Faculty of Mathematics and Computer Science

Artificial Intelligence Group

Refining the Acceptance Problem in Abstract Argumentation by selecting Extensions

Bachelor's Thesis

in partial fulfillment of the requirements for
the degree of Bachelor of Science (B.Sc.)
in Informatik

submitted by
Lucas Filipiak

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

Advisor: Kenneth Skiba
Artificial Intelligence Group

Statement

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

Essen, 09.04.2025

(Place, Date) (Signature)

Zusammenfassung

Eine Herausforderung für leichtgläubige und skeptische Akzeptanz in Abstrakter Argumentation ist die potenziell exponentielle Zahl durch eine Semantik bestimmter Extensionen, was die leichtgläubige Akzeptanz zu vieler Argumente oder die skeptische Akzeptanz zu weniger Argumente ermöglicht. In dieser Arbeit werden Methoden zur Verfeinerung des Akzeptanzproblems durch Auswahl bestimmter Extensionen auf Grundlage von Regeln aus der Wahltheorie untersucht.

Abstract

A challenge for credulous and sceptical acceptance in Abstract Argumentation is the potentially exponential number of extensions yielded by an extension-based semantics, allowing too many arguments to be credulously accepted, or too few arguments to be sceptically accepted. In this thesis, methods to refine acceptance by selecting extensions based on rules from Voting Theory are explored.

Contents

1 Introduction

Abstract Argumentation has proven to provide a powerful formalism to model argumentation processes, with many applications such as in multi-agent systems and logic programming [ABG+17]. An abstract argumentation framework (AF) consists only of a set of arguments and a binary attack relation on this set, such that the structure of the arguments is not considered [Dun95]. Extension-based semantics determine sets of arguments which, by some measure, can be accepted together and form a coherent point of view. Credulous and sceptical acceptance are two fundamental approaches to determine the overall acceptance of an argument with respect to one such semantics [KMV15]. However, these approaches pose two extremes, such that in practical applications, it can occur that too many arguments are credulously accepted or too few arguments are sceptically accepted. An alternative approach inspired by Voting Theory [BCE+16] mitigates this issue by selecting only, by some criterion, the best extensions for consideration in overall acceptance, thus potentially decreasing the number of credulously accepted arguments and potentially increasing the number of sceptically accepted arguments. In this thesis, the refinement to an acceptance in between credulous and sceptical acceptance by selecting extensions will be studied.

To determine which extensions are selected, Konieczny et al. [KMV15] have proposed pairwise comparison using the Copeland rule with one of four criteria, which selects the extensions that compare best among all extensions. However, the application of this rule might not be adequate in cases where the goal is to select the least objectionable extensions. Therefore, the application of the Simpson rule is proposed, which selects the extensions that compare best to each extension. For each semantics, this gives rise to a new semantics, called a selection semantics.

A well-established method to classify semantics are principle-based systems, such that a semantics satisfies a principle iff its set of extensions always meets certain requirements [BG07, vDTV18]. This enables the selection of semantics for practical use cases by the principles they satisfy. While extension selection might help control the number of accepted arguments, the resulting selection semantics does not necessarily inherit the principles satisfied by the underlying semantics. This is problematic because desired guarantees to the structure of the set of extensions, and consequently to the set of accepted arguments, are lost. Inheritance will be studied for the classical semantics as proposed by Dung [Dun95], the two proposed selection rules, and the four proposed comparison criteria in this thesis. Results show that principles are often not inherited by the introduced selection semantics except in trivial cases, such that many counterexamples will be presented.

The remainder of this thesis is organized as follows:

- Section 2 introduces core notions of an abstract argumentation framework and extension-based semantics.
- Section 3 explains pairwise comparison criteria and presents the Copeland-based extensions semantics.

- Section 4 introduces score functions and the Simpson rule, leading to the Simpson-based extensions semantics.
- Section 5 investigates acceptance-consistency and shows how its relation with the maintenance of non-emptiness.
- Section 6 examines the inheritance of principles from the underlying semantics by the refined semantics.
- Section 7 discusses remaining problems and mitigation techniques.
- Section 8 summarises and reflects on our main findings.
- Section 9 concludes with final remarks and directions for future research.

2 Formal Argumentation

Formal Argumentation is a field of Artificial Intelligence which seeks to formalise human argumentation processes that have originally been studied as argumentation theory in philosophy [BGGT18]. This enables the implementation of these processes in computers, with numerous practical applications including in e-governance and multi-agent systems [ABG$^+$17].

A model that considers only the relations between arguments and not their internal structure was introduced by Dung in 1995 [Dun95] and has since been subject of extensive research. Dung proposes the notion of an *abstract argumentation framework*.

Definition 2.1 (Abstract argumentation framework)**.** An *abstract argumentation framework (AF)* is a pair $(\mathcal{A}, \hookrightarrow)$ of a finite set of *arguments* $\mathcal{A}$ and a binary relation, $\hookrightarrow \subseteq \mathcal{A} \times \mathcal{A}$, on $\mathcal{A}$, called an *attack relation*.

Remark 2.1. For an AF $(\mathcal{A}, \hookrightarrow)$ with arguments $a, b \in \mathcal{A}$, we denote $(a, b) \in \hookrightarrow$ also by $a \hookrightarrow b$, say that a attacks b, and call (a, b) an attack from a (or any set of arguments $\mathcal{S} \subseteq \mathcal{A}$ where $a \in \mathcal{S}$) to b (or any set of arguments $\mathcal{S}' \subseteq \mathcal{A}$ where $b \in \mathcal{S}'$). For the negation, we write $a \not\hookrightarrow b$. The symbol of a binary relation will be used accordingly also in general.

This formalism can be used to model n-person games like the stable marriage problem as well as logic programming [Dun95]. In the following, we will see a practical example for the negotiation of ingredients in a recipe.

Example 2.1 (Smoothie AF)**.** *A real-world scenario in which three wishes for the ingredients of a smoothie are expressed, can be modelled as an AF $F = (\mathcal{A}, \hookrightarrow)$ where*

$$\mathcal{A} = \{\text{“Apple”}, \text{“Fruit”}, \text{“Vegetable”}\}$$

and

$$\hookrightarrow = \{\text{“Apple”}, \text{“Fruit”}\} \times \{\text{“Vegetable”}\} \cup \{\text{“Vegetable”}\} \times \{\text{“Apple”}, \text{“Fruit”}\},$$

such that there exist (symmetric) attacks between any two wishes iff they cannot describe the same ingredient. An apple is a fruit, therefore there exists no attack between "Apple" and "Fruit". On the other hand, no ingredient except for "Vegetable" itself qualifies as a vegetable, therefore there exist attacks between "Vegetable" and the other ingredients.

For the purposes of this thesis, the restricton to finite AFs is necessary in order to allow for quantitative comparisons between sets of arguments. Moreover, this is a common restriction in research and can be considered a natural restriction in practical use cases. This restriction enables us to visualise AFs, especially our example, by directed graphs as can be seen in Figure 1.

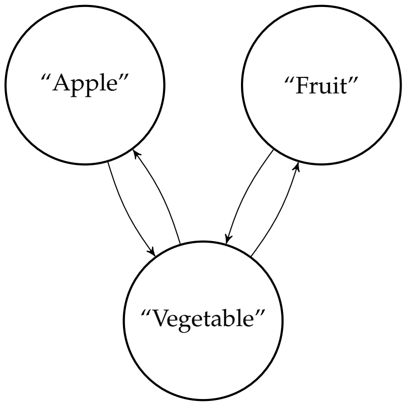

Figure 1: An AF for the primary ingredient of a smoothie

There exist many methods to extract information from an AF, called *semantics*. An important class of semantics are the *extension-based semantics*. An extension-based semantics determines sets of arguments that can be accepted together by some measure and form a coherent point of view, called *extensions* [Dun95]. As extension-based semantics are most relevant for this thesis, we consider a semantics to be extension-based unless stated otherwise.

Definition 2.2 (Power set)**.** The *power set of a set* S is denoted by $2^S = \{S' \subseteq S\}$.

Definition 2.3 (Extension-based semantics)**.** An (*extension-based*) *semantics* is a function σ on the set of all AFs, such that for all AFs $F = (\mathcal{A}, \hookrightarrow)$, it holds for the *set of extensions*, $\sigma(F)$, that

$$\sigma(F) \subseteq 2^{\mathcal{A}}.$$

Dung [Dun95] has proposed several semantics which include the *preferred*, *stable*, *complete*, and *grounded* semantics.

Remark 2.2. For an AF $F = (\mathcal{A}, \hookrightarrow)$ and semantics σ, iff a set of arguments $\mathcal{E} \subseteq \mathcal{A}$ is contained in the set of extensions $\sigma(F)$, e. g. of the preferred semantics, we also refer to it by the name of the semantics. In this example, $\mathcal{E}$ can be referred to as preferred (in F) or a preferred extension (of F). For a semantics σ in general, we say that $\mathcal{E}$ is a σ-extension (of F). As this and the following notions defined in this thesis depend on some given objects, omissions will be made for better readability where there is no ambiguity. Possible omissions are indicated by parantheses.

The semantics proposed by Dung are based on three central notions, which are those of a *conflict-free* set of arguments, an *acceptable* argument *with respect to a set of arguments*, and an *admissible* set of arguments.

Definition 2.4 (Conflict-freeness, acceptability, and admissibility)**.** For an AF $F = (\mathcal{A}, \hookrightarrow)$ and set of arguments $\mathcal{S} \subseteq \mathcal{A}$, *conflict-freeness, acceptability, and admissibility* are defined as follows:

- $\mathcal{S}$ is *conflict-free* (*in* F) iff there exist no $a, b \in \mathcal{S}$ such that $a \hookrightarrow b$.
- $a \in \mathcal{A}$ is *acceptable with respect to* $\mathcal{S}$ (*in* F) iff for all $b \in \mathcal{A}$ where $b \hookrightarrow a$, there exists $c \in \mathcal{S}$ such that $c \hookrightarrow a$.
- $\mathcal{S}$ is *admissible* (*in* F) iff $\mathcal{S}$ is conflict-free and all $a \in \mathcal{S}$ are acceptable with respect to $\mathcal{S}$.

The complete, preferred, stable, and grounded semantics are then defined such that a set of arguments is contained in a set of extensions iff it meets certain requirements.

Definition 2.5 (Minimal/maximal and least/greatest element)**.** For a set S, binary relation R, and element $x \in S$, the *minimal/maximal and least/greatest element* is defined as follows:

- Element x is *minimal with respect to* R (*among* S) iff there exists no element $y \in S$ such that $y \ R \ x$, and $x \not R \ y$ or $y \not R \ x$. Respectively, we say that x is *maximal with respect to* R (*among* S) iff the former holds in reverse.
- Element x is *the least with respect to* R (*among* S) iff for all $y \in S$, it holds that $x \ R \ y$. Respectively, we say that x is *the greatest with respect to* R (*among* S) iff the former holds in reverse.

Definition 2.6 (Complete, preferred, stable, and grounded semantics)**.** For an AF $F = (\mathcal{A}, \hookrightarrow)$ and set of arguments $\mathcal{E} \subseteq \mathcal{A}$, the *complete, preferred, stable, and grounded semantics* are defined as follows:

- $\mathcal{E}$ is *complete* (*in* F) (or $\mathcal{E} \in \mathrm{Co}(F)$) iff $\mathcal{E}$ is admissible and for all $a \in \mathcal{A}$ that are acceptable with respect to $\mathcal{E}$, it holds that $a \in \mathcal{E}$.
- $\mathcal{E}$ is *preferred* (*in* F) (or $\mathcal{E} \in \mathrm{Pr}(F)$) iff $\mathcal{E}$ is admissible and maximal with respect to set inclusion among the sets of admissible arguments. Every preferred extension is also complete. There exists at least one preferred extension.

- $\mathcal{E}$ is *stable* (*in* F) (or $\mathcal{E} \in \mathrm{St}(F)$) iff $\mathcal{E}$ is conflict-free and for all $b \in \mathcal{A} \setminus \mathcal{E}$, there exists $a \in \mathcal{E}$ such that $a \hookrightarrow b$. Every stable extension is also preferred.
- $\mathcal{E}$ is *grounded* (*in* F) (or $\mathcal{E} \in \mathrm{Gr}(F)$) iff $\mathcal{E}$ is minimal (or the least) with respect to set inclusion among the complete extensions. The grounded extension is unique.

To determine the overall acceptance of an argument under an extension-based semantics, there are two fundamental approaches which are *credulous acceptance* and *sceptical acceptance* [KMV15].

Definition 2.7 (Credulous and sceptical acceptance)**.** For an AF $F = (\mathcal{A}, \hookrightarrow)$, semantics σ, and argument $a \in \mathcal{A}$, *credulous and sceptical acceptance* are defined as follows:

- Argument a is *credulously* σ*-accepted* (*from* F) (or *credulously accepted* (*with respect to* $\sigma(F)$)) iff it is contained in *at least one* σ-extension. The set of credulously σ-accepted arguments is denoted by $\mathrm{Cr}_\sigma(F)$.
- Argument a is *sceptically* σ*-accepted* (*from* F) (or *sceptically accepted* (*with respect to* $\sigma(F)$)) iff it is contained in *every* σ-extension. The set of sceptically σ-accepted arguments is denoted by $\mathrm{Sc}_\sigma(F)$.

A problem in connection with these approaches is that the number of extensions can be exponential, possibly allowing too many arguments to be credulously accepted or too few arguments to be sceptically accepted [KMV15]. Ulbricht and Baumann [UB19] point out that the latter case can occur due to inconsistencies in the knowledge base of an agent. Besides the repair of these inconsistencies as described by Ulbricht and Baumann, one could select only certain extensions with a suitable criterion in order to mitigate this issue, decreasing the number of inconsistencies and thus potentially increasing the number of sceptically accepted arguments. A suitable criterion could be one which prefers extensions that are less prone to inconsistencies or more likely to be correct than others.

Example 2.2 (Extension selection)**.** *Consider the following example with the AF shown by Figure 1 for Example 2.1. We have two stable extensions:* $\mathcal{E}_1 = \{$*"Apple"*, *"Fruit"*$\}$ *and* $\mathcal{E}_2 = \{$*"Vegetable"*$\}$*. In a larger AF with wishes for a multitude of ingredients, credulous acceptance would not be a reasonable approach to select ingredients for a smoothie, as this could lead to the acceptance of too many ingredients. As* $\mathcal{E}_1 \cap \mathcal{E}_2 = \emptyset$*, no argument is sceptically accepted. However, it is necessary to select some ingredient, even if that means neglecting some wishes. In this case, it can be seen as desirable to consider only the larger extension* $\mathcal{E}_1$ *for credulous or sceptical acceptance [KMV15], as to fullfill the most wishes. Then, arguments "Apple" and "Fruit" are credulously and sceptically accepted.*

For this reason, alternative approaches are researched as shown in Section 3. The refinement of this acceptance problem to an acceptance in between credulous and sceptical acceptance is subject of this thesis.

3 Pairwise comparison of extensions

Konieczny et al. [KMV15] have proposed pairwise comparison of extensions by certain criteria which give rise to pairwise comparison relations on the set of extensions of an AF, which will be used in extension selection.

Definition 3.1 (Pairwise comparison criterion and relation)**.** For an AF F and semantics σ, a (*pairwise*) *comparison criterion* γ gives rise to a binary relation

$$\geq^F_{\sigma,\gamma} \subseteq \sigma(F) \times \sigma(F)$$

on $\sigma(F)$, called a (*pairwise*) *comparison relation*.

As sets of extensions shall be selected and yielded for every AF, a function to realise this is a semantics. Hence, a fitting notion is that of a *selection semantics*, a semantics which yields a subset of extensions for each AF.

Definition 3.2 (Selection semantics)**.** A *selection semantics* (*with respect to a semantics σ and comparison criterion γ*), $\mathrm{SEL}_{\sigma,\gamma}$, is a semantics, such that for all AFs F, it holds that

$$\mathrm{SEL}_{\sigma,\gamma}(F) \subseteq \sigma(F).$$

Konieczny et al. [KMV15] have proposed that extensions for an AF are then selected based on the *Copeland rule* from voting theory [BCE+16].

Definition 3.3 ($\arg\max$)**.** For a finite set S and integer-valued function f on a subset $S' \subseteq S$, if S' is non-empty, we denote by

$$\arg\max_{x \in S} f(x) = \left\{ y \in S' \;\middle|\; f(y) = \max_{z \in S'} f(z) \right\}$$

the *set of arguments in S' for which f is maximal*. If S' is empty, there exists no maximum and we set $\arg\max_{x \in S} f(x) = \varnothing$.

Definition 3.4 (Copeland score function)**.** For an AF F, semantics σ, comparison criterion γ, and extension $\mathcal{E} \in \sigma(F)$, the *Copeland score function* (*with respect to semantics σ and comparison criterion γ*), $\mathrm{Copeland}^F_{\sigma,\gamma}$, is defined by

$$\mathrm{Copeland}^F_{\sigma,\gamma}(\mathcal{E}) = \left| \left\{ \mathcal{E}' \in \sigma(F) \;\middle|\; \mathcal{E} \geq^F_{\sigma,\gamma} \mathcal{E}' \right\} \right| - \left| \left\{ \mathcal{E}'' \in \sigma(F) \;\middle|\; \mathcal{E}'' \geq^F_{\sigma,\gamma} \mathcal{E} \right\} \right| .$$

Definition 3.5 (Copeland-based extensions semantics)**.** For a Copeland score function $\mathrm{Copeland}^F_{\sigma,\gamma}$, the *Copeland-based extensions (CBE) semantics* (*with respect to semantics σ and comparison criterion γ*), $\mathrm{CBE}_{\sigma,\gamma}$, is a selection semantics and defined by

$$\mathrm{CBE}_{\sigma,\gamma}(F) = \arg\max_{\mathcal{E} \in \sigma(F)} \mathrm{Copeland}^F_{\sigma,\gamma}(\mathcal{E}).$$

Informally, this function selects the extensions that *compare best among all extensions*. It considers for every extension the number of extensions that are stronger or weaker with respect to the employed comparison criterion, and selects the extensions with the maximal difference of these "wins" and "losses".

Konieczny et al. [KMV15] have pointed out that, as a subset of extensions is selected, any argument credulously accepted with respect to $\mathrm{CBE}_{\sigma,\gamma}(F)$ is also credulously accepted with respect to $\sigma(F)$, and any argument sceptically accepted with respect to $\sigma(F)$ is also sceptically accepted with respect to $\mathrm{CBE}_{\sigma,\gamma}(F)$. This potentially *decreases the number of credulously accepted arguments* and potentially *increases the number of sceptically accepted arguments*. In fact, this *acceptance-consistency* will be achieved by any selection semantics provided that at least one extension is selected from any non-empty set of extensions. Therefore, *maintaining non-emptiness* is identified to be a desirable property of a selection semantics.

Definition 3.6 (Maintenance of non-emptiness)**.** A selection semantics $\mathrm{SEL}_{\sigma,\gamma}$ *maintains non-emptiness* iff for all AFs F such that $\sigma(F)$ is non-empty, $\mathrm{SEL}_{\sigma,\gamma}(F)$ is non-empty.

Definition 3.7 (Acceptance-consistency)**.** A selection semantics $\mathrm{SEL}_{\sigma,\gamma}$ is *acceptance-consistent* iff for all AFs F, it holds that

$$\mathrm{Cr}_{\mathrm{SEL}_{\sigma,\gamma}}(F) \subseteq \mathrm{Cr}_{\sigma}(F) \text{ and } \mathrm{Sc}_{\sigma}(F) \subseteq \mathrm{Sc}_{\mathrm{SEL}_{\sigma,\gamma}}(F).$$

Lemma 3.1 (Acceptance-consistency of $\mathrm{SEL}_{\sigma,\gamma}$)**.** *Any selection semantics that maintains non-emptiness is acceptance-consistent.*

Proof. Let F be an AF and $\mathrm{SEL}_{\sigma,\gamma}$ a selection semantics that maintains non-emptiness. Acceptance-consistency will be shown in two steps, for either requirement:

1. For any credulously $\mathrm{SEL}_{\sigma,\gamma}$-accepted argument a, there exists a selected extension $\mathcal{E} \in \mathrm{SEL}_{\sigma,\gamma}(F)$ such that $a \in \mathcal{E}$. As $\mathrm{SEL}_{\sigma,\gamma}$ is a selection semantics, $\mathcal{E}$ is a σ-extension. Therefore, a is credulously σ-accepted.

2. If $\sigma(F)$ is empty, we have $\sigma(F) = \mathrm{SEL}_{\sigma,\gamma}(F)$ as $\mathrm{SEL}_{\sigma,\gamma}$ is a selection semantics, such that we are done. If $\sigma(F)$ is non-empty, it follows from the maintenance of non-emptiness that $\mathrm{SEL}_{\sigma,\gamma}(F)$ is non-empty, such that there exists a selected extension $\mathcal{E} \in \mathrm{SEL}_{\sigma,\gamma}(F)$. As $\mathrm{SEL}_{\sigma,\gamma}$ is a selection semantics, any $\mathrm{SEL}_{\sigma,\gamma}$-extension is a σ-extension, especially $\mathcal{E}$. Any sceptically σ-accepted argument a is contained in all σ-extensions, especially $\mathcal{E}$. Therefore, a is sceptically $\mathrm{SEL}_{\sigma,\gamma}$-accepted.

□

Proposition 3.1 (Maintenance of non-emptiness by $\mathrm{CBE}_{\sigma,\gamma}$)**.** *All Copeland-based extensions semantics maintain non-emptiness.*

Proof. Let $\mathrm{CBE}_{\sigma,\gamma}$ be a Copeland-based extensions semantics. If $\sigma(F)$ is empty, we have $\sigma(F) = \mathrm{CBE}_{\sigma,\gamma}(F)$ as $\mathrm{CBE}_{\sigma,\gamma}$ is a selection semantics, such that we are done. The corresponding Simpson score function is defined on $\sigma(F)$ as for all extensions $\mathcal{E} \in \sigma(F)$, either set in its definition exists. Therefore, if $\sigma(F)$ is non-empty, a maximal Copeland score exists as $\sigma(F)$ is finite, such that at least one extension is selected. □

Corollary 3.1 (Acceptance-consistency of $\mathrm{CBE}_{\sigma,\gamma}$)**.** *All Copeland-based extensions semantics are acceptance-consistent.*

Proof. Follows from Proposition 3.1 in connection with Lemma 3.1. □

For pairwise comparison, Konieczny et al. [KMV15] have proposed the comparison criteria nonatt, strdef, delarg, and delatt.

Definition 3.8 (Strong defence)**.** For an AF $F = (\mathcal{A}, \hookrightarrow)$, sets of arguments $S, S' \subseteq \mathcal{A}$, and argument $a \in \mathcal{A}$, we say that *a is strongly defended from $\mathcal{S}$ by $\mathcal{S}'$ (in F)* iff for every attack $b \hookrightarrow a$ from $\mathcal{S}$, there exists an argument $c \in \mathcal{S}' \setminus \{a\}$, such that $c \hookrightarrow b$ and c is strongly defended from $\mathcal{S}$ by $\mathcal{S}' \setminus \{a\}$.

Definition 3.9 (Restriction)**.** The *restriction of a binary relation R to a set S* is denoted by $R_{\restriction S} = R \cap (S \times S)$.

Definition 3.10 (Comparison relations for nonatt, strdef, delarg, and delatt)**.** For an AF F, semantics σ, and extensions $\mathcal{E}, \mathcal{E}' \in \sigma(F)$, the *comparison relations for comparison criteria* nonatt, strdef, delarg, *and* delarg are defined as follows:

- $\mathcal{E} \geq^F_{\sigma,\mathrm{nonatt}} \mathcal{E}'$ iff the number of arguments in $\mathcal{E}$ not attacked by an argument in $\mathcal{E}'$ is greater than or equal to that of arguments in $\mathcal{E}'$ not attacked by an argument in $\mathcal{E}$.
- $\mathcal{E} \geq^F_{\sigma,\mathrm{strdef}} \mathcal{E}'$ iff the number of arguments in $\mathcal{E}$ strongly defended from $\mathcal{E}'$ by $\mathcal{E}$ is greater than or equal to that of arguments in $\mathcal{E}'$ strongly defended from $\mathcal{E}$ by $\mathcal{E}'$.
- $\mathcal{E} \geq^F_{\sigma,\mathrm{delarg}} \mathcal{E}'$ iff the cardinality of any largest subset $\mathcal{S}$ of $\mathcal{E}$ such that $\mathcal{E}$ is an extension of AF $(\mathcal{E} \cup \mathcal{E}', \hookrightarrow_{\restriction \mathcal{E} \cup \mathcal{E}'})$ after deleting all attacks from $\mathcal{S}$ to $\mathcal{E}'$ is greater than or equal to the cardinality of any largest subset $\mathcal{S}'$ of $\mathcal{E}'$ such that $\mathcal{E}'$ is an extension of AF $(\mathcal{E} \cup \mathcal{E}', \hookrightarrow_{\restriction \mathcal{E} \cup \mathcal{E}'})$ after deleting all attacks from $\mathcal{S}'$ to $\mathcal{E}$.
- $\mathcal{E} \geq^F_{\sigma,\mathrm{delatt}} \mathcal{E}'$ iff the maximal number of attacks from $\mathcal{E}$ to $\mathcal{E}'$ that can be deleted such that $\mathcal{E}$ is an extension of AF $(\mathcal{E} \cup \mathcal{E}', \hookrightarrow_{\restriction \mathcal{E} \cup \mathcal{E}'})$ is greater than or equal to that of attacks from $\mathcal{E}'$ to $\mathcal{E}$ that can be deleted such that $\mathcal{E}'$ is an extension of AF $(\mathcal{E} \cup \mathcal{E}', \hookrightarrow_{\restriction \mathcal{E} \cup \mathcal{E}'})$.

We can see that the relations for comparison criteria nonatt and strdef are always *total* (and thus *reflexive*) as the numbers of unattacked and strongly defended arguments always exist. However, strdef is found to be an inadequate comparison criterion for semantics based on admissibility such as the semantics proposed by Dung [Dun95], as its comparison relations coincide with those for nonatt.

Proposition 3.2 (Equality of $\geq^F_{\sigma,\mathrm{nonatt}}$ and $\geq^F_{\sigma,\mathrm{strdef}}$)**.** *For all semantics σ such that for all AFs F, $\sigma(F)$ is a set of admissible sets of arguments, comparison relations $\geq^F_{\sigma,\mathrm{nonatt}}$ and $\geq^F_{\sigma,\mathrm{strdef}}$ are equal.*

Proof. Let $F = (\mathcal{A}, \hookrightarrow)$ be an AF and σ a semantics such that for all AFs F', it holds that $\sigma(F)$ is a set of admissible sets of arguments, with extensions $\mathcal{E}, \mathcal{E}' \in \sigma(F)$ where $a_1 \in \mathcal{E}$. The equality will be shown in two steps, for either direction:

1. If a_1 is not attacked by $\mathcal{E}'$, it is trivially strongly defended from $\mathcal{E}'$ by $\mathcal{E}$.

2. Let a_1 be strongly defended from $\mathcal{E}'$ by $\mathcal{E}$. If a_1 is trivially strongly defended, it is not attacked by $\mathcal{E}'$. If a_1 is not trivially strongly defended, there exist an attack $b_1 \hookrightarrow a_1$ from $\mathcal{E}'$ and an argument $c_1 \in \mathcal{E}$, such that $c_1 \hookrightarrow b_1$ and c_1 is strongly defended from $\mathcal{E}'$ by $\mathcal{E} \setminus \{a_1\}$. We can continue this recursion with $c_1 = a_2$ and so on, until we encounter the first trivially strongly defended argument a_n, as $\mathcal{E}$ is finite such that the empty set would eventually strongly defend an argument $c_{|\mathcal{E}|-1} = a_{|\mathcal{E}|}$. Then, in contradiction to the admissibility of $\mathcal{E}'$, a_n attacks $\mathcal{E}'$ as it defends a_{n-1}, but is not attacked by $\mathcal{E}'$.

Therefore, it follows that $\geq^F_{\sigma,\mathrm{nonatt}} = \geq^F_{\sigma,\mathrm{strdef}}$. □

Despite this negative result for the effectiveness of comparison criterion strdef with semantics based on admissibility, we will continue to consider this comparison criterion as it could be used for semantics that are not based on admissibility. On the other hand, the relations for comparison criteria delarg and delatt require the existence of extensions in AFs derived from a reduced AF of any two extensions, which is generally not guaranteed as will be shown in the following.

Definition 3.11 (Totality and reflexivity)**.** Let $R \subseteq S \times S$ be a binary relation on a set S. *Totality and reflexivity* are defined as follows:

- R is *total* iff for all $x, y \in S$, it holds that $(x, y) \in R$ or $(y, x) \in R$.
- R is *reflexive* iff for all $x \in S$, it holds that $(x, x) \in R$. Totality implies reflexivity.

Proposition 3.3 (Totality and reflexivity of $\geq^F_{\sigma,\mathrm{nonatt}}$ and $\geq^F_{\sigma,\mathrm{strdef}}$)**.** *All comparison relations $\geq^F_{\sigma,\mathrm{nonatt}}$ and $\geq^F_{\sigma,\mathrm{strdef}}$ are total and reflexive.*

Proof. The numbers of unattacked and strongly defended arguments exist for any sets of arguments in an AF. □

Proposition 3.4 (Totality and reflexivity of $\geq^F_{\sigma,\mathrm{delarg}}$ and $\geq^F_{\sigma,\mathrm{delatt}}$)**.** *There exist comparison relations $\geq^F_{\sigma,\mathrm{delarg}}$ and $\geq^F_{\sigma,\mathrm{delatt}}$ that are neither total nor reflexive.*

Proof. Let $F = (\mathcal{A}, \hookrightarrow) = (\{a, b\}, \varnothing)$ be an AF and σ a semantics such that $\sigma(F) = \{\{a\}\}$ and for all other AFs $F' \neq F$, it holds that $\sigma(F) = \varnothing$. Then, the reduced AF for extensions $\mathcal{E}, \mathcal{E}' \in \sigma(F)$ according to comparison criteria delarg and delatt, $F_r = (\mathcal{E} \cup \mathcal{E}', \hookrightarrow_{\restriction \mathcal{E} \cup \mathcal{E}'}) = (\{a\}, \varnothing)$, has no extensions as $F_r \neq F$, and no attacks can be deleted. It follows for comparison relations $\geq^F_{\sigma,\mathrm{delarg}}$ and $\geq^F_{\sigma,\mathrm{delatt}}$, that $\mathcal{E} \not\geq^F_{\sigma,\mathrm{delarg}} \mathcal{E}'$ and $\mathcal{E} \not\geq^F_{\sigma,\mathrm{delatt}} \mathcal{E}'$, such that $\geq^F_{\sigma,\mathrm{delarg}}, \geq^F_{\sigma,\mathrm{delatt}} = \varnothing$. Therefore, either comparison relation is neither total nor reflexive. □

While the observation that a comparison relation is not necessarily total or reflexive is interesting by itself, so as to see what extensions are comparable, these properties seem to have no further implications for the Copeland-based extensions semantics. However, in the following section we will see that acceptance-consistency can depend on them for other selection semantics.

Example 3.1 (Copeland-based extensions semantics)**.** *Consider the following example for $\mathrm{CBE}_{\mathrm{Pr},\mathrm{nonatt}}$. In the AF shown by Figure 2, we have three preferred extensions: $\mathcal{E}_1 = \{A, E\}$, $\mathcal{E}_2 = \{B, E\}$, and $\mathcal{E}_3 = \{F\}$. Each extension is assigned a Copeland score which is iteratively calculated as follows, initially being 0. We can determine the selected extensions by counting for every two extensions $\mathcal{E}, \mathcal{E}'$ the number of arguments in $\mathcal{E}$ left unattacked by arguments in $\mathcal{E}'$ and vice versa. If the number of these arguments in $\mathcal{E}$ is greater than or equal to that in $\mathcal{E}'$, we add 1 to the Copeland score of $\mathcal{E}$. If the number of these arguments in $\mathcal{E}'$ is greater than or equal to that in $\mathcal{E}$, we add -1 to the Copeland score of $\mathcal{E}$. We can ignore cases of equality (especially if $\mathcal{E} = \mathcal{E}'$) as these are cancelled out, adding 0 to the Copeland score.*

In $\mathcal{E}_1$, argument 5 is not attacked by an argument in $\mathcal{E}_2$ and vice versa, such that we ignore this case. Argument 1 in $\mathcal{E}_1$ is not attacked by the argument in $\mathcal{E}_3$, however the argument in $\mathcal{E}_3$ is attacked by argument 5 in $\mathcal{E}_1$, such that we add 1 to the Copeland score of $\mathcal{E}_1$. This results in a Copeland score of 1 for $\mathcal{E}_1$.

In $\mathcal{E}_2$, we again ignore the equality with $\mathcal{E}_1$. Both arguments in $\mathcal{E}_2$ are attacked by the argument in $\mathcal{E}_3$ and the argument in $\mathcal{E}_3$ is attacked by argument 5 in $\mathcal{E}_2$, such that we ignore this case. This results in a Copeland score of 0 for $\mathcal{E}_2$.

The argument in $\mathcal{E}_3$ is attacked by argument 5 in $\mathcal{E}_1$, however argument 1 in $\mathcal{E}_1$ is not attacked by the argument in $\mathcal{E}_3$, such that we add -1 to the Copeland score of $\mathcal{E}_3$. The argument in $\mathcal{E}_3$ is attacked by both arguments in $\mathcal{E}_2$ and both arguments in $\mathcal{E}_2$ are attacked by the argument in $\mathcal{E}_3$, such that we ignore this case. This results in a Copeland score of -1 for $\mathcal{E}_3$.

The extensions with the maximal Copeland score of 1 are now selected. Therefore, $\mathcal{E}_1$ is selected.

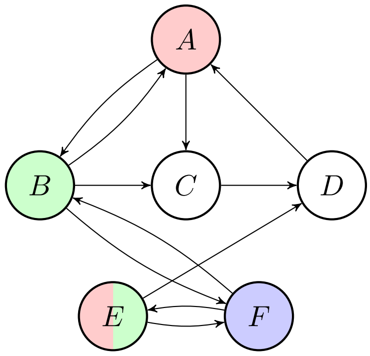

Figure 2: An AF with three preferred extensions

4 Scores and the Simpson rule

This thesis will contribute two concepts to the state of research, which are the notion of a *pairwise comparison score function* and the application of the *Simpson rule* on pairwise comparison for extension selection as outlined in Section 3. A pairwise comparison score function is an integer-valued binary partial function on the set of extensions of an AF, that is used as a means to determine not only whether an extension is stronger or weaker than another with respect to some comparison criterion, but rather a *score* of how much stronger or weaker it is. The comparison of scores is equivalent to the corresponding comparison relation.

Definition 4.1 (Pairwise comparison score function)**.** For an AF F, semantics σ, and comparison criterion γ, a (*pairwise*) *comparison score function* is an integer-valued function, $\mathrm{score}^F_{\sigma,\gamma}$, on a subset of pairs of extensions $E \subseteq \sigma(F) \times \sigma(F)$, such that for all $\mathcal{E}, \mathcal{E}' \in \sigma(F)$ and its corresponding comparison relation, $\geq^F_{\sigma,\gamma}$, it holds that

$$\{\mathcal{E}, \mathcal{E}'\} \times \{\mathcal{E}, \mathcal{E}'\} \subseteq E \text{ and } \mathrm{score}^F_{\sigma,\gamma}(\mathcal{E}, \mathcal{E}') \geq \mathrm{score}^F_{\sigma,\gamma}(\mathcal{E}', \mathcal{E}) \text{ iff } \mathcal{E} \geq^F_{\sigma,\gamma} \mathcal{E}'.$$

This is a necessary step to transfer principles from Voting Theory to extension selection and enable the application of more differentiating selection rules, especially the *Simpson rule*. Different selection rules can be considered more or less adequate depending on the context of their application. While the Copeland rule accounts for the strength of an extension among all extensions, the Simpson rule accounts for its greatest weakness compared to another extension.

4.1 Introduction of scores

The comparison relations for criteria nonatt, strdef, delarg, and delatt do not yield enough information to enable the application of more differentiating rules than the Copeland rule [BCE+16], in the sense that no "strength" is considered for the relation of any two extensions. To solve this problem, we introduce comparison score functions for these criteria, and base their definitions on the definitions of the corresponding comparison relations by considering the values used in the definitions.

Definition 4.2 (Set of attacks)**.** For an AF $F = (\mathcal{A}, \hookrightarrow)$ and sets of arguments $\mathcal{S}, \mathcal{S}' \subseteq \mathcal{A}$, we denote by

$$\hookrightarrow(\mathcal{S}, \mathcal{S}') = \hookrightarrow \cap\, (\mathcal{S} \times \mathcal{S}')$$

the *set of attacks from $\mathcal{S}$ to $\mathcal{S}'$.*

Definition 4.3 (Comparison score functions for nonatt, strdef, delarg, and delatt)**.** For an AF $F = (\mathcal{A}, \hookrightarrow)$, semantics σ, and extensions $\mathcal{E}, \mathcal{E}' \in \sigma(F)$, the *comparison score functions for criteria* nonatt, strdef, delarg, *and* delatt are defined as follows:

- $\mathrm{score}^F_{\sigma,\mathrm{nonatt}}(\mathcal{E}, \mathcal{E}') = \left|\, \{a \in \mathcal{E} \mid a \text{ is not attacked by an argument in } \mathcal{E}'\} \,\right|$
  (number of arguments in $\mathcal{E}$ not attacked by $\mathcal{E}'$)
- $\mathrm{score}^F_{\sigma,\mathrm{strdef}}(\mathcal{E}, \mathcal{E}') = \left|\, \{a \in \mathcal{E} \mid a \text{ is strongly defended from } \mathcal{E}' \text{ by } \mathcal{E}\} \,\right|$
  (number of arguments in $\mathcal{E}$ strongly defended from $\mathcal{E}'$ by $\mathcal{E}$)
- $\mathrm{score}^F_{\sigma,\mathrm{delarg}}(\mathcal{E}, \mathcal{E}') = \max\{|\, \mathcal{S} \subseteq \mathcal{E} \,| \mid \mathcal{E} \in \sigma(\mathcal{E} \cup \mathcal{E}', \hookrightarrow_{\restriction \mathcal{E} \cup \mathcal{E}'} \setminus \hookrightarrow(\mathcal{S}, \mathcal{E}'))\}$
  (maximal number of arguments in $\mathcal{E}$ from which attacks to $\mathcal{E}'$ can be deleted such that $\mathcal{E}$ is an extension after isolating $\mathcal{E}$ and $\mathcal{E}'$)
- $\mathrm{score}^F_{\sigma,\mathrm{delatt}}(\mathcal{E}, \mathcal{E}') = \max\{|\, \hookrightarrow' \subseteq \hookrightarrow(\mathcal{E}, \mathcal{E}') \,| \mid \mathcal{E} \in \sigma(\mathcal{E} \cup \mathcal{E}', \hookrightarrow_{\restriction \mathcal{E} \cup \mathcal{E}'} \setminus \hookrightarrow')\}$
  (maximal number of attacks from $\mathcal{E}$ to $\mathcal{E}'$ that can be deleted such that $\mathcal{E}$ is an extension after isolating $\mathcal{E}$ and $\mathcal{E}'$)

We can see that these functions fulfill the equivalence requirement of a comparison score function, as they represent the values used in the informal descriptions for their respective comparison relations.

4.2 Selection based on the Simpson Rule

Another rule from Voting Theory which can be used for extension selection via pairwise comparison is the *Simpson rule* [BCE+16].

Definition 4.4 (Simpson score function)**.** For an AF F, semantics σ, comparison criterion γ, and extension $\mathcal{E} \in \sigma(F)$, the *Simpson score function* (*with respect to semantics σ and comparison criterion γ*), $\mathrm{Simpson}^F_{\sigma,\gamma}$, is defined by

$$\mathrm{Simpson}^F_{\sigma,\gamma}(\mathcal{E}) = \min_{\mathcal{E}' \in \sigma(F)} \mathrm{score}^F_{\sigma,\gamma}(\mathcal{E}, \mathcal{E}') - \mathrm{score}^F_{\sigma,\gamma}(\mathcal{E}', \mathcal{E}).$$

Definition 4.5 (Simpson-based extensions semantics)**.** For a Simpson score function $\mathrm{Simpson}^F_{\sigma,\gamma}$, the *Simpson-based extensions (SBE) semantics* (*with respect to semantics σ and comparison criterion γ*), $\mathrm{SBE}_{\sigma,\gamma}$, is a selection semantics and defined by

$$\mathrm{SBE}_{\sigma,\gamma}(F) = \operatorname*{arg\,max}_{\mathcal{E} \in \sigma(F)} \mathrm{Simpson}^F_{\sigma,\gamma}(\mathcal{E}).$$

Informally, this function selects the extensions that *compare best to each extension* or are *the least objectionable*. It considers for every extension the extension it compares the worst against with respect to criterion γ, such that the difference of scores is minimal, and selects the extensions where this minimum is maximal. Unlike the Copeland rule, which considers only whether for every two extensions, one extension is stronger or weaker than the other with respect to γ, and selects the extensions with the maximal difference of "wins" and "losses", the Simpson rule accounts for the degree to which an extension is stronger or weaker. While the focus of the Copeland rule can be seen in the overall performance of an extension, the Simpson rule accounts for the one-on-one performance. In a sense, this property makes the Simpson rule more sceptical, such that universally well comparing extensions could be preferred over extensions that compare worse against at least one extension.

In some contexts, the Simpson rule might be the more adequate choice, e. g. in a debate setting where an agent has to adopt a standpoint, represented as an extension, without knowing the standpoints of the other parties, which could reflect only a small portion of extensions. While an extension selected by the Copeland rule is stronger than many extensions and weaker than few, it is not guaranteed that any weaker standpoint is adopted by an opposing party. Then, the case of an opponent adopting a standpoint expressed by an extension that is much stronger than the extension for our agent's standpoint is undesirable but possible.

A difference between both rules is that an SBE semantics selects no extension iff for no two extensions, the difference of scores is defined, whereas a CBE semantics maintains non-emptiness.

Lemma 4.1 (Emptiness of $\mathrm{SBE}_{\sigma,\gamma}(F)$)**.** *For all AFs F and Simpson-based extensions semantics $\mathrm{SBE}_{\sigma,\gamma}$, $\mathrm{SBE}_{\sigma,\gamma}(F)$ is empty iff comparison relation $\geq^F_{\sigma,\gamma}$ is empty.*

Proof. Let F be an AF and $\mathrm{SBE}_{\sigma,\gamma}(F)$ a selection semantics for comparison relation $\geq^F_{\sigma,\gamma}$. The equivalence will be shown in two steps, for either direction:

1. If $\mathrm{SBE}_{\sigma,\gamma}(F)$ is empty, only the empty binary relation exists on this set, such that $\geq^F_{\sigma,\gamma}$ is empty.
2. If $\geq^F_{\sigma,\gamma}$ is empty, it follows from Definition 4.3 of a comparison score function that comparison score function $\mathrm{score}^F_{\sigma,\gamma}$ is undefined on $\{(\mathcal{E},\mathcal{E}'),(\mathcal{E}',\mathcal{E})\}$ for any two extensions $\mathcal{E},\mathcal{E}' \in \sigma(F)$, such that the corresponding Simpson score function is undefined on $\sigma(F)$. Therefore, there exists no maximal Simpson score, such that $\mathrm{SBE}_{\sigma,\gamma}(F)$ is empty.

Therefore, it follows that $\mathrm{SBE}_{\sigma,\gamma}(F)$ is empty iff $\geq^F_{\sigma,\gamma}$ is empty. □

Corollary 4.1 (Acceptance-consistency of $\mathrm{SBE}_{\sigma,\gamma}$)**.** *There exists a Simpson-based extensions semantics that is not acceptance-consistent.*

Proof. Let F be the AF and σ the semantics from the example in the proof for Proposition 3.4, and $\mathrm{SBE}_{\sigma,\mathrm{delarg}}$ the corresponding Simpson-based extensions semantics.

Then, comparison relation $\geq^F_{\sigma,\mathrm{delarg}}$ is empty, such that it follows from Lemma 4.1 that $\mathrm{SBE}_{\sigma,\mathrm{delarg}}(F)$ is empty. As we have $\mathrm{Sc}_\sigma(F) = \{a\} \nsubseteq \mathrm{Sc}_{\mathrm{SBE}_{\sigma,\mathrm{delarg}}}(F) = \varnothing$, $\mathrm{SBE}_{\sigma,\mathrm{delarg}}$ is not acceptance-consistent. $\square$

Example 4.1 (Simpson-based extensions semantics)**.** *Consider the following example for* $\mathrm{SBE}_{\mathrm{Pr,nonatt}}$ *with the AF F shown by Figure 2 from Example 3.1. We have the three preferred extensions $\mathcal{E}_1 = \{A, E\}$, $\mathcal{E}_2 = \{B, E\}$, and $\mathcal{E}_3 = \{F\}$. Each extension is assigned a Simpson score which is iteratively calculated as follows, initially being* 0*. For other criteria, the Simpson score would initially be undefined and not guaranteed to be defined at the end of this process, however due to the reflexivity of comparison relation $\geq^F_{\mathrm{Pr,nonatt}}$ as shown by Proposition 3.3, it is possible to start with 0 here. We can determine the selected extensions by counting for every two extensions $\mathcal{E}, \mathcal{E}'$ the number of arguments in $\mathcal{E}$ left unattacked by arguments in $\mathcal{E}'$ and vice versa. If the Simpson score of $\mathcal{E}$ is undefined or greater than the difference, we update it with the difference. Because we start with 0, cases of equality (especially if $\mathcal{E} = \mathcal{E}'$) can be ignored as these do not contribute to the Simpson score.*

In $\mathcal{E}_1$, argument 5 is not attacked by an argument in $\mathcal{E}_2$ and vice versa, such that we ignore this case. Argument 1 in $\mathcal{E}_1$ is not attacked by the argument in $\mathcal{E}_3$, however the argument in $\mathcal{E}_3$ is attacked by argument 5 in $\mathcal{E}_1$, such that we have the difference

$$\mathrm{score}^F_{\mathrm{Pr,nonatt}}(\mathcal{E}_1, \mathcal{E}_3) - \mathrm{score}^F_{\mathrm{Pr,nonatt}}(\mathcal{E}_3, \mathcal{E}_1) = 1 - 0 = 1$$

and do not update the Simpson score of $\mathcal{E}_1$. This results in a Simpson score of 0 for $\mathcal{E}_1$.

In $\mathcal{E}_2$, we again ignore the equality with $\mathcal{E}_1$. Both arguments in $\mathcal{E}_2$ are attacked by the argument in $\mathcal{E}_3$ and the argument in $\mathcal{E}_3$ is attacked by argument 5 in $\mathcal{E}_2$, such that we ignore this case. This results in a Simpson score of 0 for $\mathcal{E}_2$.

The argument in $\mathcal{E}_3$ is attacked by argument 5 in $\mathcal{E}_1$, however argument 1 in $\mathcal{E}_1$ is not attacked by the argument in $\mathcal{E}_3$, such that we have the difference

$$\mathrm{score}^F_{\mathrm{Pr,nonatt}}(\mathcal{E}_3, \mathcal{E}_1) - \mathrm{score}^F_{\mathrm{Pr,nonatt}}(\mathcal{E}_1, \mathcal{E}_3) = 0 - 1 = -1$$

and update the Simpson score of $\mathcal{E}_3$ with -1. The argument in $\mathcal{E}_3$ is attacked by both arguments in $\mathcal{E}_2$ and both arguments in $\mathcal{E}_2$ are attacked by the argument in $\mathcal{E}_3$, such that we ignore this case. This results in a Simpson score of -1 for $\mathcal{E}_3$.

The extensions with the maximal Simpson score of 0 are now selected. Therefore, $\mathcal{E}_1$ and $\mathcal{E}_2$ are selected, in contrast to the analogue Example 3.1 for the Copeland-based extensions, where only $\mathcal{E}_1$ is selected.

5 Acceptance-consistency

As we have seen in the previous sections, the notions of totality, reflexivity, maintenace of non-emptiness, and acceptance-consistency are interconnected. Acceptance-consistency can be considered a desirable property of a selection semantics and is satisfied by all Copeland-based extensions semantics. Totality or reflexivity of all comparison relations for a semantics and comparison criterion implies the maintenance of non-emptiness and thus acceptance-consistency of the corresponding

Simpson-based extensions semantics, making this a useful criterion to determine the acceptance-consistency of the remaining selection semantics, especially for the semantics proposed by Dung [Dun95] and the comparison criteria proposed by Konieczny et al. [KMV15].

Lemma 5.1 (Acceptance-consistency of $\mathrm{SBE}_{\sigma,\gamma}$). *For all semantics σ and comparison criteria γ such that for all AFs F, comparison relation $\geq^F_{\sigma,\gamma}$ is total or reflexive, Simpson-based extensions semantics $\mathrm{SBE}_{\sigma,\gamma}$ maintains non-emptiness and is acceptance-consistent.*

Proof. Let σ be a semantics and γ a comparison criterion such that for all AFs F, comparison relation $\geq^F_{\sigma,\gamma}$ is total or reflexive. Let $\mathrm{SBE}_{\sigma,\gamma}$ be the respective Simpson-based extensions semantics. If $\sigma(F)$ is non-empty, it follows from Lemma 4.1 that $\mathrm{SBE}_{\sigma,\gamma}(F)$ is non-empty, as $\geq^F_{\sigma,\gamma}$ is total or reflexive and thus non-empty. Therefore, $\mathrm{SBE}_{\sigma,\gamma}$ maintains non-emptiness, and is acceptance-consistent as follows from Lemma 3.1. □

For comparison criteria nonatt and strdef, Lemma 5.1 delivers an immediate result for the acceptance-consistency of any corresponding SBE semantics, whereas for comparison criteria delarg and delarg, a distinction by the underlying semantics is required.

Corollary 5.1 (Acceptance-consistency of $\mathrm{SBE}_{\sigma,\mathrm{nonatt}}$ and $\mathrm{SBE}_{\sigma,\mathrm{strdef}}$). *For comparison criteria nonatt and strdef, all corresponding Simpson-based extensions semantics are acceptance-consistent.*

Proof. Follows from Proposition 3.3 in connection with Lemma 5.1. □

5.1 Complete semantics

Proposition 5.1 (Totality of $\geq^F_{\mathrm{Co,delarg}}$ and $\geq^F_{\mathrm{Co,delatt}}$). *There exist comparison relations $\geq^F_{\mathrm{Co,delarg}}$ and $\geq^F_{\mathrm{Co,delatt}}$ that are not total.*

Proof. Let $F = (\mathcal{A}, \hookrightarrow)$ be the AF shown by Figure 3. We have complete extensions $\mathcal{E} = \{A, D\}$ and $\mathcal{E}' = \{B\}$. The reduced AF according to comparison criteria delarg and delatt, $F_r = (\mathcal{E} \cup \mathcal{E}', \hookrightarrow_{\restriction \mathcal{E} \cup \mathcal{E}'}) = (\{A, B, D\}, \{(A, B), (B, A)\})$ has a non-empty grounded extension $\{D\} \nsubseteq \mathcal{E}'$, such that $\mathcal{E}'$ is not complete in F_r and no sets of attacks can be deleted such that it would be complete. It follows for comparison relations $\geq^F_{\mathrm{Co,delarg}}$ and $\geq^F_{\mathrm{Co,delatt}}$ that $\mathcal{E} \ngeq^F_{\mathrm{Co,delarg}} \mathcal{E}', \mathcal{E}' \ngeq^F_{\mathrm{Co,delarg}} \mathcal{E}$ and $\mathcal{E} \ngeq^F_{\mathrm{Co,delatt}} \mathcal{E}', \mathcal{E}' \ngeq^F_{\mathrm{Co,delatt}} \mathcal{E}$. □

Proposition 5.2 (Reflexivity of $\geq^F_{\mathrm{Co,delarg}}$ and $\geq^F_{\mathrm{Co,delatt}}$). *All comparison relations $\geq^F_{\mathrm{Co,delarg}}$ and $\geq^F_{\mathrm{Co,delatt}}$ are reflexive.*

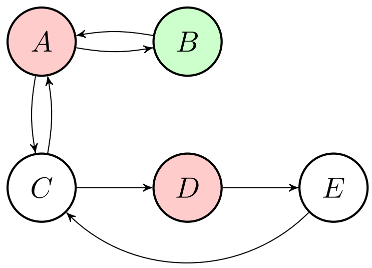

Figure 3: An AF with two complete extensions

Complete	Reflexive	Total
nonatt	✓	✓
strdef	✓	✓
delarg	✓	✗
delatt	✓	✗

Table 1: Properties of comparison relations for the complete semantics

Proof. Let $F = (\mathcal{A}, \hookrightarrow)$ be an AF. For any extension $\mathcal{E} \in \mathrm{Co}(F)$, the reduced AF according to comparison criteria delarg and delatt, $F_r = (\mathcal{E} \cup \mathcal{E}, \hookrightarrow_{\restriction \mathcal{E} \cup \mathcal{E}'}) = (\mathcal{E}, \varnothing)$, has no attacks as $\mathcal{E}$ is conflict-free. Then, $\mathcal{E}$ is an extension of F_r, such that it follows for comparison relations $\geq^F_{\mathrm{Co,delarg}}$ and $\geq^F_{\mathrm{Co,delatt}}$, that $\mathcal{E} \geq^F_{\mathrm{Co,delarg}} \mathcal{E}$ and $\mathcal{E} \geq^F_{\mathrm{Co,delatt}} \mathcal{E}$. □

Corollary 5.2 (Acceptance-consistency of $\mathrm{SBE}_{\mathrm{Co,delarg}}$ and $\mathrm{SBE}_{\mathrm{Co,delatt}}$)**.** *All Simpson-based extensions semantics* $\mathrm{SBE}_{\mathrm{Co,delarg}}$ *and* $\mathrm{SBE}_{\mathrm{Co,delatt}}$ *are acceptance-consistent.*

Proof. Follows from Proposition 5.2 in connecton with Lemma 5.1. □

5.2 Preferred semantics

Proposition 5.3 (Totality of $\geq^F_{\mathrm{Pr,delarg}}$)**.** *There exists a comparison relation* $\geq^F_{\mathrm{Pr,delarg}}$ *that is not total.*

Proof. Analogous to the proof for Proposition 5.1, as the extensions are preferred and the preferred extensions are complete. □

Proposition 5.4 (Totality of $\geq^F_{\mathrm{Pr,delatt}}$)**.** *There exists a comparison relation* $\geq^F_{\mathrm{Pr,delatt}}$ *that is not total.*

Proof. Let $F = (\mathcal{A}, \hookrightarrow)$ be the AF shown by Figure 4. We have preferred extensions $\mathcal{E} = \{A\}$ and $\mathcal{E}' = \{B, D\}$. The reduced AF according to comparison criterion delatt, $F_r = (\mathcal{E} \cup \mathcal{E}', \hookrightarrow_{\restriction \mathcal{E} \cup \mathcal{E}'}) = (\{A, B, D\}, \varnothing)$, has only the extension $\mathcal{E} \cup \mathcal{E}'$. Then,

Preferred	Reflexive	Total
nonatt	✓	✓
strdef	✓	✓
delarg	✓	✗
delatt	✓	✗

Table 2: Properties of comparison relations for the preferred semantics

there exists no subset of attacks in F_r from $\mathcal{E}$ that can be deleted such that $\mathcal{E}'$ is an extension. It follows for comparison relation $\geq^F_{\mathrm{Pr,delatt}}$, that $\mathcal{E} \not\geq^F_{\mathrm{Pr,delatt}} \mathcal{E}'$ and $\mathcal{E}' \not\geq^F_{\mathrm{Pr,delatt}} \mathcal{E}$. Therefore, $\geq^F_{\mathrm{Pr,delatt}}$ is not total. □

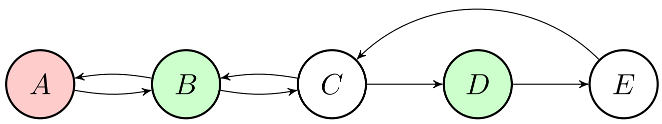

Figure 4: An AF with two preferred extensions

Proposition 5.5 (Reflexivity of $\geq^F_{\mathrm{Pr,delarg}}$ and $\geq^F_{\mathrm{Pr,delatt}}$). *All comparison relations* $\geq^F_{\mathrm{Pr,delarg}}$ *and* $\geq^F_{\mathrm{Pr,delatt}}$ *are reflexive.*

Proof. Analogous to the proof for Proposition 5.2, as the extension of the reduced AF is preferred, and the preferred extensions are conflict-free as well as complete. □

Corollary 5.3 (Acceptance-consistency of $\mathrm{SBE}_{\mathrm{Pr,delarg}}$ and $\mathrm{SBE}_{\mathrm{Pr,delatt}}$). *All Simpson-based extensions semantics* $\mathrm{SBE}_{\mathrm{Pr,delarg}}$ *and* $\mathrm{SBE}_{\mathrm{Pr,delatt}}$ *are acceptance-consistent.*

Proof. Follows from Proposition 5.5 in connection with Lemma 5.1. □

5.3 Stable semantics

Proposition 5.6 (Totality of $\geq^F_{\mathrm{St,delarg}}$ and $\geq^F_{\mathrm{St,delatt}}$). *All comparison relations* $\geq^F_{\mathrm{St,delarg}}$ *and* $\geq^F_{\mathrm{St,delatt}}$ *are total.*

Proof. Let $F = (\mathcal{A}, \hookrightarrow)$ be an AF and $\geq^F_{\mathrm{St,delarg}}, \geq^F_{\mathrm{St,delatt}}$ comparison relations. Any stable extensions $\mathcal{E}, \mathcal{E}'$ of F are extensions of the reduced AF according to comparison criteria delarg and delatt, $F_r = (\mathcal{E} \cup \mathcal{E}', \hookrightarrow_{\restriction \mathcal{E} \cup \mathcal{E}'})$, as all attacks from $\mathcal{E}$ to $\mathcal{E}'$ and vice versa are retained, and no attacks are added. It follows that $\mathcal{E} \geq^F_{\mathrm{St,delarg}} \mathcal{E}'$ or $\mathcal{E}' \geq^F_{\mathrm{St,delarg}} \mathcal{E}$, as for the comparison of either extension to the other, at least all attacks from the empty subset of arguments can be deleted such that either is an extension, and $\mathcal{E} \geq^F_{\mathrm{St,delatt}} \mathcal{E}'$ or $\mathcal{E}' \geq^F_{\mathrm{St,delatt}} \mathcal{E}$, as at least the empty set of attacks can

Stable	Reflexive	Total
nonatt	✓	✓
strdef	✓	✓
delarg	✓	✗
delatt	✓	✓

Table 3: Properties of comparison relations for the stable semantics

be deleted such that the either is an extension. Therefore, $\geq^F_{\mathrm{St,delarg}}$ and $\geq^F_{\mathrm{St,delatt}}$ are total. □

Corollary 5.4 (Reflexivity of $\geq^F_{\mathrm{St,delarg}}$ and $\geq^F_{\mathrm{St,delatt}}$)**.** *All comparison relations* $\geq^F_{\mathrm{St,delarg}}$ *and* $\geq^F_{\mathrm{St,delatt}}$ *are reflexive.*

Proof. Follows from Proposition 5.6. □

Corollary 5.5 (Acceptance-consistency of $\mathrm{SBE}_{\mathrm{St,delarg}}$ and $\mathrm{SBE}_{\mathrm{St,delatt}}$)**.** *All Simpson-based extensions semantics* $\mathrm{SBE}_{\mathrm{St,delarg}}$ *and* $\mathrm{SBE}_{\mathrm{St,delatt}}$ *are acceptance-consistent.*

Proof. Follows from Corollary 5.4 in connection with Lemma 5.1. □

5.4 Grounded semantics

Proposition 5.7 (Reflexivity of $\geq^F_{\mathrm{Gr,delarg}}$ and $\geq^F_{\mathrm{Gr,delatt}}$)**.** *All comparison relations* $\geq^F_{\mathrm{Gr,delarg}}$ *and* $\geq^F_{\mathrm{Gr,delatt}}$ *are reflexive.*

Proof. Analogous to the proof for Proposition 5.2, as the extension of the reduced AF is grounded, and the grounded extension is conflict-free as well as complete. □

Proposition 5.8 (Totality of $\geq^F_{\mathrm{Gr,delarg}}$ and $\geq^F_{\mathrm{Gr,delatt}}$)**.** *All comparison relations* $\geq^F_{\mathrm{Gr,delarg}}$ *and* $\geq^F_{\mathrm{Gr,delatt}}$ *are total.*

Proof. Follows from Proposition 5.7 as the grounded extension is unique. □

Corollary 5.6 (Acceptance-consistency of $\mathrm{SBE}_{\mathrm{Gr,delarg}}$ and $\mathrm{SBE}_{\mathrm{Gr,delatt}}$)**.** *All Simpson-based extensions semantics* $\mathrm{SBE}_{\mathrm{Gr,delarg}}$ *and* $\mathrm{SBE}_{\mathrm{Gr,delatt}}$ *are acceptance-consistent.*

Proof. Follows from Proposition 5.7 in connection with Lemma 5.1. □

Grounded	Reflexive	Total
nonatt	✓	✓
strdef	✓	✓
delarg	✓	✓
delatt	✓	✓

Table 4: Properties of comparison relations for the grounded semantics

5.5 Summary

In this section we have learned that the Copeland- and Simpson-based extensions semantics induced by the classical semantics by Dung [Dun95] and comparison criteria proposed by Konieczny et al. [KMV15] are acceptance-consistent. While acceptance-consistency is desirable, it does not guarantee that a selection semantics can be considered effective at decreasing the number of credulously accepted arguments and increasing the number of sceptically accepted arguments, it is rather a necessary condition of effectiveness. As an example, consider from the following Proposition 6.36 the selection semantics $\mathrm{CBE}_{\mathrm{St},\mathrm{nonatt}}$ and $\mathrm{SBE}_{\mathrm{St},\mathrm{nonatt}}$. Any of two stable extensions leaves only their intersection unattacked in the other, such that all extensions compare equally and the selection semantics coincide with the stable semantics. Where ineffectiveness is not trivial like in this case, determining how effective a selection semantics is would require a complex analysis, possibly driven empirically by real-world applications, that is beyond the scope of this thesis. Until further results are available, it is only possible to assume that extension selection will be reasonably effective, especially with the comparison criteria proposed by Konieczny et al. The notion of acceptance-consistency can guide the search for potentially effective selection semantics.

6 Inheritance of principles

A desirable behaviour of a selection semantics is the *inheritance of principles satisfied by the underlying semantics*, such that the choice to use a semantics in a certain context based on principles, as suggested by van der Torre and Vesic [vDTV18], is not affected by the application of the selection semantics.

Definition 6.1 (Inheritance of a principle)**.** A selection semantics $\mathrm{SEL}_{\sigma,\gamma}$ *inherits a principle* iff it is not satisfied by σ or it is satisfied by $\mathrm{SEL}_{\sigma,\gamma}$.

For some basic principles as introduced by Baroni and Giacomoin [BG07], this inheritance is independent from the employed selection rule and criterion with the exception of directionality. These and more principles which can be examined for inheritance can be found in an overview of such a *principle-based* system by van der Torre and Vesic [vDTV18], where the principles relevant for this thesis are defined.

Definition 6.2 (Principles)**.** For a semantics σ, *principles* are defined as follows:

- Semantics σ satisfies *admissibility* iff for all AFs F, all $\mathcal{E} \in \sigma(F)$ are admissible.
- Semantics σ satisfies *strong admissibility* iff for all AFs $F = (\mathcal{A}, \hookrightarrow)$, $\mathcal{E} \in \sigma(F)$, and $a \in \mathcal{E}$, a is strongly defended from $\mathcal{A}$ by $\mathcal{E}$.
- Semantics σ satisfies *naivety* iff for all AFs F and $\mathcal{E} \in \sigma(F)$, $\mathcal{E}$ is maximal with respect to set inclusion among the conflict-free sets of arguments.
- Semantics σ satisfies *indirect conflict-freeness* iff for all AFs F, $\mathcal{E} \in \sigma(F)$, and $a, b \in \mathcal{E}$, there exists no path of odd length between a and b in F with respect to the attack relation.
- Semantics σ satisfies *reinstatement* iff for all AFs $F = (\mathcal{A}, \hookrightarrow)$, $\mathcal{E} \in \sigma(F)$, and $a \in \mathcal{A}$ such that a is acceptable with respect to $\mathcal{E}$, it holds that $a \in \mathcal{E}$.
- Semantics σ satisfies *weak reinstatement* iff for all AFs $F = (\mathcal{A}, \hookrightarrow)$, $\mathcal{E} \in \sigma(F)$, and $a \in \mathcal{A}$ such that a is strongly defended from $\mathcal{A}$ by $\mathcal{E}$, it holds that $a \in \mathcal{E}$.
- Semantics σ satisfies *CF-reinstatement* iff for all AFs $F = (\mathcal{A}, \hookrightarrow)$, $\mathcal{E} \in \sigma(F)$, and $a \in \mathcal{A}$ such that a is acceptable with respect to $\mathcal{E}$ and $\mathcal{E} \cup \{a\}$ is conflict-free, it holds that $a \in \mathcal{E}$.
- Semantics σ satisfies *I-maximality* iff for all AFs F, $\mathcal{E} \in \sigma(F)$, $\mathcal{E}$ is maximal with respect to set inclusion among $\sigma(F)$.
- Semantics σ satisfies *allowing abstention* iff for all AFs F, $\mathcal{E}_1, \mathcal{E}_2 \in \sigma(F)$, and $a \in \mathcal{E}_2$ such that a is attacked by an argument in $\mathcal{E}_1$, there exists $\mathcal{E}_3 \in \sigma(F)$ such that $a \notin \mathcal{E}_3$ and a is not attacked by an argument in $\mathcal{E}_3$.
- Semantics σ satisfies *crash resistance* iff there exists no AF $F' = (\mathcal{A}', \hookrightarrow')$ such that for all AFs $F = (\mathcal{A}, \hookrightarrow)$ where $\mathcal{A} \cap \mathcal{A}' = \emptyset$, it holds that $\sigma(\mathcal{A} \cup \mathcal{A}', \hookrightarrow \cup \hookrightarrow') = \sigma(F')$.
- Semantics σ satisfies *non-interference* iff for all AFs $F = (\mathcal{A}, \hookrightarrow)$ and $\mathcal{S} \subseteq A$ such that there exists no attack from $\mathcal{A} \setminus \mathcal{S}$ to $\mathcal{S}$ and vice versa, it holds that $\sigma(\mathcal{S}, \hookrightarrow_{\restriction \mathcal{S}}) = \{\mathcal{E} \cap \mathcal{S} \mid \mathcal{E} \in \sigma(F)\}$.
- Semantics σ satisfies *weak directionality* iff for all AFs $F = (\mathcal{A}, \hookrightarrow)$ and $\mathcal{S} \subseteq \mathcal{A}$ such that there exists no attack from $\mathcal{A} \setminus \mathcal{S}$ to $\mathcal{S}$, it holds that $\sigma(\mathcal{S}, \hookrightarrow_{\restriction \mathcal{S}}) \supseteq \{\mathcal{E} \cap \mathcal{S} \mid \mathcal{E} \in \sigma(F)\}$.
- Semantics σ satisfies *semi-directionality* iff for all AFs $F = (\mathcal{A}, \hookrightarrow)$ and $\mathcal{S} \subseteq \mathcal{A}$ such that there exists no attack from $\mathcal{A} \setminus \mathcal{S}$ to $\mathcal{S}$, it holds that $\sigma(\mathcal{S}, \hookrightarrow_{\restriction \mathcal{S}}) \subseteq \{\mathcal{E} \cap \mathcal{S} \mid \mathcal{E} \in \sigma(F)\}$.

- Semantics σ satisfies *directionality* iff for all AFs $F = (\mathcal{A}, \hookrightarrow)$ and $\mathcal{S} \subseteq \mathcal{A}$ such that there exists no attack from $\mathcal{A} \backslash \mathcal{S}$ to $\mathcal{S}$, it holds that $\sigma(\mathcal{S}, \hookrightarrow_{\restriction \mathcal{S}}) = \{\mathcal{E} \cap \mathcal{S} \mid \mathcal{E} \in \sigma(F)\}$. Directionality combines weak directionality and semi-directionality.

- Semantics σ satisfies *succinctness* iff there exist no AF $F = (\mathcal{A}, \hookrightarrow)$ and $(a, b) \in \hookrightarrow$ such that for all AFs $F' = AF[']$ where $\mathcal{A} \subseteq \mathcal{A}'$ and $\hookrightarrow \subseteq \hookrightarrow'$, it holds that $\sigma(F') = \sigma(\mathcal{A}', \hookrightarrow' \setminus \{(a, b)\})$.

- Semantics σ satisfies *tightness* iff for all AFs $F = (\mathcal{A}, \hookrightarrow)$, credulously σ-accepted $a \in \mathcal{A}$, and $\mathcal{E} \in \sigma(F)$ where $\mathcal{E} \cup \{a\} \notin \sigma(F)$, there exists $b \in \mathcal{E}$ such that there exists no $\mathcal{E}' \in \sigma(F)$ where $a, b \in \mathcal{E}'$.

- Semantics σ satisfies *conflict-sensitiveness* iff for all AFs $F = (\mathcal{A}, \hookrightarrow)$ and $\mathcal{E}, \mathcal{E}' \in \sigma(F)$ where $\mathcal{E} \cup \mathcal{E}' \notin \sigma(F)$, there exist $a, b \in \mathcal{E} \cup \mathcal{E}'$ such that there exists no $\mathcal{E}'' \in \sigma(F)$ where $a, b \in \mathcal{E}''$.

- Semantics σ satisfies *com-closure* iff for all AFs $F = (\mathcal{A}, \hookrightarrow)$ and non-empty $S \subseteq \sigma(F)$ such that for all $a, b \in \bigcup_{\mathcal{E} \in S} \mathcal{E}$, it holds that $a, b \in \mathcal{E}'$ for some $\mathcal{E}' \in \sigma(F)$, there exists one and only one $\mathcal{E}'' \in \sigma(F)$, such that $\bigcup_{\mathcal{E} \in S} \mathcal{E} \subseteq \mathcal{E}''$ and $\mathcal{E}''$ is minimal with respect to set inclusion among the extensions that satisfy the former.

Other principles pertain to the semantics' relation to the strong components of the graph for an AF (*SCC-recursiveness*) or the change of the set of sceptically accepted arguments if symmetric attacks are reduced to one-sided attacks (*scepticism adequacy* and *resolution adequacy*) [vDTV18].

We can see that besides admissibility, strong admissibility, reinstatement, weak reinstatement, CF-reinstatement, and I-maximality, as Konieczny et al. [KMV15] have pointed out, also naivety and indirect conflict-freeness are always inherited by a selection semantics.

Remark 6.1 (Trivial inheritance). Any selection semantics inherits admissibility, strong admissibility, reinstatement, weak reinstatement, CF-reinstatement, and I-maximality.

Proposition 6.1 (Trivial inheritance)**.** *Any selection semantics inherits naivety and indirect conflict-freeness.*

Proof. Let F be an AF and $\mathrm{SEL}_{\sigma,\gamma}$ a selection semantics. As $\mathrm{SEL}_{\sigma,\gamma}(F) \subseteq \sigma(F)$, if σ satisfies naivety, $\mathrm{SEL}_{\sigma,\gamma}(F)$ is a subset of conflict-free sets of arguments that are maximal with respect to set inclusion, such that $\mathrm{SEL}_{\sigma,\gamma}$ inherits naivety, and if σ satisfies indirect conflict-freeness, $\mathrm{SEL}_{\sigma,\gamma}(F)$ is a subset of sets of arguments without paths of odd length between any two arguments within each extension, such that $\mathrm{SEL}_{\sigma,\gamma}$ inherits indirect conflict-freeness. □

The following results show that for the other principles, inheritance is often not achieved by the introduced selection semantics if the relevant principle is satisfied by the underlying semantics. As the relevant comparison criteria are not directly related to the relevant principles, largely negative results were to be expected and are to be expected for other comparison criteria and principles if their definitions do not suggest inheritance. Whereas the definition of more conforming principles would contradict the goal of guiding the search for fit selection semantics with existing principles and provide no added value as the search would only be influenced for the sake of influencing the search, the only reasonable method to achieve more positive results is to define more conforming, and most importantly meaningful, comparison criteria.

6.1 Complete semantics

The complete semantics satisfies the following principles which are not trivially inherited: allowing abstention, crash resistance, non-interference, weak directionality, semi-directionality, directionality, and com-closure [vDTV18].

Proposition 6.2 (Inheritance of crash resistance by $\mathrm{SEL}_{\mathrm{Co},\gamma}$)**.** *Any selection semantics* $\mathrm{SEL}_{\mathrm{Co},\gamma}$ *that maintains non-emptiness inherits crash resistance.*

Proof. Let $F' = (\mathcal{A}', \hookrightarrow')$ be an AF and $\mathrm{SEL}_{\mathrm{Co},\gamma}$ a selection semantics that maintains non-emptiness. Given a new argument $a \notin \mathcal{A}'$ and AF $F = (\{a\}, \varnothing)$, for AF $F'' = (\mathcal{A}' \cup \{a\}, \hookrightarrow' \cup \varnothing) = (\mathcal{A}' \cup \{a\}, \hookrightarrow')$, there exists a selected extension $\mathcal{E} \in \mathrm{SEL}_{\mathrm{Co},\gamma}(F'')$ as the set of complete extensions is non-empty and $\mathrm{SEL}_{\mathrm{Co},\gamma}$ maintains non-emptiness. As a is unattacked, it is contained in the grounded extension $\mathrm{Gr}(F'')$ which is included in complete extension $\mathcal{E}$. Then we have $a \in \mathcal{E}$, but $a \notin \mathcal{A}'$ such that $\mathcal{E} \notin \mathrm{SEL}_{\mathrm{Co},\gamma}(F')$ and thus $\mathrm{SEL}_{\mathrm{Co},\gamma}(F'') \neq \mathrm{SEL}_{\mathrm{Co},\gamma}(F')$. Therefore, $\mathrm{SEL}_{\mathrm{Co},\gamma}$ satisfies and thus inherits crash resistance. □

Corollary 6.1 (Inheritance of crash resistance by $\mathrm{CBE}_{\mathrm{Co},\gamma}$ and $\mathrm{SBE}_{\mathrm{Co},\gamma}$)**.** *Any Copeland-based extensions semantics* $\mathrm{CBE}_{\mathrm{Co},\gamma}$ *and Simpson-based extensions semantics* $\mathrm{SBE}_{\mathrm{Co},\gamma}$ *where* $\gamma \in \{\mathrm{nonatt}, \mathrm{strdef}, \mathrm{delarg}, \mathrm{delarg}\}$ *inherit crash resistance.*

Proof. Follows from Propositions 3.1, 3.3, and 5.7 in connection with Lemma 5.1, in connection with Proposition 6.2. □

6.1.1 Copeland-based extensions

Proposition 6.3 (Inheritance of allowing abstention by $\mathrm{CBE}_{\mathrm{Co,nonatt}}$ and $\mathrm{CBE}_{\mathrm{Co,strdef}}$)**.** *Copeland-based extensions semantics* $\mathrm{CBE}_{\mathrm{Co,nonatt}}$ *and* $\mathrm{CBE}_{\mathrm{Co,strdef}}$ *do not inherit allowing abstention.*

Proof. Let F be the AF shown by Figure 5. For Copeland-based extensions semantics $\mathrm{CBE}_{\mathrm{Co,nonatt}}$, we have two extensions: $\mathcal{E}_1 = \{A\}$ and $\mathcal{E}_2 = \{B\}$. As argument A is attacked by $\mathcal{E}_2$ but there exists no extension which neither contains A nor attacks it, F does not satisfy the condition set for all AFs by principle allowing abstention. The

Complete	Adm.	Reinst.	W. reinst.	CF-reinst.	All. abst.	Crash resist.
CBE						
nonatt	✓	✓	✓	✓	✗	✓
strdef	✓	✓	✓	✓	✗	✓
delarg	✓	✓	✓	✓	✗	✓
delatt	✓	✓	✓	✓	✗	✓
SBE						
nonatt	✓	✓	✓	✓	✗	✓
strdef	✓	✓	✓	✓	✗	✓
delarg	✓	✓	✓	✓	?	✓
delatt	✓	✓	✓	✓	?	✓

Table 5: Inheritance of principles satisfied by the complete semantics for complete-based selection semantics, part 1

Complete	Non-int.	W. dir.	S.-dir.	Dir.	Com-closure
CBE					
nonatt	?	✗	✗	✗	?
strdef	?	✗	✗	✗	?
delarg	✗	✗	✗	✗	?
delatt	✗	✗	✗	✗	?
SBE					
nonatt	?	?	✗	✗	?
strdef	?	?	✗	✗	?
delarg	?	✗	✗	✗	✗
delatt	?	✗	✗	✗	✗

Table 6: Inheritance of principles satisfied by the complete semantics for complete-based selection semantics, part 2

same holds for $\mathrm{CBE}_{\mathrm{Co,strdef}}$ as follows from Proposition 3.2. Therefore, $\mathrm{CBE}_{\mathrm{Co,nonatt}}$ and $\mathrm{CBE}_{\mathrm{Co,strdef}}$ do not satisfy and thus not inherit allowing abstention. □

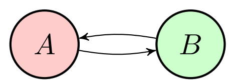

Figure 5: An AF with two $\mathrm{CBE}_{\mathrm{Co,nonatt}}$-extensions

Proposition 6.4 (Inheritance of allowing abstention by $\mathrm{CBE}_{\mathrm{Co,delarg}}$)**.** *Copeland-based extensions semantics* $\mathrm{CBE}_{\mathrm{Co,delarg}}$ *does not inherit allowing abstention.*

Proof. Let F be the AF shown by Figure 6. For Copeland-based extensions semantics $\mathrm{CBE}_{\mathrm{Co,delarg}}$, we have two extensions: $\mathcal{E}_1 = \{A, B\}$ and $\mathcal{E}_2 = \{B, C\}$. As argument C is attacked by $\mathcal{E}_1$ but there exists no extension which neither contains C nor attacks it, F does not satisfy the condition set for all AFs by principle allowing abstention. Therefore, $\mathrm{CBE}_{\mathrm{Co,delarg}}$ does not satisfy and thus not inherit allowing abstention. □

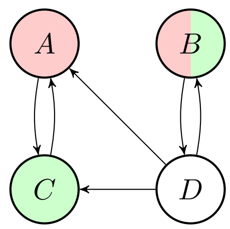

Figure 6: An AF with two $\mathrm{CBE}_{\mathrm{Co,delarg}}$-extensions

Proposition 6.5 (Inheritance of allowing abstention by $\mathrm{CBE}_{\mathrm{Co,delatt}}$)**.** *Copeland-based extensions semantics* $\mathrm{CBE}_{\mathrm{Co,delatt}}$ *does not inherit allowing abstention.*

Proof. Let F be the AF shown by Figure 7. For Copeland-based extensions semantics $\mathrm{CBE}_{\mathrm{Co,delatt}}$, we have two extensions: $\mathcal{E}_1 = \{A, C\}$ and $\mathcal{E}_2 = \{A, D\}$. As argument C is attacked by $\mathcal{E}_2$ but there exists no extension which neither contains C nor attacks it, F does not satisfy the condition set for all AFs by principle allowing abstention. Therefore, $\mathrm{CBE}_{\mathrm{Co,delatt}}$ does not satisfy and thus not inherit allowing abstention. □

Proposition 6.6 (Inheritance of non-interference by $\mathrm{CBE}_{\mathrm{Co,delarg}}$)**.** *Copeland-based extensions semantics* $\mathrm{CBE}_{\mathrm{Co,delarg}}$ *does not inherit non-interference.*

Proof. Let $F = (\mathcal{A}, \hookrightarrow)$ be the AF shown by Figure 8. For Copeland-based extensions semantics $\mathrm{CBE}_{\mathrm{Co,delarg}}$, we have two extensions: $\{B, C, D\}$ and $\{B, D, E\}$. For $\mathcal{S} = \{C, E\}$, there exists no attack from $\mathcal{A} \setminus \mathcal{S}$ to $\mathcal{S}$ and vice versa. As the restricted

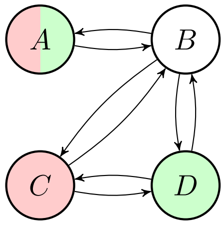

Figure 7: An AF with two $\mathrm{CBE}_{\mathrm{Co,delatt}}$-extensions

AF has extensions $\mathrm{CBE}_{\mathrm{Co,delarg}}(\mathcal{S}, \hookrightarrow_{\restriction \mathcal{S}}) = \{\varnothing, \{C\}, \{E\}\}$ not equal to the original extensions' intersections $\{\mathcal{E} \cap \mathcal{S} \mid \mathcal{E} \in \mathrm{CBE}_{\mathrm{Co,delarg}}(F)\} = \{\{C\}, \{E\}\}$, F does not satisfy the condition set for all AFs by non-interference. Therefore, $\mathrm{CBE}_{\mathrm{Co,delarg}}$ does not satisfy and thus not inherit non-interference. □

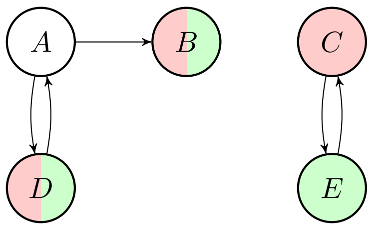

Figure 8: An AF with two $\mathrm{CBE}_{\mathrm{Co,delarg}}$-extensions

Proposition 6.7 (Inheritance of non-interference by $\mathrm{CBE}_{\mathrm{Co,delatt}}$)**.** *Copeland-based extensions semantics* $\mathrm{CBE}_{\mathrm{Co,delatt}}$ *does not inherit non-interference.*

Proof. Let $F = (\mathcal{A}, \hookrightarrow)$ be the AF shown by Figure 9. For Copeland-based extensions semantics $\mathrm{CBE}_{\mathrm{Co,delatt}}$, we have two extensions: $\{C, D, E, I\}$ and $\{D, E, F, I\}$. For $\mathcal{S} = \{C, F\}$, there exists no attack from $\mathcal{A} \setminus \mathcal{S}$ to $\mathcal{S}$ and vice versa. As the restricted AF has extensions $\mathrm{CBE}_{\mathrm{Co,delatt}}(\mathcal{S}, \hookrightarrow_{\restriction \mathcal{S}}) = \{\varnothing, \{C\}, \{F\}\}$ not equal to the original extensions' intersections $\{\mathcal{E} \cap \mathcal{S} \mid \mathcal{E} \in \mathrm{CBE}_{\mathrm{Co,delatt}}(F)\} = \{\{C\}, \{F\}\}$, F does not satisfy the condition set for all AFs by non-interference. Therefore, $\mathrm{CBE}_{\mathrm{Co,delatt}}$ does not satisfy and thus not inherit non-interference. □

Proposition 6.8 (Inheritance of weak directionality by $\mathrm{CBE}_{\mathrm{Co,nonatt}}$ and $\mathrm{CBE}_{\mathrm{Co,strdef}}$)**.** *Copeland-based extensions semantics* $\mathrm{CBE}_{\mathrm{Co,nonatt}}$ *and* $\mathrm{CBE}_{\mathrm{Co,strdef}}$ *do not inherit weak directionality.*

Proof. Let $F = (\mathcal{A}, \hookrightarrow)$ be the AF shown by Figure 10. For Copeland-based extensions semantics $\mathrm{CBE}_{\mathrm{Co,nonatt}}$, we have three extensions: $\{A, E\}$, $\{C, D\}$, and $\{B, C, E\}$. For $\mathcal{S} = \{A, B, C, D\}$, there exists no attack from $\mathcal{A} \setminus \mathcal{S}$ to $\mathcal{S}$. As the restricted AF has extensions $\mathrm{CBE}_{\mathrm{Co,nonatt}}(\mathcal{S}, \hookrightarrow_{\restriction \mathcal{S}}) = \{\{C, B\}, \{C, D\}\}$ not including

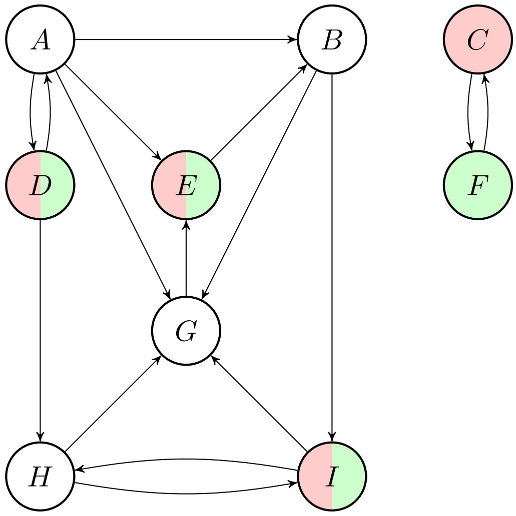

Figure 9: An AF with two $\mathrm{CBE}_{\mathrm{Co,delatt}}$-extensions

the original extensions' intersections $\{\mathcal{E} \cap \mathcal{S} \mid \mathcal{E} \in \mathrm{CBE}_{\mathrm{Co,nonatt}}(F)\} = \{\{A\}, \{C, B\}, \{C, D\}\}$, F does not satisfy the condition set for all AFs by weak directionality. The same holds for $\mathrm{CBE}_{\mathrm{Co,strdef}}$ as follows from Proposition 3.2. Therefore, $\mathrm{CBE}_{\mathrm{Co,nonatt}}$ and $\mathrm{CBE}_{\mathrm{Co,strdef}}$ do not satisfy and thus not inherit weak directionality. $\square$

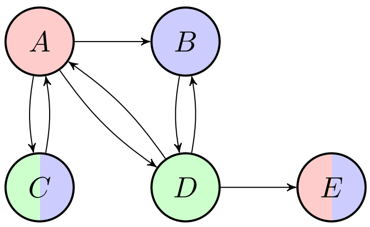

Figure 10: An AF with three $\mathrm{CBE}_{\mathrm{Co,nonatt}}$-extensions

Proposition 6.9 (Inheritance of weak directionality by $\mathrm{CBE}_{\mathrm{Co,delarg}}$)**.** *Copeland-based extensions semantics* $\mathrm{CBE}_{\mathrm{Co,delarg}}$ *does not inherit weak directionality.*

Proof. Let $F = (\mathcal{A}, \hookrightarrow)$ be the AF shown by Figure 11. For Copeland-based extensions semantics $\mathrm{CBE}_{\mathrm{Co,delarg}}$, we have three extensions: $\varnothing$, $\{A, B\}$, and $\{C, D\}$. For $\mathcal{S} = \{A, B, C\}$, there exists no attack from $\mathcal{A} \setminus \mathcal{S}$ to $\mathcal{S}$. As the restricted AF has extensions $\mathrm{CBE}_{\mathrm{Co,delarg}}(\mathcal{S}, \hookrightarrow_{\restriction \mathcal{S}}) = \{\{A, B\}\}$ not including the original extensions' intersections $\{\mathcal{E} \cap \mathcal{S} \mid \mathcal{E} \in \mathrm{CBE}_{\mathrm{Co,delarg}}(F)\} = \{\varnothing, \{C\}, \{A, B\}\}$, F does not satisfy the condition set for all AFs by weak directionality. Therefore, $\mathrm{CBE}_{\mathrm{Co,delarg}}$ does not satisfy and thus not inherit weak directionality. $\square$

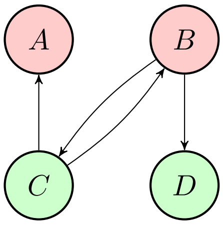

Figure 11: An AF with three $\mathrm{CBE}_{\mathrm{Co,delarg}}$-extensions

Proposition 6.10 (Inheritance of weak directionality by $\mathrm{CBE}_{\mathrm{Co,delatt}}$)**.** *Copeland-based extensions semantics* $\mathrm{CBE}_{\mathrm{Co,delatt}}$ *does not inherit weak directionality.*

Proof. Let $F = (\mathcal{A}, \hookrightarrow)$ be the AF shown by Figure 12. For Copeland-based extensions semantics $\mathrm{CBE}_{\mathrm{Co,delatt}}$, we have three extensions: $\varnothing$, $\{A, C\}$, and $\{B, D, E\}$. For $\mathcal{S} = \{A, D, E\}$, there exists no attack from $\mathcal{A} \setminus \mathcal{S}$ to $\mathcal{S}$. As the restricted AF has extensions $\mathrm{CBE}_{\mathrm{Co,delatt}}(\mathcal{S}, \hookrightarrow\restriction_{\mathcal{S}}) = \{\{D, E\}\}$ not including the original extensions' intersections $\{\mathcal{E} \cap \mathcal{S} \mid \mathcal{E} \in \mathrm{CBE}_{\mathrm{Co,delatt}}(F)\} = \{\varnothing, \{A\}, \{E, D\}\}$, F does not satisfy the condition set for all AFs by weak directionality. Therefore, $\mathrm{CBE}_{\mathrm{Co,delatt}}$ does not satisfy and thus not inherit weak directionality. □

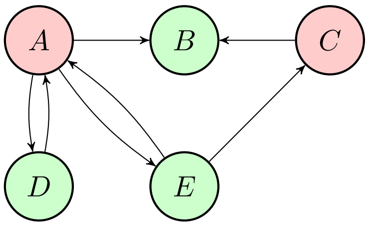

Figure 12: An AF with three $\mathrm{CBE}_{\mathrm{Co,delatt}}$-extensions

Proposition 6.11 (Inheritance of semi-directionality by $\mathrm{CBE}_{\mathrm{Co,nonatt}}$ and $\mathrm{CBE}_{\mathrm{Co,strdef}}$)**.** *Copeland-based extensions semantics* $\mathrm{CBE}_{\mathrm{Co,nonatt}}$ *and* $\mathrm{CBE}_{\mathrm{Co,strdef}}$ *do not inherit semi-directionality.*

Proof. Let $F = (\mathcal{A}, \hookrightarrow)$ be the AF shown by Figure 6. For Copeland-based extensions semantics $\mathrm{CBE}_{\mathrm{Co,nonatt}}$, we have two extensions: $\{A, B\}$ and $\{B, C\}$. For $\mathcal{S} = \{B, D\}$, there exists no attack from $\mathcal{A} \setminus \mathcal{S}$ to $\mathcal{S}$. As the restricted AF has extensions $\mathrm{CBE}_{\mathrm{Co,nonatt}}(\mathcal{S}, \hookrightarrow\restriction_{\mathcal{S}}) = \{\{B\}, \{D\}\}$ not included in the original extensions' intersections $\{\mathcal{E} \cap \mathcal{S} \mid \mathcal{E} \in \mathrm{CBE}_{\mathrm{Co,nonatt}}(F)\} = \{\{B\}\}$, F does not satisfy the condition set for all AFs by semi-directionality. The same holds for $\mathrm{CBE}_{\mathrm{Co,strdef}}$ as follows from Proposition 3.2. Therefore, $\mathrm{CBE}_{\mathrm{Co,nonatt}}$ and $\mathrm{CBE}_{\mathrm{Co,strdef}}$ do not satisfy and thus not inherit semi-directionality. □

Proposition 6.12 (Inheritance of semi-directionality by $\mathrm{CBE}_{\mathrm{Co,delarg}}$)**.** *Copeland-based extensions semantics* $\mathrm{CBE}_{\mathrm{Co,delarg}}$ *does not inherit semi-directionality.*

Proof. Let $F = (\mathcal{A}, \hookrightarrow)$ be the AF shown by Figure 13. For Copeland-based extensions semantics $\mathrm{CBE}_{\mathrm{Co,delarg}}$, we have one extension: $\{A, C\}$. For $\mathcal{S} = \{B, C\}$, there exists no attack from $\mathcal{A} \backslash \mathcal{S}$ to $\mathcal{S}$. As the restricted AF has extensions $\mathrm{CBE}_{\mathrm{Co,delarg}}(\mathcal{S}, \hookrightarrow_{\restriction \mathcal{S}}) = \{\varnothing, \{B\}, \{C\}\}$ not included in the original extensions' intersections $\{\mathcal{E} \cap \mathcal{S} \mid \mathcal{E} \in \mathrm{CBE}_{\mathrm{Co,delarg}}(F)\} = \{\{C\}\}$, F does not satisfy the condition set for all AFs by semi-directionality. Therefore, $\mathrm{CBE}_{\mathrm{Co,delarg}}$ does not satisfy and thus not inherit semi-directionality. □

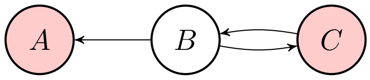

Figure 13: An AF with one $\mathrm{CBE}_{\mathrm{Co,delarg}}$-extension

Proposition 6.13 (Inheritance of semi-directionality by $\mathrm{CBE}_{\mathrm{Co,delatt}}$). *Copeland-based extensions semantics* $\mathrm{CBE}_{\mathrm{Co,delatt}}$ *does not inherit semi-directionality.*

Proof. Let $F = (\mathcal{A}, \hookrightarrow)$ be the AF shown by Figure 14. For Copeland-based extensions semantics $\mathrm{CBE}_{\mathrm{Co,delatt}}$, we have one extension: $\{B, C\}$. For $\mathcal{S} = \{C, D\}$, there exists no attack from $\mathcal{A} \backslash \mathcal{S}$ to $\mathcal{S}$. As the restricted AF has extensions $\mathrm{CBE}_{\mathrm{Co,delatt}}(\mathcal{S}, \hookrightarrow_{\restriction \mathcal{S}}) = \{\varnothing, \{C\}, \{D\}\}$ not included in the original extensions' intersections $\{\mathcal{E} \cap \mathcal{S} \mid \mathcal{E} \in \mathrm{CBE}_{\mathrm{Co,delatt}}(F)\} = \{\{C\}\}$, F does not satisfy the condition set for all AFs by semi-directionality. Therefore, $\mathrm{CBE}_{\mathrm{Co,delatt}}$ does not satisfy and thus not inherit semi-directionality. □

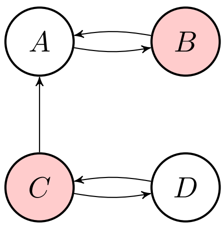

Figure 14: An AF with one $\mathrm{CBE}_{\mathrm{Co,delatt}}$-extension

6.1.2 Simpson-based extensions

Proposition 6.14 (Inheritance of allowing abstention by $\mathrm{SBE}_{\mathrm{Co,nonatt}}$ and $\mathrm{SBE}_{\mathrm{Co,strdef}}$). *Simpson-based extensions semantics* $\mathrm{SBE}_{\mathrm{Co,nonatt}}$ *and* $\mathrm{SBE}_{\mathrm{Co,strdef}}$ *do not inherit allowing abstention.*

Proof. Analogous to the proof for Proposition 6.3 as the same extensions are selected. □

Proposition 6.15 (Inheritance of weak directionality by $\mathrm{SBE}_{\mathrm{Co,delarg}}$). *Simpson-based extensions semantics* $\mathrm{SBE}_{\mathrm{Co,delarg}}$ *does not inherit weak directionality.*

Proof. Let $F = (\mathcal{A}, \hookrightarrow)$ be the AF shown by Figure 15. For Simpson-based extensions semantics $\mathrm{SBE}_{\mathrm{Co,delarg}}$, we have three extensions: $\varnothing$, $\{A, D\}$, and $\{B, C\}$. For $\mathcal{S} = \{A, B, C\}$, there exists no attack from $\mathcal{A} \setminus \mathcal{S}$ to $\mathcal{S}$. As the restricted AF has extensions $\mathrm{SBE}_{\mathrm{Co,delarg}}(\mathcal{S}, \hookrightarrow_{\restriction \mathcal{S}}) = \{\varnothing, \{B, C\}\}$ not including the original extensions' intersections $\{\mathcal{E} \cap \mathcal{S} \mid \mathcal{E} \in \mathrm{SBE}_{\mathrm{Co,delarg}}(F)\} = \{\varnothing, \{A\}, \{B, C\}\}$, F does not satisfy the condition set for all AFs by weak directionality. Therefore, $\mathrm{SBE}_{\mathrm{Co,delarg}}$ does not satisfy and thus not inherit weak directionality. □

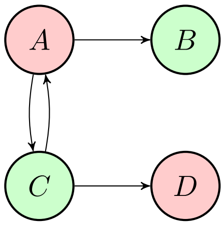

Figure 15: An AF with three $\mathrm{SBE}_{\mathrm{Co,delarg}}$-extensions

Proposition 6.16 (Inheritance of weak directionality by $\mathrm{SBE}_{\mathrm{Co,delatt}}$). *Simpson-based extensions semantics* $\mathrm{SBE}_{\mathrm{Co,delatt}}$ *does not inherit weak directionality.*

Proof. Let $F = (\mathcal{A}, \hookrightarrow)$ be the AF shown by Figure 16. For Simpson-based extensions semantics $\mathrm{SBE}_{\mathrm{Co,delatt}}$, we have three extensions: $\varnothing$, $\{A, D\}$, and $\{B, C, E\}$. For $\mathcal{S} = \{B, D, E\}$, there exists no attack from $\mathcal{A} \setminus \mathcal{S}$ to $\mathcal{S}$. As the restricted AF has extensions $\mathrm{SBE}_{\mathrm{Co,delatt}}(\mathcal{S}, \hookrightarrow_{\restriction \mathcal{S}}) = \{\varnothing, \{B, E\}\}$ not including the original extensions' intersections $\{\mathcal{E} \cap \mathcal{S} \mid \mathcal{E} \in \mathrm{SBE}_{\mathrm{Co,delatt}}(F)\} = \{\varnothing, \{D\}, \{B, E\}\}$, F does not satisfy the condition set for all AFs by weak directionality. Therefore, $\mathrm{SBE}_{\mathrm{Co,delatt}}$ does not satisfy and thus not inherit weak directionality. □

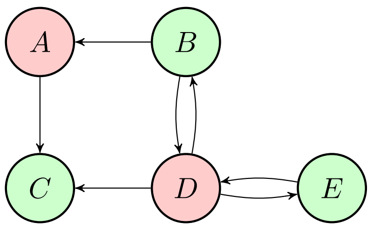

Figure 16: An AF with three $\mathrm{SBE}_{\mathrm{Co,delatt}}$-extensions

Proposition 6.17 (Inheritance of semi-directionality by $\mathrm{SBE}_{\mathrm{Co,nonatt}}$ and $\mathrm{SBE}_{\mathrm{Co,strdef}}$). *Simpson-based extensions semantics* $\mathrm{SBE}_{\mathrm{Co,nonatt}}$ *and* $\mathrm{SBE}_{\mathrm{Co,strdef}}$ *do not inherit semi-directionality.*

Proof. Let $F = (\mathcal{A}, \hookrightarrow)$ be the AF shown by Figure 17. For Simpson-based extensions semantics $\mathrm{SBE}_{\mathrm{Co,nonatt}}$, we have one extension: $\{B, D, F\}$. For $\mathcal{S} = \{B, C\}$, there exists no attack from $\mathcal{A} \backslash \mathcal{S}$ to $\mathcal{S}$. As the restricted AF has extensions $\mathrm{SBE}_{\mathrm{Co,nonatt}}(\mathcal{S}, \hookrightarrow_{\restriction \mathcal{S}}) = \{\{B\}, \{C\}\}$ not included in the original extensions' intersections $\{\mathcal{E} \cap \mathcal{S} \mid \mathcal{E} \in \mathrm{SBE}_{\mathrm{Co,nonatt}}(F)\} = \{\{B\}\}$, F does not satisfy the condition set for all AFs by semi-directionality. The same holds for $\mathrm{SBE}_{\mathrm{Co,strdef}}$ as follows from Proposition 3.2. Therefore, $\mathrm{SBE}_{\mathrm{Co,nonatt}}$ and $\mathrm{SBE}_{\mathrm{Co,strdef}}$ do not satisfy and thus not inherit semi-directionality. □

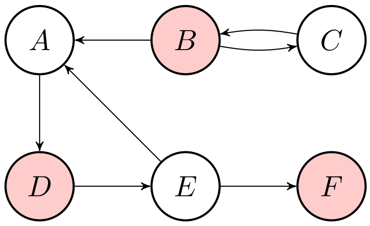

Figure 17: An AF with one $\mathrm{SBE}_{\mathrm{Co,nonatt}}$-extension

Proposition 6.18 (Inheritance of semi-directionality by $\mathrm{SBE}_{\mathrm{Co,delarg}}$). *Simpson-based extensions semantics* $\mathrm{SBE}_{\mathrm{Co,delarg}}$ *does not inherit semi-directionality.*

Proof. Let $F = (\mathcal{A}, \hookrightarrow)$ be the AF shown by Figure 13. For Simpson-based extensions semantics $\mathrm{SBE}_{\mathrm{Co,delarg}}$, we have two extensions: $\varnothing$ and $\{A, C\}$. For $\mathcal{S} = \{B, C\}$, there exists no attack from $\mathcal{A} \setminus \mathcal{S}$ to $\mathcal{S}$. As the restricted AF has extensions $\mathrm{SBE}_{\mathrm{Co,delarg}}(\mathcal{S}, \hookrightarrow_{\restriction \mathcal{S}}) = \{\varnothing, \{B\}, \{C\}\}$ not included in the original extensions' intersections $\{\mathcal{E} \cap \mathcal{S} \mid \mathcal{E} \in \mathrm{SBE}_{\mathrm{Co,delarg}}(F)\} = \{\varnothing, \{C\}\}$, F does not satisfy the condition set for all AFs by semi-directionality. Therefore, $\mathrm{SBE}_{\mathrm{Co,delarg}}$ does not satisfy and thus not inherit semi-directionality. □

Proposition 6.19 (Inheritance of semi-directionality by $\mathrm{SBE}_{\mathrm{Co,delatt}}$). *Simpson-based extensions semantics* $\mathrm{SBE}_{\mathrm{Co,delatt}}$ *does not inherit semi-directionality.*

Proof. Let $F = (\mathcal{A}, \hookrightarrow)$ be the AF shown by Figure 18. For Simpson-based extensions semantics $\mathrm{SBE}_{\mathrm{Co,delatt}}$, we have two extensions: $\varnothing$ and $\{A, D\}$. For $\mathcal{S} = \{A, C, D\}$, there exists no attack from $\mathcal{A} \setminus \mathcal{S}$ to $\mathcal{S}$. As the restricted AF has extensions $\mathrm{SBE}_{\mathrm{Co,delatt}}(\mathcal{S}, \hookrightarrow_{\restriction \mathcal{S}}) = \{\varnothing, \{C\}, \{A, D\}\}$ not included in the original extensions' intersections $\{\mathcal{E} \cap \mathcal{S} \mid \mathcal{E} \in \mathrm{SBE}_{\mathrm{Co,delatt}}(F)\} = \{\varnothing, \{A, D\}\}$, F does not satisfy the condition set for all AFs by semi-directionality. Therefore, $\mathrm{SBE}_{\mathrm{Co,delatt}}$ does not satisfy and thus not inherit semi-directionality. □

Proposition 6.20 (Inheritance of com-closure by $\mathrm{SBE}_{\mathrm{Co,delarg}}$). *Simpson-based extensions semantics* $\mathrm{SBE}_{\mathrm{Co,delarg}}$ *does not inherit com-closure.*

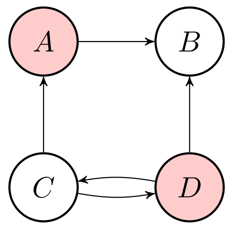

Figure 18: An AF with two $\text{SBE}_{\text{Co,delatt}}$-extensions

Proof. Let $F = (\mathcal{A}, \hookrightarrow)$ be the AF shown by Figure 19. For Simpson-based extensions semantics $\text{SBE}_{\text{Co,delarg}}$, we have nine extensions: $\emptyset$, $\{B\}$, $\{C\}$, $\{G\}$, $\{D, G\}$, $\{A, C, G\}$, $\{B, D, G\}$, $\{C, D, G\}$, and $\{C, E, F\}$. For $S = \{\{C\}, \{G\}\}$ and all $a, b \in \bigcup_{\mathcal{E} \in S} \mathcal{E}$, there exists an extension $\mathcal{E}' \in \text{SBE}_{\text{Co,delarg}}(F)$ such that $a, b \in \mathcal{E}'$. As there exist two extensions $\mathcal{E}'' \in \{\{A, C, G\}, \{C, D, G\}\}$ such that $\bigcup_{\mathcal{E} \in S} \mathcal{E} \subseteq \mathcal{E}''$ and $\mathcal{E}''$ is minimal with respect to set inclusion among the extensions that satisfy the former, F does not satisfy the condition set for all AFs by com-closure. Therefore, $\text{SBE}_{\text{Co,delarg}}$ does not satisfy and thus not inherit com-closure. □

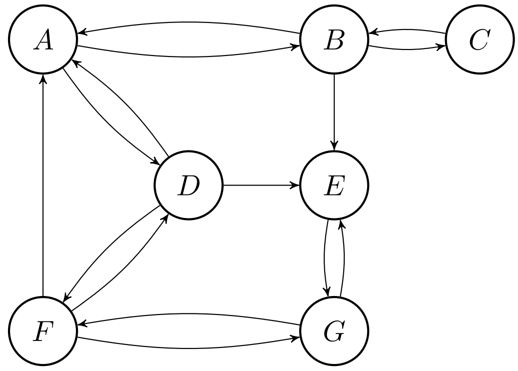

Figure 19: An AF with nine $\text{SBE}_{\text{Co,delarg}}$-extensions

Proposition 6.21 (Inheritance of com-closure by $\text{SBE}_{\text{Co,delatt}}$)**.** *Simpson-based extensions semantics* $\text{SBE}_{\text{Co,delatt}}$ *does not inherit com-closure.*

Proof. Let $F = (\mathcal{A}, \hookrightarrow)$ be the AF shown by Figure 20. For Simpson-based extensions semantics $\text{SBE}_{\text{Co,delatt}}$, we have fourteen extensions: $\{I\}$, $\{A, I\}$, $\{D, I\}$, $\{H, I\}$, $\{I, G\}$, $\{A, D, I\}$, $\{A, F, I\}$, $\{A, G, I\}$, $\{D, G, I\}$, $\{G, H, I\}$, $\{A, C, H, I\}$, $\{A, D, F, I\}$, $\{A, D, G, I\}$, and $\{A, C, F, H, I\}$. For $S = \{\{A, I\}, \{G, I\}, \{H, I\}\}$ and all $a, b \in \bigcup_{\mathcal{E} \in S} \mathcal{E}$, there exists an extension $\mathcal{E}' \in \text{SBE}_{\text{Co,delatt}}(F)$ such that $a, b \in \mathcal{E}'$. As there exists no extension $\mathcal{E}'' \in \text{SBE}_{\text{Co,delatt}}(F)$ such that $\bigcup_{\mathcal{E} \in S} \mathcal{E} \subseteq \mathcal{E}''$, F does not satisfy the condition set for all AFs by com-closure. Therefore, $\text{SBE}_{\text{Co,delatt}}$ does not satisfy and thus not inherit com-closure. □

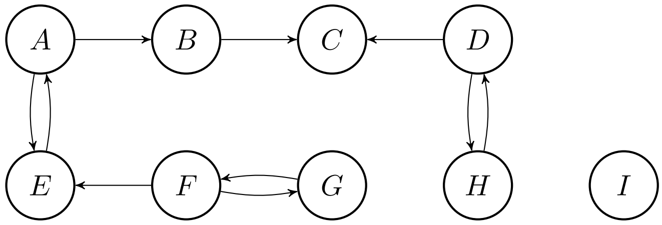

Figure 20: An AF with fourteen $\mathrm{SBE}_{\mathrm{Co,delatt}}$-extensions

6.2 Preferred semantics

The preferred semantics satisfies the following principles which are not trivially inherited: crash resistance, non-interference, weak directionality, semi-directionality, directionality, conflict-sensitiveness, and com-closure [vDTV18].

Lemma 6.1 (Attacks between preferred extensions). *For all AFs F with preferred extensions $\mathcal{E}, \mathcal{E}' \in \mathrm{Pr}(F)$ where $\mathcal{E} \neq \mathcal{E}'$, there exists an attack from $\mathcal{E}$ to $\mathcal{E}'$.*

Proof. Let $F = (\mathcal{A}, \hookrightarrow)$ be an AF with preferred extensions $\mathcal{E}, \mathcal{E}' \in \mathrm{Pr}(F)$ where $\mathcal{E} \neq \mathcal{E}'$. As the preferred extensions are maximal with respect to set inclusion among the admissible sets of arguments, the union $\mathcal{E} \cup \mathcal{E}'$ is not preferred such that it is not admissible. Then, $\mathcal{E} \cup \mathcal{E}'$ is not acceptable with respect to itself or is not conflict-free. As a union of sets that are acceptable with respect to each self, $\mathcal{E} \cup \mathcal{E}'$ is acceptable with respect to itself, such that it is not conflict-free. Then, there exists an attack from $\mathcal{E}$ to $\mathcal{E}'$ or from $\mathcal{E}'$ to $\mathcal{E}$, and as the extensions are acceptable with respect to each self, it follows that either attack exists. Therefore, there exists an attack from $\mathcal{E}$ to $\mathcal{E}'$. $\square$

Proposition 6.22 (Conflict-sensitiveness of preferred-based semantics). *Any semantics σ such that for all AFs F, it holds that $\sigma(F) \subseteq \mathrm{Pr}(F)$, satisfies conflict-sensitiveness.*

Proof. Let $F = (\mathcal{A}, \hookrightarrow)$ be an AF and σ a semantics such that for all AFs F', it holds that $\sigma(F') \subseteq \mathrm{Pr}(F')$. For any two extensions $\mathcal{E}, \mathcal{E}' \in \sigma(F)$, if their union $\mathcal{E} \cup \mathcal{E}'$ is not an extension, they are distinct and as they are preferred, it follows from Lemma 6.1 that there exists an attack $a \hookrightarrow b$ from $\mathcal{E}$ to $\mathcal{E}'$ such that $a, b \in \mathcal{E} \cup \mathcal{E}'$. Then, there exists no extension $\mathcal{E}'' \in \sigma(F)$ where $a, b \in \sigma(F)$, as the extensions are preferred and the preferred extensions are conflict-free. Therefore, σ satisfies conflict-sensitiveness. $\square$

Corollary 6.2 (Inheritance of conflict-sensitiveness by $\mathrm{CBE}_{\mathrm{Pr},\gamma}$ and $\mathrm{SBE}_{\mathrm{Pr},\gamma}$). *Any Copeland-based extensions semantics $\mathrm{CBE}_{\mathrm{Pr},\gamma}$ and Simpson-based extensions semantics $\mathrm{SBE}_{\mathrm{Pr},\gamma}$ inherit conflict-sensitiveness.*

Proof. Follows from Definition 3.2 in connection with Proposition 6.22. $\square$

Proposition 6.23 (Inheritance of crash resistance by $\mathrm{SEL}_{\mathrm{Pr},\gamma}$). *Any selection semantics $\mathrm{SEL}_{\mathrm{Pr},\gamma}$ that maintains non-emptiness inherits crash resistance.*

Preferred	Adm.	Reinst.	W. reinst.	CF-reinst.	I-max.	Crash resist.
CBE						
nonatt	✓	✓	✓	✓	✓	✓
strdef	✓	✓	✓	✓	✓	✓
delarg	✓	✓	✓	✓	✓	✓
delatt	✓	✓	✓	✓	✓	✓
SBE						
nonatt	✓	✓	✓	✓	✓	✓
strdef	✓	✓	✓	✓	✓	✓
delarg	✓	✓	✓	✓	✓	✓
delatt	✓	✓	✓	✓	✓	✓

Table 7: Inheritance of principles satisfied by the preferred semantics for preferred-based selection semantics, part 1

Proof. Analogous to the proof for Proposition 6.2 as the set of preferred extensions is non-empty and the preferred extensions are complete. □

Corollary 6.3 (Inheritance of crash resistance by $\mathrm{CBE}_{\mathrm{Pr},\gamma}$ and $\mathrm{SBE}_{\mathrm{Pr},\gamma}$)**.** *Any Copeland-based extensions semantics* $\mathrm{CBE}_{\mathrm{Pr},\gamma}$ *and Simpson-based extensions semantics* $\mathrm{SBE}_{\mathrm{Pr},\gamma}$ *where* $\gamma \in \{\mathrm{nonatt}, \mathrm{strdef}, \mathrm{delarg}, \mathrm{delarg}\}$ *inherit crash resistance.*

Proof. Follows from Propositions 3.1, 3.3, and 5.7 in connection with Lemma 5.1, in connection with Proposition 6.23. □

Proposition 6.24 (Com-closure of preferred-based semantics)**.** *Any semantics* σ *such that for all AFs* F*, it holds that* $\sigma(F) \subseteq \mathrm{Pr}(F)$*, satisfies com-closure.*

Proof. Let $F = (\mathcal{A}, \hookrightarrow)$ be an AF and σ a semantics such that for all AFs F', it holds that $\sigma(F') \subseteq \mathrm{Pr}(F')$. As the extensions are preferred, it follows from Lemma 6.1 that there exists an attack between any two distinct extensions. Then, only non-empty sets of singular extensions form unions, of each every two arguments are contained in an extension. As the extensions are preferred and the preferred extensions satisfy I-maximality which is trivially inherited by σ, for each union, there exists one and only one extension that contains its arguments, the extension that forms the union. Therefore, σ satisfies com-closure. □

Corollary 6.4 (Inheritance of com-closure by $\mathrm{SEL}_{\mathrm{Pr},\gamma}$)**.** *Any Copeland-based extensions semantics* $\mathrm{CBE}_{\mathrm{Pr},\gamma}$ *and Simpson-based extensions semantics* $\mathrm{SBE}_{\mathrm{Pr},\gamma}$ *inherit com-closure.*

Proof. Follows from Definition 3.2 in connection with Proposition 6.24. □

Preferred	Non-int.	W. dir.	S.-dir.	Dir.	Conf.-sens.	Com-closure
CBE						
nonatt	?	✗	✗	✗	✓	✓
strdef	?	✗	✗	✗	✓	✓
delarg	?	✗	✗	✗	✓	✓
delatt	?	✗	✗	✗	✓	✓
SBE						
nonatt	?	?	✗	✗	✓	✓
strdef	?	?	✗	✗	✓	✓
delarg	?	✗	✗	✗	✓	✓
delatt	?	✗	✗	✗	✓	✓

Table 8: Inheritance of principles satisfied by the preferred semantics for preferred-based selection semantics, part 2

6.2.1 Copeland-based extensions

Proposition 6.25 (Inheritance of weak directionality by $\mathrm{CBE}_{\mathrm{Pr,nonatt}}$ and $\mathrm{SBE}_{\mathrm{Pr,strdef}}$). *Copeland-based extensions semantics* $\mathrm{CBE}_{\mathrm{Pr,nonatt}}$ *and Simpson-based extensions semantics* $\mathrm{CBE}_{\mathrm{Pr,strdef}}$ *do not inherit weak directionality.*

Proof. Let $F = (\mathcal{A}, \hookrightarrow)$ be an AF where

$$\mathcal{A} = \{a, b, c, d, e, f, g, h, i, j\}$$

and

$$\hookrightarrow = \{(a, b), (a, e), (a, h), (a, j), (b, f), (b, g), (c, e), (c, f), (c, g), (c, h), (c, i), (e, b), (e, d), (e, f), (e, g), (e, i), (f, d), (g, c), (g, d), (g, i), (g, j), (h, a), (h, b), (h, c), (h, i), (i, d), (i, f), (i, g), (i, h), (j, a), (j, b), (j, c), (j, e), (j, f), (j, i)\}.$$

For Copeland-based extensions semantics $\mathrm{CBE}_{\mathrm{Pr,nonatt}}$, we have two extensions: $\{a, c, d\}$ and $\{a, f, g\}$. For $\mathcal{S} = \{a, b, c, e, f, g, h, i, j\}$, there exists no attack from $\mathcal{A} \setminus \mathcal{S}$ to $\mathcal{S}$. As the restricted AF has extensions $\mathrm{CBE}_{\mathrm{Pr,nonatt}}(\mathcal{S}, \hookrightarrow_{\upharpoonright \mathcal{S}}) = \{\{a, f, g\}\}$ not including the original extensions' intersections $\{\mathcal{E} \cap \mathcal{S} \mid \mathcal{E} \in \mathrm{CBE}_{\mathrm{Pr,nonatt}}(F)\} = \{\{a, c\}, \{a, f, g\}\}$, F does not satisfy the condition set for all AFs by weak directionality. The same holds for $\mathrm{CBE}_{\mathrm{Pr,strdef}}$ as follows from Proposition 3.2. Therefore, $\mathrm{CBE}_{\mathrm{Pr,nonatt}}$ and $\mathrm{CBE}_{\mathrm{Pr,strdef}}$ do not satisfy and thus not inherit weak directionality. □

Proposition 6.26 (Inheritance of weak directionality by $\mathrm{CBE}_{\mathrm{Pr,delarg}}$). *Copeland-based extensions semantics* $\mathrm{CBE}_{\mathrm{Pr,delarg}}$ *does not inherit weak directionality.*

Proof. Let $F = (\mathcal{A}, \hookrightarrow)$ be the AF shown by Figure 15. For Copeland-based extensions semantics $\mathrm{CBE}_{\mathrm{Pr,delarg}}$, we have two extensions: $\{A, D\}$ and $\{B, C\}$. For $\mathcal{S} = \{A, B, C\}$, there exists no attack from $\mathcal{A} \setminus \mathcal{S}$ to $\mathcal{S}$. As the restricted AF has extensions $\mathrm{CBE}_{\mathrm{Pr,delarg}}(\mathcal{S}, \hookrightarrow_{\upharpoonright \mathcal{S}}) = \{\{B, C\}\}$ not including the original extensions' intersections $\{\mathcal{E} \cap \mathcal{S} \mid \mathcal{E} \in \mathrm{CBE}_{\mathrm{Pr,delarg}}(F)\} = \{\{A\}, \{B, C\}\}$, F does not satisfy the condition set for all AFs by weak directionality. Therefore, $\mathrm{CBE}_{\mathrm{Pr,delarg}}$ does not satisfy and thus not inherit weak directionality. □

Proposition 6.27 (Inheritance of weak directionality by $\mathrm{CBE}_{\mathrm{Pr,delatt}}$)**.** *Copeland-based extensions semantics* $\mathrm{CBE}_{\mathrm{Pr,delatt}}$ *does not inherit weak directionality.*

Proof. Let $F = (\mathcal{A}, \hookrightarrow)$ be the AF shown by Figure 21. For Copeland-based extensions semantics $\mathrm{CBE}_{\mathrm{Pr,delatt}}$, we have two extensions: $\{A, D\}$ and $\{B, C, E\}$. For $\mathcal{S} = \{C, D, E\}$, there exists no attack from $\mathcal{A} \setminus \mathcal{S}$ to $\mathcal{S}$. As the restricted AF has extensions $\mathrm{CBE}_{\mathrm{Pr,delatt}}(\mathcal{S}, \hookrightarrow_{\upharpoonright \mathcal{S}}) = \{\{C, E\}\}$ not including the original extensions' intersections $\{\mathcal{E} \cap \mathcal{S} \mid \mathcal{E} \in \mathrm{CBE}_{\mathrm{Pr,delatt}}(F)\} = \{\{D\}, \{C, E\}\}$, F does not satisfy the condition set for all AFs by weak directionality. Therefore, $\mathrm{CBE}_{\mathrm{Pr,delatt}}$ does not satisfy and thus not inherit weak directionality. □

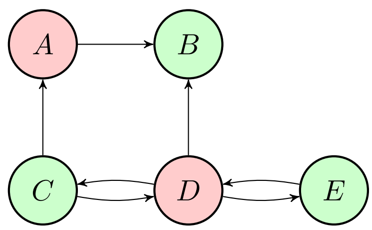

Figure 21: An AF with two $\mathrm{CBE}_{\mathrm{Pr,delatt}}$-extensions

Proposition 6.28 (Inheritance of semi-directionality by $\mathrm{CBE}_{\mathrm{Pr,nonatt}}$ and $\mathrm{CBE}_{\mathrm{Pr,strdef}}$)**.** *Copeland-based extensions semantics* $\mathrm{CBE}_{\mathrm{Pr,nonatt}}$ *and* $\mathrm{CBE}_{\mathrm{Pr,strdef}}$ *do not inherit semi-directionality.*

Proof. Let $F = (\mathcal{A}, \hookrightarrow)$ be the AF shown by Figure 22. For Copeland-based extensions semantics $\mathrm{CBE}_{\mathrm{Pr,nonatt}}$, we have one extension: $\{A, C, E\}$. For $\mathcal{S} = \{C, F\}$, there exists no attack from $\mathcal{A} \backslash \mathcal{S}$ to $\mathcal{S}$. As the restricted AF has extensions $\mathrm{CBE}_{\mathrm{Pr,nonatt}}(\mathcal{S}, \hookrightarrow_{\upharpoonright \mathcal{S}}) = \{\{C\}, \{F\}\}$ not included in the original extensions' intersections $\{\mathcal{E} \cap \mathcal{S} \mid \mathcal{E} \in \mathrm{CBE}_{\mathrm{Pr,nonatt}}(F)\} = \{\{C\}\}$, F does not satisfy the condition set for all AFs by semi-directionality. The same holds for $\mathrm{CBE}_{\mathrm{Pr,strdef}}$ as follows from Proposition 3.2. Therefore, $\mathrm{CBE}_{\mathrm{Pr,nonatt}}$ and $\mathrm{CBE}_{\mathrm{Pr,strdef}}$ do not satisfy and thus not inherit semi-directionality. □

Proposition 6.29 (Inheritance of semi-directionality by $\mathrm{CBE}_{\mathrm{Pr,delarg}}$)**.** *Copeland-based extensions semantics* $\mathrm{CBE}_{\mathrm{Pr,delarg}}$ *does not inherit semi-directionality.*

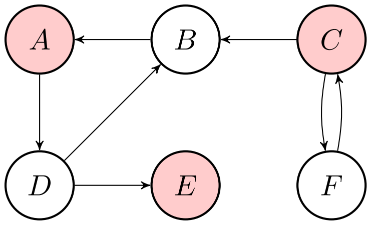

Figure 22: An AF with one $\mathrm{CBE}_{\mathrm{Pr,nonatt}}$-extension

Proof. Let $F = (\mathcal{A}, \hookrightarrow)$ be the AF shown by Figure 13. For Copeland-based extensions semantics $\mathrm{CBE}_{\mathrm{Pr,delarg}}$, we have one extension: $\{A, C\}$. For $\mathcal{S} = \{B, C\}$, there exists no attack from $\mathcal{A} \backslash \mathcal{S}$ to $\mathcal{S}$. As the restricted AF has extensions $\mathrm{CBE}_{\mathrm{Pr,delarg}}(\mathcal{S}, \hookrightarrow_{\restriction \mathcal{S}}) = \{\{B\}, \{C\}\}$ not included in the original extensions' intersections $\{\mathcal{E} \cap \mathcal{S} \mid \mathcal{E} \in \mathrm{CBE}_{\mathrm{Pr,delarg}}(F)\} = \{\{C\}\}$, F does not satisfy the condition set for all AFs by semi-directionality. Therefore, $\mathrm{CBE}_{\mathrm{Pr,delarg}}$ does not satisfy and thus not inherit semi-directionality. □

Proposition 6.30 (Inheritance of semi-directionality by $\mathrm{CBE}_{\mathrm{Pr,delatt}}$)**.** *Copeland-based extensions semantics* $\mathrm{CBE}_{\mathrm{Pr,delatt}}$ *does not inherit semi-directionality.*

Proof. Let $F = (\mathcal{A}, \hookrightarrow)$ be the AF shown by Figure 14. For Copeland-based extensions semantics $\mathrm{CBE}_{\mathrm{Pr,delatt}}$, we have one extension: $\{B, C\}$. For $\mathcal{S} = \{C, D\}$, there exists no attack from $\mathcal{A} \backslash \mathcal{S}$ to $\mathcal{S}$. As the restricted AF has extensions $\mathrm{CBE}_{\mathrm{Pr,delatt}}(\mathcal{S}, \hookrightarrow_{\restriction \mathcal{S}}) = \{\{C\}, \{D\}\}$ not included in the original extensions' intersections $\{\mathcal{E} \cap \mathcal{S} \mid \mathcal{E} \in \mathrm{CBE}_{\mathrm{Pr,delatt}}(F)\} = \{\{C\}\}$, F does not satisfy the condition set for all AFs by semi-directionality. Therefore, $\mathrm{CBE}_{\mathrm{Pr,delatt}}$ does not satisfy and thus not inherit semi-directionality. □

6.2.2 Simpson-based extensions

Proposition 6.31 (Inheritance of weak directionality by $\mathrm{SBE}_{\mathrm{Pr,delarg}}$)**.** *Simpson-based extensions semantics* $\mathrm{SBE}_{\mathrm{Pr,delarg}}$ *does not inherit weak directionality.*

Proof. Let $F = (\mathcal{A}, \hookrightarrow)$ be the AF shown by Figure 15. For Simpson-based extensions semantics $\mathrm{SBE}_{\mathrm{Pr,delarg}}$, we have two extensions: $\{A, D\}$ and $\{B, C\}$. For $\mathcal{S} = \{A, B, C\}$, there exists no attack from $\mathcal{A} \setminus \mathcal{S}$ to $\mathcal{S}$. As the restricted AF has extensions $\mathrm{SBE}_{\mathrm{Pr,delarg}}(\mathcal{S}, \hookrightarrow_{\restriction \mathcal{S}}) = \{\{B, C\}\}$ not including the original extensions' intersections $\{\mathcal{E} \cap \mathcal{S} \mid \mathcal{E} \in \mathrm{SBE}_{\mathrm{Pr,delarg}}(F)\} = \{\{A\}, \{B, C\}\}$, F does not satisfy the condition set for all AFs by weak directionality. Therefore, $\mathrm{SBE}_{\mathrm{Pr,delarg}}$ does not satisfy and thus not inherit weak directionality. □

Proposition 6.32 (Inheritance of weak directionality by $\mathrm{SBE}_{\mathrm{Pr,delatt}}$)**.** *Simpson-based extensions semantics* $\mathrm{SBE}_{\mathrm{Pr,delatt}}$ *does not inherit weak directionality.*

Proof. Analogous to the proof for Proposition 6.27 as the same extensions are selected. □

Proposition 6.33 (Inheritance of semi-directionality by $\mathrm{SBE}_{\mathrm{Pr,nonatt}}$ and $\mathrm{SBE}_{\mathrm{Pr,strdef}}$). *Simpson-based extensions semantics* $\mathrm{SBE}_{\mathrm{Pr,nonatt}}$ *and* $\mathrm{SBE}_{\mathrm{Pr,strdef}}$ *do not inherit semi-directionality.*

Proof. Analogous to the proof for Proposition 6.28 as the same extensions are selected. □

Proposition 6.34 (Inheritance of semi-directionality by $\mathrm{SBE}_{\mathrm{Pr,delarg}}$). *Simpson-based extensions semantics* $\mathrm{SBE}_{\mathrm{Pr,delarg}}$ *does not inherit semi-directionality.*

Proof. Let $F = (\mathcal{A}, \hookrightarrow)$ be the AF shown by Figure 13. For Simpson-based extensions semantics $\mathrm{SBE}_{\mathrm{Pr,delarg}}$, we have one extension: $\{A, C\}$. For $\mathcal{S} = \{B, C\}$, there exists no attack from $\mathcal{A}\backslash\mathcal{S}$ to $\mathcal{S}$. As the restricted AF has extensions $\mathrm{SBE}_{\mathrm{Pr,delarg}}(\mathcal{S}, \hookrightarrow_{\restriction\mathcal{S}}) = \{\{B\}, \{C\}\}$ not included in the original extensions' intersections $\{\mathcal{E} \cap \mathcal{S} \mid \mathcal{E} \in \mathrm{SBE}_{\mathrm{Pr,delarg}}(F)\} = \{\{C\}\}$, F does not satisfy the condition set for all AFs by semi-directionality. Therefore, $\mathrm{SBE}_{\mathrm{Pr,delarg}}$ does not satisfy and thus not inherit semi-directionality. □

Proposition 6.35 (Inheritance of semi-directionality by $\mathrm{SBE}_{\mathrm{Pr,delatt}}$). *Simpson-based extensions semantics* $\mathrm{SBE}_{\mathrm{Pr,delatt}}$ *does not inherit semi-directionality.*

Proof. Let $F = (\mathcal{A}, \hookrightarrow)$ be the AF shown by Figure 18. For Simpson-based extensions semantics $\mathrm{SBE}_{\mathrm{Pr,delatt}}$, we have one extension: $\{A, D\}$. For $\mathcal{S} = \{A, C, D\}$, there exists no attack from $\mathcal{A}\backslash\mathcal{S}$ to $\mathcal{S}$. As the restricted AF has extensions $\mathrm{SBE}_{\mathrm{Pr,delatt}}(\mathcal{S}, \hookrightarrow_{\restriction\mathcal{S}}) = \{\{C\}, \{A, D\}\}$ not included in the original extensions' intersections $\{\mathcal{E} \cap \mathcal{S} \mid \mathcal{E} \in \mathrm{SBE}_{\mathrm{Pr,delatt}}(F)\} = \{\{A, D\}\}$, F does not satisfy the condition set for all AFs by semi-directionality. Therefore, $\mathrm{SBE}_{\mathrm{Pr,delatt}}$ does not satisfy and thus not inherit semi-directionality. □

6.3 Stable semantics

The stable semantics satisfies the following principles which are not trivially inherited: weak directionality, tightness, conflict-sensitiveness, and com-closure [vDTV18].

Proposition 6.36 (Equality of St to $\mathrm{CBE}_{\mathrm{St},\gamma}$ and $\mathrm{SBE}_{\mathrm{St},\gamma}$). *The stable semantics is equal to any Copeland-based extensions semantics* $\mathrm{CBE}_{\mathrm{St},\gamma}$ *and Simpson-based extensions semantics* $\mathrm{SBE}_{\mathrm{St},\gamma}$ *where* $\gamma \in \{\mathrm{nonatt}, \mathrm{strdef}\}$.

Proof. Let F be an AF with stable extensions $\mathcal{E}, \mathcal{E}' \in \mathrm{St}(F)$. As a stable extension attacks all arguments that are not contained in it, the arguments in $\mathcal{E}$ not attacked by $\mathcal{E}'$ and vice versa are their intersection $\mathcal{E} \cap \mathcal{E}'$. Then, we have $\mathcal{E} \geq^F_{\mathrm{St,nonatt}} \mathcal{E}'$ and $\mathcal{E}' \geq^F_{\mathrm{St,nonatt}} \mathcal{E}$. It follows that $\geq^F_{\mathrm{St}} = \mathrm{St}(F) \times \mathrm{St}(F)$ such that the stable extensions are selected by Copeland-based extensions semantics $\mathrm{CBE}_{\mathrm{St,nonatt}}$ as well as Simpson-based extensions semantics $\mathrm{SBE}_{\mathrm{St,nonatt}}$. The same holds for comparison criterion strdef as follows from Proposition 3.2. Therefore, it follows that $\mathrm{St} = \mathrm{CBE}_{\mathrm{St,nonatt}} = \mathrm{SBE}_{\mathrm{St,nonatt}}$ where $\gamma \in \{\mathrm{nonatt}, \mathrm{strdef}\}$. □

Stable	Adm.	Naivety	Reinst.	W. reinst.	CF-reinst.	I-max.
CBE						
nonatt	✓	✓	✓	✓	✓	✓
strdef	✓	✓	✓	✓	✓	✓
delarg	✓	✓	✓	✓	✓	✓
delatt	✓	✓	✓	✓	✓	✓
SBE						
nonatt	✓	✓	✓	✓	✓	✓
strdef	✓	✓	✓	✓	✓	✓
delarg	✓	✓	✓	✓	✓	✓
delatt	✓	✓	✓	✓	✓	✓

Table 9: Inheritance of principles satisfied by the stable semantics for stable-based selection semantics, part 1

Corollary 6.5 (Inheritance by $\mathrm{CBE}_{\mathrm{St},\gamma}$ and $\mathrm{SBE}_{\mathrm{St},\gamma}$). *Any Copeland-based extensions semantics* $\mathrm{CBE}_{\mathrm{St},\gamma}$ *and Simpson-based extensions semantics* $\mathrm{SBE}_{\mathrm{St},\gamma}$ *where* $\gamma \in \{\mathrm{nonatt}, \mathrm{strdef}\}$ *inherit all principles.*

Proof. Follows from Proposition 6.36. □

Corollary 6.6 (Inheritance of conflict-sensitiveness by $\mathrm{CBE}_{\mathrm{St},\gamma}$ and $\mathrm{SBE}_{\mathrm{St},\gamma}$). *Any Copeland-based extensions semantics* $\mathrm{CBE}_{\mathrm{St},\gamma}$ *and Simpson-based extensions semantics* $\mathrm{SBE}_{\mathrm{St},\gamma}$ *inherit conflict-sensitiveness.*

Proof. Follows from Definition 3.2 in connection with Proposition 6.22 as the stable extensions are preferred. □

Corollary 6.7 (Inheritance of com-closure by $\mathrm{CBE}_{\mathrm{St},\gamma}$ and $\mathrm{SBE}_{\mathrm{St},\gamma}$). *Any Copeland-based extensions semantics* $\mathrm{CBE}_{\mathrm{St},\gamma}$ *and Simpson-based extensions semantics* $\mathrm{SBE}_{\mathrm{St},\gamma}$ *inherit com-closure.*

Proof. Follows from Definition 3.2 in connection with Proposition 6.24 as the stable extensions are preferred. □

6.3.1 Copeland-based extensions

Proposition 6.37 (Inheritance of weak directionality by $\mathrm{CBE}_{\mathrm{St},\mathrm{delarg}}$). *Copeland-based extensions semantics* $\mathrm{CBE}_{\mathrm{St},\mathrm{delarg}}$ *does not inherit weak directionality.*

Proof. Let $F = (\mathcal{A}, \hookrightarrow)$ be the AF shown by Figure 23. For Copeland-based extensions semantics $\mathrm{CBE}_{\mathrm{St},\mathrm{delarg}}$, we have two extensions: $\{A, C\}$ and $\{B, D\}$. For $\mathcal{S} = \{A, B, C\}$, there exists no attack from $\mathcal{A} \setminus \mathcal{S}$ to $\mathcal{S}$. As the restricted AF has

Stable	W. dir.	Tightness	Conf.-sens.	Com-closure
CBE				
nonatt	✓	✓	✓	✓
strdef	✓	✓	✓	✓
delarg	✗	✓	✓	✓
delatt	✗	✓	✓	✓
SBE				
nonatt	✓	✓	✓	✓
strdef	✓	✓	✓	✓
delarg	✗	✓	✓	✓
delatt	✗	✓	✓	✓

Table 10: Inheritance of principles satisfied by the stable semantics for stable-based selection semantics, part 2

extensions $\mathrm{CBE}_{\mathrm{St,delarg}}(\mathcal{S}, \hookrightarrow_{\restriction\mathcal{S}}) = \{\{A, C\}\}$ not including the original extensions' intersections $\{\mathcal{E} \cap \mathcal{S} \mid \mathcal{E} \in \mathrm{CBE}_{\mathrm{St,delarg}}(F)\} = \{\{B\}, \{A, C\}\}$, F does not satisfy the condition set for all AFs by weak directionality. Therefore, $\mathrm{CBE}_{\mathrm{St,delarg}}$ does not satisfy and thus not inherit weak directionality. □

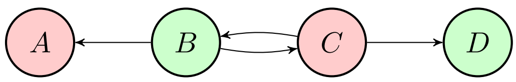

Figure 23: An AF with two $\mathrm{CBE}_{\mathrm{St,delarg}}$-extensions

Proposition 6.38 (Inheritance of weak directionality by $\mathrm{CBE}_{\mathrm{St,delatt}}$)**.** *Copeland-based extensions semantics* $\mathrm{CBE}_{\mathrm{St,delatt}}$ *does not inherit weak directionality.*

Proof. Analogous to the proof for Proposition 6.27 as the same extensions are selected. □

6.3.2 Simpson-based extensions

Proposition 6.39 (Inheritance of weak directionality by $\mathrm{SBE}_{\mathrm{St,delarg}}$)**.** *Simpson-based extensions semantics* $\mathrm{SBE}_{\mathrm{St,delarg}}$ *does not inherit weak directionality.*

Proof. Analogous to the proof for Proposition 6.37 as the same extensions are selected. □

Proposition 6.40 (Inheritance of weak directionality by $\mathrm{SBE}_{\mathrm{St,delatt}}$)**.** *Simpson-based extensions semantics* $\mathrm{SBE}_{\mathrm{St,delatt}}$ *does not inherit weak directionality.*

Grounded	Adm.	Str. adm.	Reinst.	W. reinst.	CF-reinst.	I-max.
CBE						
nonatt	✓	✓	✓	✓	✓	✓
strdef	✓	✓	✓	✓	✓	✓
delarg	✓	✓	✓	✓	✓	✓
delatt	✓	✓	✓	✓	✓	✓
SBE						
nonatt	✓	✓	✓	✓	✓	✓
strdef	✓	✓	✓	✓	✓	✓
delarg	✓	✓	✓	✓	✓	✓
delatt	✓	✓	✓	✓	✓	✓

Table 11: Inheritance of principles satisfied by the grounded semantics for grounded-based selection semantics, part 1

Proof. Analogous to the proof for Proposition 6.27 as the same extensions are selected. □

6.4 Grounded semantics

An exception to the largely negative results before are the selection semantics based on the grounded semantics, which are equal to the grounded semantics as the introduced selection rules select at least one extension: the unique grounded extension.

Corollary 6.8 (Inheritance by $\mathrm{CBE}_{\mathrm{Gr},\gamma}$ and $\mathrm{SBE}_{\mathrm{Gr},\gamma'}$)**.** *Any Copeland-based extensions semantics* $\mathrm{CBE}_{\mathrm{Gr},\gamma}$ *and Simpson-based extensions semantics* $\mathrm{SBE}_{\mathrm{Gr},\gamma'}$ *where* $\gamma' \in \{\mathrm{nonatt}, \mathrm{strdef}, \mathrm{delarg}, \mathrm{delarg}\}$ *inherit all principles.*

Proof. Follows from Propositions 3.1, 3.3, and 5.7 in connection with Lemma 5.1. □

7 Refined acceptance

In the previous sections, we have seen several selection semantics, either Copeland-based or Simpson-based extensions semantics, each of which refines a given semantics by selecting extensions and reducing the number of extensions yielded. For the classical semantics and comparison criteria, these selection semantics are *acceptance-consistent*, thereby ensuring the refined semantics do not contradict the credulous and sceptical acceptance from the extensions of the underlying semantics. However, two problems remain:

Grounded	All. abst.	Crash resist.	Non-int.	W. dir.	S.-dir.	Dir.
CBE						
nonatt	✓	✓	✓	✓	✓	✓
strdef	✓	✓	✓	✓	✓	✓
delarg	✓	✓	✓	✓	✓	✓
delatt	✓	✓	✓	✓	✓	✓
SBE						
nonatt	✓	✓	✓	✓	✓	✓
strdef	✓	✓	✓	✓	✓	✓
delarg	✓	✓	✓	✓	✓	✓
delatt	✓	✓	✓	✓	✓	✓

Table 12: Inheritance of principles satisfied by the grounded semantics for grounded-based selection semantics, part 2

Grounded	Tightness	Conf.-sens.	Com-closure
CBE			
nonatt	✓	✓	✓
strdef	✓	✓	✓
delarg	✓	✓	✓
delatt	✓	✓	✓
SBE			
nonatt	✓	✓	✓
strdef	✓	✓	✓
delarg	✓	✓	✓
delatt	✓	✓	✓

Table 13: Inheritance of principles satisfied by the grounded semantics for grounded-based selection semantics, part 3

1. *Ineffective selection.* Even if a selection semantics potentially refines the underlying semantics, it can be *practically ineffective* for some or all AFs. In particular, the selection semantics may be either
   - *deficiently selective*, selecting too few extensions (e. g. in the trivial case of the equality shown by Proposition 6.36), or
   - *overly selective*, selecting too many extensions.

   It may not always be possible to find a practically effective selection semantics among those that have been introduced in this thesis, rendering the refinement unsatisfactory.

2. *Loss of principles.* As shown in Section 6, some principles (e. g. directionality and non-interference) are often not inherited by a selection semantics unless in trivial cases.

In this section, we address the first problem by combining selection semantics, deriving new methods to refine acceptance. We then outline techniques to mitigate the second problem of principle loss.

7.1 Ineffective selection

One of three straightforward solutions to the problem of ineffective selection is a *serial selection approach.* This approach requires the concept of a *selection specifier*, which encapsulates the method by which extensions are selected. It abstracts the selection method from a fixed underlying semantics, enabling the application to any semantics.

Definition 7.1 (Selection specifier)**.** For a comparison criterion γ, *selection specifier* ϕ is a function on the set of all semantics, such that for any semantics σ, the value

$$\phi(\sigma) = \mathrm{SEL}_{\sigma,\gamma}$$

is a selection semantics with respect to σ and comparison criterion γ.

Example 7.1 (Selection specifier)**.** *From a selection specifier ϕ such that $\phi(\mathrm{Co}) = \mathrm{CBE}_{\mathrm{Co,nonatt}}$ is a Copeland-based extensions semantics, we receive for any semantics σ the Copeland-based extensions semantics*

$$\phi(\sigma) = \mathrm{CBE}_{\sigma,\mathrm{nonatt}}.$$

A *serial selection semantics* provides a way to deal with *deficiently selective* semantics, by applying selection semantics *in sequence*. Each step uses a selection specifier to reduce the set of extensions further, potentially counteracting deficiencies.

Definition 7.2 (Serial selection semantics)**.** For a semantics σ and tuple $\Phi = (\phi_1, \ldots, \phi_n)$ of n selection specifiers, the *serial selection semantics (SSS)* (*w*ith respect to semantics σ and selection specifiers Φ), $\mathrm{SSS}^{\Phi}_{\sigma}$, is a selection semantics and defined by

$$\mathrm{SSS}^{\Phi}_{\sigma} = \phi_n \circ \cdots \circ \phi_1 \circ \sigma$$

where $\circ$ denotes function composition. The final specification step $\phi_n(\cdots)$ determines the semantics and comparison criterion with respect to which $\mathrm{SSS}^{\Phi}_{\sigma}$ is a selection semantics.

Example 7.2 (Serial selection semantics)**.** *Let F be an AF, σ a semantics, and $\Phi = (\phi_1, \phi_2)$ a tuple of selection specifiers such that $\phi_1(\sigma) = \mathrm{CBE}_{\sigma,\mathrm{nonatt}}$ and $\phi_2(\sigma) = \mathrm{SBE}_{\sigma,\mathrm{delarg}}$. Then, we receive serial selection semantics*

$$\begin{aligned}\mathrm{SSS}^{\Phi}_{\sigma} &= \phi_2 \circ \phi_1 \circ \sigma \\ &= \phi_2 \circ \mathrm{CBE}_{\sigma,\mathrm{nonatt}} \\ &= \mathrm{SBE}_{\mathrm{CBE}_{\sigma,\mathrm{nonatt}},\mathrm{delarg}},\end{aligned}$$

which is defined recursively with selection semantics such that we have $(\phi_2 \circ \phi_1 \circ \sigma)\big(F\big) \subseteq (\phi_1 \circ \sigma)\big(F\big) \subseteq \sigma(F)$, potentially reducing the number of extensions.

Another way to combine ineffective selection semantics is to select the intersection or union of their extensions, applying them *in parallel*. This helps address the scenario where some semantics are deficiently selective or overly selective selective, respectively.

An *intersection-based semantics (IBS)* provides a method to deal with *deficiently selective* semantics by selecting ubiquitous extensions of selection semantics.

Definition 7.3 (Intersection-based semantics)**.** For an AF F, semantics σ and tuple $\Gamma = (\mathrm{SEL}^{1}_{\sigma,\gamma_1}, \ldots, \mathrm{SEL}^{n}_{\sigma,\gamma_n})$ of n selection semantics, the *intersection-based semantics* (*with respect to semantics σ and selection semantics* Γ), $\mathrm{IBS}^{\Gamma}_{\sigma}$, is a semantics and defined by

$$\mathrm{IBS}^{\Gamma}_{\sigma}(F) = \bigcap_{i=1}^{n} \mathrm{SEL}^{i}_{\sigma,\gamma_i}(F).$$

Remark 7.1. An argument is sceptically accepted iff it is contained in every extension (Definition 2.7). Analogously, an intersection-based semantics selects an extension iff it yielded by every considered selection semantics.

Example 7.3 (Intersection-based semantics)**.** *Let F be an AF, σ a semantics, and $\Gamma = (\mathrm{SEL}^{1}_{\sigma,\gamma_1}, \mathrm{SEL}^{2}_{\gamma_2})$ a tuple of selection semantics, such that $\sigma(F) = \{\mathcal{E}_1, \mathcal{E}_2, \mathcal{E}_3\}$, $\mathrm{SEL}^{1}_{\sigma,\gamma_1}(F) = \{\mathcal{E}_1, \mathcal{E}_2\}$, and $\mathrm{SEL}^{2}_{\sigma,\gamma_2}(F) = \{\mathcal{E}_2, \mathcal{E}_3\}$. Then, we receive intersection-based semantics*

$$\begin{aligned}\mathrm{IBS}^{\Gamma}_{\sigma} &= \mathrm{SEL}^{1}_{\sigma,\gamma_1}(F) \cap \mathrm{SEL}^{2}_{\sigma,\gamma_2}(F) \\ &= \{\mathcal{E}_1, \mathcal{E}_2\} \cap \{\mathcal{E}_2, \mathcal{E}_3\} \\ &= \{\mathcal{E}_2\},\end{aligned}$$

which is defined with selection semantics such that we have $\mathrm{IBS}^{\Gamma}_{\sigma}(F) \subseteq \mathrm{SEL}^{i}_{\sigma,\gamma_i}(F) \subseteq \sigma(F)$ *for* $1 \leq i \leq 2$*, reducing the number of extensions.*

A *union-based semantics (UBS)* provides a method to deal with *overly selective* semantics by selecting all extensions of selection semantics.

Definition 7.4 (Union-based semantics)**.** For an AF F, semantics σ and tuple $\Gamma = (\mathrm{SEL}^{1}_{\sigma,\gamma_1}, \ldots, \mathrm{SEL}^{n}_{\sigma,\gamma_n})$ of n selection semantics, the *union-based semantics* (*with respect to semantics* σ *and selection semantics* Γ), $\mathrm{UBS}^{\Gamma}_{\sigma}$, is a semantics and defined by

$$\mathrm{UBS}^{\Gamma}_{\sigma}(F) = \bigcup_{i=1}^{n} \mathrm{SEL}^{i}_{\sigma,\gamma_i}(F).$$

Remark 7.2. An argument is credulously accepted iff it is contained in at least one extension (Definition 2.7). Analogously, a union-based semantics selects an extension iff it yielded by at least one considered selection semantics.

Example 7.4 (Union-based semantics)**.** *Let* F *be an AF,* σ *a semantics, and* $\Gamma = (\mathrm{SEL}^{1}_{\sigma,\gamma_1}, \mathrm{SEL}^{2}_{\gamma_2})$ *a tuple of selection semantics, such that* $\sigma(F) = \{\mathcal{E}_1, \mathcal{E}_2, \mathcal{E}_3\}$*,* $\mathrm{SEL}^{1}_{\sigma,\gamma_1}(F) = \{\mathcal{E}_1\}$*, and* $\mathrm{SEL}^{2}_{\sigma,\gamma_2}(F) = \{\mathcal{E}_2\}$*. Then, we receive union-based semantics*

$$\begin{aligned}\mathrm{UBS}^{\Gamma}_{\sigma} &= \mathrm{SEL}^{1}_{\sigma,\gamma_1}(F) \cup \mathrm{SEL}^{2}_{\sigma,\gamma_2}(F) \\ &= \{\mathcal{E}_1\} \cup \{\mathcal{E}_2\} \\ &= \{\mathcal{E}_1, \mathcal{E}_2\},\end{aligned}$$

increasing the number of extensions.

Remark 7.3. Notably, IBS and UBS cannot be reduced to a single comparison criterion such that they are not selection semantics in the sense of Definition 3.2.

Principles satisfied by the underlying selection semantics are not generally inherited by serial selection, intersection-based, or union-based semantics. However, a property that stands out is the maintenance of non-emptiness and in consequence acceptance-consistency, which is "inherited" by the union-based semnatics.

Corollary 7.1 (Acceptance-consistency of $\mathrm{UBS}^{\Gamma}_{\sigma}$)**.** *Any union-based semantics* $\mathrm{UBS}^{\Gamma}_{\sigma}$*, such that* $\Gamma = (\mathrm{SEL}^{1}_{\gamma_1}, \ldots, \mathrm{SEL}^{n}_{\gamma_n})$ *and for all* $1 \leq i \leq n$*,* $\mathrm{SEL}^{i}_{\gamma_i}$ *maintains non-emptiness, maintains non-emptiness and is acceptance-consistent.*

Proof. Follows from Definition 7.4 in connection with Definition 3.6. □

7.2 Loss of principles

Serial selection, intersection-based, and union-based semantics can remedy ineffective selection, but as other selection semantics, they do not guarantee the inheritance of principles satisfied by the underlying semantics. Two basic approaches can be summarised:

- *Principle-aware comparison criteria.* One may define a new comparison criterion that penalises pairs of extensions that could violate a desired principle. This way, the resulting selection semantics is guided to select only conforming extensions.
- *Principle repair.* Similar to minimal repairs of inconsistencies in AFs as suggested by Ulbricht and Baumann [UB19], we can minimally modify the set of selected extensions to repair principle violations. This may require either adding otherwise unselected extensions or removing selected extensions. Trivially, inheritance can be ensured by adding all unselected extensions in cases of violation.

A design of these methods depends on the principles that are considered and is left for future work.

8 Discussion

This thesis has highlighted two central challenges in refining extension-based semantics through selection. First, even though Copeland- and Simpson-based extensions semantics can improve credulous or sceptical acceptance from the underlying semantics, their effectiveness and acceptance-consistency vary. While Copeland-based extensions semantics consistently ensure acceptance-consistency, Simpson-based extensions semantics only do so under total (or reflexive) comparison relations. Hence, important theoretical guarantees can fail in practice.

Second, principle inheritance is almost always lost after restricting a set of extensions via selection criteria, except in trivial cases such as the grounded semantics (Corollary 6.8). The counterexamples throughout Section 6 show that many principles, especially directionality and its weaker forms, are not inherited. However, it was not always possible to evaluate inheritance of a principle for the Copeland- and Simpson-based extensions semantics, since there is no established methodology to search for counterexamples. Consequently, where manual examination did not suffice, counterexamples could sometimes be found by evaluation of random and small AFs (up to 20 arguments). This automatic search limits the otherwise infinite search space. Therefore, these inheritance statuses remain unclear, although no other indicators of inheritance could be found. Generally, the literature on extension selection is sparse, offering limited theoretical basis and as such guidance.

Techniques such as serial selection, intersection-based, and union-based semantics (Section 7) can mitigate ineffective selection scenarios but do not resolve the loss of principles. Further research into more sophisticated selection methods, or repair methods that restore desirable properties, could address this gap. Furthermore, it should be noted that this thesis covers extension selection only in connection with pairwise comparison. Accounting for more than two extensions in a single comparison operation could prove to be more suitable for complex real-world applications and better capture the holistic context used by principles, especially in regard to inheritance.

In sum, selecting extensions by pairwise comparison remains promising as a way to balance credulous and sceptical acceptance. Nonetheless, without principle-conforming comparison criteria or repairs methods, inheritance is typically not given, and thus further work is needed before extension selection methods can be widely applied in a principle-based manner.

9 Conclusion

This thesis has investigated the refinement of extension-based semantics by selecting subsets of extensions according to rules from voting theory. The results demonstrate how Copeland- and Simpson-based extensions semantics can satisfy desirable properties such as acceptance-consistency under certain conditions, but also that principle inheritance often fails. Although serial selection, intersection-based, and union-based semantics (Section 7) can partially address deficiently or overly selective semantics, these methods neither guarantee inheritance.

Moreover, research on extension selection was found to be sparse, offering limited theoretical basis. This thesis seeks to fill that gap, offering several formal description of extension selections that can be studied in further research. In particular, not all introduced selection semantics could be evaluated for every principle.

Despite these remaining problems, the potential benefits of refinement remain clear. By restricting the extensions of a semantics to a smaller yet meaningful subset, it is possible to balance credulous and sceptical acceptance in a direction carefully controlled by a comparison criterion. Next steps include designing new comparison criteria tailored to ensure inheritance of principles, integrating repair methods for principle violations after selection, and establishing a foundation for real-world applications, e. g. in regard to computation of selected extensions.

Ultimately, the results here show that refining acceptance by selecting certain extensions is a promising but still evolving concept. Future work will have to balance strong theoretical guarantees like principle inheritance with practical considerations such as computational efficiency and effective implementation in real-world settings, possibly providing a more nuanced framework to choose selection semantics than the binary principle-based approach.

References

[ABG+17] Katie Atkinson, Pietro Baroni, Massimiliano Giacomin, Anthony Hunter, Henry Prakken, Chris Reed, Guillermo Simari, Matthias Thimm, and Serena Villata. Towards Artificial Argumentation. *AI Magazine*, 38(3):25–36, October 2017. Number: 3.

[ABN13] Leila Amgoud and Jonathan Ben-Naim. Ranking-Based Semantics for Argumentation Frameworks. In Weiru Liu, V. S. Subrahmanian, and Jef Wijsen, editors, *Scalable Uncertainty Management*, pages 134–147, Berlin, Heidelberg, 2013. Springer.

[BCE+16] Felix Brandt, Vincent Conitzer, Ulle Endriss, Jérôme Lang, and Ariel D. Procaccia. *Handbook of Computational Social Choice*. Cambridge University Press, New York, 2016.

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

[BGGT18] Pietro Baroni, Dov Gabbay, Massimilino Giacomin, and Leendert van der Torre, editors. *Handbook of Formal Argumentation*. College Publications, London, England, 2018.

[Dun95] Phan Minh Dung. On the acceptability of arguments and its fundamental role in nonmonotonic reasoning, logic programming and n-person games. *Artificial Intelligence*, 77(2):321–357, September 1995.

[KMV15] Sébastien Konieczny, Pierre Marquis, and Srdjan Vesic. On Supported Inference and Extension Selection in Abstract Argumentation Frameworks. In Sébastien Destercke and Thierry Denoeux, editors, *Symbolic and Quantitative Approaches to Reasoning with Uncertainty*, pages 49–59, Cham, 2015. Springer International Publishing.

[SRT+21] Kenneth Skiba, Tjitze Rienstra, Matthias Thimm, Jesse Heyninck, and Gabriele Kern-Isberner. Ranking Extensions in Abstract Argumentation. In *Proceedings of the Thirtieth International Joint Conference on Artificial Intelligence*, pages 2047–2053, Montreal, Canada, August 2021. International Joint Conferences on Artificial Intelligence Organization.

[UB19] Markus Ulbricht and Ringo Baumann. If Nothing Is Accepted – Repairing Argumentation Frameworks. *Journal of Artificial Intelligence Research*, 66:1099–1145, December 2019.

[vDTV18] Leendert van Der Torre and Srdjan Vesic. The Principle-Based Approach to Abstract Argumentation Semantics. In *Handbook of Formal Argumentation*. February 2018.