cl_{\mathcal{F}}(a) = a$ otherwise; The CAF $\mathcal{G}$ is given by $A_{\mathcal{G}} = X \cup \bar{X} \cup \{c\}$, $R_{\mathcal{G}} = \{(x, \bar{x}), (\bar{x}, x) \mid x \in X\}$, and $cl_{\mathcal{G}} = id$. Then $na(G) = na_i(\mathcal{G}) = \{X' \cup \bar{X}' \cup \{c\} \mid X' \subseteq X\}$. An example is given in Figure 6.

Let us consider the naive extensions of the underlying AF F of $\mathcal{F}$. We observe that each naive extension contains x or $\bar{x}$ for each literal x (recall that we excluded clauses containing both x and $\bar{x}$). Hence each set $X' \cup \bar{X}'$, $X' \subseteq X$, is conflict-free in F. Moreover, we can add one of the clause arguments $cl \in C$ to the set $X' \cup \bar{X}'$ iff the set X' is *not* a model of φ: by construction, X' is a model of φ iff all clause arguments are attacked. It holds that $X' \cup \bar{X}'$ is naive in F iff all clause-arguments are attacked iff X' is a model of φ.

Hence φ is unsatisfiable iff $X' \cup \bar{X}' \cup \{cl\}$ is naive for some $cl \in C$. Hence, we obtain the desired result for h-naive semantics: φ is unsatisfiable iff $na_h(\mathcal{F}) = na_h(\mathcal{G})$. Since each subset of a naive extension is conflict-free, the statement for conflict-free semantics follows. □

For i-naive semantics, we utilize a construction by Kiesel and Rapberger (2021).

Proposition A.11. $\textsc{Ver-OE}^{wf}_{na_i}$ *is* coNP-*hard.*

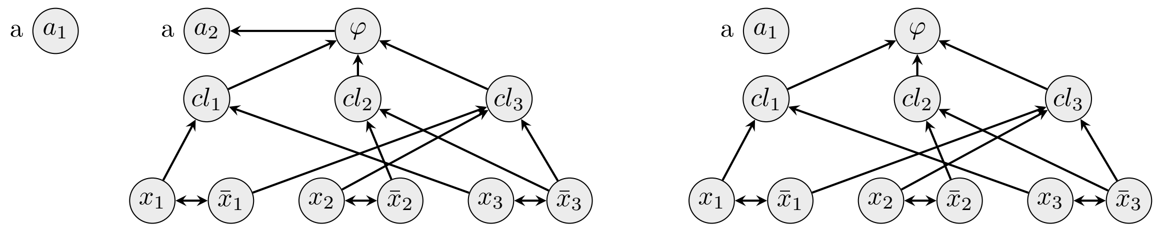

Figure 7: CAF $\mathcal{F}$ and $\mathcal{G}$ from the proof of Proposition A.11 for a formula φ which is given by the clauses $\{\{x_1, x_3\}, \{\bar{x}_3, \bar{x}_2)\}, \{\bar{x}_1, \bar{x}_3, x_2\}\}$.

Proof. Consider an UNSAT instance φ given by clauses C over variables in X. We let $\mathcal{F}$ be defined as follows:

$$\begin{aligned} A_{\mathcal{F}} =& X \cup \bar{X} \cup C \cup \{\varphi\} \cup \{a_1, a_2\}, \text{ with } \bar{X} = \{\bar{x} \mid x \in X\}, \\ R_{\mathcal{F}} =& \{(x, cl) \mid cl \in C, x \in cl\} \cup \{(\bar{x}, cl) \mid cl \in C, \neg x \in cl\} \\ & \cup \{(x, \bar{x}), (\bar{x}, x) \mid x \in X\} \cup \{(cl_i, \varphi) \mid i \leq n\} \cup \{(\varphi, a_2)\}, \end{aligned}$$

with $cl(a_1) = cl(a_2) = a$ and $cl(x) = x$ otherwise. We obtain $\mathcal{G}$ from $\mathcal{F}$ by setting $A_{\mathcal{G}} = A_{\mathcal{F}} \setminus \{a_2\}$ and $R_{\mathcal{G}} = R_{\mathcal{F}} \setminus \{(\varphi, a_2)\}$. The resulting CAF $\mathcal{G}$ has a unique claim per argument. An example of this construction is given in Figure 7.

We can compute the naive extensions of F from G as follows: We obtain $na(F)$ from $na(G)$ by (1) taking all extensions containing φ, i.e., $E \in na(F)$ for all $E \in na(G)$ with $\varphi \in E$; (2) replacing φ in each extension by a_2, i.e., $(E \setminus \{\varphi\}) \cup \{a_2\} \in na(F)$ for all $E \in na(G)$ with $\varphi \in E$; and (3) adding a_2 to all naive extensions of E not containing φ, i.e., $E \cup \{a_2\} \in na(F)$ for all $E \in na(G)$ with $\varphi \notin E$. Hence $na(F) = \{E \mid E \in na(G), \varphi \in E\} \cup \{(E \setminus \{\varphi\}) \cup \{a_2\} \mid E \in na(G), \varphi \in E\} \cup \{E \cup \{a_2\} \mid E \in na(G), \varphi \notin E\}$.

In case (1) and (3), the set E (and its modified version) has the same claims in $\mathcal{F}$ and $\mathcal{G}$ (for the latter, observe that a_1 is contained in each extension thus each set contains claim a). Now, consider a set $E \in na(G)$ with $\varphi \in E$, and let $E' = (E \setminus \{\varphi\}) \cup \{a_2\}$. It holds that $cl(E') = X' \cup \bar{X}' \cup \{a\}$ for some $X' \subseteq X$. Observe that E' is not naive in G since E is a proper superset of it. On the other hand, the set E' is naive in $\mathcal{F}$ iff it is in conflict with each clause-argument, i.e., iff E' attacks each $cl_i \in C$. By well-known results (Dvořák & Dunne, 2018), this is the case iff X' is a model of φ, i.e., iff φ is satisfiable. Hence $na_i(\mathcal{F}) = na_i(\mathcal{G})$ iff φ is unsatisfiable. □

To obtain Π_2^{P}-hardness for i-semi-stable and i-stage semantics for well-formed CAFs we proceed analogous to the proofs of Propositions 6.3 and 6.5: again, we utilize constructions used for showing hardness of the concurrence problem (Dvořák et al., 2021) for the respective semantics.

For hybrid semi-stable and stage semantics, we utilize translations Tr_1 and Tr_3 (Dvořák et al., 2021) to obtain Π_2^{P}-hardness of deciding ordinary equivalence.

We will make use of the following reduction.

Reduction A.12. Let $\Psi = \forall Y \exists Z \varphi(Y, Z)$ be an instance of $\text{QSAT}_2^{\forall}$, where φ is given by a set of clauses $C = \{cl_1, \ldots, cl_n\}$ over atoms in $X = Y \cup Z$. Let (A, R) be the AF constructed

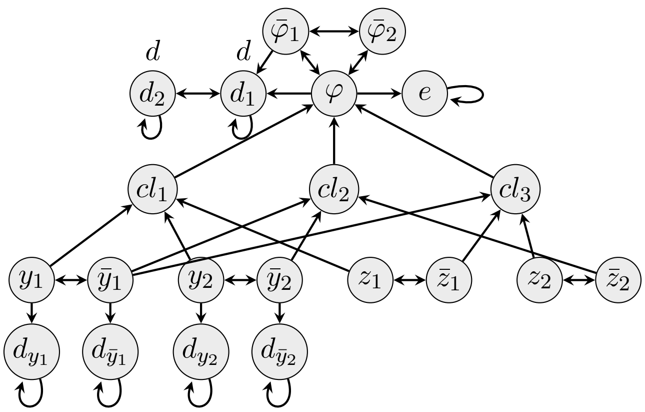

Figure 8: Reduction A.12 for the formula $\forall Y \exists Z \varphi(Y, Z)$ where $\varphi(Y, Z)$ is given by the clauses $\{\{z_1, y_1, y_2\}, \{\bar{y}_1, \bar{y}_2, \bar{z}_2)\}, \{\bar{z}_1, \bar{y}_1, z_2\}\}$. Since $cl(a) = a$ for all arguments $a \in A \setminus \{d_1, d_2\}$, we omit all claims that coincide with the arguments name.

from φ as in Reduction A.1. We define $\mathcal{F} = (A', R', cl)$ with

$$\begin{aligned} A' &= A \cup \{e, d_1, d_2, \bar{\varphi}_1, \bar{\varphi}_2\} \\ R' &= R \cup \{(a, d_a)(d_a, d_a) \mid a \in Y \cup \bar{Y}\} \cup \{(d_i, d_j) \mid i, j = 1, 2\} \cup \\ &\quad \{(a, b) \mid a, b \in \{\varphi, \bar{\varphi}_1, \bar{\varphi}_2\}, a \neq b\} \cup \{(\varphi, e), (e, e), (\varphi, d_1), (\bar{\varphi}_1, d_1)\} \end{aligned}$$

and $cl(d_1) = cl(d_2) = d$ and $cl(v) = v$ otherwise.

See Figure 8 for an example of the reduction.

Proposition A.13. Ver-OE$_\rho^{wf}$ *is* Π_2^P*-hard for* $\rho \in \{ss_i, stg_i, ss_h, stg_h\}$.

Proof. Let $\rho = ss_i$ and let $\Psi = \forall Y \exists Z \varphi(Y, Z)$ be an instance of $\text{QSAT}_2^\forall$ where φ is given by a set of clauses C over atoms in $V = Y \cup Z$. Similar as in the proof of Proposition 6.5, we make use of the complexity results for the concurrence problem: it has been shown that Ψ is valid iff $ss_i(\mathcal{F}) = ss_h(\mathcal{F})$ (Dvořák et al., 2021). It suffices to construct the CAF $\mathcal{G}$ in such a way that $ss_i(\mathcal{G}) = ss_h(\mathcal{F})$. Then it holds that Ψ is valid iff $ss_i(\mathcal{G}) = ss_h(\mathcal{F}) = ss_i(\mathcal{F})$.

To construct the CAF $\mathcal{G}$ appropriately, let us observe that the arguments d_1 and d_2 do not influence hybrid semi-stable semantics at all; since d_2 is neither contained nor attacked by any conflict-free set, they are invisible from the perspective of h-semi-stable semantics. Hence we can remove both arguments without influencing the outcome. This yields our desired CAF $\mathcal{G} = (A_\mathcal{G}, R_\mathcal{G}, id)$: we let $A_\mathcal{G} = A_\mathcal{F} \setminus \{d_1, d_2\}$ and $R_\mathcal{G} = R_\mathcal{F} \cap A_\mathcal{G}^2$. We obtain Ψ is valid iff $ss_i(\mathcal{G}) = ss_h(\mathcal{F})$.

This concludes the proof for i-semi-stable semantics. For i-stage semantics, we note that $ss_i(\mathcal{F}) = stg_i(\mathcal{F})$ and $ss_h(\mathcal{F}) = stg_h(\mathcal{F})$ (Dvořák et al., 2021). For the hybrid counter-parts, we utilize translations Tr_1 and Tr_3 (cf. Definition A.6). □

It remains to prove Π_2^P-hardness for preferred semantics. Deciding ordinary equivalence w.r.t. preferred semantics is Π_2^P-hard for the class of AFs (Baumann et al., 2019). Hence we obtain the following result.

Proposition A.14. Ver-OE$_{pr_i}^{wf}$ *is* Π_2^P*-hard.*

References

Alfano, G., Greco, S., & Parisi, F. (2021a). Incremental computation in dynamic argumentation frameworks. *IEEE Intelligent Systems*, *36*(6), 80–86.

Alfano, G., Greco, S., Parisi, F., Simari, G. I., & Simari, G. R. (2021b). Incremental computation for structured argumentation over dynamic DeLP knowledge bases. *Artificial Intelligence*, *300*, 103553.

Amgoud, L., Besnard, P., & Vesic, S. (2014). Equivalence in logic-based argumentation. *Journal of Applied Non-Classical Logics*, *24*(3), 181–208.

Baroni, P., Cerutti, F., Giacomin, M., & Guida, G. (2011). Afra: Argumentation framework with recursive attacks. *International Journal of Approximate Reasoning*, *52*(1), 19–37.

Baroni, P., & Riveret, R. (2019). Enhancing statement evaluation in argumentation via multi-labelling systems. *Journal of Artificial Intelligence Research*, *66*, 793–860.

Baumann, R. (2011). Splitting an argumentation framework. In *Proceedings of the 11th International Conference on Logic Programming and Nonmonotonic Reasoning (LPNMR 2011)*, Vol. 6645 of *Lecture Notes in Computer Science*, pp. 40–53. Springer.

Baumann, R. (2012a). Normal and strong expansion equivalence for argumentation frameworks. *Artificial Intelligence*, *193*, 18–44.

Baumann, R. (2012b). What does it take to enforce an argument? Minimal change in abstract argumentation. In *Proceedings of the 20th European Conference on Artificial Intelligence (ECAI 2012)*, Vol. 242 of *Frontiers in Artificial Intelligence and Applications*, pp. 127–132. IOS Press.

Baumann, R. (2018). On the nature of argumentation semantics: Existence and uniqueness, expressibility, and replaceability. In *Handbook of Formal Argumentation*, chap. 14. College Publications.

Baumann, R., & Brewka, G. (2015). AGM meets abstract argumentation: Expansion and revision for Dung frameworks. In *Proceedings of the 24th International Conference on Artificial Intelligence (IJCAI 2015)*, pp. 2734–2740. AAAI Press.

Baumann, R., Dvořák, W., Linsbichler, T., & Woltran, S. (2019). A general notion of equivalence for abstract argumentation. *Artificial Intelligence*, *275*, 379–410.

Baumann, R., Linsbichler, T., & Woltran, S. (2016). Verifiability of argumentation semantics. In *Proceedings of the 6th International Conference of Computational Models of Argument (COMMA 2016)*, pp. 83–94.

Baumann, R., Rapberger, A., & Ulbricht, M. (2022). Equivalence in argumentation frameworks with a claim-centric view - classical results with novel ingredients. In *Priceedings of the 36th AAAI Conference on Artificial Intelligence (AAAI 2022)*, pp. 5479–5486. AAAI Press.

Baumann, R., & Spanring, C. (2017). A study of unrestricted abstract argumentation frameworks. In *Proceedings of the 26th International Joint Conference on Artificial Intelligence (IJCAI 2017)*, pp. 807–813.

Baumann, R., & Strass, H. (2022). An abstract, logical approach to characterizing strong equivalence in non-monotonic knowledge representation formalisms. *Artificial Intelligence*, *305*, 103680.

Bernreiter, M., Dvořák, W., Rapberger, A., & Woltran, S. (2022). The effect of preferences in abstract argumentation under a claim-centric view. In *Proceedings of the 20th International Workshop on Non-Monotonic Reasoning (NMR 2022)*, Vol. 3197 of *CEUR Workshop Proceedings*, pp. 27–38. CEUR-WS.org.

Besnard, P., & Hunter, A. (2001). A logic-based theory of deductive arguments. *Artificial Intelligence*, *128*(1-2), 203–235.

Besnard, P., & Hunter, A. (2018). A review of argumentation based on deductive arguments. In *Handbook of Formal Argumentation*, chap. 9, pp. 436–484. College Publications.

Bondarenko, A., Toni, F., & Kowalski, R. A. (1993). An assumption-based framework for non-monotonic reasoning. In *Proceedings of the 2nd International Workshop on Logic Programming and Non-monotonic Reasoning (LPNMR 1993)*, pp. 171–189. MIT Press.

Caminada, M. (2006). Semi-stable semantics. In *Proceedings of the International Conference on Computational Models of Argument (COMMA 2006)*, Vol. 144 of *Frontiers in Artificial Intelligence and Applications*, pp. 121–130. IOS Press.

Caminada, M., Sá, S., Alcântara, J. F. L., & Dvořák, W. (2015). On the equivalence between logic programming semantics and argumentation semantics. *International Journal of Approximate Reasoning*, *58*, 87–111.

Cayrol, C., de Saint-Cyr, F. D., & Lagasquie-Schiex, M. (2010). Change in abstract argumentation frameworks: Adding an argument. *Journal of Artificial Intelligence Research*, *38*, 49–84.

Cayrol, C., & Lagasquie-Schiex, M.-C. (2005). On the acceptability of arguments in bipolar argumentation frameworks. In *Proceedings of the 8th European Conference on Symbolic and Quantitative Approaches to Reasoning and Uncertainty (ECSQARU 2005)*, pp. 378–389. Springer.

Coste-Marquis, S., Konieczny, S., Mailly, J., & Marquis, P. (2014). On the revision of argumentation systems: Minimal change of arguments statuses. In *Proceedings of the 14th International Conference on Principles of Knowledge Representation and Reasoning (KR 2014)*. AAAI Press.

Cyras, K., Fan, X., Schulz, C., & Toni, F. (2018). Assumption-based argumentation: Disputes, explanations, preferences. In *Handbook of Formal Argumentation*, chap. 7, pp. 365–408. College Publications. also appears in IfCoLog Journal of Logics and their Applications 4(8):2407–2456.

Cyras, K., & Toni, F. (2016). ABA+: Assumption-based argumentation with preferences. In *Proceedings of the 15th International Conference on Principles of Knowledge Representation and Reasoning (KR 2016)*, pp. 553–556. AAAI Press.

Dung, P. M. (1995). On the acceptability of arguments and its fundamental role in nonmonotonic reasoning, logic programming and n-person games. *Artificial Intelligence*, *77*(2), 321–357.

Dvořák, W., & Dunne, P. E. (2018). Computational problems in formal argumentation and their complexity. In Baroni, P., Gabbay, D., Giacomin, M., & van der Torre, L. (Eds.), *Handbook of Formal Argumentation*. College Publications.

Dvořák, W., Greßler, A., Rapberger, A., & Woltran, S. (2021). The complexity landscape of claim-augmented argumentation frameworks. In *Proceedings of hte 35th AAAI Conference on Artificial Intelligence (AAAI 2021)*, pp. 6296–6303. AAAI Press.

Dvořák, W., Greßler, A., Rapberger, A., & Woltran, S. (2023). The complexity landscape of claim-augmented argumentation frameworks. *Artificial Intelligence*, *317*, 103873.

Dvořák, W., Rapberger, A., & Woltran, S. (2020a). 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)*, pp. 341–350.

Dvořák, W., Rapberger, A., & Woltran, S. (2020b). On the different types of collective attacks in abstract argumentation: equivalence results for SETAFs. *Journal of Logic and Computation*, *30*(5), 1063–1107.

Dvořák, W., & Woltran, S. (2020). Complexity of abstract argumentation under a claim-centric view. *Artificial Intelligence*, *285*, 103290.

Falappa, M. A., Kern-Isberner, G., & Simari, G. R. (2002). Explanations, belief revision and defeasible reasoning. *Artificial Intelligence*, *141*(1/2), 1–28.

Gabbay, D., Giacomin, M., Simari, G. R., & Thimm, M. (Eds.). (2021). *Handbook of Formal Argumentation*, Vol. 2. College Publications.

Gorogiannis, N., & Hunter, A. (2011). Instantiating abstract argumentation with classical logic arguments: Postulates and properties. *Artificial Intelligence*, *175*(9-10), 1479–1497.

Kiesel, R., & Rapberger, A. (2021). Concurrence for well-formed CAFs: Naive semantics. *CoRR*, *abs/2112.10445*.

König, M., Rapberger, A., & Ulbricht, M. (2022). Just a matter of perspective. In *Proceedings of the 9th International Conference on Computational Models of Argument (COMMA 2022)*, Vol. 353 of *Frontiers in Artificial Intelligence and Applications*, pp. 212–223. IOS Press.

Liao, B. S., Jin, L., & Koons, R. C. (2011). Dynamics of argumentation systems: A division-based method. *Artificial Intelligence*, *175*(11), 1790–1814.

Lifschitz, V., Pearce, D., & Valverde, A. (2001). Strongly equivalent logic programs. *ACM Transactions on Computational Logic (TOCL)*, *2*(4), 526–541.

Lin, F. (2002). Reducing strong equivalence of logic programs to entailment in classical propositional logic. In Fensel, D., Giunchiglia, F., McGuinness, D. L., & Williams, M. (Eds.), *Proceedings of the 8th International Conference on Principles and Knowledge Representation and Reasoning (KR 2002)*, pp. 170–176. Morgan Kaufmann.

Modgil, S., & Prakken, H. (2014). The $ASPIC^+$ framework for structured argumentation: a tutorial. *Argument & Computation*, *5*(1), 31–62.

Modgil, S., & Prakken, H. (2018). Abstract rule-based argumentation. In *Handbook of Formal Argumentation*, chap. 6, pp. 287–364. College Publications. Also appears in IfCoLog Journal of Logics and their Applications 4(8):2319–2406.

Nielsen, S. H., & Parsons, S. (2006). 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*, Vol. 4766 of *Lecture Notes in Computer Science*, pp. 54–73. Springer.

Oikarinen, E., & Woltran, S. (2011). Characterizing strong equivalence for argumentation frameworks. *Artificial Intelligence*, *175*(14-15), 1985–2009.

Papadimitriou, C. (1994). Computational complexity. *Addison-Wesley*, *2*(3), 4.

Pearce, D., Tompits, H., & Woltran, S. (2001). Encodings for equilibrium logic and logic programs with nested expressions. In *Progress in Artificial Intelligence, Knowledge Extraction, Multi-agent Systems, Logic Programming and Constraint Solving, Proceedings on the 10th Portuguese Conference on Artificial Intelligence (EPIA 2001)*, Vol. 2258 of *Lecture Notes in Computer Science*, pp. 306–320. Springer.

Rapberger, A. (2020). Defining argumentation semantics under a claim-centric view. In *Proceedings of the 9th European Starting AI Researchers' Symposium (STAIRS 2020)*, Vol. 2655 of *CEUR Workshop Proceedings*. CEUR-WS.org.

Rapberger, A. (2023). *Unpacking The Argument – A Claim-Centric View On Abstract Argumentation*. Dissertation, Technische Universität Wien.

Rapberger, A., & Ulbricht, M. (2022). On dynamics in structured argumentation formalisms. In Kern-Isberner, G., Lakemeyer, G., & Meyer, T. (Eds.), *Proceedings of the 19th International Conference on Principles of Knowledge Representation and Reasoning, KR 2022, Haifa, Israel. July 31 - August 5, 2022*.

Rocha, V. H. N., & Cozman, F. G. (2022). 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)*.

Rotstein, N. D., Moguillansky, M. O., Falappa, M. A., García, A. J., & Simari, G. R. (2008). Argument theory change: Revision upon warrant. In Besnard, P., Doutre, S., & Hunter, A. (Eds.), *Computational Models of Argument: Proceedings of COMMA 2008, Toulouse, France, May 28-30, 2008*, Vol. 172 of *Frontiers in Artificial Intelligence and Applications*, pp. 336–347. IOS Press.

Rotstein, N. D., Moguillansky, M. O., García, A. J., & Simari, G. R. (2010). A dynamic argumentation framework. In *Proceedings of the 3rd International Conference on Computational Models of Argument (COMMA 2010)*, Vol. 216 of *Frontiers in Artificial Intelligence and Applications*, pp. 427–438. IOS Press.

Snaith, M., & Reed, C. (2017). Argument revision. *Journal of Logic and Computation*, *27*(7), 2089–2134.

Thimm, M. (2012). A probabilistic semantics for abstract argumentation.. In *Proceedings of the 20th European Conference on Artificial Intelligence (ECAI 2012)*, Vol. 242 of *Frontiers in Artificial Intelligence and Applications*, pp. 750–755. IOS Press.

Toni, F. (2014). A tutorial on assumption-based argumentation. *Argument & Computation*, *5*(1), 89–117.

Truszczyński, M. (2006). Strong and uniform equivalence of nonmonotonic theories – an algebraic approach. *Annals of Mathematics and Artificial Intelligence*, *48*(3–4), 245–265.

Verheij, B. (1996). Two approaches to dialectical argumentation: admissible sets and argumentation stages. *Proceedings of the 8th Duch Conference on Artificial Intelligence (NAIC 1996)*, *96*, 357–368.

Wallner, J. P., Niskanen, A., & Järvisalo, M. (2017). Complexity results and algorithms for extension enforcement in abstract argumentation. *Journal of Artificial Intelligence Research*, *60*, 1–40.

Wu, Y., Caminada, M., & Gabbay, D. M. (2009). Complete extensions in argumentation coincide with 3-valued stable models in logic programming. *Studia Logica*, *93*(2-3), 383–403.