

Characterizing Strong Equivalence for Argumentation Frameworks*

Emilia Oikarinen

Helsinki Institute for Information Technology HIIT
University of Helsinki
P.O. Box 68, FI-00014 University of Helsinki, Finland
emilia.oikarinen@cs.helsinki.fi

Stefan Woltran

Vienna University of Technology
Institute of Information Systems
Favoritenstraße 9–11, A-1040 Vienna, Austria
woltran@dbai.tuwien.ac.at

Abstract


Since argumentation is an inherently dynamic process, it is of great importance to understand the effect of incorporating new information into given argumentation frameworks. In this work, we address this issue by analyzing equivalence between argumentation frameworks under the assumption that the frameworks in question are incomplete, i.e. further information might be added later to both frameworks simultaneously. In other words, instead of the standard notion of equivalence (which holds between two frameworks, if they possess the same extensions), we require here that frameworks F and G are also equivalent when conjoined with any further framework H . Due to the nonmonotonicity of argumentation semantics, this concept is different to (but obviously implies) the standard notion of equivalence. We thus call our new notion strong equivalence and study how strong equivalence can be decided with respect to the most important semantics for abstract argumentation frameworks. We also consider variants of strong equivalence in which we define equivalence with respect to the sets of arguments credulously (or skeptically) accepted, and restrict strong equivalence to augmentations H where no new arguments are raised.

Introduction


In Artificial Intelligence, the area of argumentation (see (Bench-Capon and Dunne 2007) for an excellent summary) has become one of the central issues during the last decade with abstract argumentation frameworks (AFs, for short) as introduced by Dung (1995) being the most popular formalization on the conceptual level of argumentation. In a nutshell, such frameworks formalize statements (in general, such statements can be inferential structures themselves) together with a relation denoting conflicts between them, and the semantics gives an abstract handle to solve these inherent conflicts between statements by selecting admissible subsets of them. A number of papers compared and investigated properties of the different semantics which have been proposed for such frameworks (see, e.g. (Baroni and Giacomin 2009) and the references therein). However, the concept of equivalence between two frameworks has not received that much attention yet, although the inherent nonmonotonicity

of argumentation makes equivalence a subtle property. Let us have some motivating examples to illustrate this.


To start with, consider the following two frameworks¹ $F = (\{a, b, c\}, \{(a, b), (b, c), (c, a)\})$ and $G = (\{a, b, c\}, \{(b, a), (a, c), (c, b)\})$



which are equivalent under most of the known semantics. Let us use the preferred semantics here, which selects maximal conflict-free and self-defending sets of arguments (a formal definition is given in the next section). Then, both frameworks have the same unique preferred extension, namely the empty set. However, if we add a new argument which attacks b , the situation becomes different:



If we let $H = (\{b, d\}, \{(d, b)\})$, the AF on the left can be considered as the “union” of F and H which we will denote $F \cup H$, slightly abusing notation. The AF on the right is then $G \cup H$. However, $F \cup H$ and $G \cup H$ are not equivalent. In fact, $\{c, d\}$ is the unique preferred extension of $F \cup H$ and $\{a, d\}$ is the unique preferred extension of $G \cup H$. However, it is not even necessary to add a new argument to make this implicit difference between F and G explicit. Consider now $H' = (\{a, b\}, \{(a, b)\})$. Then $F \cup H'$ and $G \cup H'$ are



and we obtain that $F \cup H' = F$ has the empty set as preferred extension, while $\{a\}$ is the preferred extension of $G \cup H'$.

This leads us towards definitions for stronger variants of equivalence. As the central notion we want to study in this paper, we define *strong equivalence*² between two AFs F and G as the problem of deciding whether F and G remain

*This work was supported by the Vienna Science and Technology Fund (WWTF), grant ICT08-028.
Copyright © 2010, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

¹We assume a bijective relation between statements and their names.

²We follow terminology as used in other KR formalisms, in particular, logic programming (Lifschitz, Pearce, and Valverde 2001).

equivalent under any extensions of the two AFs, i.e., whether for each further AF H , $F \cup H$ and $G \cup H$ are equivalent. The study of such an equivalence notion is motivated by the following observations:

- *Implicit vs. explicit information:* as we have seen in the example above, the two frameworks store different information. However, the semantics do not make this difference visible, unless the AFs are suitably extended, i.e. some new information is added. Strong equivalence can thus be understood as a property which decides whether two argumentation frameworks provide the same implicit information.
- *Replacement property:* in classical logics, the replacement theorem states that a subformula ϕ can be faithfully replaced in any theory T by a subformula ψ , in case ϕ and ψ are equivalent. By a faithful replacement we mean here, that the models of T are not changed by this replacement, i.e. T is equivalent to $T[\phi/\psi]$. Again, our example shows that argumentation does not satisfy the replacement property (the same holds for other nonmonotonic formalisms). However, the replacement property is a central condition for simplifications. If we consider strong equivalence, we obtain a replacement property for argumentation frameworks. Note that strong equivalence thus gives us a handle on local simplifications, i.e. simplifying parts of a framework without looking at the entire framework.
- *Dynamics of argumentation:* the process of argumentation is dynamic and evolving with time (Rotstein et al. 2008b). In other words, during an argumentation process it is natural that new arguments are raised, which then attack or are attacked by already existing arguments. As well, it might be the case that we learn that two existing arguments are in an attack-relation, which we were not aware of before. This all leads to the necessity to understand the semantics also for incomplete argumentation frameworks. Strong equivalence provides new theoretical insight wrt. this aspect, which, we believe, complements results about the revision of AFs (Cayrol, Dupin de Saint-Cyr, and Lagasque-Schiex 2008; Falappa, Kern-Isberner, and Simari 2009).

Indeed, these observations are not exclusive to argumentation; the study of different equivalence notions has already received attention in other (nonmonotonic) formalisms for knowledge representation (Lifschitz, Pearce, and Valverde 2001; Turner 2004; Truszczynski 2006). However, for abstract argumentation we are not aware of such results.

Our main contributions are summarized as follows:

- We provide characterizations how to decide strong equivalence with respect to the semantics as defined by Dung (1995), as well as semi-stable (Caminada 2006) and ideal semantics (Dung, Mancarella, and Toni 2007). Our results show that strong equivalence wrt. admissible, preferred, ideal, and respectively, semi-stable semantics coincides, while the remaining semantics (stable, grounded, and complete) yield different notions of strong equivalence. In the case of self-loop free AFs, our results are even stronger; here, strong equivalence for all considered semantics reduces to syntactical equivalence (i.e., the compared AFs have to be exactly the same).

- We then study some variants of strong equivalence. In particular, we consider strong equivalence when defined over credulously and respectively skeptically accepted arguments. Interestingly, in nearly all semantics this concept reduces to the corresponding standard notion of strong equivalence.

- Finally, we weaken the concept of strong equivalence by considering only augmentation of frameworks which do not add any new arguments to the compared frameworks. While this restriction (which we call local equivalence) has no effect for admissible, preferred, ideal, and semi-stable semantics, it leads to some subtleties for the other semantics.

The paper is organized as follows. First, we introduce the necessary background for abstract argumentation frameworks and the semantics we consider here. The third section, which is the main part of this work, contains our characterization theorems for strong equivalence. Then, we look at the notion of strong equivalence when defined with respect to acceptability, and introduce the aforementioned concept of local equivalence. We conclude the paper with a discussion of related work and give pointers to open problems.

Preliminaries

We consider a fixed countable set $\mathcal{U}$ of arguments. An argumentation framework (AF) is a pair (A, R) where $A \subseteq \mathcal{U}$ is a finite set of arguments and $R \subseteq A \times A$ represents the attack-relation. For an AF $F = (B, S)$ we use $A(F)$ to refer to B and $R(F)$ to refer to S . When clear from the context, we often write $a \in F$ (instead of $a \in A(F)$) and $(a, b) \in F$ (instead of $(a, b) \in R(F)$). For two AFs F and G , we define the union $F \cup G$ and intersection $F \cap G$ as expected, i.e. $F \cup G = (A(F) \cup A(G), R(F) \cup R(G))$ and $F \cap G = (A(F) \cap A(G), R(F) \cap R(G))$. We write $F \setminus G$ to denote $(A(F), R(F) \setminus R(G))$. Likewise, we write $F \subseteq G$ in case $A(F) \subseteq A(G)$ and $R(F) \subseteq R(G)$.

For an AF $F = (A, R)$ and $S \subseteq A$, we say that (i) S is *conflict-free* in F if there are no $a, b \in S$ such that $(a, b) \in R$; (ii) $a \in A$ is *defeated* by S in F if there is $b \in S$ such that $(b, a) \in R$; and (iii) $a \in A$ is *defended* by S in F if for each $b \in A$ with $(b, a) \in R$, b is defeated by S in F . Note that an argument a with $(a, a) \in F$ cannot be defended by a conflict-free set.

Semantics for argumentation frameworks are given via a function σ which assigns to each AF $F = (A, R)$ a set $S \subseteq 2^A$ of extensions. We consider here $\sigma \in \{s, a, p, c, g, i, ss\}$ for stable, admissible, preferred, complete, grounded, ideal, and respectively, semi-stable extensions (Dung 1995; Dung, Mancarella, and Toni 2007; Caminada 2006).

Definition 1 Let $F = (A, R)$ be an AF and $S \subseteq A$.

- S is a *stable extension* of F , i.e., $S \in s(F)$, if S is conflict-free in F and each $a \in A \setminus S$ is defeated by S in F .
- S is an *admissible extension* of F , i.e., $S \in a(F)$, if S is conflict-free in F and each $a \in S$ is defended by S in F .
- S is a *preferred extension* of F , i.e., $S \in p(F)$, if $S \in a(F)$ and for each $T \in a(F)$, $S \not\subseteq T$.
- S is a *complete extension* of F , i.e., $S \in c(F)$, if $S \in a(F)$ and for each $a \in A$ defended by S in F , $a \in S$ holds.

- S is a grounded extension of F , i.e., $S \in g(F)$, if $S \in c(F)$, and for each $T \in c(F)$, $T \not\subseteq S$.
- S is an ideal extension of F , i.e., $S \in i(F)$, if $S \in a(F)$, $S \subseteq \bigcap_{T \in p(F)} T$, and for each $U \in a(F)$, such that $U \subseteq \bigcap_{T \in p(F)} T$, $S \not\subseteq U$.
- S is a semi-stable extension of F , i.e., $S \in ss(F)$, if $S \in a(F)$, and for each $T \in a(F)$, $R^+(S) \not\subseteq R^+(T)$, where $R^+(U) = U \cup \{b \mid (a, b) \in R, a \in U\}$.

It is well known that for $\sigma \in \{a, p, c, g, i, ss\}$, and any AF F , $\sigma(F) \neq \emptyset$ (recall that we deal with finite AFs); only stable semantics may yield an empty set of extensions. Moreover, for $\sigma \in \{g, i\}$, and any AF F , $\sigma(F)$ contains exactly one extension. The grounded extensions of an AF $F = (A, R)$, is given by the least fixed point of the operator $\Gamma_F : 2^A \rightarrow 2^A$, where $\Gamma_F(S) = \{a \in A \mid a \text{ is defended by } S \text{ in } F\}$. Also the following relations hold, for each AF F :

$$s(F) \subseteq ss(F) \subseteq p(F) \subseteq c(F) \subseteq a(F). \quad (1)$$

In case $s(F) \neq \emptyset$, $s(F) = ss(F)$ holds (Caminada 2006) for finite AFs. Moreover, the ideal extension of F is also a complete one (Dung, Mancarella, and Toni 2007).

When comparing frameworks, the picture becomes more opaque. Interestingly, we only have a few relations between the different semantics.

Proposition 1 For any AFs F, G , we have

- (1) $a(F) = a(G) \implies \theta(F) = \theta(G)$, for $\theta \in \{p, i\}$;
- (2) $c(F) = c(G) \implies \theta(F) = \theta(G)$, for $\theta \in \{p, i, g\}$.

Proof We show the claim for $\theta = p$; the other cases are similar. Let $\sigma \in \{a, c\}$ and assume $\sigma(F) = \sigma(G)$ but $p(F) \neq p(G)$. Wlog. let $S \in p(F) \setminus p(G)$. Then, $S \in \sigma(F)$ ($p(F) \subseteq \sigma(F)$ holds for all AFs F), and $S \in \sigma(G)$ (by assumption). Since preferred extensions are maximal, for all $S \in \sigma(G)$ there is $S' \in p(G)$ such that $S \subseteq S'$. Since $S \notin p(G)$, there is $S' \in p(G)$ such that $S \subset S'$. Recall that $p(G) \subseteq \sigma(G)$, and hence $S' \in \sigma(G)$. Since $\sigma(F) = \sigma(G)$, $S' \in \sigma(F)$. But this is in contradiction with $S \in p(F)$. $\square$

None of the other relations hold (even if we restrict to comparisons between AFs which are given over the same arguments) as witnessed by the following collection of counter examples.

The first example shows that there exist AFs F and G , such that $\sigma(F) = \sigma(G) \not\implies \theta(F) = \theta(G)$, where $\sigma \in \{s, ss, a, p, i\}$ and $\theta \in \{g, c\}$.


Example 1 Let F and G be the following AFs:



We have $a(F) = a(G) = \{\emptyset, \{b\}\}$ and $s(F) = ss(F) = s(G) = ss(G) = \{\{b\}\}$. However, $\emptyset$ is a complete extension of F (since each argument faces at least one attack), while this is not the case for G (where b is thus defended by the empty set), i.e. $\{b\}$ is the only complete extension of G . As well, we observe that $\emptyset$ is the grounded extension of F and $\{b\}$ is the grounded extension of G . Thus, we have $\sigma'(F) = \sigma'(G) \not\implies \theta(F) = \theta(G)$ for $\sigma' \in \{s, ss, a\}$ and $\theta \in \{g, c\}$. By Proposition 1, we can extend this observation to cover ideal and preferred semantics, as well.

The next example presents AFs F and G , for which $\sigma(F) = \sigma(G) \not\implies a(F) = a(G)$, $\sigma \in \{s, p, i, c, g, ss\}$.

Example 2 Let F and G be the following AFs:



Since both F and G are acyclic, we have $\sigma(F) = \sigma(G) = \{a, c, e\}$ for $\sigma \in \{s, p, i, c, g, ss\}$. However, $\{a, c\} \in a(F)$, but $\{a, c\} \notin a(G)$. Thus, $\sigma(F) = \sigma(G) \not\implies a(F) = a(G)$.

We now show that for $\sigma \in \{a, p, c, g, i\}$ and $\theta \in \{s, ss\}$, $\sigma(F) = \sigma(G) \implies \theta(F) = \theta(G)$ does not hold in general.


Example 3 Let F and G be as follows:



We have $a(F) = c(F) = a(G) = c(G) = \{\emptyset, \{a\}, \{b\}\}$. However, $s(F) = \{\{b\}\}$ and $s(G) = \{\{a\}, \{b\}\}$. Hence, via Proposition 1, we thus obtain $\sigma(F) = \sigma(G) \not\implies s(F) = s(G)$, for $\sigma \in \{a, c, p, g, i\}$. Since $s(F) \neq \emptyset \neq s(G)$, our observations immediately extend to semi-stable semantics.

Next, we present an example due to Dunne (2009) to show $\sigma(F) = \sigma(G) \not\implies \theta(F) = \theta(G)$, for $\sigma \in \{p, g, s, ss\}$ and $\theta \in \{a, c, i\}$.

Example 4 Let now F and G be as follows:



We have $g(F) = g(G) = \{\emptyset\}$, $p(F) = p(G) = s(F) = s(G) = ss(F) = ss(G) = \{\{a, d\}, \{b, d\}\}$. However, $i(F) = \{\emptyset\}$ and $i(G) = \{\{d\}\}$ since $\{d\}$ is admissible for G , and not for F . Thus $\{d\}$ is not a complete extension of F , but it is a complete extension of G (recall that the ideal extension is always a complete extension as well).

We now provide an example which shows $\sigma(F) = \sigma(G) \not\implies p(F) = p(G)$, for $\sigma \in \{s, ss, g, i\}$.


Example 5 Consider here F and G are given as follows:



We have $ss(F) = s(F) = ss(G) = s(G) = \{\{b\}\}$, but different preferred extensions, viz. $p(F) = \{\{a\}, \{b\}\}$, $p(G) = \{\{b\}, \{c\}\}$. In turn, $i(F) = i(G) = \{\emptyset\}$, and also $g(F) = g(G) = \{\emptyset\}$ is evident. Thus, the desired relations follow.

The final example shows that stable equivalence and semi-stable equivalence are incomparable.

Example 6 Let F , G , and H be as follows



We have $ss(F) = ss(G) = \{\{b\}\}$ and $ss(H) = \{\emptyset\}$; moreover, $s(F) = s(H) = \emptyset$ and $s(G) = \{\{b\}\}$. Hence, $ss(F) = ss(G) \not\implies s(F) = s(G)$ and $s(F) = s(H) \not\implies ss(F) = ss(H)$.

Strong Equivalence

As a first novel notion to compare AFs we consider the following concept which we call strong equivalence (wrt. a given semantics σ).

Definition 2 Two AFs F and G are strongly equivalent to each other wrt. a semantics σ , in symbols $F \equiv_s^\sigma G$, iff for each AF H , $\sigma(F \cup H) = \sigma(G \cup H)$ holds.

By definition, we have that $F \equiv_s^\sigma G$ implies $\sigma(F) = \sigma(G)$, i.e. standard equivalence between wrt. σ . However, the converse direction does not hold in general.

Example 7 Recall the frameworks from the introduction

$$\begin{aligned} F &= (\{a, b, c\}, \{(a, b), (b, c), (c, a)\}) \text{ and} \\ G &= (\{a, b, c\}, \{(a, c), (c, b), (b, a)\}). \end{aligned}$$

Then $\sigma(F) = \sigma(G) = \{\emptyset\}$, for $\sigma \in \{\mathbf{a}, \mathbf{p}, \mathbf{c}, \mathbf{g}, \mathbf{i}, \mathbf{ss}\}$. Moreover, $s(F) = s(G) = \emptyset$. Consider $H = (\{b, d\}, \{(d, b)\})$, with d being a fresh argument different from a, b, c . Then, $\{c, d\}$ is the only stable extension of $F \cup H$, while $\{c, d\} \notin a(G \cup H)$. Inspecting the relations in (1), we can conclude $\sigma'(F \cup H) \neq \sigma'(G \cup H)$, for $\sigma' \in \{\mathbf{s}, \mathbf{ss}, \mathbf{p}, \mathbf{c}, \mathbf{a}\}$. By definition, $F \not\equiv_{\sigma'}^{\sigma'} G$. It remains to check ideal strong-equivalence and grounded strong-equivalence. In fact, we have $i(F \cup H) = p(F \cup H) = \{\{c, d\}\} \neq \{\{a, d\}\} = p(G \cup H) = i(G \cup H)$, as well as $g(F \cup H) = c(F \cup H) = \{\{c, d\}\} \neq \{\{a, d\}\} = c(G \cup H) = g(G \cup H)$.

Let us mention at this point, that the notion of conflict-freeness is not responsible for the behavior observed above. In fact, for AFs F, H , and a set S of arguments, the following propositions are equivalent: (i) S is conflict-free in both F and H ; (ii) S is conflict-free in $F \cup H$. We thus can immediately conclude that for AFs F, G with $A(F) = A(G)$, the following holds: F and G have the same set of conflict-free sets iff $F \cup H$ and $G \cup H$ have the same set of conflict-free sets, for any AF H .

In the subsequent sections, we provide characterizations for strong equivalence wrt. the different semantics we consider. For all cases we will provide syntactical criteria which we call *kernels*. The idea is that (syntactical) equivalence of kernels characterizes strong equivalence wrt. the considered semantics.

Strong Equivalence wrt. Stable Semantics

Our first goal is to identify attacks which do not contribute when computing stable extensions of an AF F . Indeed, we need to find attacks which do not contribute in the evaluation of F , no matter how F is extended. Since stable semantics are solely based on conflict-free sets and attacks, a good candidate for such “useless” attacks are pairs (a, b) where also (a, a) is contained in the attack-relation.

Definition 3 For an AF $F = (A, R)$, we define the *s*-kernel of F as $F^{s\kappa} = (A, R^{s\kappa})$ where

$$R^{s\kappa} = R \setminus \{(a, b) \mid a \neq b, (a, a) \in R\}.$$

A few properties are clear by definition and are implicitly used later on: for each AF F and each $a \in A(F)$, $A(F) = A(F^{s\kappa})$, $R(F) \supseteq R(F^{s\kappa})$, and $(a, a) \in F$ iff $(a, a) \in F^{s\kappa}$.

Example 8 Let $F = (\{a, b\}, \{(a, a), (a, b)\})$. We have $s(F) = \emptyset$. However, also for $F^{s\kappa} = (\{a, b\}, \{(a, a)\})$, we have $s(F^{s\kappa}) = \emptyset$. As another example, consider $G = (\{a, b\}, \{(a, a), (a, b), (b, a)\})$. Here, $s(G) = \{\{b\}\}$. Note that also $G^{s\kappa} = (\{a, b\}, \{(a, a), (b, a)\})$ possesses $\{b\}$ as its only stable extension.

Indeed, the above observation follows a general principle, which we show next.

Lemma 1 For any AF F , $s(F) = s(F^{s\kappa})$.

Proof First observe that, for each set S , S is conflict-free in F iff S is conflict-free in $F^{s\kappa}$. Moreover, for each such conflict-free set S and each $a \in A$, a is defeated by S in F iff a is defeated by S in $F^{s\kappa}$ (if a is defeated by S in F , there exists $b \in S$, such that $(b, a) \in F$. Since $b \in S$ and S is conflict-free in F , we get $(b, b) \notin F$. By definition, $(b, a) \in F^{s\kappa}$. The if-direction follows from the observation that $R(F^{s\kappa}) \subseteq R(F)$). The claim follows now easily. $\square$

The next technical lemma shows that the notion of an *s*-kernel is robust wrt. composition of AFs.

Lemma 2 Let F and G be AFs, such that $F^{s\kappa} = G^{s\kappa}$. Then, $(F \cup H)^{s\kappa} = (G \cup H)^{s\kappa}$ for all AFs H .

Proof Suppose $F^{s\kappa} = G^{s\kappa}$ and let $(a, b) \in (F \cup H)^{s\kappa}$. We show $(a, b) \in (G \cup H)^{s\kappa}$. Since $(a, b) \in (F \cup H)^{s\kappa}$ we know that $(a, a) \notin F \cup H$. Thus, $(a, a) \notin F^{s\kappa}$, $(a, a) \notin G^{s\kappa}$ (by assumption $F^{s\kappa} = G^{s\kappa}$), and $(a, a) \notin H^{s\kappa}$. Now, since $(a, b) \in (F \cup H)^{s\kappa}$, $(a, b) \in F^{s\kappa}$ or $(a, b) \in H^{s\kappa}$. In case $(a, b) \in H^{s\kappa}$, $(a, b) \in (G \cup H)^{s\kappa}$ follows since $(a, a) \notin G^{s\kappa}$ (thus, $(a, a) \notin G \cup H$). In case $(a, b) \in F^{s\kappa}$, we get by assumption $F^{s\kappa} = G^{s\kappa}$, that $(a, b) \in G^{s\kappa}$, and since $(a, a) \notin H^{s\kappa}$, $(a, b) \in (G \cup H)^{s\kappa}$ follows. The other direction is symmetric. $\square$

We proceed with our first main theorem and show that syntactical equivalence of *s*-kernels characterizes strong equivalence between F and G wrt. stable semantics.

Theorem 1 For any AFs F and G : $F^{s\kappa} = G^{s\kappa}$ iff $F \equiv_s^s G$.

Proof Suppose $F^{s\kappa} = G^{s\kappa}$ and let H, S s.t. $S \in s(F \cup H)$. We show $S \in s(G \cup H)$. By Lemma 1, $S \in s((F \cup H)^{s\kappa})$ and we get from Lemma 2, $S \in s((G \cup H)^{s\kappa})$. Thus, $S \in s(G \cup H)$, again by Lemma 1. By symmetry and definition of strong equivalence, we get $F^{s\kappa} = G^{s\kappa}$ implies $F \equiv_s^s G$.

For the converse direction, suppose $F^{s\kappa} \neq G^{s\kappa}$. We show $F \not\equiv_s^s G$. First, we consider the case $A(F^{s\kappa}) \neq A(G^{s\kappa})$. This implies $A(F) \neq A(G)$ by the definition of an *s*-kernel. Wlog. let $a \in A(F) \setminus A(G)$. We use $B = (A(F) \cup A(G)) \setminus \{a\}$, and c as a fresh argument. Consider

$$H = (B \cup \{c\}, \{(c, b) \mid b \in B\}).$$

Suppose now, a is contained in some $S \in s(F \cup H)$. Then, we are done since a cannot be contained in any $S' \in s(G \cup H)$, since $a \notin A(G \cup H)$. Otherwise, we extend H to $H' = H \cup (\{a\}, \emptyset)$. Then, $\{a, c\}$ is the unique stable extension of $G \cup H'$. On the other hand, observe that $F \cup H' = F \cup H$, hence by assumption, a is not contained in any $S \in s(F \cup H')$ or $s(F \cup H')$ is empty. In both cases, we get $F \not\equiv_s^s G$.

Now suppose $A(F^{s\kappa}) = A(G^{s\kappa})$. Then, we have $R(F^{s\kappa}) \neq R(G^{s\kappa})$. Wlog. assume there exists some $(a, b) \in R(F^{s\kappa}) \setminus R(G^{s\kappa})$. We define

$$H = (A(F), \{(a, c) \mid c \in A(F) \setminus \{a, b\}\}).$$

Let $a = b$ (thus $(a, a) \in R(F)$ and $(a, a) \notin R(G)$). Then, $\{a\} \notin s(F \cup H)$ (since $\{a\}$ is not conflict-free in $F \cup H$) and $\{a\} \in s(G \cup H)$ (since a attacks all other arguments in $G \cup H$). Hence, in what follows we can assume that any self-loop is either contained in both $R(F^{s\kappa})$ and $R(G^{s\kappa})$ or in none of them. Let us thus now consider $a \neq b$. Since $(a, b) \in F^{s\kappa}$, it holds that $(a, b) \in F$, $(a, a) \notin F$, and, furthermore, we now can assume that $(a, a) \notin G$. Now, $\{a\} \in s(F \cup H)$ (since a attacks all other arguments) and $\{a\} \notin s(G \cup H)$ (since b is not defeated by $\{a\}$ in $G \cup H$; recall that $(a, b) \notin R(G^{s\kappa})$ and since $(a, a) \notin R(G)$, we also have $(a, b) \notin R(G)$). Thus, $F \not\equiv_s^s G$ follows. $\square$

Strong Equivalence wrt. Admissible, Preferred, Ideal and Semi-Stable Semantics

We now provide a slightly more restrictive notion of a kernel, which turns out to serve as a uniform characterization for strong equivalence wrt. four different semantics

Definition 4 For an AF $F = (A, R)$, we define the *a*-kernel of F as $F^{a\kappa} = (A, R^{a\kappa})$ where $R^{a\kappa}$ is given as

$$R \setminus \{(a, b) \mid a \neq b, (a, a) \in R, \{(b, a), (b, b)\} \cap R \neq \emptyset\}.$$

The following properties also hold for the notion of an *a*-kernel: For each AF F and each argument a , $A(F) = A(F^{a\kappa})$, $R(F) \supseteq R(F^{a\kappa})$, and $(a, a) \in F$ iff $(a, a) \in F^{a\kappa}$.

Example 9 We first show that the notion of *s*-kernels following Definition 3 is too weak to capture equivalence wrt. admissible extensions. Recall AF $F = (\{a, b\}, \{(a, a), (a, b)\})$ from Example 8. We clearly have $a(F) = \{\emptyset\}$. However, for $F^{s\kappa} = (\{a, b\}, \{(a, a)\})$, we now have $a(F^{s\kappa}) = \{\emptyset, \{b\}\}$. For another example, consider $G = (\{a, b\}, \{(a, a), (a, b), (b, a)\})$, which has as an *a*-kernel $G^{a\kappa} = (\{a, b\}, \{(a, a), (b, a)\})$. One can check that $a(G) = \{\emptyset, \{b\}\} = a(G^{a\kappa})$.

Concerning the relationship between an *s*-kernel and an *a*-kernel, we obviously have $F^{s\kappa} \subseteq F^{a\kappa}$ for each AF F . A stronger relation between the two notions is as follows:³

Lemma 3 For any AFs F, G , $F^{a\kappa} = G^{a\kappa}$ implies $F^{s\kappa} = G^{s\kappa}$.

We now give the two important properties a notion of a kernel has to fulfill (cf. Lemmas 1 and 2 for the stable case).

Lemma 4 For any AF F , $\sigma(F) = \sigma(F^{a\kappa})$ ($\sigma \in \{a, p, i, ss\}$).

Lemma 5 If $F^{a\kappa} = G^{a\kappa}$, then $(F \cup H)^{a\kappa} = (G \cup H)^{a\kappa}$ for all AFs H .

Interestingly, syntactical equivalence of *a*-kernels captures strong equivalence wrt. four different semantics.

Theorem 2 The following propositions are equivalent for all AFs F and G :

- (a) $F^{a\kappa} = G^{a\kappa}$
- (b) $F \equiv_s^a G$
- (c) $F \equiv_s^{ss} G$
- (d) $F \equiv_s^p G$
- (e) $F \equiv_s^i G$

Proof (a) $\Rightarrow$ (b) and (a) $\Rightarrow$ (c): Similar as in the proof of Theorem 1 using Lemmas 4 and 5.

(b) $\Rightarrow$ (d) and (b) $\Rightarrow$ (e): Suppose $F \equiv_s^a G$, and let H be any AF. By definition, $a(F \cup H) = a(G \cup H)$, and by Proposition 1, $\theta(F \cup H) = \theta(G \cup H)$, for $\theta \in \{p, i\}$. Thus $F \equiv_s^\theta G$ follows.

(c) $\Rightarrow$ (a), (d) $\Rightarrow$ (a), and (e) $\Rightarrow$ (a): Let $\theta \in \{p, i, ss\}$. Suppose $F^{a\kappa} \neq G^{a\kappa}$. We show $F \not\equiv_s^\theta G$. In case $\theta(F^{a\kappa}) \neq \theta(G^{a\kappa})$ we are done (by Lemma 4, we get $\theta(F) \neq \theta(G)$ and thus $F \not\equiv_s^\theta G$). In what follows, we thus assume $\theta(F) = \theta(G)$.

First consider the case $A(F^{a\kappa}) \neq A(G^{a\kappa})$. By definition of an *a*-kernel then $A(F) \neq A(G)$. Wlog. let $a \in A(F) \setminus A(G)$. Since $a \notin A(G)$, we have $a \notin S$ for each $S \in \theta(G)$, and thus, since $\theta(F) = \theta(G)$, $a \notin S'$ for each $S' \in \theta(F)$. Let $H = (\{a\}, \emptyset)$. Clearly, $F \cup H = F$ and thus $\theta(F \cup H) = \theta(F)$. On the other hand, $a \in S$ for any $S \in \theta(G \cup H)$. This can be seen as follows: First, $a \in S$ for any $S \in p(G \cup H)$, since a is not attacked in $G \cup H$. Hence, $a \in S$ for any $S \in ss(G \cup H)$ (since $ss(I) \subseteq p(I)$, for any AF I). Moreover, we have $\{a\} \in a(G \cup H)$ and $a \in \bigcap_{S \in p(G \cup H)} S$. Thus a has to be contained in the ideal extension of $G \cup H$, as well.

Now suppose $A(F^{a\kappa}) = A(G^{a\kappa})$, i.e. $A(F) = A(G)$. Thus wlog. there exists some $(a, b) \in R(F^{a\kappa}) \setminus R(G^{a\kappa})$. Let $B = A(F) \setminus \{a, b\}$. First, assume $a = b$. We define

$$H = (A(F), \{(a, c), (c, c) \mid c \in B\})$$

and obtain $(F \cup H)^{a\kappa} = (A(F), \{(a, a)\} \cup \{(c, c) \mid c \in B\})$ and $(G \cup H)^{a\kappa} = (A(F), \{(a, c), (c, c) \mid c \in B\})$. Thus, $\emptyset$ is the only preferred extension of $(F \cup H)^{a\kappa}$ (since it is the only conflict-free set here) and $\{a\}$ is the only preferred extension of $G \cup H$ (since a attacks all its attackers and a does not attack itself by the assumption). We obtain $p((F \cup H)^{a\kappa}) \neq p((G \cup H)^{a\kappa})$. Note that for an AF which possesses a unique preferred extension S , S has to be also the unique semi-stable extension (in case, the AF is finite, which is the case here), and the ideal extension. Hence, $\theta((F \cup H)^{a\kappa}) \neq \theta((G \cup H)^{a\kappa})$, and by Lemma 4, $\theta(F \cup H) \neq \theta(G \cup H)$ (for $\theta \in \{p, i, ss\}$).

Hence, in what follows we can assume that any self-loop is either contained in both $R(F^{a\kappa})$ and $R(G^{a\kappa})$ or in none of them. Let us thus now consider $a \neq b$. We continue our proof by different cases for the presence of attack (a, a) . If $(a, a) \notin R(F^{a\kappa})$, we define

$$H = (A(F), \{(b, a), (b, b)\} \cup \{(a, c), (c, c) \mid c \in B\}).$$

It can be checked that $p(F \cup H) = \{\{a\}\}$ and $p(G \cup H) = \{\emptyset\}$. By the same observation as above, $\theta(F \cup H) \neq \theta(G \cup H)$, for $\theta \in \{p, i, ss\}$. If $(a, a) \in R(F^{a\kappa})$, we proceed as follows. Since $(a, b) \in F^{a\kappa}$, it holds that $(b, b) \notin R(F^{a\kappa})$ and $(b, a) \notin R(F^{a\kappa})$. We define

$$H = (A(F), \{(b, c), (c, c) \mid c \in B\})$$

³From now on, we will omit some proofs due to limited space.

and obtain $\theta(F \cup H) = \{\emptyset\}$ while $\theta(G \cup H) = \{\{b\}\}$. Hence, in all cases there is an AF H , such that $\theta(F \cup H) \neq \theta(G \cup H)$. By definition of strong equivalence, we arrive at $F \not\equiv_s^\theta G$ (for $\theta \in \{p, i, ss\}$). $\square$

Strong Equivalence wrt. Grounded Semantics

We next consider the grounded semantics and require a further kernel.

Definition 5 For an AF $F = (A, R)$, we define the g-kernel of F as $F^{g\kappa} = (A, R^{g\kappa})$ where $R^{g\kappa}$ is defined as

$$R \setminus \{(a, b) \mid a \neq b, (b, b) \in R, \{(a, a), (b, a)\} \cap R \neq \emptyset\}.$$

As for our previous kernels, these properties also hold for g-kernels: $A(F) = A(F^{g\kappa})$, $R(F) \supseteq R(F^{g\kappa})$, and, $(a, a) \in F$ iff $(a, a) \in F^{g\kappa}$, for each AF F and each argument a .

As regards to the relation of the g-kernel to our other kernels, we notice that $F^{g\kappa}$ is incomparable with both $F^{a\kappa}$ and $F^{s\kappa}$ in general. For instance, for AF $G = (\{a, b\}, \{(a, a), (a, b), (b, a)\})$, we have $G^{g\kappa} = (\{a, b\}, \{(a, a), (a, b)\})$ and $G^{a\kappa} = G^{s\kappa} = (\{a, b\}, \{(a, a), (b, a)\})$. Thus both $G^{g\kappa} \not\subseteq G^{a\kappa} = G^{s\kappa}$ and $G^{a\kappa} = G^{s\kappa} \not\subseteq G^{g\kappa}$.

However, we observe an interesting symmetry between a-kernels and g-kernels. In fact, for any AF F and a, b which are not both self-attacking, $(a, b) \in F \setminus F^{a\kappa}$ iff $(b, a) \in F \setminus F^{g\kappa}$. However, in case both a and b are self-attacking, $(a, b) \in F \setminus F^{a\kappa}$ iff $(a, b) \in F \setminus F^{g\kappa}$.

Also g-kernels satisfy similar properties as we have shown for other kernels before.

Lemma 6 For any AF F , $g(F) = g(F^{g\kappa})$.

Lemma 7 If $F^{g\kappa} = G^{g\kappa}$, then $(F \cup H)^{g\kappa} = (G \cup H)^{g\kappa}$ for all AFs H .

We proceed to show that g-kernels characterize strong equivalence wrt. the grounded extensions.

Theorem 3 For any AFs F and G : $F^{g\kappa} = G^{g\kappa}$ iff $F \equiv_s^g G$.

Proof The only if-direction is similar as in the proof of Theorem 1 using Lemmas 6 and 7.

For the if-direction, suppose $F^{g\kappa} \neq G^{g\kappa}$. In case $g(F^{g\kappa}) \neq g(G^{g\kappa})$ we are done since, by Lemma 6, $g(F) \neq g(G)$ which obviously implies $F \not\equiv_s^g G$. In what follows, we thus assume $g(F) = g(G)$.

First consider the case $A(F^{g\kappa}) \neq A(G^{g\kappa})$. By definition this holds, iff, $A(F) \neq A(G)$. Wlog. let $a \in A(F) \setminus A(G)$. Since $a \notin A(G)$, we have $a \notin S$ for $S \in g(G) = g(F)$. Let $H = (\{a\}, \emptyset)$. Clearly, $F \cup H = F$ and thus $g(F \cup H) = g(F)$. On the other hand, $a \in S'$ for $S' \in g(G \cup H)$, since there is no attack on a in $G \cup H$. Consequently, $F \not\equiv_s^g G$.

Now suppose $A(F^{g\kappa}) = A(G^{g\kappa})$, i.e. $A(F) = A(G)$. Thus wlog. there exists some $(a, b) \in R(F^{g\kappa}) \setminus R(G^{g\kappa})$. Let $c \in \mathcal{U}$ be a new argument not contained in $A(F)$ and $B = A(F) \setminus \{a, b\}$. Assuming $a = b$, i.e., $(a, a) \in F$ and $(a, a) \notin G$, let $H = (B \cup \{c\}, \{(c, d) \mid d \in B\})$. Then, $\{c\} \in g(F \cup H)$ (c is defended by $\emptyset$ in $F \cup H$; no other argument is defended by $\{c\}$ in $F \cup H$) and $\{c\} \notin g(G \cup H)$ (c is defended by $\emptyset$ in $G \cup H$ and a is defended by $\{c\}$ in

$G \cup H$). Hence, we can assume that any self-loop is either contained in both F and G or in none of them.

Let $a \neq b$. Since $(a, b) \in R(F^{g\kappa})$, $(a, b) \in R(F)$ and $(b, b) \notin R(F)$; or $(a, a) \notin R(F)$ and $(b, a) \in R(F)$. If $(b, b) \notin R(F)$, then also $(b, b) \notin R(G)$. Moreover, since $(a, b) \notin R(G^{g\kappa})$, also $(a, b) \notin R(G)$. Again, let $c \in \mathcal{U}$ be a new argument not contained in $A(F)$ and $B = A(F) \setminus \{a, b\}$. We take

$$H = (A(F) \cup \{c\}, \{(a, a), (b, a)\} \cup \{(c, d) \mid d \in B\}).$$

Now, $\{c\} \in c(F \cup H)$ and $\emptyset \notin c(F \cup H)$. Thus, $g(F \cup H) = \{\{c\}\}$. On the other hand, $\{c\} \notin c(G \cup H)$ and thus $\{c\} \notin g(G \cup H)$. If $(b, b) \in R(F)$, then $(b, b) \in R(G)$, $(a, a) \notin R(F)$, $(a, a) \notin R(G)$, and $(b, a) \notin R(F)$. We take

$$H = (A(F) \cup \{c, e\}, \{(c, d) \mid d \in B\} \cup \{(b, e)\}),$$

where e is a new argument not contained in $A(F) \cup \{c\}$. Now, $g(F \cup H) = \{\{a, c, e\}\}$ while $g(G \cup H) = \{\{a, c\}\}$ if $(b, a) \notin G$, and $g(G \cup H) = \{\{c\}\}$ if $(b, a) \in G$. Thus $F \not\equiv_s^g G$ follows. $\square$

Strong Equivalence wrt. Complete Semantics

Finally, we introduce a kernel characterizing strong equivalence wrt. complete semantics. We notice that wrt. complete extensions, attacks (a, b) for which both arguments are self-attacking are irrelevant, since neither a nor b can ever be defended with any conflict-free S .

Definition 6 For an AF $F = (A, R)$, we define the c-kernel of F as $F^{c\kappa} = (A, R^{c\kappa})$ where

$$R^{c\kappa} = R \setminus \{(a, b) \mid a \neq b, (a, a), (b, b) \in R\}$$

Concerning the relationship between a c-kernel and the other notions of kernels introduced, $F^{s\kappa} \subseteq F^{a\kappa} \subseteq F^{c\kappa}$ and $F^{g\kappa} \subseteq F^{c\kappa}$ hold for any F (recall that $F^{g\kappa}$ is incomparable with both $F^{s\kappa}$ and $F^{a\kappa}$). Similarly to Lemma 3 there is also a stronger relation between the notions.

Lemma 8 For any AFs F and G , $F^{c\kappa} = G^{c\kappa}$ implies $F^\tau = G^\tau$ for $\tau \in \{s\kappa, a\kappa, g\kappa\}$.

In fact, there is stronger relationship between a-, g- and c-kernels.

Lemma 9 For any AFs F and G : $F^{c\kappa} = G^{c\kappa}$ iff jointly $F^{a\kappa} = G^{a\kappa}$ and $F^{g\kappa} = G^{g\kappa}$.

Proof The only-if direction is by Lemma 8. For the other direction, suppose $F^{c\kappa} \neq G^{c\kappa}$. Wlog. let $(a, b) \in F^{c\kappa} \setminus G^{c\kappa}$. Hence, $(a, b) \in F$ and $(a, a) \notin F$ or $(b, b) \notin F$. If $(a, a) \notin F$, $(a, b) \in F^{a\kappa}$ by construction. Moreover, we have $G^{a\kappa} \subseteq G^{c\kappa}$ and thus $(a, b) \notin G^{a\kappa}$. Hence, $F^{a\kappa} \neq G^{a\kappa}$. For the other case, i.e. if $(b, b) \notin F$, $(a, b) \in F^{g\kappa}$ follows by construction. We also have $G^{g\kappa} \subseteq G^{c\kappa}$ and thus $(a, b) \notin G^{g\kappa}$. Therefore in this case, $F^{g\kappa} \neq G^{g\kappa}$. $\square$

We continue with properties of c-kernels which we then use to show that c-kernels characterize strong equivalence wrt. complete semantics.

Lemma 10 For any AF F , $c(F) = c(F^{c\kappa})$.

Lemma 11 If $F^{c\kappa} = G^{c\kappa}$, then $(F \cup H)^{c\kappa} = (G \cup H)^{c\kappa}$ for all AFs H .

Theorem 4 For any AFs F and G : $F^{c\kappa} = G^{c\kappa}$ iff $F \equiv_s^c G$.

Proof The only-if direction can be shown via Lemmas 10 and 11 and we just sketch the if-direction. Suppose $F^{c\kappa} \neq G^{c\kappa}$. By Lemma 9, we have $F^{a\kappa} \neq G^{a\kappa}$ or $F^{g\kappa} \neq G^{g\kappa}$. In case $F^{a\kappa} \neq G^{a\kappa}$, we get by Theorem 2, $F \not\equiv_s^p G$, i.e. there exists an AF H , such that $p(F \cup H) \neq p(G \cup H)$. In case, $F^{g\kappa} \neq G^{g\kappa}$ we get by Theorem 3, $F \not\equiv_s^g G$, i.e. there exists an AF H , such that $g(F \cup H) \neq g(G \cup H)$. In both cases, Proposition 1 yields $c(F \cup H) \neq c(G \cup H)$, i.e. $F \not\equiv_s^c G$. $\square$

Summary of Results for Strong Equivalence

We summarize our results for strong equivalence.


Theorem 5 The following relations holds for all AFs F , G , and $\sigma \in \{a, p, i, ss\}$,

- (1) $F \equiv_s^\sigma G \implies F \equiv_s^s G$
- (2) $F \equiv_s^c G \iff (F \equiv_s^\sigma G \text{ and } F \equiv_s^g G)$

Proof Relation (1) is by Lemma 3 and Theorems 1 and 2. Relation (2) is by Lemma 9 and Theorems 2, 3 and 4. $\square$

None of the other relations hold, as we sketch next.

Example 10 Consider the following frameworks:



We have $F^{a\kappa} = F^{s\kappa} = G^{a\kappa} = G^{s\kappa} = G^{g\kappa} = G$ while $F^{g\kappa} = H$. Thus, $F^{a\kappa} = G^{a\kappa}$ (similarly for $F^{s\kappa} = G^{s\kappa}$) does not imply $F^{g\kappa} = G^{g\kappa}$, i.e., $F \equiv_s^\sigma G$ does not imply $F \equiv_s^g G$ for $\sigma \in \{s, a, p, i, ss\}$ (Theorems 1–3). On the other hand, $F^{c\kappa} = F$ and $G^{c\kappa} = G$, and therefore $F \equiv_s^\sigma G$ does not imply $F \equiv_s^c G$ for $\sigma \in \{s, a, p, i, ss\}$ (Theorems 1, 2 and 4). Moreover, $H^{g\kappa} = H^{a\kappa} = H^{c\kappa} = H$ and $H^{s\kappa} = I$. Using our observations on kernels of F , we get that $F \equiv_s^g H$ does not imply $F \equiv_s^\sigma H$ for $\sigma \in \{s, a, p, i, ss, c\}$ (Theorems 1–4). Finally, observe that $I^{s\kappa} = I^{a\kappa} = I$ and recall that $H^{a\kappa} = H$ and $H^{s\kappa} = I$. Thus, $H \equiv_s^s I$ does not imply $H \equiv_s^\sigma I$ for $\sigma \in \{a, p, i, ss\}$ (Theorems 1 and 2).

An inspection of the definition of kernels shows that self-loops play a crucial role. In fact, the following observation is quite straightforward:

Lemma 12 For any self-loop free AF F , $F = F^{s\kappa} = F^{a\kappa} = F^{c\kappa} = F^{g\kappa}$.

Thus for self-loop free AFs the concept of strong equivalence collapses to syntactic equivalence for all semantics.

Theorem 6 For any self-loop free AFs F and G , we have $F = G$ iff $F \equiv_s^\sigma G$, $\sigma \in \{s, a, p, c, g, i, ss\}$.

Strong Equivalence in Terms of Consequences

An alternative approach to “strong” notions of equivalence between argumentation frameworks is to define such a concept in terms of consequences. To this end, let for an AF F , $cred^\sigma(F) = \bigcup_{S \in \sigma(F)} S$ be the set of credulous consequences of F (wrt. semantics σ), and $skept^\sigma(F) = \bigcap_{S \in \sigma(F)} S$ the set of skeptical consequences of F (wrt. σ). In case $s(F) = \emptyset$, we define⁴ $skept^s(F) = A(F)$.

Definition 7 Let $\rho \in \{cred, skept\}$. We call AFs F and G strongly ρ -equivalent wrt. semantics σ , in symbols $F \equiv_{s;\rho}^\sigma G$, iff for each AF H , $\rho^\sigma(F \cup H) = \rho^\sigma(G \cup H)$ holds.

We observe that for admissible semantics strong $skept$ -equivalence is trivial, i.e. $skept^a(F) = skept^a(G)$ holds for any AFs F , G (basically since $\emptyset \in a(H)$, for every AF H). In all remaining cases, the concepts from Definition 7 coincide with strong equivalence as defined in Definition 2.

Theorem 7 For any AFs F , G , and $\sigma \in \{s, p, c, g, i, ss\}$, the following propositions are equivalent:

- (1) $F \equiv_s^\sigma G$;
- (2) $F \equiv_{s;cred}^\sigma G$;
- (3) $F \equiv_{s;skept}^\sigma G$.

Also, (1) iff (2) for $\sigma = a$.

Proof The assertion holds for $\sigma \in \{g, i\}$, since each AF possesses a unique grounded, respectively ideal, extension. Also, it is clear that (1) implies (2), and (1) implies (3), for any of the considered semantics. We now show (2) implies (1) for $\sigma \in \{s, a, p, c, ss\}$, and (3) implies (1) for $\sigma \in \{s, p, c, ss\}$.

For the case $A(F) \neq A(G)$ we can use similar arguments as in previous proofs. In what follows, let us thus assume $A(F) = A(G)$. Moreover, let $\theta \in \{s, a\}$ and $F \not\equiv_s^\theta G$, i.e. there exists an AF H , such that $\theta(F \cup H) \neq \theta(G \cup H)$. Wlog. let $S \in \theta(F \cup H) \setminus \theta(G \cup H)$. Observe that $S \neq \emptyset$ holds: for $\theta = a$, this is obvious; for $\theta = s$, this would yield that $F \cup H$ is the empty AF, i.e. $A(F \cup H) = \emptyset$. But by the assumption $A(F) = A(G)$, we get $A(F \cup H) = A(G \cup H)$, and thus $G \cup H$ had to be the empty AF, as well; a contradiction to our assumption $F \not\equiv_s^\sigma G$.

Consider now the following AF K

$$(A(F \cup H) \cup S', \bigcup_{a \in S} \{(a, a'), (a', a'), (a', b) \mid b \in S \setminus \{a\}\} \cup \{(b, b) \mid b \in A(F \cup H) \setminus S\})$$

where $S' = \{a' \mid a \in S\}$ is a set of disjoint fresh arguments. Let $F^\dagger = F \cup H \cup K$ and $G^\dagger = G \cup H \cup K$.

For the case $\theta = s$, i.e. where we had S and H such that $S \in s(F \cup H) \setminus s(G \cup H)$, one can now show $s(F^\dagger) = \{S\}$ and $s(G^\dagger) = \emptyset$. We obtain $cred^s(F^\dagger) = S \neq \emptyset = cred^s(G^\dagger)$. $F \not\equiv_{s;cred}^\sigma G$ follows by definition (since $F^\dagger = F \cup (H \cup K)$ and $G^\dagger = G \cup (H \cup K)$). This shows that (2) implies (1) for stable semantics. Moreover, we have $skept^s(F^\dagger) = S \neq A(G^\dagger) = skept^s(G^\dagger)$ since $S \subset A(G^\dagger)$

⁴Another reasonable definition for this case would be $skept^s(F) = \emptyset$.

holds by definition of K , which contains at least one fresh argument $a' \notin S$. Thus, we also have $F \not\equiv_{s;skept}^s G$. This shows that (3) implies (1) for stable semantics.

In the case $\theta = a$, we have $a(F^\dagger) = \{S, \emptyset\}$ and $a(G^\dagger) = \{\emptyset\}$. Therefore, we also have $p(F^\dagger) = ss(F^\dagger) = \{S\}$ and $S \in c(F^\dagger)$, while $p(G^\dagger) = ss(G^\dagger) = c(G^\dagger) = \{\emptyset\}$. Thus, for $\sigma' \in \{a, p, c, ss\}$, we have $cred^{\sigma'}(F^\dagger) = S \neq \emptyset = cred^{\sigma'}(G^\dagger)$, and thus $F \not\equiv_{s;cred}^{\sigma'} G$. This shows that (2) implies (1) for those semantics; in fact, since we assumed $F \not\equiv_s^a G$, we had implicitly assumed also $F \not\equiv_s^p G$ and $F \not\equiv_{ss}^s G$ (by Theorem 2) as well as $F \not\equiv_s^c G$ (by relation (2) in Theorem 5). Moreover, we have $skept^{ss}(F^\dagger) = skept^p(F^\dagger) = S \neq \emptyset = skept^p(G^\dagger) = skept^{ss}(G^\dagger)$, and thus $F \not\equiv_{s;skept}^{\sigma''} G$, for $\sigma'' \in \{p, ss\}$ showing that (3) implies (1) for those two semantics. It remains to show $F \not\equiv_{s;skept}^c G$. For the moment, our construction does not guarantee $\emptyset \notin c(F^\dagger)$ and thus $skept^c(F^\dagger) = skept^c(G^\dagger)$ still might hold. Consider a further AF $L = (S \cup \{c, d\}, \{(c, d) \mid (d, s) \mid s \in S\})$. Then, $a(G^\dagger \cup L) = \{\emptyset, \{c\}\}$, and since c is not attacked, $c(G^\dagger \cup L) = \{\{c\}\}$. On the other hand, $c(F^\dagger \cup L) = \{S \cup \{c\}\}$, since c defends all elements from S and S remains admissible in $F^\dagger \cup L$. We obtain $skept^c(F^\dagger \cup L) = S \cup \{c\} \neq \{c\} = skept^c(G^\dagger \cup L)$. $\square$

Local Equivalence

So far, we have considered an arbitrary *context* for equivalence, i.e. we put no restriction on the AFs H which are considered to be conjoined with the AFs F and G under comparison. Now we weaken this requirement by considering only AFs H which do not introduce any new arguments. We call the resulting equivalence notion local equivalence.

Definition 8 Two AFs F and G are called locally (strong) equivalent to each other wrt. a semantics σ , in symbols $F \equiv_l^\sigma G$, iff for each H with $A(H) \subseteq A(F) \cup A(G)$, $\sigma(F \cup H) = \sigma(G \cup H)$.

For some semantics, strong and local equivalence coincide.

Theorem 8 For any AFs F, G : $F \equiv_l^\sigma G$ iff $F \equiv_s^\sigma G$ for $\sigma \in \{a, p, i, ss\}$.

Proof If $F \equiv_s^\sigma G$, then $F \equiv_l^\sigma G$. Assume that $F \not\equiv_s^\sigma G$. Then, there is an AF H such that $\sigma(F \cup H) \neq \sigma(G \cup H)$. By inspecting the proof of Theorem 2 we notice that we can always find H such that $A(H) \subseteq A(F) \cup A(G)$ holds, and therefore $F \not\equiv_l^\sigma G$, for $\sigma \in \{a, p, i, ss\}$. $\square$

For the other semantics, we observe certain differences. We start with the case of stable semantics in which strong and local equivalence almost coincide. The following example illustrates the case in which two AFs are locally equivalent, but not strongly equivalent wrt. stable semantics.

Example 11 Consider $F = (\{a, b\}, \{(b, b), (b, a)\})$ and $G = (\{b\}, \{(b, b)\})$. Here, $s(F) = s(G) = \emptyset$, but from $A(F) \neq A(G)$, we have $F^{s\kappa} \neq G^{s\kappa}$, i.e., $F \not\equiv_s^s G$. On the other hand, we observe that all H such that $A(H) \subseteq \{a, b\}$ yield $s(F \cup H) = s(G \cup H)$. $F \equiv_l^s G$ follows.

Theorem 9 For any AFs F, G : $F \equiv_l^s G$ iff $F \equiv_s^s G$ or both $s(F) = s(G) = \emptyset$ and there is $a \in (A(F) \setminus A(G)) \cup$

$(A(G) \setminus A(F))$ such that $(a, a) \notin F \cup G$ and for all $b \in (A(F) \cup A(G)) \setminus \{a\}$, $(a, b) \notin F \cup G$ and $(b, b) \in F \cap G$.

Proof By definition $F \equiv_s^s G$ implies $F \equiv_l^s G$ for all AFs F and G . Let us assume $s(F) = s(G) = \emptyset$ and there is $a \in (A(F) \setminus A(G)) \cup (A(G) \setminus A(F))$ such that $(a, a) \notin F \cup G$ and for all $b \in (A(F) \cup A(G)) \setminus \{a\}$ it holds $(a, b) \notin F \cup G$ and $(b, b) \in F \cap G$. Wlog. let $a \in A(F) \setminus A(G)$. Thus $A(F) = A(G) \cup \{a\}$. Consider an arbitrary AF H such that $A(H) \subseteq A(F)$. Since $(b, b) \in F \cap G$ for all $b \in A(G)$ and H can only involve arguments from $A(F)$, the only possible conflict-free sets for $F \cup H$ and $G \cup H$ are $\emptyset$ and $\{a\}$.

If there is $b \in A(G)$ not defeated by $\{a\}$ in H , then b is not defeated by $\{a\}$ in neither $F \cup H$ nor $G \cup H$, and $s(F \cup H) = s(G \cup H) = \emptyset$. On the other hand, if each $b \in A(G)$ is defeated by $\{a\}$ in H , then each $b \in A(G)$ is defeated by $\{a\}$ in both $F \cup H$ and $G \cup H$, and $s(F \cup H) = s(G \cup H) = \{\{a\}\}$ follows.

For the other direction, assume first $A(F) = A(G)$ and $F \not\equiv_s^s G$. Wlog. let $(a, b) \in F^{s\kappa} \setminus G^{s\kappa}$. We notice that for $H = (A(F), \{(a, c) \mid c \in A(F) \setminus \{a, b\}\})$ used in the proof of Theorem 1 it holds $A(H) \subseteq A(F)$ and $s(F \cup H) \neq s(G \cup H)$.

Assume then $A(F) \neq A(G)$ Wlog. let $a \in A(F) \setminus A(G)$. If $s(F) \neq s(G)$, then $F \not\equiv_l^s G$ follows by definition. Thus, we can assume $s(F) = s(G)$. Consider first the case $s(F) \neq \emptyset$. Since $a \notin A(G)$, $a \notin S'$ for each $S' \in s(G) = s(F)$. Similarly to the proof of Theorem 1, we take $H = (\{a\}, \emptyset)$ and notice $s(F \cup H) \neq s(G \cup H)$.

Thus, let us assume $s(F) = s(G) = \emptyset$. We denote $B = (A(F) \cup A(G)) \setminus \{a\}$. If $(a, a) \in F$, we consider

$$H = (B \cup \{a\}, \{(a, b) \mid b \in B\}).$$

Now $\{a\} \in s(G \cup H)$ while $\{a\} \notin s(F \cup H)$. If $(a, a) \notin F$ and there is $b \in B$ such $(b, b) \notin F$ or $(b, b) \notin G$, let

$$H = (B, \{(b, c) \mid c \in B \setminus \{b\}\}).$$

If $(b, b) \notin R(F)$, then $\{a, b\} \in s(F \cup H)$ while $\{a, b\} \notin s(G \cup H)$. If $(b, b) \in R(F)$ and $(b, b) \notin R(G)$, then $\{b\} \in s(G \cup H)$ while $\{b\} \notin s(F \cup H)$. Finally, if $(a, a) \notin F$ and $(b, b) \in F \cap G$ for each $b \in (A(F) \cup A(G)) \setminus \{a\}$. Assume there is $b \in B$ such that $(a, b) \in F$. We take

$$H = (A(F) \cup A(G), \{(a, c) \mid c \in (A(F) \cup A(G)) \setminus \{b\}\}).$$

Now, $\{a\} \in s(F \cup H)$ and $\{a\} \notin s(G \cup H)$ (since b is not defeated by $\{a\}$ in $G \cup H$).

Thus in each case we have found an AF H such that $A(H) \subseteq A(F) \cup A(G)$ and $s(F \cup H) \neq s(G \cup H)$ and $F \not\equiv_l^s G$ follows. $\square$

We proceed with the complete semantics. Our first observation hereby is: for any AFs F and G (i) $F \equiv_l^c G$ implies $F \equiv_s^p G$; and (ii) $F \equiv_s^c G$ implies $F \equiv_l^c G$. While (ii) is clear from definition, (i) follows from Proposition 1 and Theorem 8. Moreover, both implications can be shown to be strict.

Example 12 Let $F = (\{a, b\}, \{(a, a), (a, b), (b, a)\})$ and $G = (\{a, b\}, \{(a, a), (b, a)\})$. We have $F^{a\kappa} = G^{a\kappa}$ and thus, by Theorems 2 and 8, $F \equiv_l^p G$. On the other hand, $c(F) = \{\emptyset, \{b\}\} \neq \{\{b\}\} = c(G)$, and hence, $F \not\equiv_l^c G$.

Example 13 For the other implication, let $A = \{a, b, c\}$ and $R = \{(a, a), \{(a, a), (b, a), (b, c), (c, b), (c, c)\}\}$, and consider AFs $F = (A, R \cup \{(a, b)\})$ and $G = (A, R)$. We have $F^{c\kappa} \neq G^{c\kappa}$ and thus, by Theorem 4, $F \not\equiv_s^c G$. On the other hand, $F \equiv_l^c G$ can be verified as follows. First, we have $c(F) = \{\{b\}\} = c(G)$. Since we are interested in local strong equivalence, there are only two attacks, (b, b) and (a, b) , we can properly add via an AF H with $A(H) \subseteq A(F) \cup A(G)$. If $(b, b) \in H$, $c(F \cup H) = \{\emptyset\} = c(G \cup H)$ is clear. If $(a, b) \in H$, $F \cup H = G \cup H$. In both cases, we have $c(F \cup H) = c(G \cup H)$ and $F \equiv_l^c G$ follows.

These observations suggest that we require a kernel which lies inbetween the notions of a-kernel and c-kernel.

Definition 9 For an AF $F = (A, R)$, we define the c-local-kernel of F as $F^{cl\kappa} = (A, R^{cl\kappa})$ where

$$R^{cl\kappa} = R \setminus \{(a, b) \mid a \neq b, (a, a) \in R, ((b, b) \in R \text{ or } (b, a), (c, b) \in R \text{ for some } c \neq a)\}.$$

We observe that $F^{a\kappa} \subseteq F^{cl\kappa} \subseteq F^{c\kappa}$ holds for any AF F , and provide a few lemmata along the lines as we did for other kernels.

Lemma 13 For any AF F , $c(F) = c(F^{cl\kappa})$.

Lemma 14 If $F^{cl\kappa} = G^{cl\kappa}$, then $(F \cup H)^{cl\kappa} = (G \cup H)^{cl\kappa}$ for all AFs H with $A(H) \subseteq A(F) \cup A(G)$.

As it was the case for local equivalence wrt. stable semantics, there is also a certain pattern of frameworks here, which has to be taken into account additionally. For instance, $F^{cl\kappa} \neq G^{cl\kappa}$ for F and G in Example 13, even though they are locally equivalent wrt. complete semantics.

Definition 10 An AF F is called *b-saturated* iff there exists an argument $b \in A(F)$ such that $(b, b) \notin F$, $(a, b) \in F$ for some $a \neq b$, and for each $a \in A(F) \setminus (\Gamma_F(\emptyset) \cup \{b\})$, $(a, a) \in F$ and, whenever $(a, b) \in F$, $(d, a) \in F$ for all $d \in \{b\} \cup \Gamma_F(\emptyset)$.

Now we are ready to state the characterization theorem for the complete semantics.

Theorem 10 For any AFs F, G : $F \equiv_l^c G$ iff either $F^{cl\kappa} = G^{cl\kappa}$ or F and G are both *b-saturated* with $A(F) = A(G)$, $\Gamma_F(\emptyset) = \Gamma_G(\emptyset)$, and $(d, a) \in F$ iff $(d, a) \in G$ holds, for each $a \in A(F)$ and $d \in \{b\} \cup \Gamma_F(\emptyset)$.

Finally, we consider the grounded semantics. As with stable and complete semantics, local equivalence wrt. grounded semantics is close to strong equivalence, but in addition there are certain frameworks which have to be taken into account. For instance, consider AFs F and G such that $A(F) = A(G)$ and $g(F) = g(G) = \emptyset$. Now, $g(F \cup H) = g(G \cup H) = \emptyset$ always holds for any AF H such that $A(H) \subseteq A(F)$ and thus $F \equiv_l^g G$ follows.

We start with additional notation. Given an argument a and an AF F , F_a^* is the framework resulting from F by deleting argument a , all arguments attacked by a and by deleting all attacks adjacent to some of the deleted arguments. Note that, in particular, if $\Gamma_F(\emptyset) = S$, then $g(F) = S \cup g(\bigcap_{a \in S} F_a^*)$. Also note that arguments defended by a in F are unattacked in F_a^* .

Example 14 Consider AFs

$$\begin{aligned} F &= (\{a, b, c, d\}, \{(a, b), (b, c), (c, d), (d, c)\}) \text{ and} \\ G &= (\{a, b, c, d\}, \{(a, b), (c, d), (d, c)\}). \end{aligned}$$

Now, $F^{g\kappa} = F \neq G = G^{g\kappa}$ and thus $F \not\equiv_s^g G$. We notice that $\Gamma_F(\emptyset) = \Gamma_G(\emptyset) = \{a\}$ and $F_a^* = G_a^* = (\{c, d\}, \{(c, d), (d, c)\})$. One can now check that given any AF H such that $A(H) \subseteq A(F)$, $g(F \cup H) = g(G \cup H)$ holds, and thus $F \equiv_l^g G$ holds.

There are also other special cases, which we call pathological based on their structure containing self-loops for all arguments (except potentially one).

Definition 11 An AF F is *self-loop pathological* if $(a, a) \in F$ for all $a \in A(F)$, and *b-pathological* if $b \in A(F)$ is unattacked in F and $(a, a) \in F$ for all $a \neq b$.

Finally, we are ready to state the characterization theorem for local equivalence wrt. grounded semantics. In addition to strongly equivalent AFs, we have to take into account frameworks similar to those in Example 14 and the pathological cases.

Theorem 11 For AFs F, G : $F \equiv_l^g G$ iff

- $F \equiv_s^g G$; or
- jointly $A(F) = A(G)$, $\Gamma_F(\emptyset) = \Gamma_G(\emptyset) = S$, and
 - in case $S \neq \emptyset$: $(F_a^*)^{c\kappa} = (G_a^*)^{c\kappa}$ for all $a \in S$; or
 - in case $S = \{a\}$: both F_a^* and G_a^* are self-loop pathological, or $A(F_a^*) = A(G_a^*)$ and both F_a^* and G_a^* are *b-pathological* for some b .

Summary of Results for Local Equivalence

In the following we give a summary of our results for local equivalence.


Theorem 12 The following relations holds for all AFs F, G , and $\sigma \in \{a, p, i, ss\}$,

- (1) $F \equiv_l^\sigma G \implies F \equiv_l^s G$
- (2) $F \equiv_l^c G \implies (F \equiv_l^\sigma G \text{ and } F \equiv_l^g G)$

Proof First, assume $F \equiv_l^\sigma G$. Then $F \equiv_s^\sigma G$ by Theorem 8, and $F \equiv_s^s G$ by Theorem 5. Thus $F \equiv_l^s G$ follows. Secondly, assume $F \equiv_l^c G$. Consider an arbitrary AF H such that $A(H) \subseteq A(F) \cup A(G)$. Since $F \equiv_l^c G$, we have $c(F \cup H) = c(G \cup H)$. By Proposition 1, $g(F \cup H) = g(G \cup H)$ and since H was arbitrary, $F \equiv_l^g G$ follows. Also, $p(F \cup H) = p(G \cup H)$ by Proposition 1, and $F \equiv_l^p G$ follows. By Theorems 8 and 2, $F \equiv_l^\sigma G$ for $\sigma \in \{a, p, i, ss\}$ follows. $\square$

We sketch next that none of the other relations hold.

Example 15 Recall AFs F, G, H , and I from Example 10:



By recalling that $F \equiv_s^\sigma G$ for $\sigma \in \{a, p, i, ss\}$ and $F \equiv_s^s G$, and noticing that $A(F) = A(G)$, we get $F \equiv_l^\sigma G$ and $F \equiv_l^s G$. On the other hand, $F \not\equiv_l^g G$ since $g(F) = \{\emptyset\} \neq \{\{b\}\} = g(G)$. Since all arguments in F and H are attacked, $g(F) = g(H) = \{\emptyset\}$ and $F \equiv_l^g H$ follows. However, $F \not\equiv_l^c H$ since $c(F) \neq c(H)$, $F \not\equiv_l^\sigma H$ since $F \not\equiv_s^\sigma H$, and $F \not\equiv_l^s H$ since $F \not\equiv_s^s H$ and $A(F) = A(H)$. Finally $H \equiv_l^s I$ since $H \equiv_s^s I$, whereas $H \not\equiv_l^\sigma I$ since $H \not\equiv_s^\sigma I$. The observation $H \not\equiv_l^c I$ follows by Theorem 12.

Similarly to strong equivalence, self-loops are in a crucial role with local equivalence. In fact, for self-loop free AFs, the concept of local equivalence collapses to syntactic equivalence for all but grounded semantics (see Example 14 for self-loop free AFs F and G such that $F \neq G$ and $F \equiv_l^g G$).

Theorem 13 *For any self-loop free AFs F and G , $F = G$ iff $F \equiv_l^\sigma G$ for $\sigma \in \{s, a, p, c, i, ss\}$.*

Discussion

We studied strong equivalence in the context of abstract argumentation and provided characterizations how to decide strong equivalence with respect to the most important semantics for argumentation frameworks. In particular, we showed that strong equivalence wrt. admissible, preferred, ideal, and respectively, semi-stable semantics coincides, while stable, grounded, and complete semantics yield different notions of strong equivalence. This is, however, only true for argumentation frameworks which possess self loops. For self-loop free frameworks we could show that strong equivalence wrt. all considered semantics amounts to syntactical equivalence. This strengthens the assumption that abstract argumentation frameworks provide a very “compact” KR formalism. In other words, in terms of strong equivalence there is no room for redundancy in such (self-loop free) frameworks.

We also considered strong equivalence when defined over credulously and respectively skeptically accepted arguments, and showed that in all but admissible semantics these concepts reduce to the corresponding standard notion of strong equivalence. Finally, we weakened the concept of strong equivalence by considering only augmentation of AFs which do not add any new arguments to the compared frameworks.

Understanding dynamics in argumentation is essential for revising argumentation theories (Falappa, Kern-Isberner, and Simari 2009). Indeed, we provided here a somewhat orthogonal access for understanding the semantics of dynamically evolving AFs (neither arguments nor attacks are withdrawn). Our results indicate that dynamically equivalent knowledge is very close to symbolic equivalence in terms of AFs. In our setting, “evolving” refers to the union of a given AF by a further AF which carries the additional information. It is argued in (Coste-Marquis et al. 2007) that a union of AFs does not necessarily take into account the implicit knowledge of its parts, and thus union is not considered as a suitable operation for merging AFs seen as individual agents. However, strongly equivalent AFs have the same extensions regardless of the additional knowledge

provided, and thus strong equivalence can indeed be understood as a property deciding whether two AFs provide the same implicit information.

There is further work which is related to our investigations. Rotstein et al. (2008a) add dynamics to abstract argumentation via so-called dynamic argumentation frameworks (DAFs). Given an associated set of evidence, DAFs are reduced to AFs in the classical sense. Evidence update and erasure are used to change the instance of the DAF. Attack refinement and abstraction (Boella, Kaci, and van der Torre 2009a; 2009b) is viewed as an addition or respectively removal of a single attack from the set of arguments of the original AF. Cayrol, Dupin de Saint-Cyr, and Lagasque-Schiex (2008) carefully analyzed the situation when a new argument a is introduced, limiting to cases in which this argument only has a single interaction with the original AF. In other words, given an AF F it can be combined with H such that $A(H) = A(F) \cup \{a\}$, $a \notin A(F)$, $R(H) \in \{\{(a, b)\}, \{(b, a)\}\}$. Under these assumptions, the relationship between the set of extensions of the original AF F and its revision $F \cup H$ is then analyzed. Note, however, that in our proofs we mostly used more involved AFs H to show that two AFs are not strongly equivalent. Nonetheless, there might be certain common aspects in revising AFs and strong equivalence between AFs; this analysis is ongoing work.

Recent work by Lonc and Truszczyński (2010) investigates equivalence relations on graphs in a very general setting. Their methods could provide an alternative way to derive some of the results provided here, in particular Theorem 6. For other results in this paper, the techniques provided in (Lonc and Truszczyński 2010) seem not to be applicable, since these techniques rely on self-loop free graphs and deal with unrestricted extensions of graphs.

Our future goal is to obtain deeper insight to the dynamics of AFs by considering further variants of equivalence obtained by parameterizing the augmented AFs. A general parameterization scheme would cover the cases studied in this work as special cases. Other interesting special cases can be seen to arise. For instance, one could consider generalizations of the setting as Cayrol, Dupin de Saint-Cyr, and Lagasque-Schiex (2008), i.e., allow the context AF H to attack F through new arguments only, or H to contain new arguments which can be attacked by arguments from F only.

Acknowledgements

The authors want to thank Gerd Brewka, Ringo Baumann and the anonymous referees for their comments which helped to improve this paper.

References

- Baroni, P., and Giacomin, M. 2009. Semantics of abstract argument systems. In Rahwan, I., and Simari, G., eds., *Argumentation in Artificial Intelligence*. Springer. 25–44.
- Bench-Capon, T. J. M., and Dunne, P. E. 2007. Argumentation in artificial intelligence. *Artif. Intell.* 171(10-15):619–641.

- Boella, G.; Kaci, S.; and van der Torre, L. 2009a. Dynamics in argumentation with single extensions: Abstraction principles and the grounded extension. In Sossai, C., and Chemello, G., eds., *Proceedings of the 10th European Conference on Symbolic and Quantitative Approaches to Reasoning with Uncertainty (ECSQARU 2009)*, volume 5590 of *Lecture Notes in Computer Science*, 107–118. Springer.
- Boella, G.; Kaci, S.; and van der Torre, L. 2009b. Dynamics in argumentation with single extensions: attack refinement and the grounded extension. In Sierra, C.; Castelfranchi, C.; Decker, K. S.; and Sichman, J. S., eds., *Proceedings of the 8th International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS 2009)*, Volume 2, 1213–1214. IFAAMAS.
- Caminada, M. 2006. Semi-stable semantics. In Dunne, P. E., and Bench-Capon, T. J. M., eds., *Computational Models of Argument: Proceedings of COMMA 2006*, volume 144 of *Frontiers in Artificial Intelligence and Applications*, 121–130. IOS Press.
- Cayrol, C.; Dupin de Saint-Cyr, F.; and Lagasquie-Schiex, M.-C. 2008. Revision of an argumentation system. In Brewka, G., and Lang, J., eds., *Proceedings of the 11th International Conference on Principles of Knowledge Representation and Reasoning (KR 2008)*, 124–134. AAAI Press.
- Coste-Marquis, S.; Devred, C.; Konieczny, S.; Lagasquie-Schiex, M.-C.; and Marquis, P. 2007. On the merging of Dung’s argumentation systems. *Artif. Intell.* 171(10-15):730–753.
- Dung, P. M.; Mancarella, P.; and Toni, F. 2007. Computing ideal sceptical argumentation. *Artif. Intell.* 171(10-15):642–674.
- Dung, P. M. 1995. On the acceptability of arguments and its fundamental role in nonmonotonic reasoning, logic programming and n-person games. *Artif. Intell.* 77(2):321–358.
- Dunne, P. E. 2009. The computational complexity of ideal semantics. *Artif. Intell.* 173(18):1559–1591.
- Falappa, M. A.; Kern-Isberner, G.; and Simari, G. R. 2009. Belief revision and argumentation theory. In Rahwan, I., and Simari, G., eds., *Argumentation in Artificial Intelligence*. Springer. 341–360.
- Lifschitz, V.; Pearce, D.; and Valverde, A. 2001. Strongly Equivalent Logic Programs. *ACM Transactions on Computational Logic* 2(4):526–541.
- Lonc, Z., and Truszczyński, M. 2010. On graph equivalences preserved under extensions. <http://arxiv.org/abs/1002.1749>.
- Rotstein, N. D.; Moguillansky, M. O.; Falappa, M. A.; García, A. J.; and Simari, G. R. 2008a. Argument theory change: Revision upon warrant. In Besnard, P.; Doutre, S.; and Hunter, A., eds., *Computational Models of Argument: Proceedings of COMMA 2008*, volume 172 of *Frontiers in Artificial Intelligence and Applications*, 336–347. IOS Press.
- Rotstein, N. D.; Moguillansky, M. O.; Garcia, A. J.; and Simari, G. R. 2008b. An abstract argumentation framework for handling dynamics. In *Proceedings of the 12th International Workshop on Non-Monotonic Reasoning (NMR’08)*, 131–139.
- Truszczyński, M. 2006. Strong and uniform equivalence of nonmonotonic theories - an algebraic approach. *Ann. Math. Artif. Intell.* 48(3-4):245–265.
- Turner, H. 2004. Strong equivalence for causal theories. In Lifschitz, V., and Niemelä, I., eds., *Proceedings of the 7th International Conference on Logic Programming and Non-monotonic Reasoning (LPNMR’04)*, volume 2923 of *Lecture Notes in Computer Science*, 289–301. Springer.