A Principle-Based Analysis of Weakly Admissible Semantics

Jeremie DAUPHIN [a] Tjitze RIENSTRA [b] and Leendert VAN DER TORRE [a,c]
[a] *University of Luxembourg*
[b] *University of Koblenz-Landau*
[c] *Zhejiang University*

Abstract. Baumann, Brewka and Ulbricht recently introduced weak admissibility as an alternative to Dung's notion of admissibility, and they use it to define weakly preferred, weakly complete and weakly grounded semantics of argumentation frameworks. In this paper we analyze their new semantics with respect to the principles discussed in the literature on abstract argumentation. Moreover, we introduce two variants of their new semantics, which we call *qualified* and *semi-qualified* semantics, and we check which principles they satisfy as well. Since the existing principles do not distinguish our new semantics from the ones of Baumann *et al.*, we also introduce some new principles to distinguish them. Besides selecting a semantics for an application, or for algorithmic design, our new principle-based analysis can also be used for the further search for weak admissibility semantics.

Keywords. Formal argumentation, abstract argumentation, principle-based analysis, weak admissibility

1. Introduction

There are three classes of abstract argumentation semantics, which can be illustrated on their behaviour on odd and even cycles in the three argumentation frameworks in Figure 1. Roughly, in Dung's admissibility-based semantics [8], the maximal extensions may contain arguments of even-length cycles but no arguments of odd-length cycles, unless the odd-length cycle is attacked by some accepted argument. For example, the set of preferred extensions of F_1 is $\{\emptyset\}$, of F_2 is $\{\{d,g\},\{e,g\}\}$, and of F_3 is $\{\emptyset\}$. In naive-based semantics like the CF2 semantics [2], the extensions typically include arguments that are only attacked by self-attacking arguments, such as the argument b in F_1 below. In addition, odd-length cycles and even-length cycles are treated similarly in the sense that

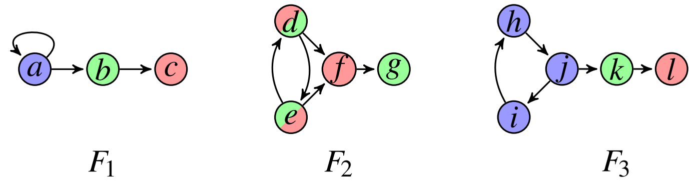

Figure 1. Three argumentation frameworks

naive extensions may also contain arguments from odd-length cycles, for example one of h, i or j in F_3. Under the *weakly admissible* semantics, recently introduced by Baumann, Brewka and Ulbricht (BBU) [3], the set of weakly preferred extensions of F_1 is $\{\{b\}\}$, of F_2 is $\{\{d,g\},\{e,g\}\}$, and of F_3 is $\{\{k\}\}$. These extensions are visualised in Figure 1: green arguments are in all the extensions, red arguments are not in the extensions and attacked by an argument in the extension (called out) and blue arguments are not in the extension and not out (called undecided). The arguments colored both red and green are in some but not all extensions.

At the moment of writing of this paper, the BBU semantics was only compared to existing semantics by their behaviour on a few examples, but a more systematic comparison was lacking. Just before sending the camera-ready version of this paper, we received a paper of the same authors [4] which will appear at a conference this year. That paper contains a table with a principle-based analysis, though most of the principles introduced and discussed in that paper are quite different from the ones in this paper, and thus that paper is complementary to this one.

The weakly admissible semantics are defined in terms of a recursive definition, which makes the analysis more difficult. Whereas many different variants of admissibility-based and naive-based semantics have been introduced and analysed, thus far only weakly complete, weakly grounded and weakly preferred semantics have been introduced from the third category. We therefore raise the following questions in this paper:

1. How do BBU's weak-admissibility based semantics compare to the existing semantics? That is, which principles does they satisfy?
2. Which other semantics can be defined along the lines of weak admissibility, giving the same results for the frameworks in Figure 1?
3. How can these new semantics be distinguished from the weak-admissibility based semantics? Which principles do these new semantics satisfy?

In general, one of the main purposes of axiomatisation in formal logic is to understand the logic with an intuitively understandable small set of principles. In proposing axioms, care should be taken to ensure that each axiom is sufficiently reasonable and sufficiently independent of others. Ideally, there should be some degree of philosophical motivation behind them. However, in the principle-based analysis of abstract argumentation, thus far the focus has been on the use of principles to differentiate semantics, and to assist computational techniques using decomposibility. Concerning the first question, Baumann *et al*. show that the weakly grounded extensions are not necessarily unique, and the principle-based analysis in this paper shows that weakly complete semantics does not satisfy directionality or SCC decomposibility.

The new semantics we define in this paper are based on SCC decomposability principles due to Baroni *et al*. [2]. This approach has previously been used to define the CF2 and Stage2 semantics. When we consider only the extensions of the framework in Figure 1, a recursive procedure comes to mind. As we show in detail later, if we use the scheme introduced by Baroni *et al*. to define CF2 (where all arguments are qualified), and we replace the base function with Dung's semantics, we get a procedure which gives the same extensions as BBU's weakly preferred semantics for the argumentation frameworks in Figure 1.

The layout of this paper is as follows. In Section 2 we introduce the *reduct admissibility* principle. We also repeat the definitions of weak admissibility and the related semantics, and we illustrate them using some new examples. In Section 3 we introduce the *semi-qualified admissibility* principle and we show that it is not satisfied by the BBU semantics. In Section 4 we introduce *weak SCC decomposability* and show it is not satisfied by the BBU semantics. In Section 5 we introduce our new two variants of semi-admissible semantics and we show which principles they satisfy. In Section 6 we discuss related and future work.

2. Weak Admissibility And The Reduct Principle

In this section we recall the definitions of the recently introduced weak-admissibility based semantics [3], and we introduce the *reduct admissibility* principle to characterise these semantics. Some notation: Given an AF $F = (A, \rightarrow)$ and $E \subseteq A$, we use E^+ to denote the set $\{b \in A \mid a \rightarrow b, a \in E\}$, use $F{\downarrow_E}$ to denote the set $(E, \rightarrow \cap E \times E)$, and use F^E to denote the E-*reduct* of F, which is the set $F{\downarrow_{E^*}}$ where $E^* = A \setminus (E \cup E^+)$. So the E-reduct of an argumentation framework F consists of the arguments that are neither in E nor attacked by E, and the attacks between these arguments.

The notion of weak admissibility weakens the requirement that every argument is defended against every attacker. Whereas an admissible extension must defend every member from every attacker, a weakly admissible extension does not require defence from attackers that do not appear in any weakly admissible set of F^E.

Definition 1. *[3] Let $F = (A, \rightarrow)$ be an AF. The set of* weakly admissible *sets of F is denoted $ad^w(F)$ and defined by $E \in ad^w(F)$ if and only if E is conflict-free (i.e., there are no $x, y \in E$ such that $x \rightarrow y$) and for every attacker y of E we have $y \notin \cup ad^w(F^E)$.*

Furthermore, a set E is said to *weakly defend* a set X if for every attacker y of X we have that either E attacks y, or y does not appear in E, does not appear in some weakly admissible set of F^E, and X is included in some weakly admissible set of F.

Definition 2. *[3] Let $F = (A, \rightarrow)$ be an AF. A set $E \subseteq A$* weakly defends *a set $X \subseteq A$ whenever, for every attacker y of X, either E attacks y, or $y \notin \cup ad^w(F^E)$, $y \notin E$ and $X \subseteq X' \in ad^w(F)$.*

Weak defense is related to weak admissibility in the sense that every conflict-free set is weakly admissible if and only if it weakly defends itself [3]. The weakly complete, preferred and grounded semantics are defined as follows.

Definition 3. *[3] Let $F = (A, \rightarrow)$ be an AF and $E \subseteq A$. We say that E is:*

- *a* weakly complete *extension of F ($E \in co^w(F)$) iff $E \in ad^w(F)$ and for every X such that $E \subseteq X$ that is w-defended by E, we have $X \subseteq E$.*
- *a* weakly preferred *extension of F ($E \in pr^w(F)$) iff E is $\subseteq$-maximal in $ad^w(F)$.*
- *a* weakly grounded *extension of F ($E \in gr^w(F)$) iff E is $\subseteq$-minimal in $co^w(F)$.*

We give some examples to illustrate these definitions.

Example 1. *The AF visualized on the left in Fig. 2 consists of a cycle of length three with one self-attacking argument. While the unique complete extension is $\emptyset$, the unique weakly complete extension is $\{a_1\}$. Intuitively, since b is self-attacking, a_1 does not need*

Figure 2. On the left: A self-attacking argument inside a 3-cycle, with a single weakly complete extension $\{a_1\}$. On the right: A 2-3 cycle, with only weakly complete extension $\{a_3\}$.

to be defended from b and can therefore be accepted. Now consider the AF visualized on the right in Fig. 2, consisting of a combination of a 2 cycle with a 3 cycle. While this AF has two complete extensions $\emptyset$ and $\{a_3\}$ it has one weakly complete extension $\{a_3\}$. Intuitively, since a_2 is not in any extension we have that the empty set weakly defends a_3, and so there is no reason not to accept it. The last AF we discuss is two connected 3-cycles, visualized in Fig. 3. This AF demonstrates that the weakly grounded extension

Figure 3. Two connected 3-cycles with one extra argument. $gr^w(F) = \{\{b\}, \{a_1, d\}\}$.

is not unique. Since the empty set defends both $\{a_1\}$ and $\{b\}$, yet not both at the same time, both of these sets are weakly grounded.

We now define a weaker version of the admissibility principle based on the definition of ad^w, which we call reduct admissibility. The motivation for the reduct admissibility principle is taken directly from BBU's motivation of weak admissibility. We quote: "It is indeed important that a set of arguments defends itself. However, [...] isnt it sufficient to counterattack those arguments which have the slightest chance of being accepted?" [3].

Definition 4 (Reduct admissibility)**.** *We say that a semantics σ satisfies* reduct admissibility *iff for any argumentation framework $F = (A, \to)$, for every extension $E \in \sigma(F)$, we have that $\forall a \in E$, $(b, a) \in R$, we have $b \notin \bigcup \sigma(F^E)$.*

Proposition 1. *co^w, gr^w and pr^w satisfy reduct admissibility.*

Proof. For $\sigma \in \{co^w, gr^w, pr^w\}$, for all $E \in \sigma(F)$ we have that E is weakly admissible. So, for every attacker y of E, $y \notin \bigcup ad^w(F^E)$, and therefore also $y \notin \bigcup \sigma(F^E)$. $\square$

As a principle, reduct admissibility is a bit complex due to the use of the reduct. In the following section we define an alternative principle that formalizes the same idea without referring to the reduct.

3. The Semi-qualified Admissibility Principle

We now define the *semi-qualified admissibility* principle and determine which semantics satisfy it. We then focus on some of the principles already found in the literature [9] and investigate whether the BBU semantics satisfy them.

When looking at the reduct admissibility principle, we may ask why the acceptability of an attacker is judged based on the reduct, and not on the original framework itself. For the definition of a semantics, assessing the acceptability of attackers on the reduct allows for a recursive definition that is guaranteed to terminate for finite AFs. When looking at a principle, this concern disappears and we therefore provide the definition of a different principle, which we call *semi-qualified admissibility*. Semi-qualified admissibility states that an extension only needs to defend itself against attackers that appear in at least one extension of the same framework.

Definition 5 (Semi-Qualified admissibility)**.** *We say that a semantics σ satisfies* semi-qualified admissibility *iff for every argumentation framework $F = (A, \rightarrow)$ and every extension $E \in \sigma(F)$ we have that $\forall a \in E$, if $b \rightarrow a$ and $b \in \bigcup \sigma(F)$ then $\exists c \in E$ s.t. $c \rightarrow b$.*

Proposition 2. *co^w, gr^w and pr^w don't satisfy semi-qualified admissibility.*

Proof. Consider the AF F shown in Figure 3. Here, the set $\{a_1, d\}$ is a co^w, gr^w and pr^w extension of F. This extension is attacked by b and we also have $b \in \cup co^w(F)$ (similarly for gr^w and pr^w). However there is no $x \in \{a_1, d\}$ such that $x \rightarrow b$. □

One can easily see that our two new principles fail for CF2 and stage2 in a 3-cycle.

Proposition 3. *CF2 and stage2 don't satisfy reduct nor semi-qualified admissibility.*

The following definition introduces a number of well-known principles from the literature [9].

Definition 6. *A semantics σ satisfies the principle of:*

- admissibility *iff for every argumentation framework F, every $E \in \sigma(F)$ is conflict-free and classically defends itself in F;*
- naivety *iff for every argumentation framework F, for every $E \in \sigma(F)$, E is a $\subseteq$-maximal conflict-free set in F;*
- reinstatement *iff for every argumentation framework $F = (A, \rightarrow)$, for every $E \in \sigma(F)$ and $a \in A$ it holds that if E classically defends a then $a \in E$;*
- I-maximality *iff for every AF F, for every $E_1, E_2 \in \sigma(F)$, if $E_1 \subseteq E_2$ then $E_1 = E_2$;*
- allowing abstention *iff for every AF $F = (A, \rightarrow)$ and $a \in A$, if there exist $E_1, E_2 \in \sigma(F)$ s.t. $a \in E_1$ and $a \in E_2^+$, then there exists $E_3 \in \sigma(F)$ s.t. $a \notin E_3 \cup E_3^+$;*
- directionality *iff for every AF $F = (A, \rightarrow)$ and $S \subseteq A$ s.t. $S \cap (A \setminus S)^+ = \emptyset$, it holds that $\sigma(F\downarrow_S) = \{E \cap S \mid E \in \sigma(F)\}$.*

We now state some results regarding these principles for the semantics based on weak admissibility. Table 1 summarises our findings. Note that one can easily see that reduct and semi-qualified admissibility follow from admissibility.

Proposition 4. *co^w, gr^w and pr^w don't satisfy admissibility.*

Proof. Consider the AF $F = (\{a,b\}, \{(a,a),(a,b)\}$. We have $co^w(F) = gr^w(F) = pr^w(F) = \{\{b\}\}$, but $\{b\}$ does not classically defend itself from a. □

Proposition 5. *co^w, gr^w and pr^w don't satisfy naivety.*

Proof. Consider the AF $F = (\{a,b,c\}, \{(a,b),(b,c),(c,a)\})$. We have $co^w(F) = gr^w(F) = pr^w(F) = \{\emptyset\}$ while e.g. $\{a\}$ is conflict-free. □

	co	pr	CF2	st2	co^w	gr^w	pr^w	q-**co**	q-**gr**	q-**pr**	sq-**co**	sq-**pr**
Admissibility	✓	✓	×	×	×	×	×	×	×	×	×	×
Naivety	×	×	✓	✓	×	×	×	×	×	×	×	×
Reinst.	✓	✓	×	×	✓	✓	✓	✓	✓	✓	✓	✓
I-Max.	×	✓	✓	✓	×	✓	✓	×	✓	✓	×	✓
Allow. abs.	✓	×	×	×	×	×	×	×	✓	×	✓	×
Directionality	✓	✓	✓	✓	×	×	?	✓	✓	✓	✓	✓
Semi-qual. adm.	✓	✓	×	×	×	×	×	×	✓	×	✓	✓
Reduct adm.	✓	✓	×	×	✓	✓	✓	×	?	?	?	?
SCC Decomp.	✓	✓	✓	✓	×	×	?	✓	✓	✓	×	×
W-SCC Decomp.	✓	✓	✓	✓	?	?	?	✓	✓	✓	✓	✓

Table 1. Principles satisfied by the weak-admissibility and qualified and semi-qualified semantics, with complete, preferred, CF2 and stage2 for comparison.

Proposition 6. *co^w, gr^w and pr^w satisfy reinstatement.*

Proof. Follows from Proposition 5.9 from Baumann *et al*. [3]. □

Proposition 7. *co^w does not satisfy I-maximality.*[1]

Proof. Consider the AF $F = (\{a,b\},\{(a,b),(b,a)\})$. $co^w = \{\{a\},\{b\},\emptyset\}$, and $\emptyset \subseteq \{a\}$ but $\emptyset \neq \{a\}$. □

Proposition 8. *gr^w and pr^w satisfy I-maximality.*

Proof. By definition, every set in pr^w is a $\subseteq$-maximal weakly admissible set, therefore none is a strict subset of the other. Similarly, every set in gr^w is by definition a $\subseteq$-minimal weakly grounded extension, and therefore none is a strict subset of another. □

Proposition 9. *co^w, gr^w and pr^w do not satisfy allowing abstention.*

Proof. Consider the AF visualized in Fig. 3. $co^w = gr^w = pr^w = \{\{a_1,d\},\{b\}\}$, and $d \in \{b\}^+$, but there is no extension E_3 where $d \notin E_3 \cup E_3^+$. □

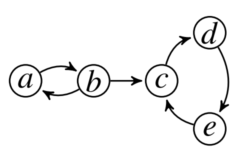

Figure 4. The weak complete and grounded semantics are not directional.

The following proposition answers two open questions of Baumann *et al.* [4].

Proposition 10. *co^w and gr^w are not directional.*

Proof. Consider the AF $F = (\{a,b,c,d,e\},\{(a,b),(b,a),(b,c),(c,d),(d,e),(e,c)\})$ visualized in Figure 4. Directionality would imply that

[1] In [4] it is mistakenly mentioned that gr^w does not satisfy I-maximality.

$$co^w(F\downarrow_{\{a,b\}}) = co^w(F)\downarrow_{\{a,b\}}.$$

However we have $co^w(F\downarrow_{\{a,b\}}) = \{\{a\},\{b\},\emptyset\}$ and $co^w(F) = \{\{a\},\{b,d\}\}$ and hence $co^w(F)\downarrow_{\{a,b\}} = \{\{a\},\{b\}\}$. Similarly, we have $gr^w(F\downarrow_{\{a,b\}}) = \{\emptyset\}$, but $gr^w(F)\downarrow_{\{a,b\}} = \{\{a\},\{b\}\}$. □

Open Question 1. *Is the weakly preferred semantics directional? Baumann et al. [4] answer this question affirmative, but they do not provide a proof.*

4. Decomposability Principles

We now discuss two additional principles. We will use these principles as a basis for the definition of two families of semantics in the next section. The first principle is *SCC decomposability*. This principle was introduced by Baroni *et al*. under the name *full decomposability w.r.t. SCC partitioning* [1].

SCC decomposability is defined for labelling-based semantics [6]. We first need to introduce some notation. A labelling L of an AF F is a function that maps each argument of F to a label $\mathtt{I}$ (in, or accepted), $\mathtt{O}$ (out, or rejected) or $\mathtt{U}$ (undecided). We use $\mathscr{L}(F)$ to denote the set of all possible labellings of F. A labelling-based semantics σ maps each AF F to a set $\mathscr{L}_\sigma(F) \subseteq \mathscr{L}(F)$. We denote the set of SCCs (strongly connected components) of F by $\mathscr{S}(F)$. Let $F = (A, \to)$ be an AF. An *outparent* of an SCC S of F is an argument $x \in A \setminus S$ such that $x \to y$ for some $y \in S$. We denote by $OP_F(S)$ the set of outparents of S. Given a labelling $L \in \mathscr{L}(F)$ we denote by $L\downarrow_S$ the restriction of L to S and, given a set $X \subseteq \mathscr{L}(F)$ of labellings, denote by $X\downarrow_S$ the set $\{L\downarrow_S \mid L \in X\}$.

The SCC decomposability principle states that the set of labellings of an AF F is decomposable into the product of the sets of labellings of each SCC of F, where the set of labellings of an SCC S is a function of the labels of the outparents of S. To formalise the principle we first define the notion of *AF with input*.

Definition 7. *An* AF with input *is a tuple* $(F, A_{in}, \to_{in}, L_{in})$ *where* $F = (A, \to)$ *is an AF;* A_{in} *a set of* input arguments *such that* $A \cap A_{in} = \emptyset$*;* $\to_{in} \subseteq A_{in} \times A$ *is an* input attack relation*; and* $L_{in} \in \mathscr{L}(A_{in})$ *is an* input labelling.

A semantics is SCC decomposable if it is *represented* by some *local function*. A local function is a function f that maps each AF with input to a set of labellings. A semantics σ is represented by a local function f if the set of σ labellings of every AF F coincides with the product of the labellings of each SCC of F as determined by f.

Definition 8. *A* local function f *assigns to every AF with input* $(F, A_{in}, \to_{in}, L_{in})$ *a set* $f(F, A_{in}, \to_{in}, L_{in}) \subseteq \mathscr{L}(F)$*. We say that* f represents *the semantics* σ *if for every AF* F*,*

$$L \in \mathscr{L}_\sigma(F) \leftrightarrow \forall S \in \mathscr{S}(F), L\downarrow S \in f(F\downarrow S, OP_F(S), \to \cap OP_F(S) \times S, L\downarrow OP_F(S)).$$

A semantics σ *is* SCC decomposable *if it is represented by some local function.*

Examples of semantics that are known to be SCC decomposable are the complete, grounded and preferred semantics. We denote by $f_{\mathbf{co}}$, $f_{\mathbf{gr}}$ and $f_{\mathbf{pr}}$ the local functions representing these semantics. Their definition can be found in [1]. As for the weak-admissibility based semantics, we observe that the weak complete and weak grounded semantics are not SCC decomposable. We first define a labelling-based version of these semantics in the usual way [6]: given an AF $F = (A, \rightarrow)$, define $Ext2Lab_F : 2^A \rightarrow \mathscr{L}(F)$ by $Ext2Lab_F(E)(x) = \mathtt{I}$, if $x \in E$; $Ext2Lab_F(E)(x) = \mathtt{O}$, if $y \rightarrow x$ for some $y \in E$; and $Ext2Lab_F(E)(x) = \mathtt{U}$, otherwise. We then define the labelling-based weak complete semantics co^w by $\mathscr{L}_{co^w}(F) = \{Ext2Lab_F(E) \mid E \in co^w(F)\}$ and define the labelling-based weak preferred pr^w and grounded gr^w similarly. It then holds that

Proposition 11. *The co^w and gr^w semantics are not SCC decomposable.*

Proof. We prove it for gr^w (co^w is similar). Consider the AFs $F_1 = (\{b,c\}, \{(b,b),(b,c)\})$ and $F_2 = (\{a,b,c\}, \{(a,b),(b,a),(b,c)\})$. We then have $\mathscr{L}_{gr^w}(F_1) = \{\{a : \mathtt{U}, b : \mathtt{I}\}\}$ and $\mathscr{L}_{gr^w}(F_2) = \{\{a_1 : \mathtt{U}, a_2 : \mathtt{U}, b : \mathtt{U}\}\}$ If the gr^w semantics is SCC decomposable then there must be a local function f_{gr^w} that represents gr^w. But then $f_{gr^w}(\{c\}, \{b\}, \{(b,c)\}, \{b\colon \mathtt{U}\})$ equals both $\{\{c\colon \mathtt{I}\}\}$ and $\{\{c\colon \mathtt{U}\}\}$, which is impossible. Hence, the gr^w semantics is not SCC decomposable. $\square$

Open Question 2. *Is the pr^w semantics SCC Decomposable?*

We now introduce a new principle called *weak SCC decomposability*. Like SCC decomposability, this principle states that the set of labellings of an AF F can be decomposed into the product of the sets of labellings of each SCC of F. The difference with SCC decomposability is that the set of labellings of an SCC S is a function not only of a particular labelling of the outparents of S, but also of the set of all other labellings that the outparents of S may receive. This provides extra information in how the labellings of an SCC are determined since, in addition to knowing the actual labels of the outparents, we also know how these arguments are labelled in other labellings. To define it we extend the notion of AF with input to that of AF with *total input* as follows.

Definition 9. *An* AF with total input *is a tuple $(F, A_{in}, \rightarrow_{in}, L_{in}, S_{in})$ where $F, A_{in}, \rightarrow_{in}$ and L_{in} are defined as in definition 7, $S_{in} \subseteq \mathscr{L}(A_{in})$, and $L_{in} \in S_{in}$. We call S_{in} the set of* total input labellings *and $L_{in} \in S_{in}$ the* actual input labelling.

We say that σ is weakly SCC decomposable if there exists a *weak local function* (i.e., a function that maps each AF with total input to a set of labellings) that represents σ. A weak local function represents a semantics σ if the set of σ labellings of every AF F coincides with the product of the labellings of each SCC $S \in \mathscr{S}(F)$ as a determined by the weak local function.

Definition 10. *A* weak local function *g assigns to every AF with total input $(F, A_{in}, \rightarrow_{in}, L_{in}, S_{in})$ a set $g(F, A_{in}, \rightarrow_{in}, L_{in}, S_{in}) \subseteq \mathscr{L}(F)$. A weak local function g represents a semantics σ whenever, for every AF F, $L \in \mathscr{L}_\sigma(F)$ if and only if*

$$\forall S \in \mathscr{S}(F), L{\downarrow}S \in g(F{\downarrow}S, OP_F(S), \rightarrow \cap OP_F(S) \times S, L{\downarrow}OP_F(S), \mathscr{L}_\sigma(F){\downarrow}OP_F(S)).$$

A semantics σ is weakly SCC decomposable *if some weak local function represents σ.*

Note that SCC decomposability implies weak SCC decomposability but that the reverse does not hold. In the next section we use the weak SCC decomposability principle to define new semantics. For the weak admissibiity-based semantics we have:

Open Question 3. *Are the weakly complete, weakly grounded and weakly preferred semantics weakly SCC decomposable? Our conjecture is that they are.*

5. Qualified and Semi-Qualified Semantics

We now define two new families of semantics. They are neither admissible nor naive and they represent two new ways to deal with propagation of undecidedness. The first family are the *qualified* semantics. A qualified semantics builds on the SCC decomposability principle and is based on applying the local function of any SCC decomposable semantics with one change: in determining the labellings of an SCC S, the label $\mathtt{U}$ for an outparent x of S is treated like the label $\mathtt{O}$. This means that, if an argument x is attacked by an $\mathtt{U}$-labelled argument y, and if x and y are elements of different SCCs, then y is still qualified for acceptance (i.e., may still be labelled $\mathtt{I}$).

Definition 11. *Let σ be an SCC decomposable semantics. Let f_σ denote the local function that represents σ. We define the* qualified σ *(or q-σ) semantics as the semantics represented by the local function $f_{q\text{-}\sigma}$ defined by*

$$f_{q\text{-}\sigma}((A,\rightarrow),A_{in},\rightarrow_{in},L_{in}) = f_\sigma((A,\rightarrow),A_{in},\rightarrow_{in},L'_{in})$$

where $L'_{in}(x) = \mathtt{I}$ if $L_{in}(x) = \mathtt{I}$, and $L'_{in}(x) = \mathtt{O}$, if $L_{in}(x) = \mathtt{O}$ or $L_{in}(x) = \mathtt{U}$.

We now focus on three examples of qualified semantics, namely the qualified complete (q-**co**), qualified grounded (q-**gr**), and qualified preferred (q-**pr**) semantics. Note that, by definition, all these semantics are SCC decomposable.

Example 2. *Consider the argumentation frameworks shown in Figure 1. The AF F_1 has a unique q-**co**, q-**gr** and q-**pr** labelling, namely $\{a\colon \mathtt{U}, b\colon \mathtt{I}, c\colon \mathtt{O}\}$. The AF F_2 has three q-**co** labellings, namely $\{d\colon \mathtt{I}, e\colon \mathtt{O}, f\colon \mathtt{O}, g\colon \mathtt{I}\}$, $\{d\colon \mathtt{O}, e\colon \mathtt{I}, f\colon \mathtt{O}, g\colon \mathtt{I}\}$, and $\{d\colon \mathtt{U}, e\colon \mathtt{U}, f\colon \mathtt{I}, g\colon \mathtt{O}\}$, where the first two are also the q-**pr** labellings and the last one is also the q-**gr** labelling. The AF F_3 has a unique q-**co**, q-**gr** and q-**pr** labelling, namely $\{h\colon \mathtt{U}, i\colon \mathtt{U}, j\colon \mathtt{U}, k\colon \mathtt{I}, l\colon \mathtt{O}\}$.*

This example shows that the qualified **co**/**gr**/**pr** and weak **co**/**gr**/**pr** semantics of the three AFs in Figure 1 coincide for the AFs F_1 and F_3 but not for F_2. In F_2, the set $\{f\}$ (which corresponds to the labelling $\{d\colon \mathtt{U}, e\colon \mathtt{U}, f\colon \mathtt{I}, g\colon \mathtt{O}\}$) is not weakly admissible because it does not defend itself from d, while d does appear in some weakly admissible set of the $\{f\}$-reduct of F_2. To capture this intuition we define a second family of semantics, which builds on the weak SCC decomposability principle. It is based on applying the local function of any SCC decomposable semantics with the following change: in determining the labellings of an SCC S, the label $\mathtt{U}$ for an outparent x of S is treated like the label $\mathtt{O}$, *but only if there is no other labelling of the outparents of S where x is labelled* $\mathtt{I}$. This means that, if an argument x is attacked by an $\mathtt{U}$-labelled argument y, and if x and y are elements of different SCCs, and there is no other labelling in which y is labelled $\mathtt{I}$, then x may still be labelled $\mathtt{I}$. We call the resulting semantics *semi-qualified*.

Definition 12. *Let σ be an SCC decomposable semantics. Let f_σ denote the local function that represents σ. We define the* semi-qualified σ *(or sq-σ) semantics as the semantics represented by the weak local function $g_{sq-\sigma}$ defined by*

$$g_{sq-\sigma}((A,\rightarrow),A_{in},\rightarrow_{in},L_{in},S_{in}) = g_\sigma((A,\rightarrow),A_{in},\rightarrow_{in},L'_{in})$$

where $L'_{in}(x) = \mathtt{I}$, if $L_{in}(x) = \mathtt{I}$; $L'_{in}(x) = \mathtt{O}$, if $L_{in}(x) = \mathtt{O}$; $L'_{in}(x) = \mathtt{O}$, if $L_{in}(x) = \mathtt{U}$ and there is no $L \in S_{in}$ such that $L(x) = \mathtt{I}$; and $L'_{in}(x) = \mathtt{U}$, if $L_{in}(x) = \mathtt{U}$ and there is some $L \in S_{in}$ such that $L(x) = \mathtt{I}$.

Any semi-qualified semantics is, by definition, weakly SCC decomposable. Furthermore, note that, for a unique status semantics σ (such as the grounded semantics) the qualified σ and semi-qualified σ semantics coincide.

Example 3. *The sq-**co**, sq-**gr** and sq-**pr** labellings of the AFs F_1 and F_3 shown in Figure 1 are the same as the q-**co**, q-**gr** and q-**pr** labellings (see Example 2). The sq-**co** labellings of F_2 are different from the q-**co** labellings. The sq-**co** labellings of F_2 are $\{d\colon \mathtt{I}, e\colon \mathtt{O}, f\colon \mathtt{O}, g\colon \mathtt{I}\}$, $\{d\colon \mathtt{O}, e\colon \mathtt{I}, f\colon \mathtt{O}, g\colon \mathtt{I}\}$, and $\{d\colon \mathtt{U}, e\colon \mathtt{U}, f\colon \mathtt{U}, g\colon \mathtt{U}\}$, where the first two are also the q-**pr** labellings and the last one is also the q-**gr** labelling.*

Note that the semi-qualified labellings in the example above coincide with the weak-admissibility based extensions. Thus, they provide an alternative approach to achieve weak-admissibility like behaviour. They are not equivalent, however. In particular, the semi-qualified complete, grounded and preferred semantics are different in how they evaluate isolated SCCs. For instance, consider the AF shown in Figure 3 but without the argument d. This AF consists of a single SCC and has only one semi-qualified complete (and hence grounded and preferred) labelling in which all arguments are undecided.

Table 1 includes an overview of principles satisfied by the (semi)-qualified complete, grounded and preferred semantics. We omit the *sq*-**gr** semantics, which is equivalent to the *q*-**gr** semantics. Failure of admissibility and naivety is demonstrated by Examples 2 and 3. The same holds for failure of allowing abstention under the *q*-**pr** and *sq*-**pr** semantics and I-maximality under the *q*-**co** and *sq*-**co** semantics. Satisfaction of reinstatement under all semantics follows easily, and so does satisfaction of I-maximality under the *q*-**pr** and *q*-**pr** semantics. The *q*-**gr** semantics trivially satisfies allowing abstention and I-maximality. Non-interference and Directionality follow from weak SCC decomposability together with the property that a local function returns a non-empty set of labellings for all possible inputs, which holds for the local functions that we use. Finally, allowing abstention does not hold under the *q*-**co** semantics (see the argument b in the AF F_2 in Example 2). We now consider the remaining principles and state a number of open questions at the end.

Proposition 12. *sq-**co** satisfies allowing abstention.*

Proof. (Sketch) We show that we can transform an AF F into an AF F' such that $\mathscr{L}_{sq\text{-}\mathbf{co}}(F) = \mathscr{L}_{\mathbf{co}}(F')$. Let $S_1,\ldots,S_n$ be an ordering SCCs of F such that if a directed path from S_i to S_j exists, then $i < j$. Define A_i and $\rightarrow_i$ by $A_0 = \emptyset$, $\rightarrow_0 = \emptyset$, and for $i > 0$, $A_i = A_{i-1} \cup S_i$ and $\rightarrow_i = \rightarrow_{i-1} \cup (\rightarrow \cap (S_i \times S_i)) \cup (\rightarrow \cap (X_i \times S_i))$, where $X_i = \{x \in A_{i-1} \mid \exists L \in \mathscr{L}_{sq\text{-}\mathbf{co}}((A_{i-1},\rightarrow_{i-1})), L(x) = \mathtt{I}\}$. We then have $\mathscr{L}_{sq\text{-}\mathbf{co}}(F) = \mathscr{L}_{\mathbf{co}}((A_n,\rightarrow_n))$. Since **co** satisfies allowing abstention it thus follows that *sq*-**co** does too. $\square$

Proposition 13. *q-**pr** does not satisfy semi-qualified admissibility.*

Proof. The AF $F = (\{a,b,c,d,e,f\},\{(a,b),(b,a),(b,c),(c,d),(d,e),(e,c),(d,f)\})$ has a q-**pr** labelling $L = \{(a,\mathtt{I}),(b,\mathtt{O}),(c,\mathtt{U}),(d,\mathtt{U}),(e,\mathtt{U}),(f,\mathtt{I})\}$, which corresponds to the extension $E = \{a,f\}$. We have $d \to E$. Therefore, according to semi-qualified admissibility, since there is no $x \in E$ such that $x \to d$ it must hold that d is not in any q-**pr** extension of F. However, this is false, because F has a s-**pr** labelling $\{(a,\mathtt{O}),(b,\mathtt{I}),(c,\mathtt{O}),(d,\mathtt{I}),(e,\mathtt{O}),(f,\mathtt{O})\}$, which corresponds to the extension $\{b,d\}$. $\square$

Proposition 14. *q-**co** does not satisfy reduct or semi-qualified admissibility.*

Proof. Consider the AF $F = (\{a,b,c\},\{(a,b),(b,a),(b,c)\})$. This AF has a q-**co** extension $E = \{c\}$. Since $b \to E$, according to reduct admissibility, b may not be in any q-**co** extension of F^E. But $F^E = \{(\{a,b\},\{(a,b),(b,a)\})$ has a q-**co** extension $\{b\}$. This violates reduct admissibility. Semi-qualified admissibility is violated similarly. $\square$

Proposition 15. *sq-**co**, sq-**gr**, sq-**pr** and q-**gr** satisfy semi-qualified admissibility.*

Proof. Let $F = (A,\to)$ be an AF and let $L \in \mathscr{L}_{sq\text{-}\mathbf{co}}(F)$. Let $E = \{x \in A \mid L(x) = \mathtt{I}\}$. Suppose $x \to y$ for some $y \in E$. Then either $L(x) = \mathtt{O}$, which implies that there is a z such that $z \to x$ and $x \in E$; or $L(x) = \mathtt{U}$, which implies (via Definition 12) that there is no $L' \in \mathscr{L}_{sq\text{-}\mathbf{co}}(F)$ such that $L'(x) = \mathtt{I}$ and hence no sq-**co** extension E' of F such that $x \in E'$. Hence the sq-**co** semantics, and thus also the sq-**co** and sq-**co** semantics, satisfy semi-qualified admissibility, and so does q-**gr**, which coincides with sq-**gr**. $\square$

Open Question 4. *Do q-**pr**, q-**gr**, sq-**co** and sq-**pr** satisfy reduct admissibility?*

6. Related and Future Work

The principle-based approach was initiated by Baroni *et al*. to distinguish argumentation semantics, and then taken up by various researchers widening the scope of the "principle-based approach." For example, Doutre and colleagues have been promoting a principle-based approach to abstract argumentation, and the SESAME software [5] is an achievement in this respect. Motivated by empirical cognitive studies on argumentation semantics, Cramer and van der Torre [7] have introduced a new naive-based argumentation semantics called SCF2. A principle- based analysis shows that it has two distinguishing features:

1. If an argument is attacked by all extensions, then it can never be used in a dialogue and therefore it has no effect on the acceptance of other arguments. They call it Irrelevance of Necessarily Rejected Arguments.
2. Within each extension, if none of the attackers of an argument is accepted and the argument is not involved in a paradoxical relation, then the argument is accepted. They define paradoxicality as being part of an odd cycle, and they call this principle Strong Completeness Outside Odd Cycles.

They argue that these features together with the findings from empirical cognitive studies make SCF2 a good candidate for an argumentation semantics that corresponds well to what humans consider a rational judgment on the acceptability of arguments.

As mentioned in the introduction, just before sending the camera-ready version of this paper, we received a paper [4] with another principle-based analysis for weak admissibility, though most of the principles introduced and discussed in that paper are quite different from the ones in this paper, and thus that paper is complementary to this one.

A topic for further research is the development of a labeling-based semantics for weak admissibility, and the weakly complete, weakly grounded and weakly preferred semantics. We are also looking for labeling-based definitions of the new semantics introduced in this paper. We believe that labeling-based semantics can also be instrumental in the search for new argumentation semantics.

Acknowledgement

The work was partially supported by the German Research Foundation (DFG) project CAML (grant KE 1686/3-1).

References

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

[2] Pietro Baroni, Massimiliano Giacomin, and Giovanni Guida. Scc-recursiveness: a general schema for argumentation semantics. *Artificial Intelligence*, 168(1-2):162–210, 2005.

[3] Ringo Baumann, Gerhard Brewka, and Markus Ulbricht. Revisiting the foundations of abstract argumentation - semantics based on weak admissibility and weak defense. In *The Thirty-Fourth AAAI Conference on Artificial Intelligence, AAAI 2020, New York, USA, February 7 - February 12, 2020*. AAAI Press, 2020.

[4] Ringo Baumann, Gerhard Brewka, and Markus Ulbricht. Comparing weak admissibility semantics to their dung-style counterparts reduct, modularization, and strong equivalence in abstract argumentation. In *Principles of Knowledge Representation and Reasoning: Proceedings of the 17th International Conference, KR 2020, Rhodes, Greece, September 12-18, 2020*, 2020 (To Appear).

[5] Philippe Besnard, Sylvie Doutre, Van Hieu Ho, and Dominique Longin. SESAME - A system for specifying semantics in abstract argumentation. In Matthias Thimm, Federico Cerutti, Hannes Strass, and Mauro Vallati, editors, *Proceedings of the First International Workshop on Systems and Algorithms for Formal Argumentation (SAFA), Potsdam, Germany, September 13, 2016*, volume 1672 of *CEUR Workshop Proceedings*, pages 40–51. CEUR-WS.org, 2016.

[6] Martin W. A. Caminada and Dov M. Gabbay. A logical account of formal argumentation. *Studia Logica*, 93(2-3):109–145, 2009.

[7] Marcos Cramer and Leendert van der Torre. SCF2 - an argumentation semantics for rational human judgments on argument acceptability. In Christoph Beierle, Marco Ragni, Frieder Stolzenburg, and Matthias Thimm, editors, *Proceedings of the 8th Workshop on Dynamics of Knowledge and Belief (DKB-2019) and the 7th Workshop KI & Kognition (KIK-2019), Kassel, Germany, September 23, 2019*, volume 2445 of *CEUR Workshop Proceedings*, pages 24–35. CEUR-WS.org, 2019.

[8] 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):321–358, 1995.

[9] Leendert van der Torre and Srdjan Vesic. The Principle-Based Approach to Abstract Argumentation Semantics. In Pietro Baroni, Dov Gabbay, Massimiliano Giacomin, and Leendert van der Torre, editors, *Handbook of Formal Argumentation*, chapter 12, pages 2735–2778. College Publications, London, 2018.