4

Abstract Argumentation Frameworks and Their Semantics

Pietro Baroni, Martin Caminada, Massimiliano Giacomin

ABSTRACT. The current chapter presents an overview on the state of the art of Dung's abstract argumentation frameworks and their semantics, covering both some of the most influential literature proposals and some general issues concerning semantics definition and evaluation. As to the former point the chapter reviews Dung's original notions of complete, grounded, preferred, and stable semantics, as well as a variety of notions subsequently proposed in the literature namely, naïve, semi-stable, ideal, eager, stage, $\mathcal{CF}2$, and stage2 semantics, considering both the extension-based and the labelling-based approaches with respect to their definitions. As to the latter point the chapter analyzes the notions of argument justification and skepticism comparison and discusses semantics agreement.

1 Introduction

This chapter is devoted to the formalism of *abstract argumentation frameworks* introduced by Dung [1995]. This formalism is based on the idea that arguments are defeasible entities which may attack each other and whose acceptance is subject to evaluation. In presence of conflicts, an argument cannot be accepted just because it exists: its acceptance depends on the existence of possible counter arguments, that can then themselves be attacked by counter arguments, and so on. Formally, an argumentation framework is represented as a directed graph in which the arguments are represented as nodes and the attack relation is represented by the arrows. Given such a graph, one is naturally led to examine the question of which set(s) of arguments can be accepted: answering this question corresponds to defining an *argumentation semantics*. Various proposals have been formulated in this respect, and in the current chapter we will describe some of the mainstream approaches. Before entering the technical presentation, however, some important general considerations are worth introducing. As sketched above, the formalism of argumentation frameworks is exclusively centered on the notion of attack between arguments and on the evaluation of argument acceptability, based on the intuition that the existence of attacks prevents all arguments to be accepted together. In this formal context, arguments are deprived of all their features apart their identity: their origin, structure and any other characteristics differentiating them are abstracted away, leaving room only to the property of attacking (or being attacked by) their homogeneous (from the abstract point of view) peers. This extreme simplification of an otherwise complex and articulated phenomenon like argumentation is a key factor for understanding the huge interest that the study of this formalism has attracted, as well as its boundaries. Theoretical cleanness and wide applicability are among the formidable strengths following from the simplicity of the formalism. Argumentation frameworks allow the investigation of notions and properties (far from being trivially simple, by the way) which are purely related to the existence of attacks without being obfuscated or complicated by the many accidental or complementary properties of the entities involved in the attack themselves. In this sense Dung's theory can be regarded as an attack (or conflict) calculus, which has been initially formulated for the need of dealing with attacks between arguments but then stands on its own feet, even independently of the original interpretation in argumentative terms. This yields a powerful generality: as shown in Dung's original paper, several more specific formal settings can be regarded (and better understood) as special cases of argumentation frameworks, in areas ranging from nonmonotonic reasoning to game theory. Remarkably, some of these settings are only loosely related to the notion of argumentation. This has shown since the beginning that theoretical investigations of Dung's formalism enjoy a wide applicability across a variety of domains. It must also be observed, however, that dealing with attacks, while being crucial, is by no means sufficient to provide a formal counterpart of actual argumentation processes. In this sense, the temptation of considering abstract argumentation theory as a self-sufficient tool for formal argumentation should be regarded as an oversimplification. To avoid this risk, it is important to keep in mind that assessing argument acceptability in presence of attacks, which is the essence of Dung's theory, is only one specific (although important) aspect in formalizing argument-based reasoning and needs to be integrated and bridged with other formal components. As an example of how to apply Dung's theory within a broadened setting, consider the use of abstract argumentation theory for the purpose of nonmonotonic inference from a knowledge base.

In this context, one can distinguish three steps (see Figure 1). First of all, one would use an underlying knowledge base to generate a set of arguments and determine in which ways these arguments attack each other (step 1). The result is then an argumentation framework, to be represented as a directed graph in which the internal structure of the arguments, as well as the nature of the attack relation have been abstracted away. Based on this argumentation framework, the next step is to determine the sets of arguments that can be accepted, using a pre-defined criterion corresponding to an argumentation semantics (step 2). After the set(s) of accepted arguments have been identified, one then has to identify the set(s) of accepted conclusions (step 3), for which there exist various approaches.

As an example of how things work, suppose one applies the argumentation process in the context of logic programming. In particular, suppose the knowledge base consists of a logic program P of the following form.

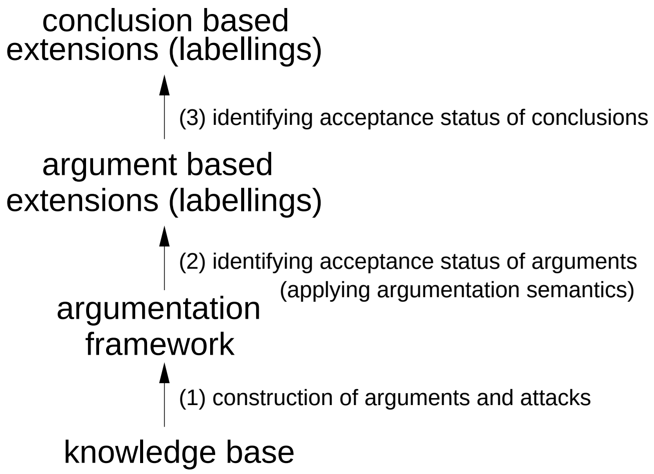

Figure 1. Argumentation for inference

$$
\begin{array}{ll}
b \leftarrow c, \texttt{not}\ a & a \leftarrow \texttt{not}\ b \\
p \leftarrow c, d, \texttt{not}\ p & p \leftarrow \texttt{not}\ a \\
c \leftarrow d & d \leftarrow
\end{array}
$$

In that case, following for instance the approach of [Wu *et al.*, 2009] where arguments consist of trees of rules and attack is based on weak negation, one can construct the argumentation framework shown in Figure 2.

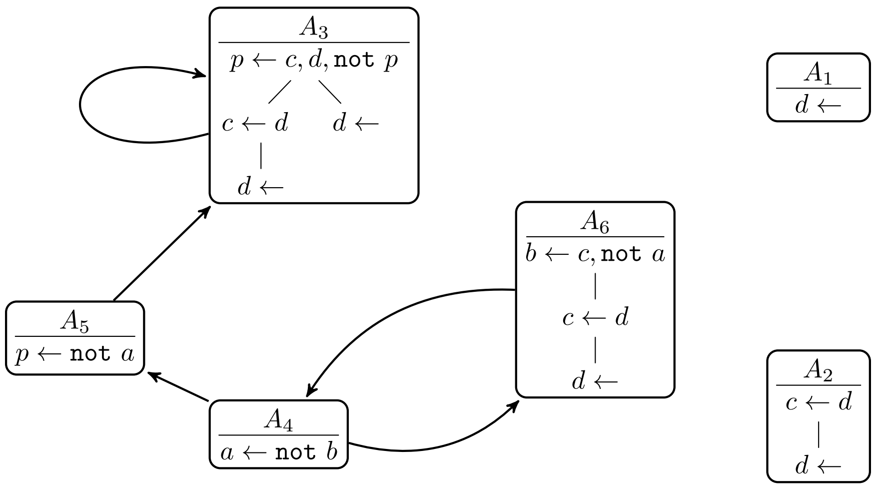

Figure 2. The argumentation framework built from the logic program P

This argumentation framework has exactly one stable extension[1] of arguments: $\{A_1, A_2, A_5, A_6\}$. This extension of arguments yields the following extension of conclusions: $\{d, c, p, b\}$. It has been proved [Dung, 1995] that when applying stable semantics at step 2, the overall result (extensions of conclusions) is precisely the same as when applying stable model semantics [Gelfond and Lifschitz, 1988; Gelfond and Lifschitz, 1991] to the original logic program. In a similar way, the three-step argumentation process can also be applied to simulate other logic programming semantics. We refer to [Dung, 1995; Wu *et al.*, 2009; Caminada *et al.*, 2015] for an overview.

As mentioned before, one of the strengths of the argumentation approach is that it turns out to be powerful enough to model not just logic programming but a whole range of formalisms, including Default Logic [Dung, 1995; Caminada *et al.*, 2012] and Nute's Defeasible Logic [Governatori *et al.*, 2004]. Other scholars have subsequently used the argumentation approach to specify their own formalisms for non-monotonic entailment, like ASPIC+ [Modgil and Prakken, 2014], ABA [Toni, 2014] and logic-based argumentation [Gorogiannis and Hunter, 2011].

Overall, the argumentation process described above leads to a number of questions:

1. *What is the content of the knowledge base with which arguments are constructed?* Different argument-based formalisms start with different types of knowledge bases. In the case of argument-based logic programming [Dung, 1995; Wu *et al.*, 2009; Caminada *et al.*, 2015], like described above, the knowledge base is simply a logic program. In the case of logic-based argumentation [Gorogiannis and Hunter, 2011] the knowledge base consists of propositions. In the case of ABA [Toni, 2014], it contains rules and assumptions. In the case of ASPIC+ [Modgil and Prakken, 2014], it contains rules and formulas, as well as a preference ordering, to determine argument strength. In spite of their differences, what unifies these formalisms is that each of them can be seen as applying the argumentation process of Figure 1. The reader can find details in chapters 6, 7, 8 and 9 of this volume.

2. *Given a knowledge base, how to precisely construct the associated argumentation framework?* Even for the same type of knowledge base, there can be several ways of constructing the associated argumentation framework, each with their own advantages and disadvantages. Details will be provided in chapters 6, 7, and 9.

3. *Once the argumentation framework has been constructed, how to determine which arguments to accept and reject?* This is the key question to be studied in the current chapter. Our aim is to provide an overview of

[1] A stable extension attacks precisely those arguments that are not in it. We refer to Section 3.6 for details.

criteria for argument acceptance, as have been stated in the literature. Hence, the current chapter focuses on step 2 of the argumentation process in Figure 1.

4. *How to make sure the overall outcome makes sense?* Abstract argumentation theory selects arguments purely on the basis of the topology of the graph, without looking at their actual contents. In particular, these contents can have a logical form (as is the case for instance in ABA [Toni, 2014], ASPIC+ [Modgil and Prakken, 2014] and logic-based argumentation [Gorogiannis and Hunter, 2011]). The question is how to make sure that the overall conclusions yielded by the formalism are consistent or satisfy any other desirable property. This is a crucial research issue: chapter 15 will present some key desirable properties as well as some of the approaches for satisfying them.

To complete this introduction, it is worth mentioning again that while inference from a knowledge base is an important domain for abstract argumentation theory, a wider range of applications can be found in the literature, ranging from decision making [Amgoud, 2009] to topics like coalition formation and the stable marriage problem [Dung, 1995]. It is also fair to mention that in the literature there are also formalisms for argument-based reasoning, like *DeLP* [García and Simari, 2004], which adopt alternative approaches with respect to abstract argumentation theory for the assessment of argument acceptance.

The remaining part of this chapter is structured as follows. First, in Section 2 we formally describe the notion of an argumentation framework and present some relevant basic concepts.

Then in Section 3 we present some relatively well-known and well-established argumentation semantics, both in terms of argument extensions and in terms of argument labellings. In Section 4 we provide a comprehensive treatment of the notions of argument justification and skepticism, including skepticism comparison between the reviewed semantics, while in Section 5 we discuss the issue of semantics agreement. Finally Section 6 provides a brief summary and concludes the chapter.

2 Basic concepts

Central to the theory of abstract argumentation is the notion of an *argumentation framework*, which, as mentioned in Section 1, is essentially a directed graph in which the arguments are represented by the nodes and the attack relation is represented by the arrows[2]. Given the tutorial nature of this chapter, we keep the presentation simple by restricting ourselves to finite argumentation frameworks, while briefly mentioning non-finite frameworks where appropriate. The reader may refer to chapter 17 for a coverage of properties of infinite argumentation frameworks.

[2]In Dung's theory, attack is a one-to-one relationship, which deviates from earlier work of, for instance, [Vreeswijk, 1993] which is centered around the notion of *collective attack*, meaning that a set of arguments is collectively attacking another argument.

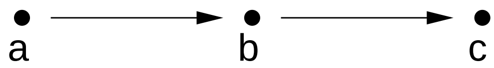

Figure 3. A simple argumentation framework

Definition 2.1 *An* argumentation framework *is a pair* $\langle Ar, att \rangle$ *in which* Ar *is a finite set of arguments and* $att \subseteq Ar \times Ar$.

We say that argument $a \in Ar$ *attacks* argument $b \in Ar$ (or that a is an *attacker* of b) iff $(a, b) \in att$. If $\mathcal{A}rgs \subseteq Ar$ and $a \in Ar$ then we say that a attacks $\mathcal{A}rgs$ iff there exists $b \in \mathcal{A}rgs$ such that a attacks b. Likewise, we say that $\mathcal{A}rgs$ attacks a iff there exists $b \in \mathcal{A}rgs$ such that b attacks a. For $a \in Ar$ then we write a^- for $\{b \mid (b, a) \in att\}$ and a^+ for $\{b \mid (a, b) \in att\}$. Likewise, for $\mathcal{A}rgs \subseteq Ar$ we write $\mathcal{A}rgs^-$ for $\{b \mid \exists a \in \mathcal{A}rgs : (b, a) \in att\}$ and $\mathcal{A}rgs^+$ for $\{b \mid \exists a \in \mathcal{A}rgs : (a, b) \in att\}$. All these notions refer to a given argumentation framework, which is left implicit in the relevant notation for the sake of simplicity and conciseness.

We will also need to consider the restriction of an argumentation framework to a subset of its arguments.

Definition 2.2 *Given an argumentation framework* $AF = \langle Ar, att \rangle$ *and a set of arguments* $\mathcal{A}rgs \subseteq Ar$*, the* restriction *of* AF *to* $\mathcal{A}rgs$*, denoted as* $AF{\downarrow}_{\mathcal{A}rgs}$ *is the argumentation framework* $\langle \mathcal{A}rgs, att \cap (\mathcal{A}rgs \times \mathcal{A}rgs) \rangle$.

An argumentation framework encodes, through the attack relation, the existing conflicts within a set of arguments. It is then interesting to identify the conflict outcomes, which, roughly speaking, means determining which arguments should be accepted (let's say, "survive the conflict") and which arguments should be rejected (let's say, "are defeated in the conflict"), according to some reasonable criterion.

Consider for instance the argumentation framework depicted in Figure 3. Which arguments are able to survive the conflict? Is there only one possibility or are there several solutions available? While the reader may resort to her/his personal intuition to devise a specific answer in this simple case, it appears that a well-defined systematic method is needed to deal with the case of arbitrarily complex argumentation frameworks: such a formal method to identify conflict outcomes for any argumentation framework is called *argumentation semantics*.

Two main approaches to the definition of argumentation semantics are available in the literature: the *labelling*-based approach and the *extension*-based approach.

The idea underlying the *labelling*-based approach is to give each argument a label. A sensible (though not the only possible) choice for the set of labels is: `in`, `out` or `undec`, where the label `in` means the argument is accepted, the label `out` means the argument is rejected and the label `undec` means one abstains from an opinion on whether the argument is accepted or rejected. Each argument then gets exactly one label. In Figure 3, one might start assigning the label `in` to argument a, as it does not receive attacks, then derive that the argument b should be `out`. Then the attack from b to c can be considered ineffective or, in other words, it can be said that a defends c against b and one can assume that c should be `in` in turn. While this labeling may sound reasonable, other choices are, at least in principle, available: e.g. one might assign all arguments the label `in`, but this seems incompatible with the existence of conflicts among them, or one might assign all arguments the label `undec`, but this seems excessively cautious at least as far as the unattacked argument a is concerned. Thus, a specific labelling-based argumentation semantics provides a way to select "reasonable" labellings among all the possible ones, according to some criterion embedded in its definition.

The idea underlying the *extension*-based approach is to identify sets of arguments, called extensions, which can survive the conflict together and thus represent collectively a reasonable position an autonomous reasoner might take. To illustrate how one could use an incremental procedure for extension construction, in Figure 3 one might start including the argument a, as it does not receive attacks, then exclude the argument b, and then assume that c should be included in turn, ending up with the extension $\{a, c\}$. Also in this case other choices are available, at least in principle: e.g. one might consider the extension $\{a, b, c\}$, but (again) this seems incompatible with the existing conflicts among arguments, or one might consider the empty set as extension, but this seems excessively cautious since at least a seems to deserve inclusion in any extension. Thus, a specific extension-based argumentation semantics provides a way to select "reasonable" sets of arguments among all the possible ones, according to some criterion embedded in its definition.

Let us now turn to the formal counterpart of the notions exemplified above.

A generic labelling assigns to each argument of an argumentation framework a label taken from a predefined set.

Definition 2.3 *Let $AF = \langle Ar, att \rangle$ be an argumentation framework and Λ a set of labels. A* $\Lambda-$labelling *is a total function $\mathcal{L}ab : Ar \longrightarrow \Lambda$. The set of all* $\Lambda-$labellings *of AF is denoted as $\mathfrak{L}(\Lambda, AF)$.*

A labelling-based semantics prescribes a set of labellings for any argumentation framework.

Definition 2.4 *Given an argumentation framework $AF = \langle Ar, att \rangle$ and a set of labels Λ, a* labelling-based semantics *σ associates with AF a subset of $\mathfrak{L}(\Lambda, AF)$, denoted as $\mathcal{L}_\sigma(AF)$.*

We will also need the notion of restriction of a labelling to a set of arguments.

Definition 2.5 *Given an argumentation framework $AF = \langle Ar, att \rangle$, a set of labels Λ, a $\Lambda-$*labelling *$\mathcal{L}ab$, and a set of arguments $\mathcal{A}rgs \subseteq Ar$, the* restriction *of $\mathcal{L}ab$ to $\mathcal{A}rgs$, denoted as $\mathcal{L}ab\downarrow_{\mathcal{A}rgs}$ is defined as $\mathcal{L}ab \cap (\mathcal{A}rgs \times \Lambda)$.*

In this chapter we focus on the case $\Lambda = \{\mathtt{in}, \mathtt{out}, \mathtt{undec}\}$, a sensible choice for Λ which has received considerable attention in the literature [Caminada, 2006a; Caminada, 2007a; Rahwan and Larson, 2008; Caminada and Gabbay, 2009; Caminada and Pigozzi, 2011; Rahwan and Tohmé, 2010]. An alternative approach can be found in [Jakobovits and Vermeir, 1999], where a four-valued labelling is considered. The idea of labelling can also be put in correspondence with the notion of status assignment in inference graphs [Pollock, 1995]. A first investigation of the connections between defeat status assignments and extensions in Dung's argumentation frameworks was provided in [Verheij, 1996].

We will implicitly assume the use of $\Lambda = \{\mathtt{in}, \mathtt{out}, \mathtt{undec}\}$, when the reference to the label set is omitted. In particular, given a labelling $\mathcal{L}ab$, we write $\mathtt{in}(\mathcal{L}ab)$ for $\{a \mid \mathcal{L}ab(a) = \mathtt{in}\}$, $\mathtt{out}(\mathcal{L}ab)$ for $\{a \mid \mathcal{L}ab(a) = \mathtt{out}\}$ and $\mathtt{undec}(\mathcal{L}ab)$ for $\{a \mid \mathcal{L}ab(a) = \mathtt{undec}\}$. A labelling can be represented as a set of pairs. For instance, the first labelling exemplified above for Figure 3 can be described as $\{(a, \mathtt{in}), (b, \mathtt{out}), (c, \mathtt{in})\}$. Sometimes we will also represent a labelling $\mathcal{L}ab$ as the triple $(\mathtt{in}(\mathcal{L}ab), \mathtt{out}(\mathcal{L}ab), \mathtt{undec}(\mathcal{L}ab))$. The same labelling for Figure 3 can thus be represented as $(\{a, c\}, \{b\}, \emptyset)$.

Turning to the extension-based approach, since an extension is just a set of arguments, the definition of extension-based semantics is quite simple and does not require preliminary notions.

Definition 2.6 *An* extension-based semantics σ *associates with any argumentation framework $AF = \langle Ar, att \rangle$ a subset of 2^{Ar}, denoted as $\mathcal{E}_\sigma(AF)$.*

Some observations about the relations between the labelling and extension-based approaches are worth remarking. First, as set membership can be formulated in terms of a simple binary labelling, e.g. with $\Lambda = \{\in, \notin\}$, the extension-based approach can be regarded as a special case of the general labelling-based approach. The latter is therefore more general, while the former, probably due to its simplicity, has received by far more attention in previous literature.

Considering the three-valued labelling we focus on in this chapter, a correspondence with the extension-based approach can be drawn, so that a semantics based on this labelling can be turned into an extension-based one through a simple mapping. In fact, given a labelling of an AF, the labels `in` can be understood as identifying the members of an extension. This kind of correspondence can be easily identified in the example concerning Figure 3 described above and is formally expressed by the following definitions.

Definition 2.7 *Given an argumentation framework $AF = \langle Ar, att \rangle$ and a labelling $\mathcal{L}ab$ the corresponding set of arguments $\mathtt{Lab2Ext}(\mathcal{L}ab)$ is defined as $\mathtt{Lab2Ext}(\mathcal{L}ab) = \mathtt{in}(\mathcal{L}ab)$.*

Definition 2.8 *Given an argumentation framework $AF = \langle Ar, att \rangle$ and a labelling-based semantics σ, the set of extensions corresponding to $\mathcal{L}_\sigma(AF)$ is given by $\mathcal{E}_\sigma(AF) = \{\mathtt{Lab2Ext}(\mathcal{L}ab) \mid \mathcal{L}ab \in \mathcal{L}_\sigma(AF)\}$.*

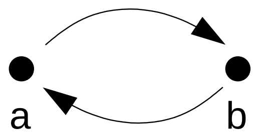

Figure 4. An argumentation framework with mutual attack

On the other hand, given a set of arguments E it is possible to define a corresponding three-valued labelling by distinguishing the arguments belonging to E, those attacked by some member of E, and those which neither belong to E nor are attacked by E. As this correspondence is well-defined only if E satisfies some basic conditions, we defer its formal definition to Section 3.1 (Definition 3.6).

We now introduce some notions which are common to both approaches.

First it can be noted that both approaches encompass (in general) the existence of a set of alternatives (either labellings or extensions) for a single argumentation framework. It may be the case, however, that a semantics σ is defined so that a univocal outcome is prescribed for each argumentation framework (formally for any argumentation framework AF, $|\mathcal{L}_\sigma(AF)| = 1$ or $|\mathcal{E}_\sigma(AF)| = 1$). In this case, the semantics is said to belong to the *unique-status* (or *single-status*) approach, while in the general case it is said to belong to the *multiple-status* approach.

Consider the argumentation framework of Figure 4 representing a mutual attack. A unique-status approach may prescribe the $\{(a, \mathtt{undec}), (b, \mathtt{undec})\}$ labelling or analogously a single empty extension, corresponding to an explicit abstention from decision. On the other hand, a multiple-status approach may encompass the two alternative labellings $\{(a, \mathtt{in}), (b, \mathtt{out})\}$ and $\{(a, \mathtt{out}), (b, \mathtt{in})\}$ or analogously the extensions $\{a\}$ and $\{b\}$, corresponding to two opposite ways of solving the conflict.

As evident from the previous example, a semantics σ does not provide, in general, the "last word" about the status of an argument a. In fact σ may prescribe both a labelling where a is labelled `in` and another where a is labelled `out` (or, analogously, an extension including a and another one not). In the view of producing a synthetic evaluation for each argument, one has then to consider questions like "Is being `in` in all labellings significantly different from being `in` only in some of them?" or "If an argument is not `in` in all labellings should it being labelled `out` or `undec` in the remaining labellings make some difference?". Analogous questions may arise for the extension-based approach. It emerges that the assessment of a synthetic *justification status* for each argument of an argumentation framework is a further distinct (and not trivial) step after the identification of labellings or extensions. This will be dealt with in Section 4.

3 An Overview of Argumentation Semantics

In this section we provide an overview of some well-known argumentation semantics, starting from the very basic notion of "naïve semantics" and then discussing Dung's original concepts of complete, stable, preferred and grounded semantics [Dung, 1995], as well as the subsequently introduced ideal [Dung *et al.*, 2007], semi-stable [Verheij, 1996; Caminada, 2006b] and eager [Caminada, 2007b] semantics.[3] These semantics can be considered as mainstream, since they share a basic property called *admissibility* and have been subject to much study, including the specification of proof procedures and of properties regarding computational complexity. We also treat three additional semantics, namely stage [Verheij, 1996], $\mathcal{CF}2$ [Baroni *et al.*, 2005b] and stage2 [Dvořák and Gaggl, 2012b; Dvořák and Gaggl, 2012a; Dvořák and Gaggl, 2016] semantics. Unlike the other semantics considered in this chapter, stage, $\mathcal{CF}2$ and stage2 semantics are not admissibility-based, but they have quite unique characteristics that make them worthwhile to examine.

The presentations of the various semantics roughly follow a common line. First, the underlying intuitive idea is introduced, then the semantics formal definition is given according to both the labelling and the extension-based approach, and finally the presentation is completed by discussing illustrative examples and examining additional important formal properties and inter-semantics relationships. We do not deal with algorithmic or implementation issues in this chapter, as this matter is extensively treated in chapter 14. However, without any claim of providing an adequate coverage of the state of the art, in some places we will mention some algorithms and the relevant literature sources which, in our opinion, represent useful readings to get a further insight on the nature and behaviour of the considered semantics. Each semantics is denoted by a short abbreviation for easy reference. As for examples, the relatively simple ones provided in Figures 5-7 will be used as a common reference throughout this section, adding other more specific and/or complex ones where necessary. We invite the reader to give a look to Figures 5-7 in order to set up a "personal view" on how the conflict they encode might be resolved, and then comparing this view with those emerging from the various semantics proposals analyzed in the following. Before dealing directly with semantics we need however to examine (in the next subsection) two general properties, which underlie most of them, namely admissibility[4] and conflict-freeness.

3.1 Admissibility and conflict-freeness

To introduce the notion of admissibility let us start from a very simple principle: for every argument a one accepts (or rejects) one should be able to explain why

[3]Please notice that terms like preferred semantics or ideal semantics correspond to existing terminology in the literature and do not imply any value judgements.

[4]Admissibility has been introduced as a semantic property, not as a semantics in [Dung, 1995]. In the subsequent literature, however, the term admissible semantics has often been used. We will also refer to admissible semantics later in the chapter where convenient for presentation purposes.

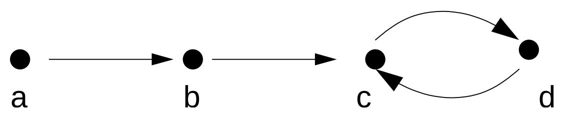

Figure 5. An argumentation framework with unidirectional and mutual attacks

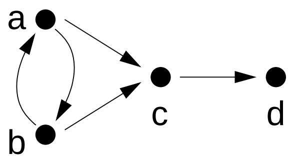

Figure 6. The case of “floating” acceptance

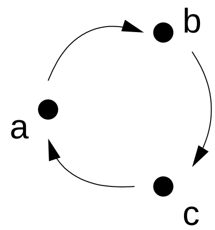

Figure 7. Cycle of three attacking arguments

it is accepted (or rejected), taking into account the acceptance or rejection of other arguments connected to a through the attack relation. This concept lends itself to slightly different, though converging, realisations in the labelling and in the extension-based approach.

In the labelling-based approach, assigning the `in` label to an argument a can be explained by having assigned the `out` label to all its attackers (or by a being attacked by no argument) so that a is not affected by any attack, while assigning the `out` label to a can be explained by having assigned the `in` label to one of its attackers, which enables a to be rejected.

This is expressed by the following definitions.

Definition 3.1 *Let $\mathcal{L}ab$ be a labelling of argumentation framework $\langle Ar, att\rangle$.*

- *An* `in`*-labelled argument is said to be* legally `in` *iff all its attackers are labelled* `out`*.*
- *An* `out`*-labelled argument is said to be* legally `out` *iff it has at least one attacker that is labelled* `in`

Definition 3.2 *Let $AF = \langle Ar, att\rangle$ be an argumentation framework. An* admissible labelling *is a labelling $\mathcal{L}ab$ where each* `in`*-labelled argument is legally* `in` *and each* `out`*-labelled argument is legally* `out`*.*

Note that, according to this definition, for any argumentation framework a labelling where all arguments are `undec` is admissible. Let us now examine admissible labellings in the reference examples. Considering Figure 5, it is evident that a, having no attackers, can only be labelled legally `in` or `undec`. Considering the latter case, b can only be labelled `undec`, which implies that c cannot be legally `in`. If c is labelled `undec` then d is `undec` too, otherwise c is labelled `out` entailing that d is labelled `in`. This yields two admissible labellings, the trivial one $(\emptyset, \emptyset, \{a, b, c, d\})$ and $(\{d\}, \{c\}, \{a, b\})$. The case where a is labelled `in` leaves two alternatives for b. If b is labelled `undec` we have the same options as above for c and d yielding the two additional admissible labellings $(\{a\}, \emptyset, \{b, c, d\})$ and $(\{a, d\}, \{c\}, \{b\})$. Finally if b is labelled `out`, three alternatives are left open for c and d: they can be both labelled `undec` or c can be legally labelled `in` if d is labelled `out` and vice versa, yielding other three labellings: $(\{a\}, \{b\}, \{c, d\})$, $(\{a, c\}, \{b, d\}, \emptyset)$, $(\{a, d\}, \{b, c\}, \emptyset)$.

In Figure 6, with a similar reasoning as in the previous example it can be noted that a and b can be both labelled `undec` or one `in` and the other `out`. The first case yields only the trivial labelling $(\emptyset, \emptyset, \{a, b, c, d\})$, in the other cases c may be labelled `undec`, yielding d `undec`, or `out` leaving for d both the options `undec` and `in`. Altogether there are seven admissible labellings whose enumeration is left to the reader (see Table 2 later).

In Figure 7 no admissible labellings besides the trivial one $(\emptyset, \emptyset, \{a, b, c\})$ are possible, since no argument defends itself and every argument attacks the

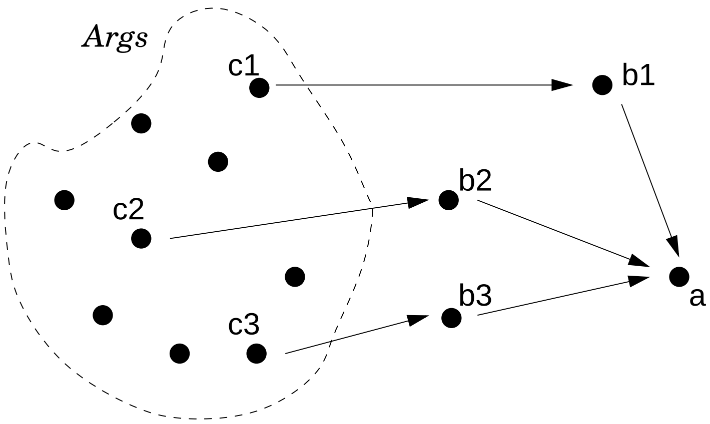

Figure 8. $\mathcal{A}rgs$ defends argument a

argument which would defend it: for instance a would need a defense against c but attacks its potential defender b, and similarly for the other two arguments.

Turning now to the extension-based approach, the inclusion of an argument a in an extension E can be supported by the fact that E rules out all the attackers of a by in turn attacking them (if any). To put it in other words, E "defends" a. This is formalized in the following definitions.

Definition 3.3 *Let $AF = \langle Ar, att \rangle$ be an argumentation framework and $\mathcal{A}rgs \subseteq Ar$. The set $\mathcal{A}rgs$* defends[5] *$a \in Ar$ iff $\forall b \in a^- \exists c \in \mathcal{A}rgs : c$ attacks b. The function $\mathrm{F}_{AF} : 2^{Ar} \longrightarrow 2^{Ar}$ such that $\mathrm{F}_{AF}(\mathcal{A}rgs) = \{a \mid \mathcal{A}rgs$ defends $a\}$ is called the characteristic function of AF.*

An example of defense is given in Figure 8. Here we have an argument a that has three attackers: b_1, b_2 and b_3. $\mathcal{A}rgs$ defends a because it attacks all these attackers.

Having introduced the notion of defense, a basic requirement for a set of arguments is the capability to defend all its elements. It is however natural to also require that the set of arguments features some sort of "internal coherence": no conflict should be allowed within a set of arguments which are considered able to survive the conflict *together*. This leads to the definition of conflict-free set.

Definition 3.4 *Let $AF = \langle Ar, att \rangle$ be an argumentation framework and $\mathcal{A}rgs \subseteq Ar$. The set $\mathcal{A}rgs$ is* conflict-free *iff $\neg\exists a, b \in \mathcal{A}rgs : a$ attacks b.*

Note that this definition also rules out sets containing *self-attacking* (also called *self-defeating*) arguments (in the case $a = b$).

[5]The original terminology in [Dung, 1995] was that an argument a is *acceptable* w.r.t. a set of arguments $\mathcal{A}rgs$. However, the more intuitive expression that an argument a is *defended* by a set of arguments $\mathcal{A}rgs$ is commonly used in the literature and we prefer it.

Admissible labellings	Admissible sets
$\{(a,\mathtt{undec}),(b,\mathtt{undec}),(c,\mathtt{undec}),(d,\mathtt{undec})\}$	$\emptyset$
$\{(a,\mathtt{undec}),(b,\mathtt{undec}),(c,\mathtt{out}),(d,\mathtt{in})\}$	$\{d\}$
$\{(a,\mathtt{in}),(b,\mathtt{undec}),(c,\mathtt{undec}),(d,\mathtt{undec})\}$	$\{a\}$
$\{(a,\mathtt{in}),(b,\mathtt{undec}),(c,\mathtt{out}),(d,\mathtt{in})\}$	$\{a,d\}$
$\{(a,\mathtt{in}),(b,\mathtt{out}),(c,\mathtt{undec}),(d,\mathtt{undec})\}$	$\{a\}$
$\{(a,\mathtt{in}),(b,\mathtt{out}),(c,\mathtt{in}),(d,\mathtt{out})\}$	$\{a,c\}$
$\{(a,\mathtt{in}),(b,\mathtt{out}),(c,\mathtt{out}),(d,\mathtt{in})\}$	$\{a,d\}$

Table 1. Admissibile labellings and sets in the example of Figure 5

An admissible set [Dung, 1995] is required to be both internally coherent and able to defend its elements.

Definition 3.5 *Let $AF = \langle Ar, att\rangle$ be an argumentation framework. A set $\mathcal{A}rgs \subseteq Ar$ is called an* admissible set *iff $\mathcal{A}rgs$ is conflict-free and $\mathcal{A}rgs \subseteq \mathrm{F}_{AF}(\mathcal{A}rgs)$. The set of admissible sets of AF is denoted as $\mathcal{AS}(AF)$.*

As evident from this definition, the empty set is admissible for any argumentation framework. Apart from this trivial case, let us examine conflict-free and admissible sets in the reference examples. Considering Figure 5, one can observe that the non empty conflict-free sets are $\{a\}$, $\{b\}$, $\{c\}$, $\{d\}$, $\{a,c\}$, $\{a,d\}$, $\{b,d\}$. Among them, $\{a\}$, having no attackers, is admissible (actually $\mathrm{F}_{AF}(\{a\}) = \{a\}$). The sets $\{b\}$ and $\{c\}$ are not admissible (b does not defend itself from a and c does not defend itself from b), while $\{d\}$ is, as it defends itself from c (in particular $\mathrm{F}_{AF}(\{d\}) = \{a,d\}$). Moreover the sets $\{a,c\}$ and $\{a,d\}$ are admissible (in the former case c defends itself from the attack by d and is defended by a against b, in the latter both a and d are able to defend themselves), while $\{b,d\}$ is not (a defense for b against a is lacking). Applying analogous considerations in Figure 6, it can be seen that the non empty admissible sets are $\{a\}$, $\{b\}$, $\{a,d\}$ and $\{b,d\}$. On the other hand, in Figure 7 only the empty set is admissible since the non empty conflict-free sets are just the singletons $\{a\}$, $\{b\}$, $\{c\}$ but no argument defends itself from the attack it receives.

As probably noticed by the reader, the above examples point out a correspondence between admissible labellings and admissible sets (see an overview in Tables 1 and 2 for the two more articulated examples).

Before stating this correspondence in the general case, we need to provide the mapping from sets of arguments to labellings that was not introduced in previous section since it is well-defined only for conflict-free sets of arguments[6].

[6]If a set $\mathcal{A}rgs$ of arguments is not con ict-free $\mathcal{A}rgs \cap \mathcal{A}rgs^+$ is not empty, i.e. some argument would be labelled both `in` and `out` according to $\mathtt{Ext2Lab}(\mathcal{A}rgs)$.

Admissible labellings	Admissible sets
$\{(a, \texttt{undec}), (b, \texttt{undec}), (c, \texttt{undec}), (d, \texttt{undec})\}$	$\emptyset$
$\{(a, \texttt{in}), (b, \texttt{out}), (c, \texttt{undec}), (d, \texttt{undec})\}$	$\{a\}$
$\{(a, \texttt{in}), (b, \texttt{out}), (c, \texttt{out}), (d, \texttt{undec})\}$	$\{a\}$
$\{(a, \texttt{in}), (b, \texttt{out}), (c, \texttt{out}), (d, \texttt{in})\}$	$\{a, d\}$
$\{(a, \texttt{out}), (b, \texttt{in}), (c, \texttt{undec}), (d, \texttt{undec})\}$	$\{b\}$
$\{(a, \texttt{out}), (b, \texttt{in}), (c, \texttt{out}), (d, \texttt{undec})\}$	$\{b\}$
$\{(a, \texttt{out}), (b, \texttt{in}), (c, \texttt{out}), (d, \texttt{in})\}$	$\{b, d\}$

Table 2. Admissibile labellings and sets in the example of Figure 6

The idea is that the members of the set are labelled `in`, the arguments attacked by the set are labelled `out` and the remaining arguments are labelled `undec`.

Definition 3.6 *Given an argumentation framework* $AF = \langle Ar, att\rangle$ *and a conflict-free set* $\mathcal{A}rgs \subseteq Ar$ *the corresponding labelling* $\texttt{Ext2Lab}(\mathcal{A}rgs)$ *is defined as* $\texttt{Ext2Lab}(\mathcal{A}rgs) = (\mathcal{A}rgs, \mathcal{A}rgs^+, Ar \setminus (\mathcal{A}rgs \cup \mathcal{A}rgs^+))$.

Let us call an extension-based semantics conflict-free if each of its extensions is a conflict-free set. We can then extend the above definition to sets of extensions.

Definition 3.7 *Given an argumentation framework* $AF = \langle Ar, att\rangle$ *and a conflict-free extension-based semantics* σ*, the set of labellings corresponding to* $\mathcal{E}_\sigma(AF)$ *is given by* $\mathcal{L}_\sigma(AF) = \{\texttt{Ext2Lab}(E) \mid E \in \mathcal{E}_\sigma(AF)\}$.

The correspondence between admissible labellings and admissible sets stated by Proposition 3.8 has been proved in [Caminada and Gabbay, 2009].

Proposition 3.8 *For any argumentation framework* $AF = \langle Ar, att\rangle$

- *if* $\mathcal{A}rgs$ *is an admissible set then* $\texttt{Ext2Lab}(\mathcal{A}rgs)$ *is an admissible labelling;*
- *if* $\mathcal{L}ab$ *is an admissible labelling then* $\texttt{Lab2Ext}(\mathcal{L}ab)$ *is an admissible set.*

It can be noted that the correspondence is not bijective, since different admissible labellings may give rise to the same admissible set. For instance, in the argumentation framework of Figure 5 both $(\{a\}, \{b\}, \{c, d\})$ and $(\{a\}, \emptyset, \{b, c, d\})$ are admissible labellings, whose set of `in`-labelled arguments yields the same admissible set $\{a\}$.

To complete the correspondence it is also possible to define a notion of a conflict-free labelling which parallels the one of conflict-free set[7].

[7]We use the De nition of [Caminada, 2011]. Note that clause 2. is needed for de ning stage labellings (see Section 3.9).

Definition 3.9 *Let $\mathcal{L}ab$ be a labelling of an argumentation framework $AF = (Ar, att)$. $\mathcal{L}ab$ is* conflict-free *iff for each $a \in Ar$ it holds that:*

1. *if a is labelled* `in` *then it does not have an attacker that is labelled* `in`
2. *if a is labelled* `out` *then it has at least one attacker that is labelled* `in`

When comparing a conflict-free labelling to an admissible labelling it can be noticed that the condition on `out` labelled arguments (second bullet) is essentially the same. However, the condition for `in`-labelled arguments (first bullet) is weaker for conflict-free labellings than for admissible labellings. It then follows that every admissible labelling is also a conflict-free labelling (just like every admissible set is also a conflict-free set by definition).

Finally, it is worth recalling that admissibility and defense are related by a basic property. In terms of extensions, if an admissible sets defends an argument, it is possible to add the argument to the set while preserving its admissibility and its capability to defend any other argument. This was proved in the so called Dung's Fundamental Lemma [Dung, 1995] recalled below.

Lemma 3.10 *For any argumentation framework $AF = \langle Ar, att \rangle$, let $\mathcal{A}rgs$ be an admissible set and a, b be arguments defended by $\mathcal{A}rgs$. Then*

1. *$\mathcal{A}rgs' = \mathcal{A}rgs \cup \{a\}$ is an admissible set;*
2. *b is defended by $\mathcal{A}rgs'$.*

Apart from admissibility as commonly applied in the literature, there also exists the related concept of *strong admissibility* [Baroni and Giacomin, 2007b; Caminada, 2014]. In order to describe this, we first need to introduce the concept of a min-max numbering. Basically, what a min-max numbering does is to assign to each `in` or `out`-labelled argument a natural number (or ∞) such that the min-max number of each `in`-labelled argument becomes the maximal value of its `out`-labelled attackers, plus one, and the min-max value of each `out`-labelled argument becomes the minimal value of its `in`-labelled attackers, plus one.

Definition 3.11 *Let $\mathcal{L}ab$ be an admissible labelling of argumentation framework $\langle Ar, att \rangle$. A* min-max numbering *is a total function $\mathcal{MM}_{\mathcal{L}ab} : \mathtt{in}(\mathcal{L}ab) \cup \mathtt{out}(\mathcal{L}ab) \rightarrow \mathbb{N} \cup \{\infty\}$ such that for each $a \in \mathtt{in}(\mathcal{L}ab) \cup \mathtt{out}(\mathcal{L}ab)$ it holds that:*

- *if $\mathcal{L}ab(a) = \mathtt{in}$ then $\mathcal{MM}_{\mathcal{L}ab}(a) = max(\{\mathcal{MM}_{\mathcal{L}ab}(b) \mid b \text{ attacks } a \text{ and } \mathcal{L}ab(b) = \mathtt{out}\}) + 1$ (with $max(\emptyset)$ defined as 0)*
- *if $\mathcal{L}ab(a) = \mathtt{out}$ then $\mathcal{MM}_{\mathcal{L}ab}(a) = min(\{\mathcal{MM}_{\mathcal{L}ab}(b) \mid b \text{ attacks } a \text{ and } \mathcal{L}ab(b) = \mathtt{in}\}) + 1$ (with $min(\emptyset)$ defined as ∞)*

It can be proved that each admissible labelling $\mathcal{L}ab$ has a unique min-max labelling $\mathcal{MM}_{\mathcal{L}ab}$ [Caminada, 2014]. As an example of a min-max numbering, take the argumentation framework of Figure 5 and admissible labelling $\mathcal{L}ab = (\{a, c\}, \{b, d\}, \emptyset)$. Argument a is labelled `in` and does not have any attackers. Hence, $\mathcal{MM}_{\mathcal{L}ab}(a) = 1$ (as $max(\emptyset) = 0$). Argument b is labelled `out` and has only one attacker that is labelled `in` (a), whose min-max value we have just determined to be 1. Hence, $\mathcal{MM}_{\mathcal{L}ab}(b) = 2$. For argument c the situation is more complex, as it has two attackers (b and d) and we have determined the min-max value of only one of these ($\mathcal{MM}_{\mathcal{L}ab}(b) = 2$). This means we first need to "guess" the min-max number of d in order to determine the min-max number of c. If we would for instance guess the min-max number of d as 2, then c will be assigned the min-max number of 3, which then implies that d should actually have been assigned a min-max number of 4 (contradiction). It can be verified that whatever natural number we initially assign as a min-max value to d, the number will later turn out to be too small. The only solution is to assign both c and d not with a natural number but with ∞. In that case, $\mathcal{MM}_{\mathcal{L}ab}(c) = max(\{2, \infty\}) + 1 = \infty$ and $\mathcal{MM}_{\mathcal{L}ab}(d) = max(\{\infty\}) + 1 = \infty$, thus satisfying the constraints of a min-max numbering.

We are now ready to provide the definition of a strongly admissible labelling.

Definition 3.12 *A* strongly admissible labelling *is an admissible labelling whose min-max numbering yields natural numbers only (so no argument is numbered ∞).*

From Definition 3.12 it directly follows that every strongly admissible labelling is also an admissible labelling. Apart from applying the labelling-based approach, it is also possible to express strong admissibility using the extension-based approach [Baroni and Giacomin, 2007b; Caminada, 2014].

Definition 3.13 *Let $\langle Ar, att\rangle$ be an argumentation framework. $\mathcal{A}rgs \subseteq Ar$ is* strongly admissible *iff every $a \in \mathcal{A}rgs$ is defended by some $\mathcal{A}rgs' \subseteq \mathcal{A}rgs \setminus \{a\}$ which is strongly admissible.*

The basis of this recursive definition is given by the facts that the empty set is strongly admissible and that unattacked arguments are defended by the empty set. Intuitively, the defense of every argument in a strongly admissible set is "rooted" in an unattacked argument (a notion called *strong defense* in [Baroni and Giacomin, 2007b]). In case there are no unattacked arguments in a framework, the empty set is its only strongly admissible set.

It can be proved that each strongly admissible set is conflict free and admissible [Baroni and Giacomin, 2007b; Caminada, 2014]. As an example, consider again the argumentation framework of Figure 5. The set $\{a\}$ is strongly admissible, because a is defended by $\emptyset$ (and $\emptyset \subseteq \{a\} \setminus \{a\}$) which is trivially strongly admissible since it has no elements. Also, the set $\{a, c\}$, although admissible, is not strongly admissible. This is because no subset of $\{a, c\} \setminus \{c\}$ can defend c against d. Correspondence between strongly admissible sets and strongly admissible labellings can be established through the functions `Ext2Lab` and `Lab2Ext`.

Proposition 3.14 *For any argumentation framework $AF = \langle Ar, att \rangle$*

- *if $\mathcal{A}rgs$ is a strongly admissible set then* `Ext2Lab`$(\mathcal{A}rgs)$ *is a strongly admissible labelling;*
- *if $\mathcal{L}ab$ is a strongly admissible labelling then* `Lab2Ext`$(\mathcal{L}ab)$ *is a strongly admissible set.*

It can be noted that the correspondence is not bijective, since different strongly admissible labellings may give rise to the same strongly admissible set. For instance, in the argumentation framework of Figure 5 both $(\{a\}, \{b\}, \{c, d\})$ and $(\{a\}, \emptyset, \{b, c, d\})$ are strongly admissible labellings, whose set of `in`-labelled arguments yields the same strongly admissible set $\{a\}$.

Finally, an equivalent non-recursive characterisation for finite strongly admissible sets has been provided in [Baumann *et al.*, 2016] and is recalled in the Proposition 3.15.

Proposition 3.15 *Let $\langle Ar, att \rangle$ be an argumentation framework and $\mathcal{A}rgs \subseteq Ar$ a finite set of arguments. $\mathcal{A}rgs$ is* strongly admissible *iff there are finitely many and pairwise disjoint sets $\mathcal{A}rgs_1, \ldots \mathcal{A}rgs_n$ such that $\mathcal{A}rgs = \bigcup_{1 \leq i \leq n} \mathcal{A}rgs_i$, $\mathcal{A}rgs_1 \subseteq \mathrm{F}_{AF}(\emptyset)$ and for every $1 \leq j \leq n-1$ $\mathcal{A}rgs_{j+1} \subseteq \mathrm{F}_{AF}(\bigcup_{1 \leq i \leq j} \mathcal{A}rgs_i)$.*

In words, a strongly admissible set can be constructed starting from a set of unattacked arguments $\mathcal{A}rgs_1$ and then adding iteratively further arguments which are defended by those already included in the set.

3.2 Naïve semantics

Naïve semantics (denoted as $\mathcal{NA}$) corresponds to selecting as many arguments as possible, provided that there are no attacks among them. It is a sort of greedy strategy, driven by the only criterion of avoiding conflicts. Formally it corresponds to requiring conflict-freeness together with a maximality property and can be easily expressed in both the labelling-based and extension-based approach.

Definition 3.16 *Let $\mathcal{L}ab$ be a labelling of an argumentation framework $\langle Ar, att \rangle$. $\mathcal{L}ab$ is a* naïve labelling *iff it is a conflict-free labelling whose set of* `in`*-labelled arguments is maximal (w.r.t. set inclusion) with respect to all conflict-free labellings.*

Definition 3.17 *Let $\langle Ar, att \rangle$ be an argumentation framework. A set $\mathcal{A}rgs \subseteq Ar$ is called a* naïve extension *iff $\mathcal{A}rgs$ is a maximal conflict-free set.*

It can immediately be observed from the definitions above that naïve semantics ignores the direction of attacks, which makes it very simple but at the same time rather poor, since it overlooks an essential information carried by the formalism.

Let us illustrate naïve semantics with an example[8].

In Figure 5, there are nineteen conflict-free labellings, including for instance $(\emptyset, \emptyset, \{a, b, c, d\})$, $(\{a\}, \emptyset, \{b, c, d\})$, $(\{d\}, \{c\}, \{a, b\})$ and many others. Among them, there are four naïve labellings with set of `in`-labelled arguments $\{a, c\}$, namely $(\{a, c\}, \emptyset, \{b, d\})$, $(\{a, c\}, \{b\}, \{d\})$, $(\{a, c\}, \{d\}, \{b\})$, $(\{a, c\}, \{b, d\}, \emptyset)$, four naïve labellings where the set of `in`-labelled arguments is $\{a, d\}$, namely $(\{a, d\}, \emptyset, \{b, c\})$, $(\{a, d\}, \{b\}, \{c\})$, $(\{a, d\}, \{c\}, \{b\})$, $(\{a, d\}, \{b, c\}, \emptyset)$, and two naïve labellings with set of `in`-labelled arguments $\{b, d\}$, namely $(\{b, d\}, \emptyset, \{a, c\})$ and $(\{b, d\}, \{c\}, \{a\})$.

In the same example, there are eight conflict-free sets of arguments, namely $\emptyset$, $\{a\}$, $\{b\}$, $\{c\}$, $\{d\}$, $\{a, c\}$, $\{a, d\}$, $\{b, d\}$. Clearly only $\{a, c\}$, $\{a, d\}$, $\{b, d\}$ are maximal, i.e. naïve extensions.

From the example it emerges clearly that the relationships between naïve labellings and naïve extensions is in general many-to-one: many naïve labellings may correspond to the same naïve extension and each naïve extension corresponds to at least one, and in general many, naïve labellings.

Proposition 3.18 *For any argumentation framework $\langle Ar, att \rangle$, if $\mathcal{L}ab$ is a naïve labelling there is a naïve extension $\mathcal{A}rgs$ such that $\mathcal{A}rgs = \mathtt{Lab2Ext}(\mathcal{L}ab)$. If $\mathcal{A}rgs$ is a naïve extension then $\mathtt{Ext2Lab}(\mathcal{A}rgs)$ is a naïve labelling.*

3.3 Complete Semantics

Complete semantics ($\mathcal{CO}$) can be regarded as a strengthening of the basic requirements enforced by the idea of admissibility. Intuitively, while admissibility requires one to be able to give reasons for accepted and rejected arguments but leaves one free to abstain on any argument, complete semantics requires one to abstain only if there are no good reasons to do otherwise. That is, if one abstains from having an opinion on whether the argument is accepted or rejected, then one should have insufficient grounds to accept the argument (meaning that not all its attackers are rejected) and insufficient grounds to reject the argument (meaning that it does not have an attacker that is accepted). Note in particular that, while the trivial solution of leaving anything undecided is always admissible, it is not always complete since there can be arguments one has good reasons not to abstain about.

In the labelling-based approach, the intuition described above corresponds to extending Definition 3.1 in order to encompass a notion of an argument being *legally undecided*.

[8] A summary of the outcomes produced by all the semantics considered in this chapter on all the examples presented in this section is given in Tables 5 and 6 at the end of the section. For the sake of compactness, the summary is given in terms of extensions, the corresponding labellings being derivable with the `Ext2Lab` function.

Definition 3.19 *Let $\mathcal{L}ab$ be a labelling of an argumentation framework $\langle Ar, att \rangle$.*

- *An* `undec`*-labelled argument is said to be* legally `undec` *iff not all its attackers are labelled* `out` *and it doesn't have an attacker that is labelled* `in`*.*

Definition 3.20 *A* complete labelling *is a labelling where every* `in`*-labelled argument is legally* `in`*, every* `out`*-labelled argument is legally* `out` *and every* `undec` *labelled argument is legally* `undec`*.*

It is clear from Definitions 3.20 and 3.2 that every complete labelling is an admissible labelling (but the reverse does not hold in general).

An alternative characterisation of a complete labelling can be provided (a formal proof can be found in [Caminada and Gabbay, 2009]).

Proposition 3.21 *A labelling $\mathcal{L}ab$ of an argumentation framework (Ar, att) is a complete labelling iff for each argument $a \in Ar$ it holds that:*

1. *a is labelled* `in` *iff all its attackers are labelled* `out`*, and*
2. *a is labelled* `out` *iff it has at least one attacker that is labelled* `in`*.*

Although Proposition 3.21 does not explicitly mention `undec`, it follows that each argument that is labelled `undec` does not have all its attackers `out` (otherwise it would have to be labelled `in` by point 1) and it does not have an attacker that is labelled `in` (otherwise it would have to be labelled `out` by point 2). Therefore, each `undec`-labelled argument is legally `undec`. Comparing Proposition 3.21 with Definition 3.5 one can appreciate the difference between admissible and complete labellings from another perspective: in an admissible labelling an argument can be labelled `in` *only if* all its attackers are labelled `out`, but need not to be labelled `in` if this condition holds. In a complete labelling an argument is labelled `in` *if and only if* all its attackers are labelled `out`: thus complete labellings are more constrained and this corresponds to the lesser freedom of abstaining mentioned at the beginning of this section.

Turning to the extension-based approach, a complete extension is a conflict-free set which includes precisely those arguments it defends. That is, if an argument is defended by the set it should be in the set, and if an argument is not defended by the set, it should not be in the set. Technically this means that a complete extension is a conflict-free fixed point of the characteristic function, as stated in the following definition [Dung, 1995].

Definition 3.22 *Let $AF = \langle Ar, att \rangle$ be an argumentation framework. A set $\mathcal{A}rgs \subseteq Ar$ is called a* complete extension *iff $\mathcal{A}rgs$ is conflict-free and $\mathcal{A}rgs = \mathrm{F}_{AF}(\mathcal{A}rgs)$.*

Fig. #	Members of $\mathcal{L}_{\mathcal{CO}}(AF)$	Members of $\mathcal{E}_{\mathcal{CO}}(AF)$
Fig. 5	$\{(a, \texttt{in}), (b, \texttt{out}), (c, \texttt{undec}), (d, \texttt{undec})\}$	$\{a\}$
	$\{(a, \texttt{in}), (b, \texttt{out}), (c, \texttt{in}), (d, \texttt{out})\}$	$\{a, c\}$
	$\{(a, \texttt{in}), (b, \texttt{out}), (c, \texttt{out}), (d, \texttt{in})\}$	$\{a, d\}$
Fig. 6	$\{(a, \texttt{undec}), (b, \texttt{undec}), (c, \texttt{undec}), (d, \texttt{undec})\}$	$\emptyset$
	$\{(a, \texttt{in}), (b, \texttt{out}), (c, \texttt{out}), (d, \texttt{in})\}$	$\{a, d\}$
	$\{(a, \texttt{out}), (b, \texttt{in}), (c, \texttt{out}), (d, \texttt{in})\}$	$\{b, d\}$
Fig. 7	$\{(a, \texttt{undec}), (b, \texttt{undec}), (c, \texttt{undec})\}$	$\{\emptyset\}$

Table 3. Complete labellings and extensions in the examples of Figures 5- 7

It is clear from Definitions 3.22 and 3.5 that every complete extension is an admissible set (but the reverse does not hold in general).

Let us now provide some examples to illustrate the notion of complete semantics. In Figure 5, one can observe that, among the seven admissible labellings, only $(\{a\}, \{b\}, \{c, d\})$, $(\{a, c\}, \{b, d\}, \emptyset)$, and $(\{a, d\}, \{b, c\}, \emptyset)$ are complete. In particular, note that a is legally `in` in all labellings because all its attackers are `out` (trivially, because it has no attackers). b is legally `out` in all labellings because it has an attacker (a) that is `in`. On the other hand, c and d can be both legally `undec`, or one legally `in` and the other legally `out`. Analogously, in the same figure it can be noted that $\{a\}$ is a complete extension (a has no attackers and is therefore trivially defended by any set, a defends c from b but not from d), and $\{a, c\}$ and $\{a, d\}$ are complete extensions too.

In Figure 6, the trivial labelling $(\emptyset, \emptyset, \{a, b, c, d\})$ is complete, as well as $(\{a, d\}, \{b, c\}, \emptyset)$ and $(\{b, d\}, \{a, c\}, \emptyset)$. Analogously, $\emptyset$ is a complete extension (no unattacked arguments exist, which would be the only arguments defended by the empty set) as well as $\{a, d\}$ and $\{b, d\}$, while $\{a\}$ and $\{b\}$ are not complete extensions since they both defend also argument d.

In Figure 7 the only complete labelling is the trivial one $(\emptyset, \emptyset, \{a, b, c\})$ and analogously the only complete extension is $\emptyset$ (as it was the case for admissible labellings/sets).

As the above examples also show, there is a direct mapping between complete labellings and complete extensions: it has been proved in [Caminada and Gabbay, 2009] that this correspondence is bijective, as stated in the following proposition.

Proposition 3.23 *For any argumentation framework $AF = \langle Ar, att \rangle$, $\mathcal{L}ab$ is a complete labelling iff there is a complete extension $\mathcal{A}rgs$ such that $\mathcal{L}ab =$ $\texttt{Ext2Lab}(\mathcal{A}rgs)$.*

Table 3 shows this correspondence on the examples discussed above.

3.4 Grounded Semantics

If one regards each complete labelling (or complete extension) as a reasonable position one can take in the presence of the conflicting information expressed in the argumentation framework, then a possible question is to examine what is the most "grounded" position one can take, namely the position which is least questionable. The idea is then to accept only the arguments that one cannot avoid to accept, to reject only the arguments that one cannot avoid to reject, and abstaining as much as possible. This gives rise to the most skeptical (or least committed) semantics among those based on complete extensions, namely the *grounded* semantics ($\mathcal{GR}$).

This idea has a straightforward formal counterpart in terms of a minimality requirement[9].

Definition 3.24 *Let $AF = \langle Ar, att \rangle$ be an argumentation framework. The* grounded labelling *of AF is a complete labelling $\mathcal{L}ab$ where $\mathtt{in}(\mathcal{L}ab)$ is minimal (w.r.t. set inclusion), i.e. there is no complete labelling $\mathcal{L}ab'$ such that $\mathtt{in}(\mathcal{L}ab') \subsetneq \mathtt{in}(\mathcal{L}ab)$.*

Definition 3.25 *Let $AF = \langle Ar, att \rangle$ be an argumentation framework. The* grounded extension *of AF is a minimal (w.r.t. set inclusion) complete extension of AF (i.e. a minimal conflict-free fixed point of the characteristic function F_{AF}).*

As we have already seen complete labellings and extensions in the examples of Figures 5-7, one can easily identify those featuring the minimality property required by the above definitions. In the example of Figure 5, the grounded labelling is $(\{a\}, \{b\}, \{c, d\})$ while the grounded extension is $\{a\}$. In both Figures 6 and 7 the grounded labelling is the trivial one $((\emptyset, \emptyset, \{a, b, c, d\})$ and $(\emptyset, \emptyset, \{a, b, c\})$ respectively), and analogously the grounded extension is the empty set in both cases.

The uniqueness of the grounded labelling and extension in these examples is not accidental. Considering the grounded extension, since F_{AF} is monotonic it follows from the Knaster-Tarski theorem that F_{AF} has a unique smallest fixed point. It can then be proved that this fixed point is also conflict-free [Dung, 1995].

Proposition 3.26 *For any argumentation framework $AF = \langle Ar, att \rangle$, the following statements are equivalent:*

1. *$\mathcal{A}rgs$ is a minimal conflict-free fixed point of F_{AF}*
2. *$\mathcal{A}rgs$ is the smallest fixed point of F_{AF}*

[9]De nition 3.25 is not literally the same as the one originally given by [Dung, 1995]. We provide this equivalent version as it is more coherent with our presentation line.

It follows that:

- the grounded extension is unique (i.e. grounded semantics belongs to the unique-status approach);
- the grounded extension is the least complete extension, in particular it is included in any complete extension.

The grounded extension of an argumentation framework AF will be denoted as $GE(AF)$.

In virtue of the one-to-one correspondence between complete extensions and complete labellings established in Section 3.3, it can be proved that the grounded labelling is unique and coincides with $\mathtt{Ext2Lab}(\mathcal{A}rgs)$ where $\mathcal{A}rgs$ is the grounded extension. Similarly, if $\mathcal{L}ab$ is the grounded labelling, then $\mathtt{Lab2Ext}(\mathcal{L}ab)$ is the grounded extension.

As a confirmation of the intuitive meaning stated at the beginning of the section, it turns out that the grounded semantics can be described not only in terms of minimizing acceptance. In fact, the complete labelling where $\mathtt{in}(\mathcal{L}ab)$ is minimal is also the complete labelling $\mathcal{L}ab$ where $\mathtt{out}(\mathcal{L}ab)$ is minimal, and the complete labelling $\mathcal{L}ab$ where $\mathtt{undec}(\mathcal{L}ab)$ is maximal. This is stated in Proposition 3.28, whose proof is based on Lemma 3.27 (see [Caminada, 2006a; Caminada and Gabbay, 2009] for details).

Lemma 3.27 *Given two complete labellings $\mathcal{L}ab_1$ and $\mathcal{L}ab_2$ of an argumentation framework $\langle Ar, att\rangle$, it holds that $\mathtt{in}(\mathcal{L}ab_1) \subseteq \mathtt{in}(\mathcal{L}ab_2)$ iff $\mathtt{out}(\mathcal{L}ab_1) \subseteq \mathtt{out}(\mathcal{L}ab_2)$.*

Proposition 3.28 *Let $\mathcal{L}ab$ be a complete labelling of an argumentation framework $\langle Ar, att\rangle$. The following statements are equivalent.*

1. *$\mathcal{L}ab$ is the complete labelling where $\mathtt{in}(\mathcal{L}ab)$ is minimal (w.r.t. set inclusion)*
2. *$\mathcal{L}ab$ is the complete labelling where $\mathtt{out}(\mathcal{L}ab)$ is minimal (w.r.t. set inclusion)*
3. *$\mathcal{L}ab$ is the complete labelling where $\mathtt{undec}(\mathcal{L}ab)$ is maximal (w.r.t. set inclusion)*

Given the bijective correspondence between complete labellings and complete extensions, the above proposition can be equivalently formulated for the extension-based approach.

Proposition 3.29 *Let E be a complete extension of an argumentation framework $\langle Ar, att\rangle$. The following statements are equivalent.*

1. *E is the least (w.r.t. set inclusion) complete extension*

2. *E is the complete extension such that E^+ is minimal (w.r.t. set inclusion)*
3. *E is the complete extension such that $Ar \setminus (E \cup E^+)$ is maximal (w.r.t. set inclusion)*

There also exists a connection between grounded semantics and the concept of strong admissibility [Baroni and Giacomin, 2007b; Caminada, 2014].

Theorem 3.30 *Let $AF = \langle Ar, att \rangle$ be an argumentation framework. The grounded extension of AF is the unique maximal (w.r.t set inclusion) strongly admissible set of AF.*

Finally, an interesting property proved in [Dung, 1995] provides a useful "constructive" characterisation of grounded semantics[10] for finite (and more generally finitary[11]) argumentation frameworks.

Proposition 3.31 *The grounded extension of any finitary argumentation framework AF is equal to $\bigcup_{i=1,\dots,\infty} \mathrm{F}^i_{AF}(\emptyset)$, where $\mathrm{F}^1_{AF}(\emptyset) = \mathrm{F}_{AF}(\emptyset)$ and for $i > 1$ $\mathrm{F}^i_{AF}(\emptyset) = \mathrm{F}_{AF}(\mathrm{F}^{i-1}_{AF}(\emptyset))$.*

On the basis of Proposition 3.31 the grounded labelling (or equivalently extension) can be obtained incrementally by first labelling `in` those arguments which do not receive attacks. Then the arguments attacked by those labelled `in` are labelled `out`. The same steps are iterated considering only those arguments which have not been labelled yet, namely repeating the procedure on an argumentation framework obtained by suppressing the already labelled arguments. In particular, this corresponds to labelling `in` those unlabelled arguments which only receive attacks from arguments labelled `out`, and then labelling `out` those attacked by the newly labelled `in` arguments. The procedure is then iterated until an iteration does not produce any newly `in` or `out` labelled argument. Then, any still unlabelled arguments are labelled `undec`.

It can be noted that the first iteration corresponds to labelling `in` the arguments in $\mathrm{F}^1_{AF}(\emptyset)$ and `out` the arguments attacked by $\mathrm{F}^1_{AF}(\emptyset)$, the second iteration to labelling `in` the arguments in $\mathrm{F}^2_{AF}(\emptyset)$ and `out` the arguments attacked by $\mathrm{F}^2_{AF}(\emptyset)$, and so on. This procedure can be applied to the examples and provides another way to see that the grounded extension includes those and only those arguments whose defense is "rooted" in unattacked arguments and is the maximal strongly admissible set.

If the aim is not so much to compute the entire grounded extension (labelling) but merely to examine whether or not an argument is in the grounded extension (labelled `in` by the grounded labelling) then one could also use the proof procedures described in [Modgil and Caminada, 2009; Caminada, 2015].

[10] Note that the characterisation of strongly admissible sets in Proposition 3.15 can be seen as a generalisation of the intuition underlying this traditional result in the case of nite frameworks.

[11] An argumentation framework is nitary if every argument receives a nite number of attacks.

3.5 Preferred Semantics

While grounded semantics takes a skeptical, or least-commitment, standpoint, one can also consider the alternative view oriented at accepting as many arguments as reasonably possible. This may give rise to mutually exclusive alternatives for acceptance: for instance a mutual attack can be reasonably resolved by accepting either of the conflicting arguments, but clearly not both (these alternatives are called non-skeptical solutions in the examples below).

The idea of maximizing accepted arguments is expressed by *preferred* semantics ($\mathcal{PR}$) whose description in the labelling-based and extension-based approaches is given in the following definitions.

Definition 3.32 *Let $AF = \langle Ar, att \rangle$ be an argumentation framework. A* preferred labelling *of AF is a complete labelling $\mathcal{L}ab$ where* $\mathtt{in}(\mathcal{L}ab)$ *is maximal (w.r.t. set-inclusion) among all complete labellings, i.e. there is no complete labelling $\mathcal{L}ab'$ such that* $\mathtt{in}(\mathcal{L}ab') \supsetneq \mathtt{in}(\mathcal{L}ab)$.

Definition 3.33 *Let $AF = \langle Ar, att \rangle$ be an argumentation framework. A* preferred extension *is a maximal (w.r.t. set-inclusion) admissible set of AF.*

Considering the examples of Figures 5-7, the existence of multiple preferred labellings (or extensions) immediately emerges. For instance, in Figure 5 two non-skeptical solutions exist for the mutual attack between c and d, giving rise to the preferred labellings $(\{a, c\}, \{b, d\}, \emptyset)$ and $(\{a, d\}, \{b, c\}, \emptyset)$. Similarly, two preferred extensions exist, namely $\{a, c\}$ and $\{a, d\}$.

In Figure 6 again two alternative non-skeptical solutions exist for the mutual attack between a and b. In both cases, c is then rejected and d accepted. This intuitive description corresponds to the two preferred labellings $(\{a, d\}, \{b, c\}, \emptyset)$ and $(\{b, d\}, \{a, c\}, \emptyset)$ and, analogously, to the preferred extensions $\{a, d\}$ and $\{b, d\}$.

In Figure 7 instead, no non-trivial solutions to the conflict are available under the constraint of admissibility, as the reader may remember from previous subsections. It then follows that the unique preferred labelling in this case is $(\emptyset, \emptyset, \{a, b, c\})$ and, similarly, the only preferred extension is $\emptyset$.

As usual, the correspondences in the above examples are not accidental: it can be proved that an analogous version of Proposition 3.23 holds for preferred semantics, i.e. there is a bijective correspondence between preferred labellings and preferred extensions through the `Ext2Lab` (and `Lab2Ext`) functions.

It turns out that the complete labellings with maximal `in` are the same as the complete labellings with maximal `out`, as stated in Proposition 3.34 whose proof is based on Lemma 3.27.

Proposition 3.34 *Given an argumentation framework $AF = (Ar, att)$ the following statements are equivalent.*

1. *$\mathcal{L}ab$ is a complete labelling where* $\mathtt{in}(\mathcal{L}ab)$ *is maximal (w.r.t. set inclusion) among all complete labellings.*

2. *$\mathcal{L}ab$ is a complete labelling where* $\texttt{out}(\mathcal{L}ab)$ *is maximal (w.r.t. set inclusion) among all complete labellings.*

An analogous formulation of Proposition 3.34 for the extension-based approach could be provided in a straightforward way.

Relationships of preferred extensions with other semantics notions have been analyzed in [Dung, 1995]. Preferred extensions can for instance equivalently be characterized as maximal complete extensions.

Proposition 3.35 *Let $AF = \langle Ar, att \rangle$ be an argumentation framework and let $\mathcal{A}rgs \subseteq Ar$. The following statements are equivalent.*

1. *$\mathcal{A}rgs$ is a maximal (w.r.t. set inclusion) admissible set of AF*
2. *$\mathcal{A}rgs$ is a maximal (w.r.t. set inclusion) complete extension of AF*

This in particular implies that the grounded extension is included in any preferred extension, as it is in any complete extension. By definition, the grounded extension coincides with the intersection of all complete extensions: one may then wonder whether this holds also for preferred extensions. The answer is negative, as shown for instance by the example of Figure 6 where the grounded extension is $\emptyset$ while the intersection of the preferred extensions is $\{d\}$. Again, this fact can be easily translated to the labelling-based approach referring to the `in`-labelled arguments.

An algorithm that produces all preferred labellings (and therefore also produces all preferred extensions) is described in [Caminada, 2007a; Modgil and Caminada, 2009; Nofal *et al.*, 2014]. If the aim is merely to determine whether an argument is in at least one preferred extension (labelled `in` by at least one preferred labelling) then one could also use the proof procedures described in [Vreeswijk and Prakken, 2000; Vreeswijk, 2006; Verheij, 2007; Modgil and Caminada, 2009; Caminada *et al.*, 2016]. Proof procedures for determining whether an argument is in every preferred extension (labelled `in` by every preferred labelling) are provided in [Cayrol *et al.*, 2003; Modgil and Caminada, 2009].

3.6 Stable Semantics

So far we have discussed semantics according to the intuitive idea that an argument can be accepted, rejected or left undecided. One can however prefer more committed evaluations, in which there is no room for neutrality or shades of gray and everything is just black or white. This means that undecided arguments are simply "forbidden" as in statements like "you're either with us or against us."

This clear-and-strong view corresponds to *stable* semantics ($\mathcal{ST}$) and has a direct formulation in both the labelling-based and extension-based approach.

Definition 3.36 *Let $\mathcal{L}ab$ be a labelling of argumentation framework $AF = (Ar, att)$. $\mathcal{L}ab$ is a* stable labelling *of AF iff it is a complete labelling with* $\texttt{undec}(\mathcal{L}ab) = \emptyset$.

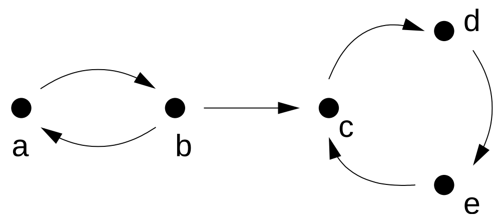

Figure 9. An argumentation framework where preferred and stable semantics differ

Definition 3.37 *Let $AF = \langle Ar, att \rangle$ be an argumentation framework. A* stable extension *of AF is a conflict-free set $\mathcal{A}rgs$ such that $\mathcal{A}rgs \cup \mathcal{A}rgs^+ = Ar$.*

In the example shown in Figure 5 there are two stable labellings, namely $(\{a,c\},\{b,d\},\emptyset)$ and $(\{a,d\},\{b,c\},\emptyset)$. Similarly, two stable extensions exist, namely $\{a,c\}$ and $\{a,d\}$. In Figure 6 the labellings $(\{a,d\},\{b,c\},\emptyset)$ and $(\{b,d\},\{a,c\},\emptyset)$ are stable and, analogously, there are two stable extensions, namely $\{a,d\}$ and $\{b,d\}$.

Figure 7 shows that the strong view underlying stable semantics cannot be universally applied. In fact, no labelling nor extension complying with the definition can be identified (the requirements of conflict-freeness and ability to attack all other arguments are incompatible in this case). This can be regarded as a limitation of stable semantics as "stable extensions do not capture the intuitive semantics of every meaningful argumentation system" [Dung, 1995]. Looking at this fact from another perspective, differently from other semantics reviewed so far, in the case of stable semantics the trivial labelling (or extension) does not represent the "default" conflict resolution one can resort to when nothing else is reasonable. It follows that, using a terminology from [Baroni and Giacomin, 2009a], stable semantics is not *universally defined*, since there are argumentation frameworks where it is intrinsically impossible to apply its "in-or-out" view. No other argumentation semantics considered in the literature has this limitation.

Apart from this critical case, the reader may have noticed that the stable labellings (extensions) coincide with the preferred ones in the other two cases. One may then wonder whether stable semantics (leaving apart critical cases) coincides with preferred semantics in general. The answer is negative, as shown by the argumentation framework of Figure 9. Here one can verify that there are three complete labellings, namely $(\emptyset,\emptyset,\{a,b,c,d,e\})$, $(\{a\},\{b\},\{c,d,e\})$ and $(\{b,d\},\{a,c,e\},\emptyset)$, and, correspondingly, three complete extensions. Two of the three labellings (extensions) are preferred, namely $(\{a\},\{b\},\{c,d,e\})$ and $(\{b,d\},\{a,c,e\},\emptyset)$, but clearly only the last one is stable.

Let us now generalize this and possibly related observations, examining properties of stable semantics in general.

First it is possible to characterize the concept of a stable labelling in other terms. In particular note that the difference between a complete labelling and an admissible labelling is that a complete labelling has the additional requirement that every `undec`-labelled argument is legally `undec`. However, if, as in Definition 3.36, there are no `undec`-labelled arguments in the first place, then this extra requirement becomes superfluous. Moreover, the fact that anything that is not labelled `in` is labelled `out` ensures that every stable labelling is also preferred (but not viceversa, as we have already seen). These considerations are summarized in Proposition 3.38 (notice that point 3 of Proposition 3.38 coincides with Definition 3.36).

Proposition 3.38 *Let $\mathcal{L}ab$ be a labelling of an argumentation framework $AF = (Ar, att)$. The following statements are equivalent:*

1. *$\mathcal{L}ab$ is a conflict-free labelling with* $\texttt{undec}(\mathcal{L}ab) = \emptyset$
2. *$\mathcal{L}ab$ is an admissible labelling with* $\texttt{undec}(\mathcal{L}ab) = \emptyset$
3. *$\mathcal{L}ab$ is a complete labelling with* $\texttt{undec}(\mathcal{L}ab) = \emptyset$
4. *$\mathcal{L}ab$ is a preferred labelling with* $\texttt{undec}(\mathcal{L}ab) = \emptyset$

On the other hand, it is immediate to see that a stable extension is an admissible set, hence the equivalent characterisations given in Proposition 3.39 (again, note that point 1 of Proposition 3.39 coincides with Definition 3.37).

Proposition 3.39 *Let $AF = \langle Ar, att \rangle$ be an argumentation framework and $\mathcal{A}rgs \subseteq Ar$ a set of arguments. The following statements are equivalent:*

1. *$\mathcal{A}rgs$ is a conflict-free set with $\mathcal{A}rgs \cup \mathcal{A}rgs^+ = Ar$*
2. *$\mathcal{A}rgs$ is an admissible set such that $\mathcal{A}rgs \cup \mathcal{A}rgs^+ = Ar$*
3. *$\mathcal{A}rgs$ is a complete extension such that $\mathcal{A}rgs \cup \mathcal{A}rgs^+ = Ar$*
4. *$\mathcal{A}rgs$ is a preferred extension such that $\mathcal{A}rgs \cup \mathcal{A}rgs^+ = Ar$*
5. *$\mathcal{A}rgs^+ = Ar \backslash \mathcal{A}rgs$*

As probably evident from above, the bijective labellings-extensions correspondence through `Ext2Lab` (and `Lab2Ext`) holds for stable semantics too as proved in [Caminada and Gabbay, 2009].

An algorithm that produces all stable labellings (and therefore also all stable extensions) is described in [Caminada, 2007a; Modgil and Caminada, 2009]. If the aim is merely to determine whether an argument is in at least one stable extension (labelled `in` by at least one stable labelling) then one could also use the proof procedures described in [Caminada and Wu, 2009]. Proof procedures for determining whether an argument is in every stable extension (labelled `in` by every stable labelling) are also provided in [Caminada and Wu, 2009].

3.7 Semi-Stable Semantics

As illustrated in the previous section, the requirement of "forbidding" undecided arguments turns out to yield no results in some cases. A more sophisticated idea consists in expressing a definite opinion on the largest possible set of arguments, while restricting as much as possible (but not necessarily avoiding) those which are left undecided. This intuition lies at the basis of *semi-stable* semantics ($\mathcal{SST}$), which can be defined as follows.

Definition 3.40 *Let $\mathcal{L}ab$ be a labelling of an argumentation framework $AF = (Ar, att)$. $\mathcal{L}ab$ is a* semi-stable labelling *of AF iff $\mathcal{L}ab$ is a complete labelling where $\mathtt{undec}(\mathcal{L}ab)$ is minimal (w.r.t. set inclusion) among all complete labellings, i.e. there is no complete labelling $\mathcal{L}ab'$ such that $\mathtt{undec}(\mathcal{L}ab') \subsetneq \mathtt{undec}(\mathcal{L}ab)$.*

Definition 3.41 *Let $AF = \langle Ar, att \rangle$ be an argumentation framework. A* semi-stable extension *of AF is a complete extension $\mathcal{A}rgs$ where $\mathcal{A}rgs \cup \mathcal{A}rgs^{+}$ is maximal (w.r.t. set inclusion) among all complete extensions, i.e there is no complete extension $\mathcal{A}rgs'$ such that $(\mathcal{A}rgs' \cup \mathcal{A}rgs'^{+}) \supsetneq (\mathcal{A}rgs \cup \mathcal{A}rgs^{+})$.*

It follows directly that each stable labelling is also a semi-stable labelling and that semi-stable labellings coincide with stable labellings when the latter exist. This is because a stable labelling is a complete labelling with an empty set of undec-labelled arguments. Hence, it is a complete labelling where the set of undec-labelled arguments is minimal (so a semi-stable labelling). Furthermore, if there exists at least one stable labelling then the set of undec-labelled arguments has to be empty in any complete labelling with a minimal set of undec-labelled arguments (semi-stable labelling) and hence any such a labelling has to be stable. The same relationship holds between stable and semi-stable extensions: each stable extension is a semi-stable extension, and semi-stable extensions coincide with stable extensions when the latter exist. Accordingly, we already know, from previous section, the behaviour of semi-stable semantics in the examples of Figures 5 and 6.

Even in situations where stable extensions/labellings do not exist, the existence of semi-stable labellings (or extensions) is anyway guaranteed, since they are selected among the (always existing) complete ones. In particular, in the example of Figure 7 the only semi-stable labelling (extension) is (again) the trivial one.

The maximization requirement imposed by semi-stable semantics is intuitively similar, but clearly different, from the maximization requirement in the definition of preferred semantics. One may wonder whether these different maximizations actually lead to the same results. The answer is negative, see also [Verheij, 2003], as shown by the example of Figure 9, where there are two preferred labellings (and then two corresponding extensions) namely $(\{a\}, \{b\}, \{c, d, e\})$ and $(\{b, d\}, \{a, c, e\}, \emptyset)$, but only the latter is semi-stable (as well as stable).

Equivalent characterisations of semi-stable semantics in terms of admissible labellings/sets and of preferred labellings/extensions are available, see e.g. [Caminada and Gabbay, 2009], as summarized in the following propositions.

Proposition 3.42 *Let $\mathcal{L}ab$ be a labelling of an argumentation framework $AF = (Ar, att)$. The following statements are equivalent.*

1. *$\mathcal{L}ab$ is a complete labelling where* $\texttt{undec}(\mathcal{L}ab)$ *is minimal (w.r.t. set inclusion) among all complete labellings*
2. *$\mathcal{L}ab$ is an admissible labelling where* $\texttt{undec}(\mathcal{L}ab)$ *is minimal (w.r.t. set inclusion) among all admissible labellings*
3. *$\mathcal{L}ab$ is a preferred labelling where* $\texttt{undec}(\mathcal{L}ab)$ *is minimal (w.r.t. set inclusion) among all preferred labellings*

Proposition 3.43 *Let $AF = \langle Ar, att\rangle$ be an argumentation framework, and let $\mathcal{A}rgs \subseteq Ar$. The following statements are equivalent.*

1. *$\mathcal{A}rgs$ is a complete extension where $\mathcal{A}rgs \cup \mathcal{A}rgs^+$ is maximal (w.r.t. set inclusion) among all complete extensions*
2. *$\mathcal{A}rgs$ is an admissible set where $\mathcal{A}rgs \cup \mathcal{A}rgs^+$ is maximal (w.r.t. set inclusion) among all admissible sets*
3. *$\mathcal{A}rgs$ is a preferred extension where $\mathcal{A}rgs \cup \mathcal{A}rgs^+$ is maximal (w.r.t. set inclusion) among all preferred extensions*

Finally, the usual bijective labellings-extension correspondence holds for semi-stable semantics too, see [Caminada, 2007a; Caminada and Gabbay, 2009]. An algorithm that produces all semi-stable labellings (and therefore also all semi-stable extensions) is described in [Caminada, 2007a; Modgil and Caminada, 2009].

The concept of semi-stable semantics can be traced back to the notion of admissible stage extensions (see Section 3.9) introduced by [Verheij, 1996]. Although there are differences in the basic formalisation (Verheij for instance does not use the standard extension-based approach) it can be proved that Verheij's approach is equivalent to that of Caminada, who, independently from Verheij, rediscovered the same concept under the name of semi-stable semantics [Caminada, 2006b].

3.8 Ideal and Eager Semantics

The notion of *ideal* semantics ($\mathcal{ID}$) can perhaps be best explained using a description concerning a judgment aggregation context [Caminada and Pigozzi, 2011], where different people have different opinions on a set of arguments, each opinion being expressed as a labelling, and an aggregated opinion, namely an aggregated labelling, has to be produced. In particular the *ideal* labelling/extension results from the following assumption on the aggregation procedure:

- each participant tries to accept as many arguments as possible, that is, the opinions to be aggregated correspond to the set of the preferred labellings/extensions;
- each argument is (tentatively) accepted or rejected only if there is unanimity on it by all participants, otherwise it is regarded as undecided;
- the resulting labelling/extension may or may not correspond to a defensible position, namely may or may not be admissible: if it is not, water it down (by abstaining about some tentatively accepted or rejected arguments in the aggregated judgment) until it becomes defensible

In order to formally define the concept of the ideal labelling, according to the intuition outlined above, we first need to treat some preliminaries (see [Caminada and Pigozzi, 2011]).

Definition 3.44 *Let $\mathcal{L}ab_1$ and $\mathcal{L}ab_2$ be labellings of an argumentation framework $AF = \langle Ar, att \rangle$. We say that $\mathcal{L}ab_2$ is* more or equally committed *than $\mathcal{L}ab_1$ ($\mathcal{L}ab_1 \sqsubseteq \mathcal{L}ab_2$) iff $\mathtt{in}(\mathcal{L}ab_1) \subseteq \mathtt{in}(\mathcal{L}ab_2)$ and $\mathtt{out}(\mathcal{L}ab_1) \subseteq \mathtt{out}(\mathcal{L}ab_2)$. We say that $\mathcal{L}ab_2$ is* compatible *with $\mathcal{L}ab_1$ ($\mathcal{L}ab_1 \approx \mathcal{L}ab_2$) iff $\mathtt{in}(\mathcal{L}ab_1) \cap \mathtt{out}(\mathcal{L}ab_2) = \emptyset$ and $\mathtt{out}(\mathcal{L}ab_1) \cap \mathtt{in}(\mathcal{L}ab_2) = \emptyset$.*

It holds that "$\sqsubseteq$" defines a partial order (reflexive, anti-symmetric, transitive) on the labellings of an argumentation framework. We can therefore talk about a labelling being "bigger" or "smaller" than another labelling with respect to "$\sqsubseteq$". The relation "$\approx$", although reflexive and symmetric, is not an equivalence relation, since it does not satisfy transitivity.[12] It holds that "$\sqsubseteq$" is at least as strong as "$\approx$"; that is, if $\mathcal{L}ab_1 \sqsubseteq \mathcal{L}ab_2$ then $\mathcal{L}ab_1 \approx \mathcal{L}ab_2$.[13]

The idea of "$\sqsubseteq$" is to define what it means for a labelling to be more committed than another labelling (this is a special case of skepticism comparison, an issue which will be dealt with systematically in Section 4). For instance, the grounded labelling is the least committed labelling among all complete labellings. The idea of "$\approx$" is to define when a labelling of one person might still be acceptable to another person. To see this, first consider that by requiring that $\mathtt{in}(\mathcal{L}ab_1) \cap \mathtt{out}(\mathcal{L}ab_2) = \emptyset$ and $\mathtt{out}(\mathcal{L}ab_1) \cap \mathtt{in}(\mathcal{L}ab_2) = \emptyset$, the relation "$\approx$" does not allow for conflicts between `in` and `out`. That is, if there is an argument that is accepted by agent Ag_1 but rejected by agent Ag_2 (or vice versa) then their labellings are not compatible. However, it is less problematic to have conflicts between `in` and `undec`, or between `out` and `undec`. Thus, compatibility provides an indication of how easy or difficult it is to share a position that is not one's own. It is easier to do this for a labelling that is compatible than for a labelling that is not compatible. In the former case the worst that can happen

[12] As a counterexample, consider $AF = (\{a, b\}, \{(a, b), (b, a)\})$. Let $\mathcal{L}ab_1 = (\{a\}, \{b\}, \emptyset)$, $\mathcal{L}ab_2 = (\emptyset, \emptyset, \{a, b\})$ and $\mathcal{L}ab_3 = (\{b\}, \{a\}, \emptyset)$. It holds that $\mathcal{L}ab_1 \approx \mathcal{L}ab_2$ and $\mathcal{L}ab_2 \approx \mathcal{L}ab_3$ but $\mathcal{L}ab_1 \not\approx \mathcal{L}ab_3$.

[13] This is because $\mathcal{L}ab_1 \approx \mathcal{L}ab_2$ i $\mathtt{in}(\mathcal{L}ab_1) \subseteq \mathtt{in}(\mathcal{L}ab_2) \cup \mathtt{undec}(\mathcal{L}ab_2)$ and $\mathtt{out}(\mathcal{L}ab_1) \subseteq \mathtt{out}(\mathcal{L}ab_2) \cup \mathtt{undec}(\mathcal{L}ab_2)$.

is that one has to abstain from something one accepts or rejects (or have to accept or reject something where one did not have an explicit opinion about). In the latter case, however, one has to make statements that go directly against one's private position.

To come back to the informal description of ideal semantics, we assume a meeting in which every preferred labelling is represented. The meeting then discusses each argument, one by one, with the aim to define an *initial labelling*. If everybody agrees that the argument is labelled `in` (that is, the argument is labelled `in` in every preferred labelling) then the argument is also labelled `in` in the initial tentative labelling. If everybody agrees that the argument is labelled `out` (that is, the argument is labelled `out` in every preferred labelling) then the argument is labelled `out` in the tentative labelling. In all other cases, the argument is labelled `undec` in the tentative labelling. After this process is over, and the tentative labelling has been finished, the meeting goes to the second phase, in which the initial labelling is "watered down" in order to become an admissible labelling. This is done by iteratively relabelling each argument that is illegally `in` or illegally `out` to `undec`. When there are no more arguments left that are illegally `in` or illegally `out`, the result is the *ideal labelling*. It was proved in [Caminada and Pigozzi, 2011] that this process results in constructing the most committed ("biggest") admissible labelling that is less or equally committed than each preferred labelling. This leads to the following definition of ideal semantics.

Definition 3.45 *Let $AF = \langle Ar, att \rangle$ be an argumentation framework. The* ideal labelling *of AF is the biggest admissible labelling that is smaller or equal to each preferred labelling.*

The uniqueness of the ideal labelling and the fact that the ideal labelling is a complete labelling have been proved in [Caminada and Pigozzi, 2011]. Since the grounded labelling is the smallest complete labelling (w.r.t. "$\sqsubseteq$") it directly follows that the ideal labelling is bigger or equal to the grounded labelling.

Proposition 3.46 *Let $AF = \langle Ar, att \rangle$ be an argumentation framework, let $\mathcal{L}ab_{grounded}$ be its grounded labelling and $\mathcal{L}ab_{ideal}$ be its ideal labelling. It holds that $\mathcal{L}ab_{grounded} \sqsubseteq \mathcal{L}ab_{ideal}$.*

There are several ways of describing the ideal labelling [Caminada, 2011].

Proposition 3.47 *Let $\mathcal{L}ab$ be a labelling of an argumentation framework $AF = \langle Ar, att \rangle$. The following statements are equivalent.*

1. *$\mathcal{L}ab$ is the biggest admissible labelling that is smaller or equal to each preferred labelling*

2. *$\mathcal{L}ab$ is the biggest admissible labelling that is compatible with each admissible labelling*

3. *$\mathcal{L}ab$ is the biggest admissible labelling that is compatible with each complete labelling*

4. *$\mathcal{L}ab$ is the biggest admissible labelling that is compatible with each preferred labelling*

The concept of ideal semantics was originally introduced in terms of extensions in [Dung *et al.*, 2007], drawing inspiration from the analogous concept of ideal sceptical semantics in extended logic programs [Alferes *et al.*, 1993].

Definition 3.48 *Let $AF = \langle Ar, att \rangle$ be an argumentation framework. An admissible set $\mathcal{A}rgs$ is called* ideal *iff it is a subset of each preferred extension. The* ideal extension *of AF is a maximal (w.r.t. set-inclusion) ideal set.*

It turns out that the ideal extension of any argumentation framework AF, denoted in the following as $ID(AF)$, is unique (which implies that it is also the biggest ideal set) and that it is also a complete extension [Dung *et al.*, 2007]. It then follows directly that the ideal extension is a superset of the grounded extension.

Proposition 3.49 *Let $AF = \langle Ar, att \rangle$ be an argumentation framework. It holds that $GE(AF) \subseteq ID(AF)$.*

There are several ways of describing the ideal extension.

Proposition 3.50 *Let $AF = \langle Ar, att \rangle$ be an argumentation framework, and let $\mathcal{A}rgs \subseteq Ar$. The following statements are equivalent.*

1. *$\mathcal{A}rgs$ is the biggest admissible set that is a subset of each preferred extension*

2. *$\mathcal{A}rgs$ is the biggest admissible set that is not attacked by any admissible set*

3. *$\mathcal{A}rgs$ is the biggest admissible set that is not attacked by any complete extension*

4. *$\mathcal{A}rgs$ is the biggest admissible set that is not attacked by any preferred extension*

In Proposition 3.50 the equivalence between points 1 and 2 follows from [Dung *et al.*, 2007, Theorem 3.3]. The equivalence between points 2, 3 and 4 follows from the fact that an argument (or set) is attacked by an admissible set iff it is attacked by a complete extension iff it is attacked by a preferred extension.

The bijective labellings-extensions correspondence through `Ext2Lab` (and `Lab2Ext`) also holds for ideal semantics [Caminada, 2011].

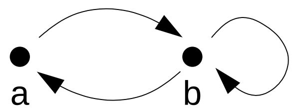

Figure 10. The ideal labelling can be less skeptical than the grounded labelling

Ideal semantics is similar to grounded semantics in the sense that it always yields a unique labelling (extension). Actually it can be seen that the ideal labelling (extension) coincides with the grounded labelling (extension) in the examples of Figures 5, 6 and 7. In particular, referring to extensions, in Figure 5 the intersection of preferred extensions $\{a\}$ coincides with the grounded extension; in Figure 6 the intersection of preferred extensions $\{d\}$ is not admissible and its only admissible subset is the empty set (coinciding with the grounded extension); in Figure 7 there is only one (empty) preferred extension, which coincides with the grounded and ideal extension.

However, as shown in Propositions 3.46 and 3.49, in general ideal semantics tends to be less skeptical than grounded semantics. As an example, in the argumentation framework of Figure 10 the grounded labelling is $(\emptyset, \emptyset, \{a, b\})$ (the grounded extension is $\emptyset$) whereas the ideal labeling is $(\{a\}, \{b\}, \emptyset)$ (the ideal extension is $\{a\}$).

To determine whether an argument is an element of the ideal extension, point 2 of Proposition 3.50 implies that it is sufficient to determine whether it is an element of an admissible set that is not attacked by any admissible set. Proof procedures for this are straightforward and have been described in [Dung *et al.*, 2007].

An alternative approach that is very close to ideal semantics is that of *eager* semantics ($\mathcal{EAG}$) [Caminada, 2007b]. Where the ideal extension is the (unique) biggest admissible (and complete) subset of each preferred extension, the eager extension is the (unique[14]) biggest admissible (and complete) subset of each semi-stable extension. This of course admits an equivalent formulation in terms of labellings.

Definition 3.51 *Let $AF = \langle Ar, att \rangle$ be an argumentation framework. The* eager extension *of AF is the maximal (w.r.t. set-inclusion) admissible set which is included in every semi-stable extension.*

Definition 3.52 *Let $AF = \langle Ar, att \rangle$ be an argumentation framework. The* eager labelling *of AF is the biggest admissible labelling that is smaller or equal to each semi-stable labelling.*

The eager extension is a superset of the ideal extension, making eager semantics (to the best of our knowledge) the most credulous unique status semantics

[14]Uniqueness of the eager extension is guaranteed for nite and nitary frameworks, while it may not hold in general for in nite frameworks.

that has been proposed in the literature. The eager extension and the associated eager labelling can be computed by first calculating all semi-stable labellings (using for instance the algorithm of [Caminada, 2007a]) and subsequently applying the judgement aggregation operators specified in [Caminada and Pigozzi, 2011].

3.9 Stage Semantics

The concept of *stage* semantics ($\mathcal{STG}$) has been introduced in [Verheij, 1996] and further developed in [Verheij, 2003] in different formal settings with respect to the ones considered in this chapter. Precise (and rather straightforward) correspondences can nevertheless be drawn so that we can describe stage semantics in terms of labellings and extensions, as was done for all other semantics in this chapter. In essence, a stage labelling is a conflict-free labelling where `undec` is minimal, while a stage extension is a conflict-free set of arguments $\mathcal{A}rgs$, where $\mathcal{A}rgs \cup \mathcal{A}rgs^{+}$ is maximal.

Definition 3.53 *Let $AF = \langle Ar, att \rangle$ be an argumentation framework. A labelling $\mathcal{L}ab$ is called a* stage labelling *of AF iff it is a conflict-free labelling where* $\texttt{undec}(\mathcal{L}ab)$ *is minimal (w.r.t. set-inclusion) among all conflict-free labellings, i.e. there is no conflict-free labelling $\mathcal{L}ab'$ such that* $\texttt{undec}(\mathcal{L}ab') \subsetneq \texttt{undec}(\mathcal{L}ab)$.

Definition 3.54 *Let $AF = \langle Ar, att \rangle$ be an argumentation framework. A* stage extension *of AF is a conflict-free set $\mathcal{A}rgs \subseteq Ar$ where $\mathcal{A}rgs \cup \mathcal{A}rgs^{+}$ is maximal (w.r.t. set inclusion) among all conflict-free sets, i.e there is no conflict-free set $\mathcal{A}rgs'$ such that $(\mathcal{A}rgs' \cup \mathcal{A}rgs'^{+}) \supsetneq (\mathcal{A}rgs \cup \mathcal{A}rgs^{+})$.*

It holds that every stable labelling (extension) is also a stage labelling (extension).

Theorem 3.55 *Let $\mathcal{L}ab$ be a labelling of an argumentation framework $AF = (Ar, att)$. If $\mathcal{L}ab$ is a stable labelling of AF then $\mathcal{L}ab$ is also a stage labelling of AF.*

Theorem 3.56 *Let $AF = \langle Ar, att \rangle$ be an argumentation framework and $\mathcal{A}rgs \subseteq Ar$. If $\mathcal{A}rgs$ is a stable extension of AF then $\mathcal{A}rgs$ is also a stage extension of AF.*

If there exists at least one stable labelling (extension), then each stage labelling (extension) is also a stable labelling (extension).

Theorem 3.57 *Let $AF = \langle Ar, att \rangle$ be an argumentation framework. If there exists at least one stable labelling of AF then every stage labelling is also a stable labelling.*

Theorem 3.58 *Let $AF = \langle Ar, att \rangle$ be an argumentation framework. If there exists at least one stable extension of AF then every stage extension is also a stable extension.*

There also exists an alternative way to describe the concept of stage semantics. In essence, a stage labelling is a stable labelling of a maximal subgraph of the argumentation framework that has at least one stable labelling, augmented with `undec` labels for the arguments that did not make their way into the subgraph. Similarly, what a stage extension does is taking a maximal subgraph of the argumentation framework that has at least one stable extension. A stage extension is then a stable extension of such a maximal subgraph.

Theorem 3.59 *Let $\mathcal{L}ab$ be a labelling of an argumentation framework $AF = \langle Ar, att \rangle$. The following two statements are equivalent.*

1. *$\mathcal{L}ab$ is a conflict-free labelling where $\mathtt{undec}(\mathcal{L}ab)$ is minimal (w.r.t. set inclusion) among all conflict-free labellings*
2. *$\mathcal{A}rgs = \mathtt{in}(\mathcal{L}ab) \cup \mathtt{out}(\mathcal{L}ab)$ is a maximal subset of Ar such that $AF\!\downarrow_{\mathcal{A}rgs}$ has a stable labelling, and $\mathcal{L}ab\!\downarrow_{\mathcal{A}rgs}$ is a stable labelling of $AF\!\downarrow_{\mathcal{A}rgs}$.*

Theorem 3.60 *Let $AF = \langle Ar, att \rangle$ be an argumentation framework and $\mathcal{A}rgs \subseteq Ar$. The following two statements are equivalent.*

1. *$\mathcal{A}rgs$ is a conflict-free set where $\mathcal{A}rgs \cup \mathcal{A}rgs^+$ is maximal (w.r.t. set inclusion) among all conflict-free sets.*
2. *$\mathcal{A}rgs \cup \mathcal{A}rgs^+$ is a maximal subset of Ar such that $AF\!\downarrow_{\mathcal{A}rgs \cup \mathcal{A}rgs^+}$ has a stable extension, and $\mathcal{A}rgs$ is a stable extension of $AF\!\downarrow_{\mathcal{A}rgs \cup \mathcal{A}rgs^+}$.*

The bijective labellings-extensions correspondence through `Ext2Lab` (and `Lab2Ext`) also holds for stage semantics, as proved in [Caminada, 2011]. An algorithm that produces all stage labellings (and therefore also all stage extensions) is described in [Caminada, 2010].

To exemplify stage labellings (extensions) let us refer as usual to the examples of Figures 5-7. Stage labellings (extensions) coincide with stable labellings (extensions), when the latter exist, as in the case of Figures 5 and 6. On the other hand, in the case of Figure 7, differently from all other semantics examined so far, stage semantics prescribes three non-trivial labellings, namely $(\{a\}, \{b\}, \{c\})$, $(\{b\}, \{c\}, \{a\})$, $(\{c\}, \{a\}, \{b\})$ (and of course the corresponding three non-empty extensions, $\{a\}$, $\{b\}$, and $\{c\}$).

Using the technical properties and the examples described above, we are now ready to describe the intuition behind stage semantics. In essence, stage semantics shares with stable semantics some sort of preference for strongly committed evaluations with respect to the undecided ones. As already seen, such an attitude is not universally applicable: the solution of stage semantics is to consider the maximal restrictions where this attitude is still applicable. In other terms, stage semantics can be interpreted as the attempt to identify and then ignore the minimal amounts of information that prevent the application of a black-and-white view of the world. Note that different information can

Figure 11. Stage semantics differs from semi-stable semantics

be ignored in different labellings (extensions), for instance in the example of Figure 7 arguments a, b, and c are alternatively ignored.

The idea of minimizing the set of `undec`-labelled arguments or, alternatively, of maximizing the range ($\mathcal{A}rgs \cup \mathcal{A}rgs^+$) of extensions is common to stage and semi-stable semantics. However, where semi-stable semantics aims to maximize the range under the condition of admissibility, stage semantics tries to maximize the range under the weaker condition of conflict-freeness. As shown above, this amounts to taking the stable labellings (extensions) of the biggest subframework that has at least one stable labelling (extension). Hence, the approach of stage semantics is comparable with the approach of handling inconsistent knowledge bases, where one can select maximal consistent subsets of the knowledge base, and then examine what holds in all of them (in the intersection of all their models). That is, it is as if stage semantics interprets the absence of stable labellings/extensions as some form of "inconsistency", which needs to be handled taking the "maximal consistent subframeworks". On the other hand, in semi-stable semantics as well as in most other semantics all arguments play a role in all extensions/labellings. In particular, an undecided argument keeps the capability to cause other arguments to be undecided, while this is not the case in stage semantics. An example is shown in Figure 11. Here, the other semantics considered up to now in the current chapter (with the exception of naïve semantics) yield a single labelling $(\emptyset, \emptyset, \{a, b\})$ corresponding to the extension $\emptyset$, whereas stage semantics yields a single labelling $(\{b\}, \emptyset, \{a\})$ corresponding to the extension $\{b\}$. In essence, what stage semantics does is to ignore argument a, since this argument causes the framework not to have any stable labelling/extension. $\mathcal{CF}2$ and stage2 semantics, examined in the next section, show the same behaviour in this example.

Another example to illustrate the difference between stage semantics and semi-stable semantics is given in Figure 12. Here, semi-stable semantics yields a single extension $\{a\}$, corresponding to a labelling $(\{a\}, \{b\}, \{c\})$. Stage semantics yields two extensions, the first one being equivalent to the one yielded by semi-stable semantics, the second one being $\{b\}$, corresponding to a labelling $(\{b\}, \{c\}, \{a\})$. The first stage extension (labelling) is the result of ignoring argument c, the second stage extension (labelling) is the result of ignoring argument a. For both possibilities, the remaining argumentation framework is a maximal one that has at least one stable extension (labelling). It can therefore be observed that under stage semantics, even an argument without any attackers (like argument a in Figure 12) is not always labelled `in`. With any

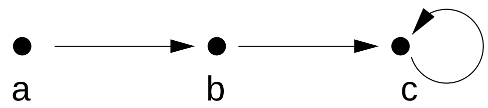

Figure 12. A peculiar case for stage semantics

other semantics considered in this chapter[15], however, an argument without any attackers is *always* labelled `in`.

3.10 $\mathcal{CF}2$ and stage2 semantics

With the exception of naïve and stage semantics, all semantics reviewed so far are admissibility-based, i.e. the labellings (extensions) they prescribe are admissible. Moreover they are compatible with the basic skeptical view represented by grounded semantics, in the sense that in any of their labellings (extensions) the accepted arguments are a superset of those accepted by the grounded semantics. Focusing now on those of these semantics which are multiple-status (namely complete, preferred, stable and semi-stable), one can notice that odd-length unidirectional attack cycles cause a sort of singularity in their behaviour. For instance, considering the example of Figure 7 only the trivial labelling (extension) is prescribed and, in the case of stable semantics, no labelling (extension) at all exists. This gives rise to a sort of unbalanced treatment of even-length and odd-length unidirectional attack cycles: non-trivial labellings (extensions) exist for the former ones, while they do not exist for the latter. This has been regarded as problematic by [Pollock, 2001], since in some contexts an "equal" treatment of cycles, independently of their length, can be more appropriate[16]. It is evident that this requires giving up the property of admissibility, as no non-trivial admissible labellings (extensions) exist for the example of Figure 7. In fact, the behaviour of stage semantics goes in that direction, since in the example of Figure 7 it prescribes three non-trivial labellings, namely $(\{a\}, \{b\}, \{c\})$, $(\{b\}, \{c\}, \{a\})$, $(\{c\}, \{a\}, \{b\})$, or, analogously, three non-empty extensions, namely $\{a\}$, $\{b\}$, $\{c\}$. Stage semantics however shows a peculiar behaviour and strongly departs from grounded semantics in some cases. As already commented in Section 3.9 a stage labelling (or extension) may even exclude from acceptance an unattacked argument (a in the example of Figure 12) while including an argument attacked by it (b in the same example). This kind of behaviour has no parallel in all other semantics considered in this chapter and, as such, appears rather hard to justify. Then the question arises as to whether it is possible to define a multiple-status semantics which is not admissibility-based, treats in an "equal" way odd and even-length

[15] With the exception of na ve semantics, which ignores the direction of attacks.

[16] [Pollock, 2001] discusses odd-length attack cycles in the context of a set of reference inference graphs for testing the intuitive validity of justi cation status assignments. Actually, the paper where the problem is raised [Pollock, 2001] is mainly focused on an approach to reasoning with variable degrees of justi cation and does not provide an explicit solution to this problematic example.

unidirectional attack cycles, while preserving compatibility with the grounded semantics in any case.

$\mathcal{CF}2$ [Baroni and Giacomin, 2003; Baroni *et al.*, 2005b] and *stage2* ($\mathcal{STG}2$) [Dvořák and Gaggl, 2012b; Dvořák and Gaggl, 2012a; Dvořák and Gaggl, 2016] semantics satisfy the above requirements. In fact, to achieve this objective a relatively sophisticated semantics definition scheme has been devised called *SCC-recursiveness*. The SCC-recursive scheme is based on the graph theoretical notion of a strongly connected component (SCC). In short, strongly connected components provide a unique partition of a directed graph into disjoint parts where all nodes are mutually reachable (it is assumed that reachability is a reflexive relation). Formally, strongly connected components are the equivalence classes induced by the path equivalence (i.e. mutual reachability) relation between nodes. To illustrate this notion, in the example of Figure 5 there are three SCCs, namely $\{a\}$, $\{b\}$, and $\{c,d\}$, in Figure 6 there are three SCCs too, namely $\{a,b\}$, $\{c\}$, and $\{d\}$, while the argumentation framework of Figure 7 consists of a unique SCC, namely $\{a,b,c\}$. As another example, in Figure 9 there are two SCCs, namely $\{a,b\}$ and $\{c,d,e\}$.

An important property of the SCC decomposition is that the graph obtained considering SCCs as single nodes is acyclic, i.e. the attack relation induces a partial order between the SCCs. The SCC-recursive scheme exploits this property and can be intuitively regarded as a constructive procedure to incrementally build extensions (or labellings) following the partial order of SCCs. In a nutshell, one "locally" applies some extension selection criterion to the *initial* SCCs, i.e. those not receiving attacks from other ones. Then, for each possible choice identified in the initial SCCs, one accordingly suppresses some arguments from the initial argumentation framework and the procedure is recursively applied to the new argumentation framework resulting from this modification, until no remaining arguments are left to process. In the case of $\mathcal{CF}2$ semantics, the "local" selection criterion[17] applied to SCCs is quite simple and corresponds to the intuition underlying naïve semantics: all maximal conflict free sets are selected. In the case of stage2 semantics the criterion corresponds to stage semantics, namely the conflict free sets with maximal range are selected. However embedding these criteria within the SCC recursive scheme gives rise to different results with respect to the original naïve and stage semantics.

We now provide a formal definition of $\mathcal{CF}2$ and stage2 semantics in terms of extensions (as this is their original and easier to follow formulation), exemplify their behaviour and review their properties. For further details and more extensive explanations of the SCC-recursive scheme the reader may refer to the original source [Baroni *et al.*, 2005b] and to chapter 18 in this volume. A labelling-based formulation of $\mathcal{CF}2$ and stage2 semantics will be examined at the end of the section.

[17] It can be remarked that all Dung's original semantics can be equivalently characterized using SCC-recursive de nitions similar to De nition 3.61, as proved in [Baroni *et al.*, 2005b].

Definition 3.61 *Given an argumentation framework* $AF = \langle Ar, att \rangle$*, a set* $\mathcal{A}rgs \subseteq Ar$ *is an* extension of $\mathcal{CF}2$ semantics *if and only if*

- $\mathcal{A}rgs$ *is a naïve extension of* AF *if* $|\mathrm{SCCS}_{AF}| = 1$
- $\forall S \in \mathrm{SCCS}_{AF}$ $(\mathcal{A}rgs \cap S) \in \mathcal{E}_{CF2}(AF{\downarrow}_{UP_{AF}(S,\mathcal{A}rgs)})$ *otherwise*

where

- SCCS_{AF} *denotes the set of strongly connected components of* AF
- *for any* $\mathcal{A}rgs, S \subseteq Ar$, $UP_{AF}(S, \mathcal{A}rgs) = \{a \in S \mid \nexists b \in \mathcal{A}rgs \setminus S : (b, a) \in att\}$.

Definition 3.62 *Given an argumentation framework* $AF = \langle Ar, att \rangle$*, a set* $\mathcal{A}rgs \subseteq Ar$ *is an* extension of stage2 semantics *if and only if*

- $\mathcal{A}rgs$ *is a stage extension of* AF *if* $|\mathrm{SCCS}_{AF}| = 1$
- $\forall S \in \mathrm{SCCS}_{AF}$ $(\mathcal{A}rgs \cap S) \in \mathcal{E}_{\mathcal{STG}2}(AF{\downarrow}_{UP_{AF}(S,\mathcal{A}rgs)})$ *otherwise*

Definitions 3.61 and 3.62 are quite complicated and their detailed illustration is beyond the scope of the chapter. We remark only that the recursion is well-founded since, in their second branch, the semantics itself is applied to a set of restricted (and disjoint) argumentation frameworks, each including a strictly lesser number of arguments with respect to the original one. This ensures that the base case, namely the application of $\mathcal{CF}2$ or stage2 semantics to an argumentation framework consisting of a single SCC (first branch of Definitions 3.61 and 3.62) is reached in a finite number of steps. Note in particular that an argumentation framework including 0 or 1 arguments necessarily consists of a single SCC.

In spite of technical complications, the idea underlying $\mathcal{CF}2$ and stage2 semantics is relatively simple and can be illustrated with reference to our examples (since in these examples their extensions coincide, we refer to $\mathcal{CF}2$ only in the following description). In Figure 5 there is one initial SCC, namely $\{a\}$, and of course it contains only one maximal conflict-free set, namely $\{a\}$ itself, which is selected for extension building. The subsequent (according to the partial order induced by the attack relation) SCC, namely $\{b\}$, is suppressed as its only element is attacked by the already selected argument a. The last SCC, namely $\{c, d\}$, then remains unaffected by previously selected elements and we can select its maximal conflict-free subsets $\{c\}$ and $\{d\}$ to be combined with the previous selection, leading to the $\mathcal{CF}2$ extensions $\{a, c\}$ and $\{a, d\}$.

In Figure 6 there is one initial SCC, namely $\{a, b\}$, whose maximal conflict-free sets are $\{a\}$ and $\{b\}$, each representing a starting point for further extension construction. As a matter of fact, in both cases the subsequent SCC, namely $\{c\}$ is suppressed, leaving the remaining SCC, $\{d\}$, unaffected and providing

$\{d\}$ itself as maximal conflict-free subset. It turns out that there are two $\mathcal{CF}2$ extensions, namely $\{a, d\}$ and $\{b, d\}$.

The argumentation framework of Figure 7 consists of only one SCC and therefore its $\mathcal{CF}2$ extensions coincide with its maximal conflict-free subsets $\{a\}$, $\{b\}$ and $\{c\}$.

In the example of Figure 9, the application of $\mathcal{CF}2$ semantics definition is more articulated. The (again unique) initial SCC is $\{a, b\}$, which, as in the previous case, yields $\{a\}$ and $\{b\}$ as starting points for further extension construction. Considering $\{a\}$, we have that b is attacked by the extension and the subsequent SCC $\{c, d, e\}$ is left unaffected. As a consequence, all its maximal conflict-free subsets $\{c\}$, $\{d\}$ and $\{e\}$ are available, yielding the three $\mathcal{CF}2$ extensions $\{a, c\}$, $\{a, d\}$ and $\{a, e\}$. Considering $\{b\}$, both a and c are attacked by the extension and therefore suppressed. The restriction of the argumentation framework to the set $\{d, e\}$ then remains to be evaluated. As $\{d\}$ is the initial SCC of this restricted argumentation framework, it is selected and then the subsequent SCC $\{e\}$ is entirely suppressed, yielding a further $\mathcal{CF}2$ extension $\{b, d\}$.

Finally, in the example of Figure 12 a unique $\mathcal{CF}2$ extension is identified, namely $\{a\}$, yielding agreement with grounded semantics.

While in the examples above the extensions of $\mathcal{CF}2$ and stage2 semantics coincide, they can be different in some cases. A simple example is an argumentation framework consisting of six arguments $a_1, \ldots, a_6$ arranged into an attack cycle, i.e. such that a_i attacks a_{i+1} for $i = 1, \ldots, 5$ and a_6 attacks a_1. In this case, consisting of a unique SCC, there are five naïve and $\mathcal{CF}2$ extensions, namely $\{a_1, a_3, a_5\}$, $\{a_2, a_4, a_6\}$, $\{a_1, a_4\}$, $\{a_2, a_5\}$, $\{a_3, a_6\}$, while only two of them (clearly the first ones) are also stage and stage2 extensions.

Having exemplified the behaviour of $\mathcal{CF}2$ and stage2 semantics, we summarize in Proposition 3.63 some of their known properties in relation to other semantics notions (in particular naïve, grounded, stable and preferred).

Proposition 3.63 *For any argumentation framework $AF = \langle Ar, att \rangle$*

- $\mathcal{E}_{\mathcal{STG}2}(AF) \subseteq \mathcal{E}_{CF2}(AF) \subseteq \mathcal{E}_{\mathcal{NA}}(AF)$*;*
- *the grounded extension is included in any $\mathcal{CF}2$ and stage2 extension;*
- *any stable extension is also a $\mathcal{CF}2$ and stage2 extension;*
- *for any preferred extension E there is a $\mathcal{CF}2$ extension E' such that $E \subseteq E'$ (note that this property does not hold for stage2 semantics)*

As mentioned above, $\mathcal{CF}2$ and stage2 semantics have been conceived and defined in the extension-based setting. The same semantic notion can however be expressed using the SCC-recursive scheme in the labelling context.

Definition 3.64 *Given an argumentation framework $AF = \langle Ar, att \rangle$, a labelling $\mathcal{L}ab$ is a $\mathcal{CF}2$* labelling *if and only if*

- *if* $|\text{SCCS}_{AF}| = 1$*,* $\mathcal{L}ab$ *is a conflict-free labelling with maximal* $\mathtt{in}(\mathcal{L}ab)$ *among conflict-free labellings and such that* $a \in \mathtt{in}(\mathcal{L}ab) \Rightarrow a^+ \subseteq \mathtt{out}(\mathcal{L}ab)$*;*

- *otherwise,* $\forall S \in \text{SCCS}_{AF}$ $\mathcal{L}ab{\downarrow}_{UP_{AF}(S,\mathcal{A}rgs)}$ *is a* $\mathcal{CF}2$ *labelling of* $AF{\downarrow}_{UP_{AF}(S,\mathcal{A}rgs)}$ *and all arguments in* $S \setminus UP_{AF}(S, \mathcal{A}rgs)$ *are labelled* $\mathtt{out}$*.*

where all notations are as in Definition 3.61.

Definition 3.65 *Given an argumentation framework* $AF = \langle Ar, att \rangle$*, a labelling* $\mathcal{L}ab$ *is a* stage2 labelling *if and only if*

- *if* $|\text{SCCS}_{AF}| = 1$*,* $\mathcal{L}ab$ *is a stage labelling;*

- *otherwise,* $\forall S \in \text{SCCS}_{AF}$ $\mathcal{L}ab{\downarrow}_{UP_{AF}(S,\mathcal{A}rgs)}$ *is a stage2 labelling of* $AF{\downarrow}_{UP_{AF}(S,\mathcal{A}rgs)}$ *and all arguments in* $S \setminus UP_{AF}(S, \mathcal{A}rgs)$ *are labelled* $\mathtt{out}$*.*

where all notations are as in Definition 3.61.

By inspection of Definitions 3.61 vs. 3.64 and 3.62 vs. 3.65, it can be seen that the bijective labellings-extensions correspondence through `Ext2Lab` (and `Lab2Ext`) holds for $\mathcal{CF}2$ and stage2 semantics.

3.11 Roundup

We now provide an overview of how the semantics that have been treated until now are related. In Figure 13 we graphically depict what can be seen as an ontology of argumentation semantics. The figure shows for instance that every stable labelling is also a stage labelling, a semi-stable labelling and a $\mathcal{CF}2$ labelling, that every semi-stable labelling is also a preferred labelling, etc.

The same relations of Figure 13 also hold for the extension-based approach. In Table 4 we provide an overview of how the admissibility-based semantics can be expressed in terms of complete labellings.

As explicitly stated in Section 2, this chapter is focused on finite argumentation frameworks and the analysis of semantics properties we have carried out relies on this assumption. One may wonder what is the impact of this restriction and what would be the implications of considering also infinite frameworks. While providing a full answer to this question is beyond the scope of this chapter, we observe in particular that in infinite frameworks the notion of maximality w.r.t set inclusion is less immediate than in finite frameworks and the existence of maximal sets of arguments respecting some criterion, which is guaranteed in finite frameworks, may fail to be achieved in infinite ones. As an example of the consequences of this fact, a semantics which is universally defined or unique status in the context of finite frameworks may not be so when considering also infinite ones, implying (among other consequences) that the skepticism comparison of Section 4 does not extend directly to the infinite case. The reader may refer to chapter 17 for a treatment of these issues.

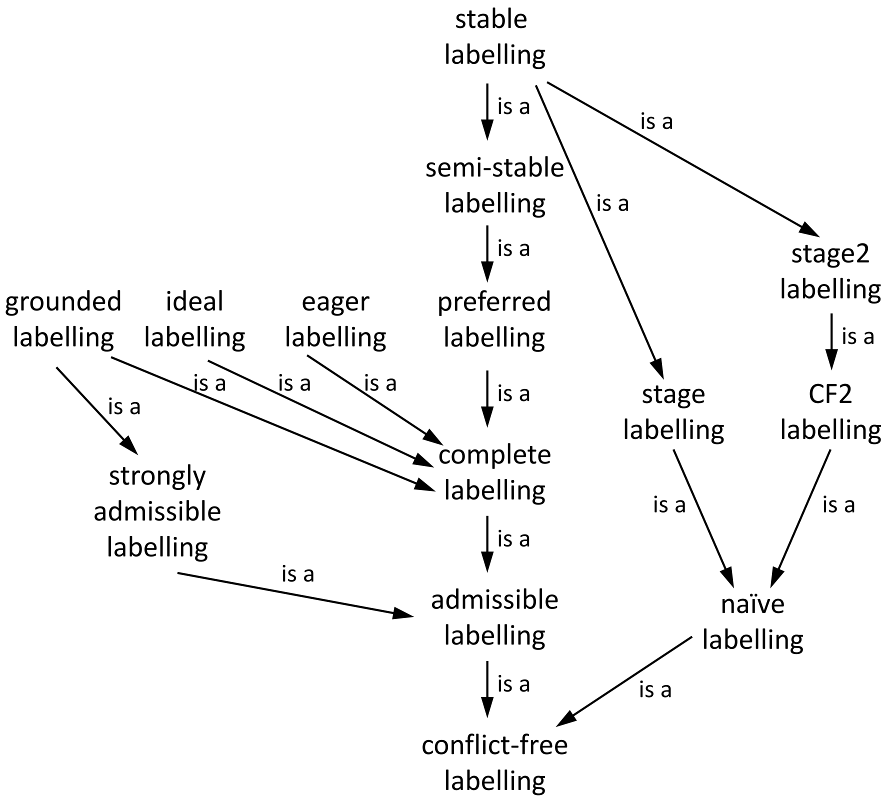

Figure 13. Relations among alternative labelling and extension notions

Table 4. Describing admissibility based semantics in terms of complete labellings

restriction on complete labelling	resulting semantics
no restrictions	complete semantics
empty `undec`	stable semantics
maximal `in`	preferred semantics
maximal `out`	preferred semantics
maximal `undec`	grounded semantics
minimal `in`	grounded semantics
minimal `out`	grounded semantics
minimal `undec`	semi-stable semantics
maximal w.r.t. $\sqsubseteq$ while compatible with each complete labelling	ideal semantics

Fig. #	$\mathcal{E}_{\mathcal{NA}}(AF)$	$\mathcal{E}_{\mathcal{CO}}(AF)$	$\mathcal{E}_{\mathcal{GR}}(AF)$	$\mathcal{E}_{\mathcal{PR}}(AF)$	$\mathcal{E}_{\mathcal{ST}}(AF)$
Fig. 3	$\{\{a,c\},\{b\}\}$	$\{\{a,c\}\}$	$\{\{a,c\}\}$	$\{\{a,c\}\}$	$\{\{a,c\}\}$
Fig. 4	$\{\{a\},\{b\}\}$	$\{\emptyset,\{a\},\{b\}\}$	$\{\emptyset\}$	$\{\{a\},\{b\}\}$	$\{\{a\},\{b\}\}$
Fig. 5	$\{\{a,c\},\{a,d\},\{b,d\}\}$	$\{\{a\},\{a,c\},\{a,d\}\}$	$\{\{a\}\}$	$\{\{a,c\},\{a,d\}\}$	$\{\{a,c\},\{a,d\}\}$
Fig. 6	$\{\{a,d\},\{b,d\},\{c\}\}$	$\{\emptyset,\{a,d\},\{b,d\}\}$	$\{\emptyset\}$	$\{\{a,d\},\{b,d\}\}$	$\{\{a,d\},\{b,d\}\}$
Fig. 7	$\{\{a\},\{b\},\{c\}\}$	$\{\emptyset\}$	$\{\emptyset\}$	$\{\emptyset\}$	$\emptyset$
Fig. 9	$\{\{a,c\},\{a,d\},\{a,e\},\{b,d\},\{b,e\}\}$	$\{\emptyset,\{a\},\{b,d\}\}$	$\{\emptyset\}$	$\{\{a\},\{b,d\}\}$	$\{\{b,d\}\}$
Fig. 10	$\{\{a\}\}$	$\{\emptyset,\{a\}\}$	$\{\emptyset\}$	$\{\{a\}\}$	$\{\{a\}\}$
Fig. 11	$\{\emptyset\}$	$\{\emptyset\}$	$\{\emptyset\}$	$\{\emptyset\}$	$\emptyset$
Fig. 12	$\{\{a\}\}$	$\{\{a\}\}$	$\{\{a\}\}$	$\{\{a\}\}$	$\{\{a\}\}$

Table 5. A summary of the examples: part I

Fig. #	$\mathcal{E}_{\mathcal{SST}}(AF)$	$\mathcal{E}_{\mathcal{ID}}(AF)$	$\mathcal{E}_{\mathcal{EAG}}(AF)$	$\mathcal{E}_{\mathcal{STG}}(AF)$	$\mathcal{E}_{\mathcal{CF}2}(AF)$	$\mathcal{E}_{\mathcal{STG}2}(AF)$
Fig. 3	$\{\{a,c\}\}$	$\{\{a,c\}\}$	$\{\{a,c\}\}$	$\{\{a,c\}\}$	$\{\{a,c\}\}$	$\{\{a,c\}\}$
Fig. 4	$\{\{a\},\{b\}\}$	$\{\emptyset\}$	$\{\emptyset\}$	$\{\{a\},\{b\}\}$	$\{\{a\},\{b\}\}$	$\{\{a\},\{b\}\}$
Fig. 5	$\{\{a,c\},\{a,d\}\}$	$\{\{a\}\}$	$\{\{a\}\}$	$\{\{a,c\},\{a,d\}\}$	$\{\{a,c\},\{a,d\}\}$	$\{\{a,c\},\{a,d\}\}$
Fig. 6	$\{\{a,d\},\{b,d\}\}$	$\{\emptyset\}$	$\{\emptyset\}$	$\{\{a,d\},\{b,d\}\}$	$\{\{a,d\},\{b,d\}\}$	$\{\{a,d\},\{b,d\}\}$
Fig. 7	$\{\emptyset\}$	$\{\emptyset\}$	$\{\emptyset\}$	$\{\{a\},\{b\},\{c\}\}$	$\{\{a\},\{b\},\{c\}\}$	$\{\{a\},\{b\},\{c\}\}$
Fig. 9	$\{\{b,d\}\}$	$\{\emptyset\}$	$\{\{b,d\}\}$	$\{\{b,d\}\}$	$\{\{a,c\},\{a,d\},\{a,e\},\{b,d\}\}$	$\{\{a,c\},\{a,d\},\{a,e\},\{b,d\}\}$
Fig. 10	$\{\{a\}\}$	$\{\{a\}\}$	$\{\{a\}\}$	$\{\{a\}\}$	$\{\{a\}\}$	$\{\{a\}\}$
Fig. 11	$\{\emptyset\}$	$\{\emptyset\}$	$\{\emptyset\}$	$\{\{b\}\}$	$\{\{b\}\}$	$\{\{b\}\}$
Fig. 12	$\{\{a\}\}$	$\{\{a\}\}$	$\{\{a\}\}$	$\{\{a\},\{b\}\}$	$\{\{a\}\}$	$\{\{a\}\}$

Table 6. A summary of the examples: part II

4 Argument Justi cation and Skepticism

4.1 The notion of justification status

Argumentation semantics can allow for the presence of more than one extension (or labelling) of arguments, following a tradition in nonmonotonic reasoning [Reiter, 1980; Pollock, 1995]. Hence, when one is interested in the overall status of a particular argument, one needs to have a way of taking the multiplicity of extensions (or labellings) into account. At a basic level, two very simple (and, in a sense, extreme) alternatives for the notion of justification status can be considered: *skeptical justification* requires that an argument is accepted in all labellings (or extensions), while *credulous justification* requires that an argument is accepted in at least one labelling (or extension). This is formalized in Definitions 4.2 and 4.3. Note that we assume, as in previous literature [Baroni and Giacomin, 2007b; Baroni and Giacomin, 2009b], that using these simple notions is meaningful only when the set of labellings or extensions is not empty, otherwise the basis for evaluation is lacking. To express this in a concise way, we introduce a specific notation in Definition 4.1.

Definition 4.1 *Given a labelling-based semantics* σ, $\mathcal{DL}_\sigma = \{AF : \mathcal{L}_\sigma(AF) \neq \emptyset\}$. *Given an extension-based semantics* σ, $\mathcal{DE}_\sigma = \{AF : \mathcal{E}_\sigma(AF) \neq \emptyset\}$.

Definition 4.2 (course-grained justification status, labelling-based) *Given a labelling-based semantics* σ *and an argumentation framework* $AF \in \mathcal{DL}_\sigma$, *an argument* a *is* skeptically justified *(or* skeptically accepted*) if* $\forall \mathcal{L}ab \in \mathcal{L}_\sigma(AF)\ \mathcal{L}ab(a) = \mathtt{in}$; *an argument* a *is* credulously justified *(or* credulously accepted*) if* $\exists \mathcal{L}ab \in \mathcal{L}_\sigma(AF) : \mathcal{L}ab(a) = \mathtt{in}$.

Definition 4.3 (course-grained justification status, extension-based) *Given an extension-based semantics* σ *and an argumentation framework* $AF \in \mathcal{DE}_\sigma$, *an argument* a *is* skeptically justified *(or* skeptically accepted*) if* $\forall E \in \mathcal{E}_\sigma(AF)\ a \in E$; *an argument* a *is* credulously justified *(or* credulously accepted*) if* $\exists E \in \mathcal{E}_\sigma(AF) : a \in E$.

Skeptical justification implies credulous justification, as long as the set of extensions (or labellings) is not empty. Also, a third justification status can be derived: an argument is *not justified* (or *rejected*) if it is not credulously justified (and hence also not skeptically justified, assuming that the set of extensions (or labellings) is not empty).

It can be noted that in any unique-status semantics skeptical and credulous acceptance coincide, so that an argument can only be accepted or rejected. In this context it is possible, however, to consider two levels of rejection, in fact a rejected argument can be attacked or not by the unique extension (or, analogously, can be labelled `out` or `undec` in the unique labelling). The former case corresponds to a stronger form of rejection (these arguments have been sometimes called *defeated outright* in the literature [Pollock, 1992]) while in the latter case rejection is clearly weaker (these arguments being called *provisionally defeated* according to the same terminology).

While the brief remarks above correspond to the prevailing approaches to the notion of justification status in the literature, one may observe that a more systematic treatment is possible, by combining the ideas concerning the status of an argument with respect to a single labelling (or extension) and those referring to a plurality of them. In fact, an argument can be in one of three possible states with respect to a single labelling (namely, `in`, `out` or `undec`) and correspondingly can be accepted, defeated outright or provisionally defeated with respect to a single extension. If a plurality of labellings (or extensions) is considered, we can ask ourselves the following three questions: is the argument accepted (labelled `in`) by at least one extension (labelling), is the argument rejected outright (labelled `out`) by at least one extension (labelling), and is the argument provisionally rejected (labelled `undec`) by at least one extension (labelling)? The answer to these three questions gives rise to a concept we will sometimes refer to as the *fine-grained justification status*, to distinguish it from the earlier introduced concepts of sceptical justification and credulous justification, which we will sometimes refer to as the *course-grained justification status*.

Fine-grained justification is also able to properly treat with a distinct status the case where no labellings or extensions at all exist. For the labellings approach, fine-grained justification status can be defined as follows.

Definition 4.4 (fine-grained justification status, labelling-based) *Given a labelling-based semantics σ and an argumentation framework $AF = \langle Ar, att \rangle$, the possible justification states of an argument a are defined by a function $\mathcal{JS}_\sigma^{AF} : Ar \to 2^{\{\mathtt{in},\mathtt{out},\mathtt{undec}\}}$ such that $\mathcal{JS}_\sigma^{AF}(a) = \{\mathcal{L}ab(a) \mid \mathcal{L}ab \in \mathcal{L}_\sigma(AF)\}$.*

If we assume a labelling-based semantics to specify the reasonable positions (labellings) one can take in the presence of the conflicting information specified in the argumentation framework, then one can give an intuitive interpretation of the concept of a justification status. For instance, the justification status of $\{\mathtt{in}\}$ means the argument has to be accepted (labelled `in`) in every reasonable position. Similarly, the justification status $\{\mathtt{in}, \mathtt{undec}\}$ means that in every reasonable position the argument is either accepted (labelled `in`) or abstained from having an explicit opinion on (labelled `undec`), but the argument cannot be rejected (labelled `out`). Such an interpretation of the notion of justification status is for instance used in [Wu and Caminada, 2010; Dvořák, 2012].

It is also possible to define the notion of fine-grained justification status in terms of the extensions approach, as is done below.

Definition 4.5 (fine-grained justification status, extension-based) *Given an extension-based semantics σ and an argumentation framework $AF = \langle Ar, att \rangle$, the possible justification states of an argument a are defined by the following mutually exclusive conditions:*

Argument	$\mathcal{JS}_{\mathcal{CO}}^{AF}$	$\mathcal{JS}_{\mathcal{GR}}^{AF}$	$\mathcal{JS}_{\mathcal{PR}}^{AF}$	$\mathcal{JS}_{\mathcal{NA}}^{AF}$
a	$\{\mathtt{in}, \mathtt{out}, \mathtt{undec}\}$	$\{\mathtt{undec}\}$	$\{\mathtt{in}, \mathtt{out}\}$	$\{\mathtt{in}, \mathtt{out}, \mathtt{undec}\}$
b	$\{\mathtt{in}, \mathtt{out}, \mathtt{undec}\}$	$\{\mathtt{undec}\}$	$\{\mathtt{in}, \mathtt{out}\}$	$\{\mathtt{in}, \mathtt{out}, \mathtt{undec}\}$
c	$\{\mathtt{out}, \mathtt{undec}\}$	$\{\mathtt{undec}\}$	$\{\mathtt{out}\}$	$\{\mathtt{in}, \mathtt{out}, \mathtt{undec}\}$
d	$\{\mathtt{in}, \mathtt{undec}\}$	$\{\mathtt{undec}\}$	$\{\mathtt{in}\}$	$\{\mathtt{in}, \mathtt{out}, \mathtt{undec}\}$

Table 7. Argument justification statuses in the example of Figure 6

- $\mathcal{E}_\sigma(AF) = \emptyset$
- $\mathcal{E}_\sigma(AF) \neq \emptyset$ *and* $\forall E \in \mathcal{E}_\sigma(AF)\ a \in E$;
- $\mathcal{E}_\sigma(AF) \neq \emptyset$ *and* $\forall E \in \mathcal{E}_\sigma(AF)\ a \in E^+$;
- $\mathcal{E}_\sigma(AF) \neq \emptyset$ *and* $\forall E \in \mathcal{E}_\sigma(AF)\ a \notin (E \cup E^+)$;
- $\exists E \in \mathcal{E}_\sigma(AF) : a \in E^+$, $\exists E \in \mathcal{E}_\sigma(AF) : a \notin (E \cup E^+)$, *and* $\nexists E \in \mathcal{E}_\sigma(AF) : a \in E$;
- $\exists E \in \mathcal{E}_\sigma(AF) : a \in E$, $\exists E \in \mathcal{E}_\sigma(AF) : a \notin (E \cup E^+)$, *and* $\nexists E \in \mathcal{E}_\sigma(AF) : a \in E^+$;
- $\exists E \in \mathcal{E}_\sigma(AF) : a \in E$, $\exists E \in \mathcal{E}_\sigma(AF) : a \in E^+$, *and* $\nexists E \in \mathcal{E}_\sigma(AF) : a \notin (E \cup E^+)$;
- $\exists E \in \mathcal{E}_\sigma(AF) : a \in E$, $\exists E \in \mathcal{E}_\sigma(AF) : a \in E^+$, *and* $\exists E \in \mathcal{E}_\sigma(AF) : a \notin (E \cup E^+)$.

Intuitively, for an argument a, each item of Definition 4.5 corresponds to a possible value $\mathcal{JS}_\sigma^{AF}(a)$ in Definition 4.4, i.e. to a subset of $\{\mathtt{in}, \mathtt{out}, \mathtt{undec}\}$. For instance the first item corresponds to $\emptyset$, the following three items correspond to $\{\mathtt{in}\}$, $\{\mathtt{out}\}$, $\{\mathtt{undec}\}$ respectively, the fifth item corresponds to $\{\mathtt{out}, \mathtt{undec}\}$ and so on.

As an example of how labelling-based justification status works, consider again the argumentation framework of Figure 6. If one applies complete semantics, three complete labellings are yielded: $(\{a, d\}, \{b, c\}, \emptyset)$, $(\{b, d\}, \{a, c\}, \emptyset)$, $(\emptyset, \emptyset, \{a, b, c, d\})$. This implies that $\mathcal{JS}_{\mathcal{CO}}^{AF}(a) = \{\mathtt{in}, \mathtt{out}, \mathtt{undec}\}$, $\mathcal{JS}_{\mathcal{CO}}^{AF}(b) = \{\mathtt{in}, \mathtt{out}, \mathtt{undec}\}$, $\mathcal{JS}_{\mathcal{CO}}^{AF}(c) = \{\mathtt{out}, \mathtt{undec}\}$ and $\mathcal{JS}_{\mathcal{CO}}^{AF}(d) = \{\mathtt{in}, \mathtt{undec}\}$.

Clearly, the justification status depends on the adopted semantics, for instance with grounded semantics only the third labelling is considered, yielding the status $\{\mathtt{undec}\}$ for all arguments, while with preferred semantics only the first two labellings are considered so that, for each argument, the label `undec` is "removed" from the statuses listed above for complete semantics. Table 7 summarizes the justification statuses of arguments for the various

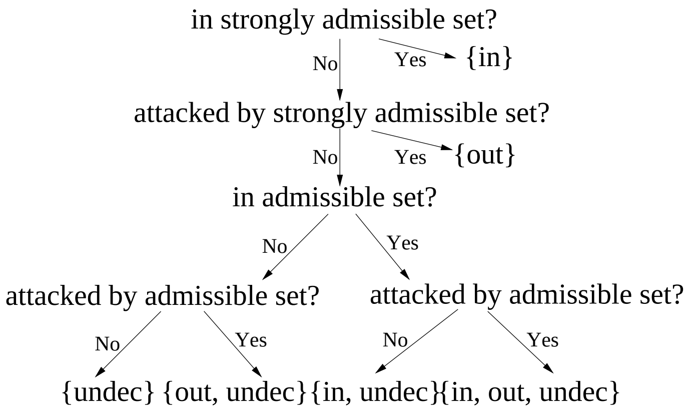

Figure 14. The justification status under complete semantics can be determined using admissibility and strong admissibility.

labelling-based[18] semantics in the example of Figure 6 (note that in this example $\mathcal{JS}^{AF}_{\mathcal{GR}} = \mathcal{JS}^{AF}_{\mathcal{ID}} = \mathcal{JS}^{AF}_{\mathcal{EAG}}$ and $\mathcal{JS}^{AF}_{\mathcal{PR}} = \mathcal{JS}^{AF}_{\mathcal{ST}} = \mathcal{JS}^{AF}_{\mathcal{SST}} = \mathcal{JS}^{AF}_{\mathcal{STG}} = \mathcal{JS}^{AF}_{\mathcal{CF}2} = \mathcal{JS}^{AF}_{\mathcal{STG}2}$).

In general, complete semantics gives rise to the following justification statuses [Caminada and Wu, 2010]:

- $\{\mathtt{in}\}$, called *strong accept*
- $\{\mathtt{out}\}$, called *strong reject*
- $\{\mathtt{in}, \mathtt{undec}\}$, called *weak accept*
- $\{\mathtt{out}, \mathtt{undec}\}$, called *weak reject*
- $\{\mathtt{undec}\}$, called *determined borderline*
- $\{\mathtt{in}, \mathtt{out}, \mathtt{undec}\}$, called *undetermined borderline*

For complete semantics, the justification status of an argument can also be determined using the concepts of admissible and strongly admissible sets, as is indicated in Figure 14.

The reader may have noticed that complete semantics yields only six different justification statuses, although theoretically it seems like eight (2^3) are possible. This is because some combinations of labels cannot occur under complete semantics. For instance, there cannot be an argumentation framework $AF = \langle Ar, att \rangle$ and argument $a \in Ar$ such that $\mathcal{JS}^{AF}_{\mathcal{CO}}(a) = \{\mathtt{in}, \mathtt{out}\}$. This is because if a is labelled `in` by a complete labelling and `out` by another, then

[18] These justi cation statuses have a direct correspondence in the extension-based approach, with the exception of na ve semantics, since in this case di erent labellings may correspond to the same extension.

there exists a third complete labelling (the grounded one) where a is labelled `undec` [Caminada and Wu, 2010] so the justification status should have been $\{\texttt{in}, \texttt{out}, \texttt{undec}\}$ instead. Similarly, $\mathcal{JS}^{AF}_{\mathcal{CO}}(a)$ cannot be $\emptyset$ because there always exists at least one complete labelling.

It can be observed that different semantics yield a different range of justification statuses [Dvořák, 2012].

Proposition 4.6 *Let $AF = \langle Ar, att \rangle$ be an argumentation framework and let $a \in Ar$.*

1. $\mathcal{JS}^{AF}_{\mathcal{GR}}(a) \in \{\{\texttt{in}\}, \{\texttt{out}\}, \{\texttt{undec}\}\}$
2. $\mathcal{JS}^{AF}_{\mathcal{AD}}(a) \in \{\{\texttt{undec}\}, \{\texttt{in}, \texttt{undec}\}, \{\texttt{out}, \texttt{undec}\}, \{\texttt{in}, \texttt{out}, \texttt{undec}\}\}$
3. $\mathcal{JS}^{AF}_{\mathcal{CO}}(a) \in 2^{\{\texttt{in},\texttt{out},\texttt{undec}\}} \setminus \{\{\texttt{in}, \texttt{out}\}, \emptyset\}$
4. $\mathcal{JS}^{AF}_{\mathcal{ST}}(a) \in \{\{\texttt{in}\}, \{\texttt{out}\}, \{\texttt{in}, \texttt{out}\}, \emptyset\}$
5. $\mathcal{JS}^{AF}_{\mathcal{PR}}(a) \in 2^{\{\texttt{in},\texttt{out},\texttt{undec}\}} \setminus \{\emptyset\}$
6. $\mathcal{JS}^{AF}_{\mathcal{SST}}(a) \in 2^{\{\texttt{in},\texttt{out},\texttt{undec}\}} \setminus \{\emptyset\}$
7. $\mathcal{JS}^{AF}_{\mathcal{STG}}(a) \in 2^{\{\texttt{in},\texttt{out},\texttt{undec}\}} \setminus \{\emptyset\}$
8. $\mathcal{JS}^{AF}_{\mathcal{ID}}(a) \in \{\{\texttt{in}\}, \{\texttt{out}\}, \{\texttt{undec}\}\}$
9. $\mathcal{JS}^{AF}_{\mathcal{EAG}}(a) \in \{\{\texttt{in}\}, \{\texttt{out}\}, \{\texttt{undec}\}\}$

As observed in [Dvořák, 2012], the justification status w.r.t. admissible, complete and preferred semantics only differs on the `undec` labels.

Proposition 4.7 *Let $AF = \langle Ar, att \rangle$ be an argumentation framework and let $a \in Ar$.*

1. $\mathcal{JS}^{AF}_{\mathcal{AD}}(a) = \mathcal{JS}^{AF}_{\mathcal{CO}}(a) \cup \{\texttt{undec}\}$
2. $\mathcal{JS}^{AF}_{\mathcal{AD}}(a) = \mathcal{JS}^{AF}_{\mathcal{PR}}(a) \cup \{\texttt{undec}\}$
3. $\mathcal{JS}^{AF}_{\mathcal{CO}}(a) = \begin{cases} \mathcal{JS}^{AF}_{\mathcal{GR}}(a) & \textit{if } a \in GE(AF) \cup GE(AF)^{+} \\ \mathcal{JS}^{AF}_{\mathcal{AD}}(a) & \textit{otherwise} \end{cases}$
4. $\mathcal{JS}^{AF}_{\mathcal{CO}}(a) = \begin{cases} \mathcal{JS}^{AF}_{\mathcal{GR}}(a) & \textit{if } a \in GE(AF) \cup GE(AF)^{+} \\ \mathcal{JS}^{AF}_{\mathcal{PR}}(a) \cup \{\texttt{undec}\} & \textit{otherwise} \end{cases}$

Fine-grained justification status, as described above, offers a number of advantages compared to the more traditional course-grained justification status that is based on sceptical and credulous acceptance. For instance, the computational complexity of *weak acceptance* under preferred semantics is lower than that of traditional sceptical preferred [Dvořák, 2012]. Furthermore, labelling-based justification status offers a subtle way of dealing with the notion of *floating conclusions* [Caminada and Wu, 2010]. Briefly, a conclusion is "floating" when it is supported by different arguments which are labelled `in` in different labellings, so that, even if there is no individual argument for the conclusion which is labelled `in` in all labellings, there is at least an argument labelled `in` supporting the conclusion in every labelling. In other words, the conclusion has different supports in different labellings, but has at least one in every labelling.

Still, labelling-based justification status has not yet received the same level of attention as the more traditional extension-based justification status, and some research questions are still open.[19]

4.2 Skepticism and skepticism relations

The term *skepticism* has been used in the literature (often in an informal way) to discuss argumentation semantics behaviour, e.g. by observing that a semantics is "more skeptical" than another one. Intuitively, a skeptical attitude tends to make less committed choices about the justification of the arguments, as well exemplified by the traditional notions of skeptical and credulous acceptance recalled in Section 4.1. In other words, a skeptical attitude tends to leave arguments in an "undecided" justification state and to accept (or reject) as least arguments as possible, while a less skeptical (or more credulous) attitude corresponds to more extensive acceptance (or rejection) of arguments. Note, in particular, that the notion of commitment (or decidedness) of a justification state must be clearly distinguished from the notion of acceptance: two justification states corresponding to definite acceptance and definite rejection, though reflecting antithetical choices about the state of an argument, have both the same highest level of commitment.

Which are the formal counterparts of these basic intuitions?

Starting from basic elements, we first need to define a criterion to compare extensions and labellings with respect to skepticism. As to extensions, this is quite simple: an extension E_1 is "more skeptical" than (to be precise, at least as skeptical as) an extension E_2 if $E_1 \subseteq E_2$, since then E_1 supports the acceptance of no more arguments than E_2. As to labellings, we have to consider both the `in` and `out` labels as being both more committed choices than `undec`. We can then state that a labelling $\mathcal{L}ab_1$ is at least as skeptical as a labelling $\mathcal{L}ab_2$ according to the inclusion of both the sets of `in` and `out` labelled arguments. These intuitions are formalized in Definition 4.8.

Definition 4.8 *Given two extensions E_1 and E_2 of an argumentation framework AF, E_1 is* at least as skeptical *as E_2, denoted as $E_1 \preceq E_2$ if and only if $E_1 \subseteq E_2$. Given two labellings $\mathcal{L}ab_1$ and $\mathcal{L}ab_2$ of an argumentation framework AF, $\mathcal{L}ab_1$ is* at least as skeptical *as $\mathcal{L}ab_2$, denoted as $\mathcal{L}ab_1 \preceq \mathcal{L}ab_2$, if and only if $\mathcal{L}ab_1 \sqsubseteq \mathcal{L}ab_2$ (see Definition 3.44).*

[19] For instance, how to nd e cient proof procedures for determining the ne-grained justi cation status with respect to preferred, semi-stable and stage semantics.

While the above relations are sufficient to compare unique-status semantics, the next step is to introduce skepticism relations between non-empty[20] sets of extensions or labellings in order to compare multiple-status semantics. As more extensively discussed in [Baroni and Giacomin, 2009b], several alternatives can be considered for this issue.

As a first basic step, one can consider a comparison method based on inclusion of the sets of accepted arguments, either according to skeptical or credulous acceptance. This gives rise to the skepticism relations stated in the following definitions.

Definition 4.9 *Given two non-empty sets of extensions $\mathcal{E}_1$ and $\mathcal{E}_2$ of an argumentation framework AF, $\mathcal{E}_1 \preceq^E_\cap \mathcal{E}_2$ if and only if $\bigcap_{E_1 \in \mathcal{E}_1} E_1 \subseteq \bigcap_{E_2 \in \mathcal{E}_2} E_2$.*

Definition 4.10 *Given two non-empty sets of extensions $\mathcal{E}_1$ and $\mathcal{E}_2$ of an argumentation framework AF, $\mathcal{E}_1 \preceq^E_\cup \mathcal{E}_2$ if and only if $\bigcup_{E_1 \in \mathcal{E}_1} E_1 \subseteq \bigcup_{E_2 \in \mathcal{E}_2} E_2$.*

Definition 4.11 *Given two non-empty sets of labellings $\mathcal{L}_1$ and $\mathcal{L}_2$ of an argumentation framework AF, $\mathcal{L}_1 \preceq^L_\cap \mathcal{L}_2$ if and only if $\bigcap_{\mathcal{L}ab_1 \in \mathcal{L}_1} \mathtt{in}(\mathcal{L}ab_1) \subseteq \bigcap_{\mathcal{L}ab_2 \in \mathcal{L}_2} \mathtt{in}(\mathcal{L}ab_2)$.*

Definition 4.12 *Given two non-empty sets of labellings $\mathcal{L}_1$ and $\mathcal{L}_2$ of an argumentation framework AF, $\mathcal{L}_1 \preceq^L_\cup \mathcal{L}_2$ if and only if $\bigcup_{\mathcal{L}ab_1 \in \mathcal{L}_1} \mathtt{in}(\mathcal{L}ab_1) \subseteq \bigcup_{\mathcal{L}ab_2 \in \mathcal{L}_2} \mathtt{in}(\mathcal{L}ab_2)$.*

To exemplify the above notions, consider first the example of Figure 4. In the extension-based approach, the grounded and ideal semantics prescribe the set of extensions $\mathcal{E}_1 = \{\emptyset\}$ while all other semantics prescribe $\mathcal{E}_2 = \{\{a\}, \{b\}\}$. Clearly, $\mathcal{E}_1 \preceq^E_\cap \mathcal{E}_2$, $\mathcal{E}_2 \preceq^E_\cap \mathcal{E}_1$, $\mathcal{E}_1 \preceq^E_\cup \mathcal{E}_2$, while it is not the case that $\mathcal{E}_2 \preceq^E_\cup \mathcal{E}_1$ (denoted as $\mathcal{E}_2 \not\preceq^E_\cup \mathcal{E}_1$). For the same example in the labelling-based approach grounded and ideal semantics prescribe the set of labellings $\mathcal{L}_1 = \{(\emptyset, \emptyset, \{a, b\})\}$ while all other semantics prescribe $\mathcal{L}_2 = \{(\{a\}, \{b\}, \emptyset), (\{b\}, \{a\}, \emptyset)\}$. Again, it can be seen that $\mathcal{L}_1 \preceq^L_\cap \mathcal{L}_2$, $\mathcal{L}_2 \preceq^L_\cap \mathcal{L}_1$, $\mathcal{L}_1 \preceq^L_\cup \mathcal{L}_2$, while $\mathcal{L}_2 \not\preceq^L_\cup \mathcal{L}_1$.

Considering the example of Figure 6, in the extension-based approach the grounded and ideal semantics prescribe the set of extensions $\mathcal{E}_1 = \{\emptyset\}$ while all other semantics prescribe $\mathcal{E}_2 = \{\{a, d\}, \{b, d\}\}$. It turns out that $\mathcal{E}_1 \preceq^E_\cap \mathcal{E}_2$ and $\mathcal{E}_1 \preceq^E_\cup \mathcal{E}_2$, while $\mathcal{E}_2 \not\preceq^E_\cap \mathcal{E}_1$ and $\mathcal{E}_2 \not\preceq^E_\cup \mathcal{E}_1$. The case of labellings is perfectly analogous with $\mathcal{L}_1 = \{(\emptyset, \emptyset, \{a, b, c, d\})\}$ for grounded and ideal semantics and $\mathcal{L}_2 = \{(\{a, d\}, \{b, c\}, \emptyset), (\{b, d\}, \{a, c\}, \emptyset)\}$ for other semantics yielding $\mathcal{L}_1 \preceq^L_\cap \mathcal{L}_2$ and $\mathcal{L}_1 \preceq^L_\cup \mathcal{L}_2$, while $\mathcal{L}_2 \not\preceq^L_\cap \mathcal{L}_1$ and $\mathcal{L}_2 \not\preceq^L_\cup \mathcal{L}_1$.

Figure 9 provides a more articulated case for comparison.

In the extension-based approach grounded and ideal semantics prescribe $\mathcal{E}_1 = \{\emptyset\}$, preferred semantics prescribes $\mathcal{E}_2 = \{\{a\}, \{b, d\}\}$, $\mathcal{CF}2$ semantics prescribes $\mathcal{E}_3 = \{\{a, c\}, \{a, d\}, \{a, e\}, \{b, d\}\}$, while stable, semi-stable and

[20] We assume that an empty set of extensions/labellings corresponds to a peculiar case and cannot be involved in skepticism comparison.

stage semantics prescribe $\mathcal{E}_4 = \{\{b, d\}\}$. It follows that for any $i, j \in \{1, 2, 3\}$ $\mathcal{E}_i \preceq^E_\cap \mathcal{E}_j$, while for any $i \in \{1, 2, 3\}$ $\mathcal{E}_i \preceq^E_\cap \mathcal{E}_4$ and $\mathcal{E}_4 \not\preceq^E_\cap \mathcal{E}_i$. On the other hand, these sets are completely ordered according to $\preceq^E_\cup$ since $\mathcal{E}_1 \preceq^E_\cup \mathcal{E}_4 \preceq^E_\cup \mathcal{E}_2 \preceq^E_\cup \mathcal{E}_3$. Again, the case of labellings is perfectly analogous.

As a further step in the analysis of skepticism relations, one may observe that also explicitly rejected arguments should be taken into account in a similar way as accepted arguments: this gives rise to the following definitions.

Definition 4.13 *Given two non-empty sets of extensions $\mathcal{E}_1$ and $\mathcal{E}_2$ of an argumentation framework AF, $\mathcal{E}_1 \preceq^E_{\cap\rightarrow} \mathcal{E}_2$ if and only if $\mathcal{E}_1 \preceq^E_\cap \mathcal{E}_2$ and $\bigcap_{E_1 \in \mathcal{E}_1} E_1^+ \subseteq \bigcap_{E_2 \in \mathcal{E}_2} E_2^+$.*

Definition 4.14 *Given two non-empty sets of extensions $\mathcal{E}_1$ and $\mathcal{E}_2$ of an argumentation framework AF, $\mathcal{E}_1 \preceq^E_{\cup\rightarrow} \mathcal{E}_2$ if and only if $\mathcal{E}_1 \preceq^E_\cup \mathcal{E}_2$ and $\bigcup_{E_1 \in \mathcal{E}_1} E_1^+ \subseteq \bigcup_{E_2 \in \mathcal{E}_2} E_2^+$.*

Definition 4.15 *Given two non-empty sets of labellings $\mathcal{L}_1$ and $\mathcal{L}_2$ of an argumentation framework AF, $\mathcal{L}_1 \preceq^L_{\cap\rightarrow} \mathcal{L}_2$ if and only if $\mathcal{L}_1 \preceq^L_\cap \mathcal{L}_2$ and $\bigcap_{\mathcal{L}ab_1 \in \mathcal{L}_1} \mathtt{out}(\mathcal{L}ab_1) \subseteq \bigcap_{\mathcal{L}ab_2 \in \mathcal{L}_2} \mathtt{out}(\mathcal{L}ab_2)$.*

Definition 4.16 *Given two non-empty sets of labellings $\mathcal{L}_1$ and $\mathcal{L}_2$ of an argumentation framework AF, $\mathcal{L}_1 \preceq^L_{\cup\rightarrow} \mathcal{L}_2$ if and only if $\mathcal{L}_1 \preceq^L_\cup \mathcal{L}_2$ and $\bigcup_{\mathcal{L}ab_1 \in \mathcal{L}_1} \mathtt{out}(\mathcal{L}ab_1) \subseteq \bigcup_{\mathcal{L}ab_2 \in \mathcal{L}_2} \mathtt{out}(\mathcal{L}ab_2)$.*

Consider again the example of Figure 4, and for a set of extensions $\mathcal{E}$ let us denote in the following $\mathcal{E}^+ = \{E^+ \mid E \in \mathcal{E}\}$. Then, referring to the already mentioned sets of extensions $\mathcal{E}_1 = \{\emptyset\}$ and $\mathcal{E}_2 = \{\{a\}, \{b\}\}$ we have $\mathcal{E}_1^+ = \{\emptyset\}$ and $\mathcal{E}_2^+ = \{\{b\}, \{a\}\}$. Clearly, $\mathcal{E}_1 \preceq^E_{\cap\rightarrow} \mathcal{E}_2$, $\mathcal{E}_2 \preceq^E_{\cap\rightarrow} \mathcal{E}_1$, $\mathcal{E}_1 \preceq^E_{\cup\rightarrow} \mathcal{E}_2$ while $\mathcal{E}_2 \not\preceq^E_{\cup\rightarrow} \mathcal{E}_1$. For the same example in the labelling-based approach, it can analogously be seen that $\mathcal{L}_1 \preceq^L_{\cap\rightarrow} \mathcal{L}_2$, $\mathcal{L}_2 \preceq^L_{\cap\rightarrow} \mathcal{L}_1$, $\mathcal{L}_1 \preceq^L_{\cup\rightarrow} \mathcal{L}_2$ and $\mathcal{L}_2 \not\preceq^L_{\cup\rightarrow} \mathcal{L}_1$.

In the example of Figure 6, we refer again to $\mathcal{E}_1 = \{\emptyset\}$ and $\mathcal{E}_2 = \{\{a, d\}, \{b, d\}\}$, yielding $\mathcal{E}_1^+ = \{\emptyset\}$ and $\mathcal{E}_2^+ = \{\{b, c\}, \{a, c\}\}$. Then $\mathcal{E}_1 \preceq^E_{\cap\rightarrow} \mathcal{E}_2$ and $\mathcal{E}_1 \preceq^E_{\cup\rightarrow} \mathcal{E}_2$, while $\mathcal{E}_2 \not\preceq^E_{\cap\rightarrow} \mathcal{E}_1$ and $\mathcal{E}_2 \not\preceq^E_{\cup\rightarrow} \mathcal{E}_1$. The case of labellings is perfectly analogous with $\mathcal{L}_1 \preceq^L_{\cap\rightarrow} \mathcal{L}_2$ and $\mathcal{L}_1 \preceq^L_{\cup\rightarrow} \mathcal{L}_2$, while $\mathcal{L}_2 \not\preceq^L_{\cap\rightarrow} \mathcal{L}_1$ and $\mathcal{L}_2 \not\preceq^L_{\cup\rightarrow} \mathcal{L}_1$.

Figure 9 provides again a more articulated case.

Considering the sets of extensions $\mathcal{E}_1 = \{\emptyset\}$, $\mathcal{E}_2 = \{\{a\}, \{b, d\}\}$, $\mathcal{E}_3 = \{\{a, c\}, \{a, d\}, \{a, e\}, \{b, d\}\}$, and $\mathcal{E}_4 = \{\{b, d\}\}$ we have $\mathcal{E}_1^+ = \{\emptyset\}$, $\mathcal{E}_2^+ = \{\{b\}, \{a, c, e\}\}$, $\mathcal{E}_3^+ = \{\{b, d\}, \{b, e\}, \{b, c\}, \{a, c, e\}\}$, $\mathcal{E}_4^+ = \{\{a, c, e\}\}$. It follows that for any $i, j \in \{1, 2, 3\}$ $\mathcal{E}_i \preceq^E_{\cap\rightarrow} \mathcal{E}_j$, while for any $i \in \{1, 2, 3\}$ $\mathcal{E}_i \preceq^E_{\cap\rightarrow} \mathcal{E}_4$ and $\mathcal{E}_4 \not\preceq^E_{\cap\rightarrow} \mathcal{E}_i$. On the other hand, $\mathcal{E}_1 \preceq^E_{\cup\rightarrow} \mathcal{E}_4 \preceq^E_{\cup\rightarrow} \mathcal{E}_2 \preceq^E_{\cup\rightarrow} \mathcal{E}_3$. Again, the case of labellings is perfectly analogous.

To have an example where the relations of the $\preceq_\cap$ kind differ from those of the $\preceq_{\cap\rightarrow}$ kind consider the example of Figure 15. In the extension-based approach all semantics but stable[21], stage, $\mathcal{CF}2$, and stage2 semantics prescribe the set

[21] The set of stable extensions is empty in this case.

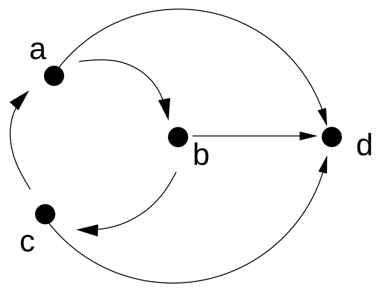

Figure 15. Cycle of three attacking arguments in turn attacking another argument

of extensions $\mathcal{E}_1 = \{\emptyset\}$ with $\mathcal{E}_1^+ = \{\emptyset\}$, while stage, $\mathcal{CF}2$, and stage2 semantics prescribe $\mathcal{E}_2 = \{\{a\}, \{b\}, \{c\}\}$ with $\mathcal{E}_2^+ = \{\{d\}\}$. It follows that $\mathcal{E}_1 \preceq^E_\cap \mathcal{E}_2$ and $\mathcal{E}_2 \preceq^E_\cap \mathcal{E}_1$, while $\mathcal{E}_1 \preceq^E_{\cap\rightarrow} \mathcal{E}_2$ but $\mathcal{E}_2 \not\preceq^E_{\cap\rightarrow} \mathcal{E}_1$. Similar considerations apply in the labelling-based approach.

Definitions 4.9-4.16 treat sets of extensions or labellings "as a whole" by simply considering their intersection or union: for instance, very different sets of extensions are treated in the same way if they have an empty intersection. In order to take account of how single extensions or labellings are defined, a different kind of definition is needed: the skepticism relation between two sets (let say $\mathcal{X}_1$ and $\mathcal{X}_2$) of extensions or labellings should be based on some comparison between their individual elements. In particular, according to a skeptical approach to argument justification, in order to state that $\mathcal{X}_1$ is at least as skeptical as $\mathcal{X}_2$, one may require that every element in $\mathcal{X}_2$ has a more skeptical counterpart in $\mathcal{X}_1$, while, according to a credulous approach, one may require dually that every element in $\mathcal{X}_1$ has a less skeptical counterpart in $\mathcal{X}_2$. This general idea is formalized by the following definitions, which resort to the basic comparisons between single extensions and labellings identified in Definition 4.8.

Definition 4.17 *Given two non-empty sets of extensions $\mathcal{E}_1$ and $\mathcal{E}_2$ of an argumentation framework AF, $\mathcal{E}_1 \preceq^E_{\cap+} \mathcal{E}_2$ if and only if $\forall E_2 \in \mathcal{E}_2\ \exists E_1 \in \mathcal{E}_1 : E_1 \preceq E_2$.*

Definition 4.18 *Given two non-empty sets of extensions $\mathcal{E}_1$ and $\mathcal{E}_2$ of an argumentation framework AF, $\mathcal{E}_1 \preceq^E_{\cup+} \mathcal{E}_2$ if and only if $\forall E_1 \in \mathcal{E}_1\ \exists E_2 \in \mathcal{E}_2 : E_1 \preceq E_2$.*

Definition 4.19 *Given two non-empty sets of labellings $\mathcal{L}_1$ and $\mathcal{L}_2$ of an argumentation framework AF, $\mathcal{L}_1 \preceq^L_{\cap+} \mathcal{L}_2$ if and only if $\forall \mathcal{L}ab_2 \in \mathcal{L}_2\ \exists \mathcal{L}ab_1 \in \mathcal{L}_1 : \mathcal{L}ab_1 \preceq \mathcal{L}ab_2$.*

Definition 4.20 *Given two non-empty sets of labellings $\mathcal{L}_1$ and $\mathcal{L}_2$ of an argumentation framework AF, $\mathcal{L}_1 \preceq^L_{\cup+} \mathcal{L}_2$ if and only if $\forall \mathcal{L}ab_1 \in \mathcal{L}_1\ \exists \mathcal{L}ab_2 \in \mathcal{L}_2 : \mathcal{L}ab_1 \preceq \mathcal{L}ab_2$.*

Let us exemplify these relations.

In the example of Figure 4, referring to the already mentioned sets of extensions $\mathcal{E}_1 = \{\emptyset\}$ and $\mathcal{E}_2 = \{\{a\}, \{b\}\}$ we have $\mathcal{E}_1 \preceq^E_{\cap+} \mathcal{E}_2$ but, differently from the previously considered relations, $\mathcal{E}_2 \not\preceq^E_{\cap+} \mathcal{E}_1$. On the other hand, $\mathcal{E}_1 \preceq^E_{\cup+} \mathcal{E}_2$ and $\mathcal{E}_2 \not\preceq^E_{\cup+} \mathcal{E}_1$. As usual, analogous relations hold for the labelling-based approach.

Similarly, in the example of Figure 6, with $\mathcal{E}_1 = \{\emptyset\}$ and $\mathcal{E}_2 = \{\{a, d\}, \{b, d\}\}$ it holds $\mathcal{E}_1 \preceq^E_{\cap+} \mathcal{E}_2$ and $\mathcal{E}_1 \preceq^E_{\cup+} \mathcal{E}_2$, while $\mathcal{E}_2 \not\preceq^E_{\cap+} \mathcal{E}_1$ and $\mathcal{E}_2 \not\preceq^E_{\cup+} \mathcal{E}_1$. It goes without saying that the same holds in the labelling-based approach.

Finally consider the case of Figure 9 with sets of extensions $\mathcal{E}_1 = \{\emptyset\}$, $\mathcal{E}_2 = \{\{a\}, \{b, d\}\}$, $\mathcal{E}_3 = \{\{a, c\}, \{a, d\}, \{a, e\}, \{b, d\}\}$, and $\mathcal{E}_4 = \{\{b, d\}\}$. We can first observe that for $i \in \{2, 3, 4\}$ $\mathcal{E}_1 \preceq^E_{\cap+} \mathcal{E}_i$ and (differently from previous relations) $\mathcal{E}_i \not\preceq^E_{\cap+} \mathcal{E}_1$. Then we can note that $\mathcal{E}_2 \preceq^E_{\cap+} \mathcal{E}_4$ and $\mathcal{E}_3 \preceq^E_{\cap+} \mathcal{E}_4$ since the only element of $\mathcal{E}_4$ (namely $\{b, d\}$) is a superset of (actually coincides with) an element of either $\mathcal{E}_2$ or $\mathcal{E}_3$. Also $\mathcal{E}_2 \preceq^E_{\cap+} \mathcal{E}_3$ since the elements $\{a, c\}$, $\{a, d\}$, and $\{a, e\}$ of $\mathcal{E}_3$ are supersets of $\{a\}$ in $\mathcal{E}_2$ and $\{b, d\}$ is present both in $\mathcal{E}_3$ and $\mathcal{E}_2$. With similar observations it can be seen that $\mathcal{E}_3 \not\preceq^E_{\cap+} \mathcal{E}_2$, $\mathcal{E}_4 \not\preceq^E_{\cap+} \mathcal{E}_2$, and $\mathcal{E}_4 \not\preceq^E_{\cap+} \mathcal{E}_3$. Turning to the relation corresponding to the credulous perspective, it can immediately be observed that for $i \in \{2, 3, 4\}$ $\mathcal{E}_1 \preceq^E_{\cup+} \mathcal{E}_i$ and $\mathcal{E}_i \not\preceq^E_{\cup+} \mathcal{E}_1$. Also, $\mathcal{E}_2 \preceq^E_{\cup+} \mathcal{E}_3$ since $\{a\}$ is included in some elements of $\mathcal{E}_3$ and $\{b, d\}$ is present both in $\mathcal{E}_2$ and $\mathcal{E}_3$. On the other hand, $\mathcal{E}_3 \not\preceq^E_{\cup+} \mathcal{E}_2$. Differently from the skeptical perspective, $\mathcal{E}_4 \preceq^E_{\cup+} \mathcal{E}_2$ and $\mathcal{E}_4 \preceq^E_{\cup+} \mathcal{E}_3$ (the only element of $\mathcal{E}_4$, namely $\{b, d\}$ is present both in $\mathcal{E}_2$ and $\mathcal{E}_3$) while it can be easily seen that $\mathcal{E}_2 \not\preceq^E_{\cup+} \mathcal{E}_4$ ($\{a\}$ is not included in any element of $\mathcal{E}_4$) and $\mathcal{E}_3 \not\preceq^E_{\cup+} \mathcal{E}_4$ (as above for sets $\{a, c\}$, $\{a, d\}$, $\{a, e\}$). Again, the case of labellings is perfectly analogous.

A stronger skepticism relation, unifying the skeptical and credulous perspectives, can be obtained by combining together the relations $\preceq_{\cap+}$ and $\preceq_{\cup+}$.

Definition 4.21 *Given two non-empty sets of extensions $\mathcal{E}_1$ and $\mathcal{E}_2$ of an argumentation framework AF, $\mathcal{E}_1 \preceq^E_{\oplus} \mathcal{E}_2$ if and only if $\mathcal{E}_1 \preceq^E_{\cap+} \mathcal{E}_2$ and $\mathcal{E}_1 \preceq^E_{\cup+} \mathcal{E}_2$.*

Definition 4.22 *Given two non-empty sets of labellings $\mathcal{L}_1$ and $\mathcal{L}_2$ of an argumentation framework AF, $\mathcal{L}_1 \preceq^L_{\oplus} \mathcal{L}_2$ if and only if $\mathcal{L}_1 \preceq^L_{\cap+} \mathcal{L}_2$ and $\mathcal{L}_1 \preceq^L_{\cup+} \mathcal{L}_2$.*

As also evident from their definitions, the various skepticism relations introduced above are related each other by implication. In particular, two implications chains can be identified in correspondence with the skeptical or credulous perspective. In fact, given two sets of extensions $\mathcal{E}_1$ and $\mathcal{E}_2$ of an argumentation framework AF, it holds that:

(1) $$\mathcal{E}_1 \preceq^E_{\oplus} \mathcal{E}_2 \Rightarrow \mathcal{E}_1 \preceq^E_{\cap+} \mathcal{E}_2 \Rightarrow \mathcal{E}_1 \preceq^E_{\cap\rightarrow} \mathcal{E}_2 \Rightarrow \mathcal{E}_1 \preceq^E_{\cap} \mathcal{E}_2$$

(2) $$\mathcal{E}_1 \preceq^E_{\oplus} \mathcal{E}_2 \Rightarrow \mathcal{E}_1 \preceq^E_{\cup+} \mathcal{E}_2 \Rightarrow \mathcal{E}_1 \preceq^E_{\cup\rightarrow} \mathcal{E}_2 \Rightarrow \mathcal{E}_1 \preceq^E_{\cup} \mathcal{E}_2$$

The only nontrivial implications in (1) and (2) concern the fact that $\preceq^E_{\cap+}$ implies $\preceq^E_{\cap\rightarrow}$, and, similarly, $\preceq^E_{\cup+}$ implies $\preceq^E_{\cup\rightarrow}$: they have been proved in [Baroni and Giacomin, 2009b].

Using Definitions 4.11, 4.12, 4.15, 4.16, 4.19, 4.20, 4.22, and the same kind of reasoning it is possible to prove that the analogous relations hold in the labelling based approach. In fact, given two sets of labellings $\mathcal{L}_1$ and $\mathcal{L}_2$ of an argumentation framework AF, it holds that:

(3) $$\mathcal{L}_1 \preceq^L_{\oplus} \mathcal{L}_2 \Rightarrow \mathcal{L}_1 \preceq^L_{\cap+} \mathcal{L}_2 \Rightarrow \mathcal{L}_1 \preceq^L_{\cap\rightarrow} \mathcal{L}_2 \Rightarrow \mathcal{L}_1 \preceq^L_{\cap} \mathcal{L}_2$$

(4) $$\mathcal{L}_1 \preceq^L_{\oplus} \mathcal{L}_2 \Rightarrow \mathcal{L}_1 \preceq^L_{\cup+} \mathcal{L}_2 \Rightarrow \mathcal{L}_1 \preceq^L_{\cup\rightarrow} \mathcal{L}_2 \Rightarrow \mathcal{L}_1 \preceq^L_{\cup} \mathcal{L}_2$$

Turning to the comparison between semantics, for a given generic relation $\preceq$ concerning either extensions or labellings it is quite natural to define an induced relation of skepticism between two semantics σ_1 and σ_2, by requiring that $\preceq$ holds for their sets of extensions or labellings. As it may happen that either σ_1 or σ_2 prescribes an empty set of extensions (or labellings) in some cases, the induced relation has to refer to a set of argumentation frameworks where both σ_1 and σ_2 prescribe non-empty sets of extensions (or labellings), namely to $\mathcal{DE}_{\sigma_1} \cap \mathcal{DE}_{\sigma_2}$ (or $\mathcal{DL}_{\sigma_1} \cap \mathcal{DL}_{\sigma_2}$) using the notation of Definition 4.1.

Definition 4.23 *Let $\preceq^E$ be a skepticism relation between sets of extensions, σ_1 and σ_2 be extension-based argumentation semantics, and $\mathcal{A}$ be a set of argumentation frameworks with $\mathcal{A} \subseteq (\mathcal{DE}_{\sigma_1} \cap \mathcal{DE}_{\sigma_2})$. The skepticism relation $\preceq^{SE}$ induced by $\preceq^E$ between σ_1 and σ_2 with reference to $\mathcal{A}$ is defined as follows: $\sigma_1 \preceq^{SE} \sigma_2$ if and only if $\forall AF \in \mathcal{A}$ $\mathcal{E}_{\sigma_1}(AF) \preceq^E \mathcal{E}_{\sigma_2}(AF)$.*

Definition 4.24 *Let $\preceq^L$ be a skepticism relation between sets of labellings, σ_1 and σ_2 be labelling-based argumentation semantics, and $\mathcal{A}$ be a set of argumentation frameworks with $\mathcal{A} \subseteq (\mathcal{DL}_{\sigma_1} \cap \mathcal{DL}_{\sigma_2})$. The skepticism relation $\preceq^{SL}$ induced by $\preceq^L$ between σ_1 and σ_2 with reference to $\mathcal{A}$ is defined as follows: $\sigma_1 \preceq^{SL} \sigma_2$ if and only if $\forall AF \in \mathcal{A}$ $\mathcal{L}_{\sigma_1}(AF) \preceq^L \mathcal{L}_{\sigma_2}(AF)$.*

Focusing on the extension-based approach, while Definition 4.23 considers a generic set of argumentation frameworks $\mathcal{A} \subseteq (\mathcal{DE}_{\sigma_1} \cap \mathcal{DE}_{\sigma_2})$ to define a skepticism relation, clearly the most interesting skepticism relation is the one corresponding to the case $\mathcal{A} = (\mathcal{DE}_{\sigma_1} \cap \mathcal{DE}_{\sigma_2})$. Then, when considering a skepticism comparison concerning more than two semantics $\sigma_1, \sigma_2, \ldots, \sigma_N$ it is reasonable to consider a common reference $\mathcal{A} = \bigcap_{i=1\ldots N} \mathcal{DE}_{\sigma_i}$. As to the semantics discussed in this chapter, only stable semantics may prescribe an empty set of extensions/labellings. Therefore two reference sets can be considered: the universe of all argumentation frameworks if stable semantics is not involved in the comparison, or $\mathcal{DE}_{\mathcal{ST}}$ otherwise. Clearly the same considerations hold in the labelling-based approach by replacing $\mathcal{DE}$ with $\mathcal{DL}$.

It is worth noting that, in general, two semantics σ_1 and σ_2 may not be comparable with respect to skepticism. For instance, it may be the case that

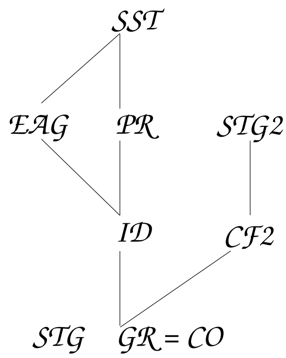

Figure 16. $\preceq^{SE}_{\cap+}$, $\preceq^{SE}_{\cap\rightarrow}$ and $\preceq^{SE}_{\cap}$ relations for any argumentation framework

there are two argumentation frameworks AF_1 and AF_2 such that σ_1 behaves more skeptically than σ_2 in the case of AF_1 but σ_2 behaves more skeptically than σ_1 in the case of AF_2, or that the two semantics yield incomparable sets of extensions for some given argumentation framework. Furthermore, the order between two semantics may be different according to the credulous or skeptical perspective.

A detailed analysis of skepticism relations between most extension-based semantics has been carried out in [Baroni and Giacomin, 2009b] to which the reader may refer for details: we report here only the resulting partial orders, graphically presented as Hasse diagrams (with the addition of eager, stage, and stage2 semantics[22] with respect to [Baroni and Giacomin, 2009b]). As mentioned above, distinct Hasse diagrams are presented for the case where stable extensions exist and for the general one.

The partial orders[23] induced by all the relations corresponding to the skeptical perspective, namely $\preceq^{SE}_{\cap+}$, $\preceq^{SE}_{\cap\rightarrow}$ and $\preceq^{SE}_{\cap}$ coincide.

The Hasse diagram corresponding to the general case is shown in Figure 16: grounded semantics is the most skeptical one and since the grounded extension is the least complete extension it turns out that $\mathcal{GR} \preceq^{SE}_{\cap+} \mathcal{CO}$ and $\mathcal{CO} \preceq^{SE}_{\cap+} \mathcal{GR}$. Ideal, preferred, and semi-stable semantics are all comparable among them and orderly less skeptical, with eager semantics in between ideal and semi-stable, while not comparable with preferred. Stage2 semantics is less skeptical than $\mathcal{CF}2$, which in turn is comparable with $\mathcal{GR}$ and $\mathcal{CO}$, while no other relation holds with the other semantics. Stage semantics is not comparable with any other, also due to its peculiar behaviour in some cases, exemplified in the

[22] The authors are grateful to Wolfgang Dvorak for suggesting the extension of these results to eager and stage2 semantics.

[23] The skepticism relations described in the following have been analyzed in [Baroni and Giacomin, 2009b] for the extension-based approach. Due to the one-to-one correspondence between extensions and labellings holding for all the semantics involved in the comparison, it is possible to prove that the skepticism relations hold also in the labelling-based approach.

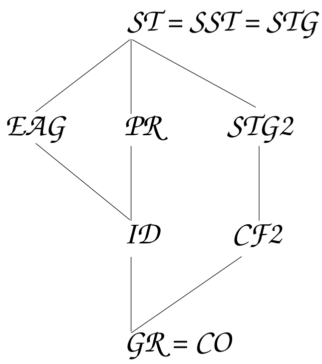

Figure 17. $\preceq^{SE}_{\cap+}$, $\preceq^{SE}_{\cap\rightarrow}$ and $\preceq^{SE}_{\cap}$ relations for argumentation frameworks in $\mathcal{DE}_{\mathcal{ST}}$ ($\mathcal{DL}_{\mathcal{ST}}$)

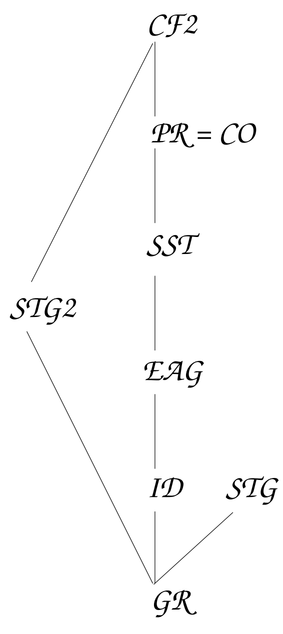

Figure 18. $\preceq^{SE}_{\cup+}$, $\preceq^{SE}_{\cup\rightarrow}$ and $\preceq^{SE}_{\cup}$ relations for any argumentation framework

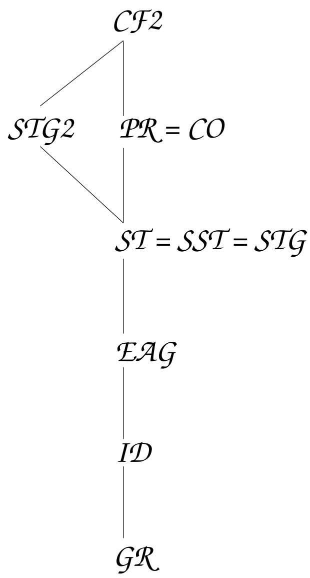

Figure 19. $\preceq^{SE}_{\cup+}$, $\preceq^{SE}_{\cup\rightarrow}$ and $\preceq^{SE}_{\cup}$ relations for argumentation frameworks in $\mathcal{DE}_{\mathcal{ST}}$ ($\mathcal{DL}_{\mathcal{ST}}$)

argumentation framework of Figure 12.

The Hasse diagram for $\preceq^{SE}_{\cap+}$, $\preceq^{SE}_{\cap\rightarrow}$ and $\preceq^{SE}_{\cap}$ considering only the argumentation frameworks where stable extensions exist (and then coincide with semi-stable and stage extensions) is shown in Figure 17. It can be noted that in this context stage2 semantics is comparable with (and more skeptical than) stable semantics.

Turning to skepticism relations based on the credulous perspective, namely $\preceq^{SE}_{\cup+}$, $\preceq^{SE}_{\cup\rightarrow}$ and $\preceq^{SE}_{\cup}$, the Hasse diagram corresponding to the general case is shown in Figure 18. An almost complete ordering is achieved due to the change of perspective. In particular, complete semantics is in mutual relation with preferred semantics: $\mathcal{PR} \preceq^{SE}_{\cup+} \mathcal{CO}$ and $\mathcal{CO} \preceq^{SE}_{\cup+} \mathcal{PR}$ since preferred extensions are maximal complete extensions. Moreover one can note that $\mathcal{CF}2$ is now comparable with any other one (and is actually the least skeptical semantics) and that the orderings between $\mathcal{PR}$ and $\mathcal{SST}$ and between $\mathcal{CF}2$ and stage2 are inverted with respect to Figure 16.

The Hasse diagram for $\preceq^{SE}_{\cup+}$, $\preceq^{SE}_{\cup\rightarrow}$ and $\preceq^{SE}_{\cup}$ considering only the argumentation frameworks where stable extensions exist is shown in Figure 19: here an almost total order is achieved, which obeys the same relations as the general case but where stable, semi-stable and stage semantics coincide and stage2 is between stable and $\mathcal{CF}2$ semantics.

Finally, the Hasse diagrams for the relations arising from the conjunction of

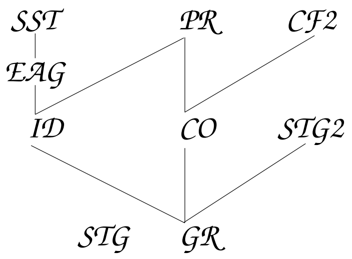

Figure 20. $\preceq_{\oplus}^{SE}$ relation for any argumentation framework.

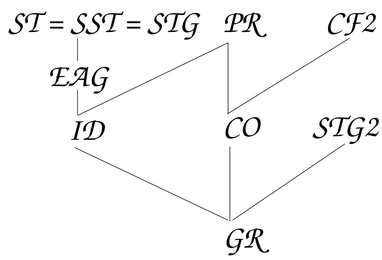

Figure 21. $\preceq_{\oplus}^{SE}$ relation for argumentation frameworks in $\mathcal{DE}_{\mathcal{ST}}$ ($\mathcal{DL}_{\mathcal{ST}}$)

the skeptical and credulous perspective are shown in Figures 20 and 21, for the general case and for argumentation frameworks where stable extensions exist respectively. As obvious, stronger relations entail lesser comparability between semantics, but one can note in particular that the role of $\mathcal{GR}$ as "bottom" skeptical reference with respect to all other semantics (but $\mathcal{STG}$) is confirmed.

5 Semantics agreement

The study of many different argumentation semantics in the literature indicates the richness and inherent complexity of the problem of "solving conflicts among arguments". In general terms, the introduction of a new semantics can be motivated by the need to achieve a particular desired outcome in some specific example considered particularly relevant and/or by the objective of satisfying some properties, at an intuitive or formal level.

While semantics differences obviously attract attention in analyses and comparisons, and have been discussed in Section 3, it is also important to characterize situations where argumentation semantics agree, i.e. exhibit the same behaviour. This is useful from several viewpoints. On one hand, situations where "most" (or even all) existing semantics agree can be regarded as providing a sort of reference behaviour against which further proposals should be confronted. On the other hand, it may be the case that in a specific application domain there are some restrictions on the structure of the argumentation frameworks that need to be considered. It is then surely interesting to know whether these restrictions lead to semantics agreement, since in this case the choice of the semantics to be adopted turns out to have no influence and, in a sense, the outcomes are universally supported in a semantics-independent way.

In this section we review the existing results about semantics agreement covering two distinct, though related, issues.

The first issue concerns the identification of the classes of argumentation frameworks where a given set of semantics agree. As will be shown in Section 5.1, using general properties of argumentation semantics only, it is possible to prove that there is a limited number of distinct agreement classes and identify them regardless of the topological properties of the argumentation frameworks belonging to them.

The second issue, dealt with in Section 5.2, concerns, in a complementary way, the analysis of argumentation frameworks where semantics agree from a topological viewpoint.

Before we enter the matter, some basic definitions and properties on semantics agreement need to be introduced.

Semantics agreement concerns comparing the set of extensions or, equivalently, of labellings prescribed by different semantics for the same framework. The analysis will be developed focusing on the extension-based approach, as this allows a more compact presentation. Furthermore, to avoid to deal with spurious situations, whenever we will discuss the properties of agreement of a given semantics σ we will implicitly refer, if not differently specified, only to argumentation frameworks belonging to $\mathcal{DE}_\sigma$, i.e. such that the set of extensions prescribed by σ is not empty.

The notion of agreement for a set of semantics on an argumentation framework is defined in the obvious way.

Definition 5.1 *Let $\mathbb{S}$ be a set of argumentation semantics and AF an argumentation framework such that $AF \in \bigcap_{\sigma \in \mathbb{S}} \mathcal{DE}_\sigma$. We say that AF is an* agreement framework *for $\mathbb{S}$ (or, equivalently, that the semantics in $\mathbb{S}$ agree on AF) iff $\forall \sigma_1, \sigma_2 \in \mathbb{S}$ it holds that $\mathcal{E}_{\sigma_1}(AF) = \mathcal{E}_{\sigma_2}(AF)$. The set of the agreement frameworks for $\mathbb{S}$ is denoted as $\mathcal{AGR}(\mathbb{S})$ and called the* agreement class *of $\mathbb{S}$.*

In general, it may be the case that $\mathcal{AGR}(\mathbb{S}_1) = \mathcal{AGR}(\mathbb{S}_2)$ for different sets of semantics $\mathbb{S}_1$ and $\mathbb{S}_2$: this in fact motivates the analysis carried out in Section 5.1. Moreover it is immediate to note that $\mathbb{S}_1 \subseteq \mathbb{S}_2 \Rightarrow \mathcal{AGR}(\mathbb{S}_2) \subseteq \mathcal{AGR}(\mathbb{S}_1)$.

5.1 Agreement classes

Systematic results on the identification of agreement classes have been obtained up to now for the semantics in the set $\Omega = \{\mathcal{GR}, \mathcal{ID}, \mathcal{CO}, \mathcal{PR}, \mathcal{ST}, \mathcal{CF}2, \mathcal{SST}\}$.

It can be noted that $\mathcal{AGR}(\mathbb{S}) \neq \emptyset$ for any set of semantics $\mathbb{S} \subseteq \Omega$ since all semantics belonging to Ω are obviously in agreement on the empty argumentation framework $AF_\emptyset = (\emptyset, \emptyset)$. Moreover, all semantics in Ω (and probably every

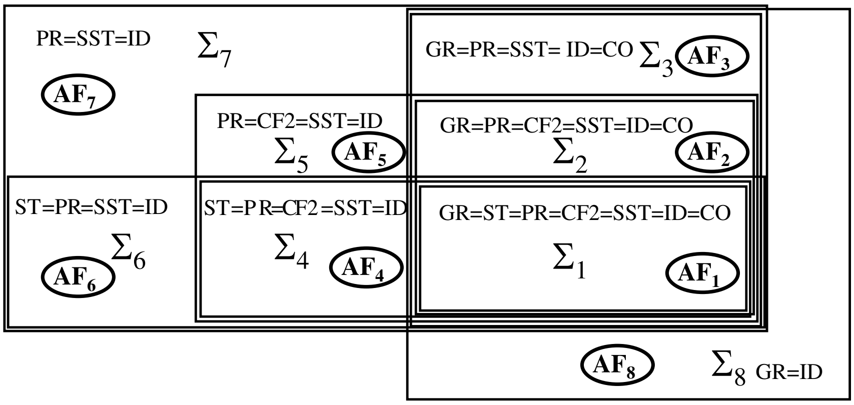

Figure 22. Venn diagram of unique-status agreement classes.

other reasonable argumentation semantics one can conceive) are in agreement also on attack-free argumentation frameworks. Namely, for an argumentation framework $AF = (\mathcal{A}rgs, \emptyset)$ for every $\sigma \in \Omega$ it holds that $\mathcal{E}_\sigma(AF) = \{\mathcal{A}rgs\}$.

As for the semantics not included in Ω, the following comments can be made. First, naïve semantics differs substantially from other semantics since, basically, it ignores the direction of attacks. Thus agreement with some other semantics can be achieved only in peculiar situations where the direction of attacks is not relevant (this is the case for instance of symmetric argumentation frameworks, mentioned in Section 5.2). As to stage semantics it must be recalled, first of all, that when stable extensions exist, they coincide with stage extensions. Thus agreement classes including stable semantics also implicitly include stage semantics. When stable extensions do not exist the peculiar behaviour of stage semantics in some cases (in particular the possibility that unattacked arguments are not included in an extension) implies that stage semantics may disagree with any other set of semantics. The identification and study of agreement classes involving eager or stage2 semantics appears to be potentially more interesting, but has not been developed in the literature yet.

Considering all subsets $\mathbb{S}$ of Ω such that $|\mathbb{S}| \geq 2$ gives rise, in principle, to 120 classes $\mathcal{AGR}(\mathbb{S})$ to be evaluated. Hovewer it has been proved in [Baroni and Giacomin, 2008] that the distinct agreement classes are only 14.

We will start by analyzing in subsection 5.1.1 unique-status agreement, namely the classes $\mathcal{AGR}(\mathbb{S})$ where $\mathbb{S}$ includes $\mathcal{GR}$ or $\mathcal{ID}$, and we will then examine agreement between multiple-status semantics in subsection 5.1.2.

5.1.1 Unique-status Agreement

The Venn diagram concerning unique-status agreement classes is shown in Figure 22 where rectangular boxes represent classes of agreement and small ellipses represent single argumentation frameworks to be used as specific examples.

The diagram will be illustrated in two main steps: first, we show that the set-theoretical relationships between the agreement classes $\Sigma_1, \ldots, \Sigma_8$ depicted in Figure 22 actually hold, then we explain why the classes $\Sigma_1, \ldots, \Sigma_8$ are the only meaningful ones in the context of unique-status agreement.

As to the first step, we proceed by following the partial order induced by inclusion (namely if $\Sigma_i \subsetneq \Sigma_j$ then $j > i$). While introducing each class Σ_i it will be necessary:

1. to identify which classes Σ_k, $k < i$, are included in Σ_i;
2. for each of these classes Σ_k to show that $\Sigma_i \setminus \Sigma_k \neq \emptyset$;
3. for any Σ_h such that $h < i$ and $\Sigma_h \nsubseteq \Sigma_i$ to examine $\Sigma_i \cap \Sigma_h$.

Point 1 will not be stressed since inclusion relationships can be directly derived from the inclusion of the relevant sets of semantics ($\mathbb{S}_1 \subseteq \mathbb{S}_2 \Rightarrow \mathcal{AGR}(\mathbb{S}_2) \subseteq \mathcal{AGR}(\mathbb{S}_1)$). As to point 2, examples of argumentation frameworks belonging to non-empty set differences will be given to prove that inclusion relationships are strict, while point 3 will be dealt with case by case.

Let us start from $\mathcal{AGR}(\{\mathcal{GR}, \mathcal{ID}, \mathcal{CO}, \mathcal{PR}, \mathcal{ST}, \mathcal{CF}2, \mathcal{SST}\})$, denoted as Σ_1. This is the class of argumentation frameworks where all semantics show a uniform single status behaviour in agreement with grounded semantics. Σ_1 includes, for instance, attack-free argumentation frameworks like the very simple $AF_1 = \langle \{a\}, \emptyset \rangle$.

$\mathcal{AGR}(\{\mathcal{GR}, \mathcal{ID}, \mathcal{CO}, \mathcal{PR}, \mathcal{CF}2, \mathcal{SST}\})$, denoted as Σ_2, corresponds to a uniform single status behaviour in agreement with grounded semantics by all but stable semantics. As shown in Figure 22, Σ_2 strictly includes Σ_1 since $(\Sigma_2 \setminus \Sigma_1)$ includes in particular $AF_2 = \langle \{a, b\}, \{(b, b)\} \rangle$.

$\Sigma_3 \triangleq \mathcal{AGR}(\{\mathcal{GR}, \mathcal{ID}, \mathcal{CO}, \mathcal{PR}, \mathcal{SST}\})$ is the last class in Figure 22 concerning agreement with grounded semantics. $(\Sigma_3 \setminus \Sigma_2) \neq \emptyset$ since it includes for instance $AF_3 = \langle \{a, b, c\}, \{(a, b), (b, c), (c, a)\} \rangle$. It is now worth noting that $\Sigma_3 \cap \mathcal{DE}_{\mathcal{ST}} = \Sigma_2 \cap \mathcal{DE}_{\mathcal{ST}} = \Sigma_1$. In fact, since semi-stable extensions coincide with stable extensions when the latter exist, whenever $AF \in \mathcal{DE}_{\mathcal{ST}}$ it must be the case that $AF \in \mathcal{AGR}(\mathbb{S})$ where $\mathbb{S}$ includes both $\mathcal{ST}$ and $\mathcal{SST}$.

On the left of Σ_1, Σ_2 and Σ_3 the diagram of Figure 22 shows four classes where several multiple-status semantics exhibit a unique-status behaviour in agreement with ideal semantics, but not necessarily also with grounded semantics. The smallest of these classes is $\Sigma_4 \triangleq \mathcal{AGR}(\{\mathcal{ID}, \mathcal{CF}2, \mathcal{ST}, \mathcal{PR}, \mathcal{SST}\})$. $(\Sigma_4 \setminus \Sigma_1) \neq \emptyset$ since it includes $AF_4 = \langle \{a, b\}, \{(a, b), (b, a), (b, b)\} \rangle$. Moreover, since $\Sigma_4 \subset \mathcal{DE}_{\mathcal{ST}}$, it is the case that $\Sigma_4 \cap (\Sigma_3 \setminus \Sigma_1) = \emptyset$. Not requiring agreement with stable semantics leads to $\Sigma_5 \triangleq \mathcal{AGR}(\{\mathcal{ID}, \mathcal{CF}2, \mathcal{PR}, \mathcal{SST}\})$. $(\Sigma_5 \setminus (\Sigma_4 \cup \Sigma_2)) \neq \emptyset$, since it includes $AF_5 = \langle \{a, b, c\}, \{(a, b), (b, a), (b, b), (c, c)\} \rangle$. Notice also that $\Sigma_5 \cap (\Sigma_3 \setminus \Sigma_2) = \emptyset$ since if $AF \in \Sigma_5 \cap \Sigma_3$ then, by definition of these classes, it also holds $AF \in \Sigma_2$. It is also clear that $\Sigma_5 \cap \mathcal{DE}_{\mathcal{ST}} = \Sigma_4$. Not requiring agreement with $\mathcal{CF}2$ semantics leads to $\Sigma_6 \triangleq \mathcal{AGR}(\{\mathcal{ID}, \mathcal{ST}, \mathcal{PR}, \mathcal{SST}\})$. $(\Sigma_6 \setminus \Sigma_4) \neq \emptyset$ since it includes $AF_6 = \langle \{a, b, c\}, \{(a, b), (a, c), (b, c), (c, a)\} \rangle$.

Again, since $\Sigma_6 \subset \mathcal{DE}_{\mathcal{ST}}$ it holds $\Sigma_6 \cap (\Sigma_3 \setminus \Sigma_1) = \emptyset$ and $\Sigma_6 \cap (\Sigma_5 \setminus \Sigma_4) = \emptyset$. Finally, excluding both stable and $\mathcal{CF}2$ semantics from the required agreement corresponds to $\Sigma_7 \triangleq \mathcal{AGR}(\{\mathcal{ID}, \mathcal{PR}, \mathcal{SST}\})$. $(\Sigma_7 \setminus (\Sigma_6 \cup \Sigma_5 \cup \Sigma_3)) \neq \emptyset$, since it includes $AF_7 = \langle \{a, b, c, d, e\}, \{(a, b), (b, a), (b, b), (c, d), (d, e), (e, c)\}\rangle$.

The last class concerning unique-status agreement is $\Sigma_8 \triangleq \mathcal{AGR}(\{\mathcal{ID}, \mathcal{GR}\})$. By definition of the relevant sets of semantics (and since no other distinct agreement classes exist, as recalled below) it holds that $\Sigma_8 \cap \Sigma_7 = \Sigma_3$, moreover it is easy to see that $\Sigma_8 \setminus \Sigma_3 \neq \emptyset$, since it includes $AF_8 = \langle \{a, b\}, \{(a, b), (b, a)\}\rangle$.

In [Baroni and Giacomin, 2008] it is shown that no other classes than $\Sigma_1, \ldots, \Sigma_8$ are meaningful in the context of unique-status agreement. The proofs are based on a set of basic properties and relationships, which we recall together in Proposition 5.2 (see Lemmata 2-10 in [Baroni and Giacomin, 2008]) as they are of general interest in the analysis and understanding of argumentation semantics.

Proposition 5.2 *Given an argumentation framework AF the following statements hold:*

1. *if* $|\mathcal{E}_{\mathcal{PR}}(AF)| = 1$*, then*
   - $\mathcal{E}_{\mathcal{PR}}(AF) = \mathcal{E}_{\mathcal{ID}}(AF)$*;*
   - $\mathcal{E}_{\mathcal{PR}}(AF) = \mathcal{E}_{\mathcal{SST}}(AF)$*;*
   - *if* $AF \in \mathcal{DE}_{\mathcal{ST}}$, $\mathcal{E}_{\mathcal{PR}}(AF) = \mathcal{E}_{\mathcal{ST}}(AF)$*.*
2. *if* $|\mathcal{E}_{\mathcal{CO}}(AF)| = 1$*, then* $\mathcal{E}_{\mathcal{CO}}(AF) = \mathcal{E}_{\mathcal{GR}}(AF) = \mathcal{E}_{\mathcal{PR}}(AF)$*;*
3. *let* $\sigma \in \{\mathcal{PR}, \mathcal{ST}, \mathcal{SST}, \mathcal{CF}2\}$*, if* $GE(AF) \in \mathcal{E}_\sigma(AF)$ *then* $\mathcal{E}_\sigma(AF) = \{GE(AF)\}$*;*
4. *if* $\mathcal{E}_{\mathcal{PR}}(AF) = \{GE(AF)\}$ *then* $\mathcal{E}_{\mathcal{CO}}(AF) = \{GE(AF)\}$*;*
5. *if* $\mathcal{E}_{\mathcal{ST}}(AF) = \{GE(AF)\}$ *then* $\mathcal{E}_{\mathcal{CF}2}(AF) = \{GE(AF)\}$*;*
6. *if* $\mathcal{E}_{\mathcal{CF}2}(AF) = \{GE(AF)\}$ *then* $\mathcal{E}_{\mathcal{CF}2}(AF) = \mathcal{E}_{\mathcal{PR}}(AF)$*;*
7. *if* $\mathcal{E}_{\mathcal{CF}2}(AF) = \{ID(AF)\}$ *then* $\mathcal{E}_{\mathcal{CF}2}(AF) = \mathcal{E}_{\mathcal{PR}}(AF)$*;*
8. *if* $\mathcal{E}_{\mathcal{SST}}(AF) = \{ID(AF)\}$ *then* $\mathcal{E}_{\mathcal{SST}}(AF) = \mathcal{E}_{\mathcal{PR}}(AF)$*;*
9. *if* $\mathcal{E}_{\mathcal{CF}2}(AF) \subseteq \mathcal{AS}(AF)$ *then* $\mathcal{E}_{\mathcal{CF}2}(AF) = \mathcal{E}_{\mathcal{PR}}(AF)$*.*

These results can be briefly expressed in words as follows. Item 1 states that when there is a unique preferred extension, it is also the unique ideal, semi-stable, and (possibly) stable extension, while item 2 states that when there is a unique complete extension it is also the unique preferred extension and the grounded extension. Item 3 says that the set of preferred extensions can include the grounded extension only if it contains only the grounded extension (and the same holds for stable, semi-stable and $\mathcal{CF}2$ extensions). If the grounded extension is the only preferred extension it is also the only complete extension (item 4), while if it is the only stable extension it is also the only $\mathcal{CF}2$ extension (item 5) and the only preferred extension (item 6). Finally, if the set of $\mathcal{CF}2$ extensions or of semi-stable extensions contains only the ideal extension this is also the unique preferred extension (items 7 and 8) and if all $\mathcal{CF}2$ extensions are admissible they coincide with the preferred extensions (item 9).

On the basis of these results, it is shown in [Baroni and Giacomin, 2008] that any agreement class $\mathcal{AGR}(\mathbb{S})$ where $\mathbb{S}$ includes either $\mathcal{GR}$ or $\mathcal{ID}$ coincides with one of the classes $\Sigma_1, \ldots, \Sigma_8$. Note in particular that Σ_1, Σ_2 and Σ_3 are the only classes where $\mathcal{CO}$ appears: this is due to the fact that $GE(AF) \in \mathcal{E}_{\mathcal{CO}}(AF)$ which implies that agreement between complete semantics and other semantics can only occur when the set of extensions consists of the grounded extension only. For this reason $\mathcal{CO}$ will not need to be considered any more in next subsection concerning multiple-status agreement.

5.1.2 Multiple-status Agreement

The complete Venn diagram concerning all agreement classes is shown in Figure 23, where the rectangle with bold lines including all the others represents the universe of all finite argumentation frameworks and again rectangular boxes represent classes of agreement and small ellipses represent single argumentation frameworks to be used as specific examples. As in the previous subsection, the diagram will be illustrated by examining first the set-theoretical relationships between the agreement classes depicted in Figure 23 and then stating that no other meaningful classes exist.

The first step will encompass the same three points as in the previous subsection. In particular, as to point 3, it can be noticed from Figure 23 that most intersections between classes correspond to unions of previously identified classes and/or differences between classes and can be easily determined by considering the sets of semantics involved. For this reason, only intersections requiring specific explanations (in particular all those concerning Σ_8) will be explicitly discussed. Our analysis will now concern agreement classes involving only multiple-status semantics except $\mathcal{CO}$.

The smallest one includes all of them: $\Sigma_9 \triangleq \mathcal{AGR}(\{\mathcal{PR}, \mathcal{CF}2, \mathcal{ST}, \mathcal{SST}\})$. $(\Sigma_9 \setminus \Sigma_4) \neq \emptyset$ as it includes $AF_9 = \langle \{a, b, c, d\}, \{(a, b), (b, a), (a, c), (b, c), (c, d)\} \rangle$. Note that $AF_9 \in ((\Sigma_9 \setminus \Sigma_4) \cap \Sigma_8)$. Also $\Sigma_9 \setminus (\Sigma_4 \cup \Sigma_8)$ is not empty since it includes, for example, $AF_{9'} = \langle \{a, b, c, d\}, \{(a, b), (b, a), (a, c), (b, c), (c, d), (d, c)\} \rangle$.

$\Sigma_{10} \triangleq \mathcal{AGR}(\{\mathcal{PR}, \mathcal{CF}2, \mathcal{SST}\})$ covers the case where stable extensions do not exist, while all other multiple-status semantics agree. $\Sigma_{10} \setminus (\Sigma_9 \cup \Sigma_5) \neq \emptyset$ since it includes, for instance, $AF_{10} = \langle \{a, b, c\}, \{(a, b), (b, a), (c, c)\} \rangle$. Note that $AF_{10} \in ((\Sigma_{10} \setminus (\Sigma_9 \cup \Sigma_5)) \cap \Sigma_8)$. Also $(\Sigma_{10} \setminus (\Sigma_9 \cup \Sigma_5 \cup \Sigma_8))$ is not empty since it includes, for example, $AF_{10'} = AF_{10} \uplus AF_4$, where, given two argumentation frameworks $AF_1 = (Ar_1, att_1)$, $AF_2 = (Ar_2, att_2)$ such that $Ar_1 \cap Ar_2 = \emptyset$, we define[24] $AF_1 \uplus AF_2 \triangleq (Ar_1 \cup Ar_2, att_1 \cup att_2)$.

[24] While we have used the same labels $a, b, \ldots$ to denote arguments of our sample argumentation frameworks, we implicitly assume that arguments with the same label in di erent

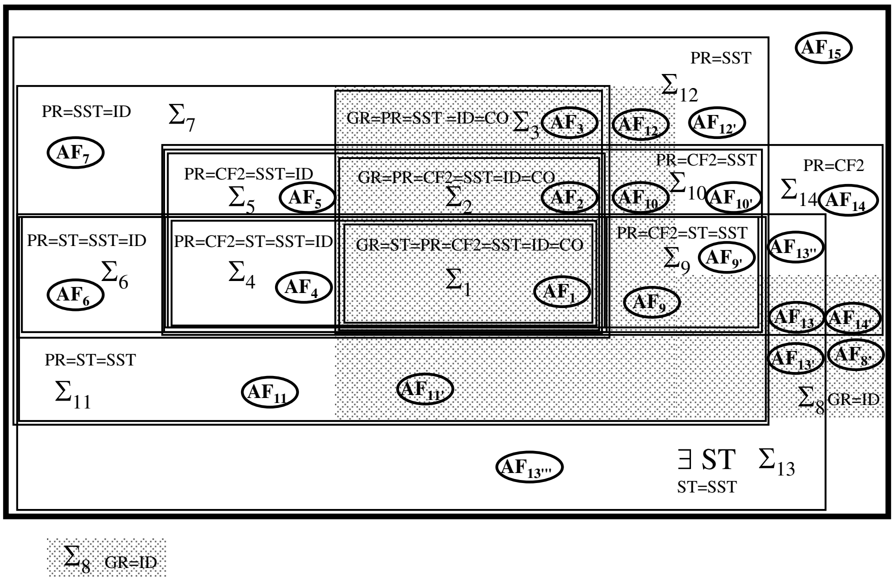

Figure 23. Venn diagram of agreement classes.

$\Sigma_{11} \triangleq \mathcal{AGR}(\{\mathcal{PR}, \mathcal{ST}, \mathcal{SST}\})$ coincides with the class of *coherent* argumentation frameworks considered in [Dung, 1995]. $(\Sigma_{11} \setminus (\Sigma_6 \cup \Sigma_9)) \neq \emptyset$ since it includes $AF_{11} = \langle \{a, b, c, d, e\}, \{(a, b), (a, c), (b, c), (c, a), (d, e), (e, d)\} \rangle$. It can be noted that $AF_{11} \notin \Sigma_8$. Also $(\Sigma_{11} \setminus (\Sigma_6 \cup \Sigma_9)) \cap \Sigma_8 \neq \emptyset$ since it includes $AF_{11'} = \langle \{a, b, c, d\}, \{(a, b), (a, c), (a, d), (b, c), (b, d), (c, a), (d, a), (d, b), (d, c)\} \rangle$.

We are now left with three classes where only a pair of multiple-status semantics are in agreement.

Let us start by considering $\Sigma_{12} \triangleq \mathcal{AGR}(\mathcal{PR}, \mathcal{SST})$. $\Sigma_{12} \setminus (\Sigma_7 \cup \Sigma_{10} \cup \Sigma_{11}) \neq \emptyset$ as it includes $AF_{12} = \langle \{a, b, c, d, e\}, \{(a, b), (b, a), (c, d), (d, e), (e, c)\} \rangle$. It can be noted that $AF_{12} \in \Sigma_8$. For an example of argumentation framework included in $\Sigma_{12} \setminus (\Sigma_7 \cup \Sigma_8 \cup \Sigma_{10} \cup \Sigma_{11})$ consider $AF_{12'} = AF_{12} \uplus AF_4$.

Finally, as already remarked, $\Sigma_{13} \triangleq \mathcal{AGR}(\mathcal{ST}, \mathcal{SST})$ coincides with the class $\mathcal{DE}_{\mathcal{ST}}$ of argumentation frameworks where stable extensions exist, while the last pair to be considered corresponds to $\Sigma_{14} \triangleq \mathcal{AGR}(\mathcal{PR}, \mathcal{CF}2)$. The part of the diagram still to be illustrated involves argumentation frameworks outside Σ_{12} and requires an articulated treatment, since the intersections $\Sigma_{13} \cap \Sigma_{14}$, $\Sigma_{13} \cap \Sigma_8$, and $\Sigma_{14} \cap \Sigma_8$ do not allow a simple characterisation in terms of the other identified classes. First the set difference $\Sigma_{13} \setminus \Sigma_{12}$ can be partitioned

argumentation frameworks are actually distinct. This assumption allows us to apply the combination operator $\uplus$ to any pair of sample argumentation frameworks while keeping a simple notation. To make the distinction explicit, argument a of AF_4 should be actually labeled a_4, argument a of AF_{10} should be labeled a_{10}, and so on, but we regard this as an unnecessary notational burden.

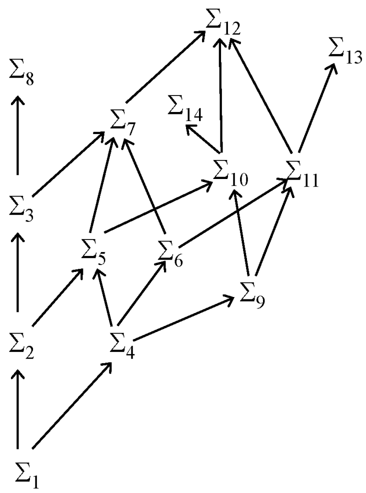

Figure 24. Inclusion relations between agreement classes.

into four non-empty subsets:

- $((\Sigma_{13} \setminus \Sigma_{12}) \cap \Sigma_{14} \cap \Sigma_8) \ni AF_{13} = \langle \{a, b, c\}, \{(a, b), (b, a), (a, c), (c, c)\} \rangle$;
- $((\Sigma_{13} \setminus \Sigma_{12}) \cap (\Sigma_8 \setminus \Sigma_{14})) \ni AF_{13'} = AF_{13} \uplus AF_{11'}$;
- $((\Sigma_{13} \setminus \Sigma_{12}) \cap (\Sigma_{14} \setminus \Sigma_8)) \ni AF_{13''} = AF_{13} \uplus AF_4$;
- $(\Sigma_{13} \setminus (\Sigma_{12} \cup \Sigma_{14} \cup \Sigma_8)) \ni AF_{13'''} = AF_{13} \uplus AF_4 \uplus AF_{11'}$.

Then, $\Sigma_{14} \setminus (\Sigma_{12} \cup \Sigma_{13})$ can be partitioned into two non-empty subsets:

- $(\Sigma_{14} \setminus (\Sigma_{12} \cup \Sigma_{13} \cup \Sigma_8)) \ni AF_{14} = AF_{13} \uplus AF_4 \uplus AF_2$;
- $((\Sigma_{14} \setminus (\Sigma_{12} \cup \Sigma_{13})) \cap \Sigma_8) \ni AF_{14'} = AF_{13} \uplus AF_2$.

Only the characterisation of Σ_8 remains to be completed, in fact $\Sigma_8 \setminus (\Sigma_{12} \cup \Sigma_{13} \cup \Sigma_{14})$ is not empty since it includes $AF_{8'} = AF_{13} \uplus AF_2 \uplus AF_{11'}$. Finally, argumentation frameworks where all semantics are in mutual disagreement also exist, like $AF_{15} = AF_{13} \uplus AF_2 \uplus AF_4 \uplus AF_{11'}$. The fact that no other agreement classes involving $\mathcal{CF}2$, $\mathcal{ST}$, $\mathcal{SST}$ and $\mathcal{PR}$ are meaningful is proved in [Baroni and Giacomin, 2008] on the basis of item 9 of Proposition 5.2.

The Hasse diagram corresponding to inclusion relationships between the agreement classes described above is shown in Figure 24 (where arrows point from subsets to supersets).

5.2 Topological properties and semantics agreement

Section 5.1 identifies the distinct agreement classes concerning a significant set of semantics and provides examples of argumentation frameworks belonging to them. It does not give any indication on how, given an argumentation framework, to answer the question of which agreement class(es) it belongs to (apart the obvious method of computing the sets of extensions prescribed by the various semantics considered). As a matter of fact, there are some significant relationships between agreement classes and some topological properties of argumentation frameworks: presenting them is the subject of this subsection.

5.2.1 Well-founded frameworks

The issue of single-status agreement has been considered as early as in the seminal paper by [Dung, 1995] where it is shown that a sufficient condition for agreement among grounded, preferred and stable semantics is that the argumentation framework is well-founded.

Definition 5.3 (Definition 29 of [Dung, 1995]) *An argumentation framework is well-founded iff there exists no infinite sequence $a_0, a_1, \ldots, a_n, \ldots$ of (not necessarily distinct) arguments such that for each i, a_{i+1} attacks a_i.*

In the case of a finite argumentation framework, well-foundedness coincides with acyclicity of the attack relation. In the light of the results presented in Section 5.1 acyclicity turns out to be a sufficient condition for membership in the agreement class Σ_1.

5.2.2 Determined argumentation frameworks

We consider now a more general conditions for agreement with grounded semantics. To this purpose we introduce the notion of *determined* argumentation framework.

Definition 5.4 *An argumentation framework $AF = \langle Ar, att \rangle$ is* determined *if and only if $\nexists a \in Ar : a \notin GE(AF) \wedge GE(AF) \cap a^- = \emptyset$.*

In words, an argumentation framework AF is determined if and only the grounded extension is also a stable extension. On the basis of the results of Section 5.1 the set of determined argumentation frameworks, denoted as $\mathcal{DET}$, coincides with the agreement class Σ_1.

Well-founded argumentation frameworks are a special case of determined argumentation frameworks: for finite frameworks, the absence of cycles is a sufficient but not necessary topological condition for $AF \in \mathcal{DET}$. A simple example of argumentation framework which is determined without being acyclic is the following: $\langle \{a, b, c\}, \{(a, b), (b, c), (c, b)\} \rangle$.

Actually, as shown in [Baroni and Giacomin, 2007a] the absence of cycles is necessary only in the initial SCCs, and then recursively in the initial SCCs of the restricted argumentation frameworks obtained by taking into account that the nodes corresponding to the initial SCCs are necessarily included in any extension. This observation gives rise to a full topological characterisation of determined argumentation frameworks, i.e. of Σ_1.

Definition 5.5 *An argumentation framework* $AF = \langle Ar, att \rangle$ *is initial-acyclic if* $AF = \langle \emptyset, \emptyset \rangle$ *or the following condition holds:* $\forall S \in \mathcal{IS}(AF)$ S *is monadic and* $AF{\downarrow}_{UP_{AF}((Ar \setminus IN(AF)), IN(AF))}$ *is initial-acyclic, where*

- $\mathcal{IS}(AF)$ *is the set of strongly connected components of* AF *which are not attacked by any other strongly connected component;*
- $IN(AF) \triangleq \bigcup_{S \in \mathcal{IS}(AF)} S$*;*
- *an argumentation framework is monadic iff it consists of a single non-self-defeating argument;*
- *for any* $\mathcal{A}rgs, S \subseteq Ar$*,* $UP_{AF}(S, \mathcal{A}rgs) = \{a \in S \mid \nexists b \in \mathcal{A}rgs \setminus S : (b, a) \in att\}$ *(as in Definition 3.61).*

The base of this recursive definition is represented by the empty argumentation framework. The recursion is well-founded as the set $IN(AF)$ is non-empty for a non-empty argumentation framework, which means that at each recursive step an argumentation framework with a strictly lesser number of nodes is considered. The set of initial-acyclic argumentation frameworks is denoted by $\mathcal{IAA}$ and it is proved in [Baroni and Giacomin, 2007a] that $\mathcal{IAA} = \mathcal{DET}$.

Proposition 5.6 *For any argumentation framework* $AF = \langle Ar, att \rangle$*,* $AF \in \mathcal{IAA}$ *if and only if* $AF \in \mathcal{DET}$*.*

5.2.3 Almost determined argumentation frameworks

While determined argumentation frameworks ensure agreement among all the semantics belonging to the set Ω, it can be observed that there is a larger class of argumentation frameworks where an almost total agreement is reached. Consider for instance the case of an argumentation framework consisting just of a self-defeating argument, namely $AF = \langle \{a\}, \{(a, a)\} \rangle$. In this case we have that $\mathcal{E}_{\mathcal{GR}}(AF) = \{\emptyset\}$ and, in virtue of the conflict-free property, for any semantics σ which admits extensions on AF it must also hold that $\mathcal{E}_\sigma(AF) = \{\emptyset\}$. However, since stable semantics is unable to prescribe extensions in this case, $\mathcal{E}_{\mathcal{ST}}(AF) = \emptyset \neq \{\emptyset\}$. In this case, disagreement arises from the non-existence of stable extensions rather than from the existence of extensions different from $GE(AF)$ and clearly AF belongs to Σ_2.

On the basis of this observation, it is useful to consider a further class of argumentation frameworks, called *almost determined.*

Definition 5.7 *An argumentation framework* $AF = \langle Ar, att \rangle$ *is* almost determined *if and only if for any* $a \in Ar$*,* $(a \notin GE(AF) \land GE(AF) \cap a^- = \emptyset) \Rightarrow (a, a) \in att$*.*

In words, an argumentation framework is almost determined if all the nodes which are not attacked nor included in the grounded extension are self-defeating. The set of almost determined argumentation frameworks is denoted as $\mathcal{AD}$. It is proved in [Baroni and Giacomin, 2007a] that frameworks in $\mathcal{AD}$ ensure agreement for all multiple status semantics in Ω but stable semantics (and actually for every SCC-recursive semantics satisfying some basic properties) and that outside $\mathcal{AD}$ there cannot be agreement with $\mathcal{CF}2$ semantics.

Proposition 5.8 *For any argumentation framework* $AF = \langle Ar, att \rangle \in \mathcal{AD}$ *it holds that* $\mathcal{E}_\sigma(AF) = \{GE(AF)\}$ *for* $\sigma \in \{\mathcal{CO}, \mathcal{PR}, \mathcal{SST}, \mathcal{CF}2\}$.

Proposition 5.9 *For any argumentation framework* $AF = \langle Ar, att \rangle \notin \mathcal{AD}$ $\mathcal{E}_{\mathcal{CF}2}(AF) \neq \{GE(AF)\}$.

The propositions above shows that $\mathcal{AD} = \Sigma_2$.

5.2.4 Limited controversial frameworks

Sections 5.2.1 - 5.2.3 deal with cases of single status agreement, turning now to multiple status agreement, a first basic result, concerning preferred and stable semantics, was introduced, again, in Dung's seminal paper with reference to the notion of *limited controversial* frameworks, based on the one of *controversial arguments*.

Definition 5.10 *Given an argumentation framework* $AF = \langle Ar, att \rangle$ *an argument* a indirectly attacks *an argument* b *iff there exists a finite sequence* $a_0, \ldots a_{2n+1}$ *such that* $a = a_0$, $b = a_{2n+1}$, *and* $\forall i \in \{0, \ldots, 2n\}$ $a_i \in a^-_{i+1}$. *An argument* a indirectly defends *an argument* b *iff there exists a finite sequence* $a_0, \ldots a_{2n}$ *such that* $a = a_0$, $b = a_{2n}$, *and* $\forall i \in \{0, \ldots, 2n-1\}$ $a_i \in a^-_{i+1}$. *An argument* a *is controversial with respect to an argument* b *if* a *indirectly attacks and indirectly defends* b. *An argument* a *is controversial if it is controversial with respect to an argument* b.

Definition 5.11 *An argumentation framework* $AF = \langle Ar, att \rangle$ *is* limited controversial *if and only if there is no infinite sequence of arguments* $a_0, \ldots, a_n, \ldots$ *such that* a_{i+1} *is controversial with respect to* a_i.

In the finite case, an argumentation framework is limited controversial if it does not include any odd-length cycle. In [Dung, 1995] it is shown that being limited controversial is a sufficient condition for being coherent, i.e. for membership in the agreement class Σ_{11}.

5.2.5 Agreement with stable semantics

The basic result on multiple status agreement recalled in Section 5.2.4 has been extended by [Baroni and Giacomin, 2007a] by characterizing a family of argumentation frameworks, called SCC-symmetric, where agreement is ensured for a class of multiple-status semantics including stable, preferred and $\mathcal{CF}2$ semantics.

First we need to introduce the notion of symmetric argumentation frameworks, noting also that symmetry is preserved in all the restrictions of a symmetric framework.

Definition 5.12 *An argumentation framework $AF = \langle Ar, att \rangle$ is* symmetric *if for any $a, b \in Ar$, $a \in b^- \Leftrightarrow b \in a^-$.*

Lemma 5.13 *Given a symmetric argumentation framework $AF = \langle Ar, att \rangle$ and a set $S \subseteq Ar$, $AF{\downarrow_S}$ is symmetric.*

As it will be more evident from Proposition 5.15, it is quite natural that extensions of a symmetric argumentation framework free of self-defeating arguments coincide with its maximal conflict free sets, if the multiple-status approach is adopted. Argumentation semantics satisfying this requirement will be called *-symmetric.

Definition 5.14 *An argumentation semantics σ is *-symmetric if for any argumentation framework AF which is symmetric and free of self-defeating arguments $\mathcal{E}_\sigma(AF) = \mathcal{E}_{\mathcal{NA}}(AF)$.*

Several significant multiple-status semantics, though their definition is based on quite different principles, share the property of being *-symmetric.

Proposition 5.15 *Stable semantics, preferred semantics, semi-stable semantics, and $\mathcal{CF}2$ semantics are *-symmetric.*

In symmetric argumentation frameworks non-mutual attacks cannot exist: this seriously limits their applicability for modeling practical situations. Their properties however provide the basis for analyzing a more interesting family of argumentation frameworks called *SCC-symmetric.*

Definition 5.16 *An argumentation framework AF is SCC-symmetric if $\forall S \in \mathrm{SCCS}_{AF}$ $AF{\downarrow_S}$ is symmetric.*

Definition 5.16 is equivalent to impose that all attacks participating in an attack cycle are mutual, while non-mutual attacks are allowed outside cycles.

Proposition 5.17 *An argumentation framework AF is SCC-symmetric if and only if for every attack cycle, i.e. for every sequence $a_0, a_1, \ldots, a_n$ such that $a_0 = a_n$ and $\forall i \in \{0, \ldots, n-1\}$ $a_i \in a_{i+1}^-$, it holds that $\forall i \in \{1, \ldots, n\}$ $a_i \in a_{i-1}^-$.*

Theorem 5.18 provides the main result about agreement in SCC-symmetric argumentation frameworks.

Theorem 5.18 *In any argumentation framework which is SCC-symmetric and free of self-defeating arguments all SCC-recursive *-symmetric semantics are in agreement, i.e. they prescribe the same set of extensions.*

The following result immediately follows from the previous theorem and Proposition 5.15.

Corollary 5.19 *If an argumentation framework AF is SCC-symmetric and free of self-defeating arguments then* $\mathcal{E}_{\mathcal{PR}}(AF) = \mathcal{E}_{CF2}(AF) = \mathcal{E}_{\mathcal{ST}}(AF) = \mathcal{E}_{\mathcal{SST}}(AF)$*, thus in particular AF is coherent.*

Summing up, being SCC-symmetric is a sufficient condition for membership in the agreement class Σ_9.

It may be noted that the classes of SCC-symmetric and limited controversial argumentation frameworks are non-disjoint but distinct. In fact, an SCC-symmetric argumentation framework may contain cycles of any length, while a limited controversial argumentation framework may consist, for instance, of an even-length cycle which is not symmetric.

It is interesting to note that the property of SCC-symmetry may be recovered from assumptions on the underlying method of argument and attack construction, which have been considered in the literature on structured argumentation and are not directly related to decomposition into SCCs. For instance in [Baroni *et al.*, 2005a] the case is considered where conflicts among arguments arise only from contradicting conclusions, namely only the *rebutting* kind of defeat is allowed while *undercutting* defeat is not (we follow here the terminology of [Pollock, 1992]). Briefly, rebutting defeats concern arguments with contradictory conclusions, while undercutting defeats concern questioning the applicability of the rule used to build an argument rather than its conclusion (see for instance [Modgil and Prakken, 2014]). This distinction is relevant because undercutting defeats can give rise to arbitrary attack cycles, while rebutting defeats cannot, since they are non mutual only if some preference ordering between arguments is in place and this ordering prevents the existence of attack cycles which are not symmetric. Formally, it is shown in Proposition 26 of [Baroni *et al.*, 2005a] that if only rebutting defeats are allowed, the corresponding argumentation framework is SCC-symmetric (such a framework is called *r-type* in [Baroni *et al.*, 2005a]). From another perspective, in [Kaci *et al.*, 2006] it is shown that when the attack relation results from a symmetric conflict relation and a transitive preference relation between arguments the argumentation framework satisfies a property called *strict acyclicity*, which is actually equivalent to SCC-symmetry through the characterisation given in Proposition 5.17.

6 Conclusions

Starting from Dung's seminal paper [Dung, 1995] abstract argumentation has received a growing interest by the research community, witnessed by a large corpus of scientific literature where an increasing variety of alternative semantics proposals is complemented by studies on general principles and properties for their assessment and comparison. The current book chapter, which revises and updates a previous tutorial paper [Baroni *et al.*, 2011a] is meant to provide a reasonably complete and up-to-date introductory survey on these aspects. In particular it provides a side-by-side treatment of the extension-based and labelling-based approaches and a coverage of skepticism-based semantics comparison and semantics agreement. While this can be sufficient to get a first impression of the subject, there are many other core aspects that are relevant for obtaining a comprehensive view, ranging from postulates and principles for argumentation semantics to notions of equivalence and locality and modularity issues: these additional aspects are covered in Part E of this volume.

Furthermore, some of the relevant aspects have only been briefly mentioned in the current chapter, for instance the procedures for computing labellings and extensions or for checking other semantics-related properties. The relevant theoretical and implementation issues are extensively treated in part D of this volume.

Many lines of further development have sprung up from Dung's core theory and are the subject of active investigation.

First, the basic model can be enriched considering other kinds of relationships between arguments in addition to attacks, like, for instance, support considered in bipolar argumentation frameworks [Cayrol and Lagasquie-Schiex, 2013; Cohen *et al.*, 2014]. Abstract Dialectical Frameworks, presented in chapter 5, are a generalized graph-based abstract formalism able to capture different kinds of interactions between the graph nodes. Other extensions of the model include considering attacks to attacks [Modgil, 2009; Baroni *et al.*, 2011b] and taking into account values and audiences in argument evaluation as in value-based argumentation frameworks [Bench-Capon, 2003].

The survey we presented in this chapter does not exhaustively treat the whole range of argumentation semantics in the literature. In particular, it is worth mentioning the notion of parametric semantics, namely semantics definition schemes which are generic with respect to the choice of an argumentation semantics, playing the role of a parameter in the context of a scheme.

Resolution-based semantics [Baroni *et al.*, 2011c] is an example of parametric semantics. Here the idea is that, given an argumentation framework AF, a set of argumentation frameworks $\mathcal{RES}(AF)$ is generated, where each member of $\mathcal{RES}(AF)$ corresponds to a resolution of all the mutual attacks in AF. Then the semantics σ adopted as parameter is applied to the elements of $\mathcal{RES}(AF)$ and among the resulting extensions those which are minimal with respect to set inclusion are selected as extensions of the resolution-based semantics for σ. The reader is referred to [Baroni *et al.*, 2011c] for all details, we only recall here that the instance of the resolution-based definition scheme based on grounded semantics shows unparalleled features in terms of principle-based evaluation, while also enjoying good computational complexity properties among multiple-status semantics [Baroni *et al.*, 2009].

Another example of parametric semantics concerns the generalisation of the notion of ideal extension [Dunne *et al.*, 2013]: basically this amounts to make Definitions 3.45 and 3.48 parametric by replacing the reference to preferred semantics with the reference to a generic semantics. This generalisation turns out to be suitably applicable to value-based argumentation frameworks.

While traditional abstract argumentation semantics produces qualitative assessments, a variety of quantitative approaches have been considered. Their study is quickly evolving, ranging from weighted argumentation frameworks [Dunne *et al.*, 2011] to various flavors of probabilistic argumentation frameworks [Hunter, 2013; Hunter, 2014] to the equational approach to argumentation semantics [Gabbay, 2012], just to mention a few. The coverage of this kind of extensions is planned for the second volume of this handbook.

To conclude, we note that while the ongoing developments listed above promise to overcome some of the restrictive assumptions embedded in the original theory of abstract argumentation frameworks, they also witness the extraordinary interest and fertility of this formalism and ensure that it will continue to represent an active research subject for many years to come.

Acknowledgments

The authors are grateful to the reviewers for their careful analysis of the chapter and their insightful suggestions.

BIBLIOGRAPHY

[Alferes *et al.*, 1993] Jose Julio Alferes, Phan Minh Dung, and Lu s Pereira. Scenario semantics of extended logic programs. In A. Nerode and L. Pereira, editors, *Proc. of the 2nd Int. Workshop on Logic Programming and Non-monotonic Reasoning (LPNMR 93)*, pages 334 348. MIT Press, 1993.

[Amgoud, 2009] Leila Amgoud. Argumentation for decision making. In I. Rahwan and G.R. Simari, editors, *Argumentation in Artificial Intelligence*, pages 301 320. Springer, 2009.

[Baroni and Giacomin, 2003] Pietro Baroni and Massimiliano Giacomin. Solving semantic problems with odd-length cycles in argumentation. In *Proc. of the 7th European Conf. on Symbolic and Quantitative Approaches to Reasoning with Uncertainty (ECSQARU 03)*, number 2711 in Springer Lecture Notes in AI, pages 440 451, Berlin, 2003. Springer Verlag.

[Baroni and Giacomin, 2007a] Pietro Baroni and Massimiliano Giacomin. Characterizing defeat graphs where argumentation semantics agree. In *Proc. of 1st Int. Workshop on Argumentation and Non-Monotonic Reasoning (ARGNMR07)*, pages 33 48, 2007.

[Baroni and Giacomin, 2007b] Pietro Baroni and Massimiliano Giacomin. On principle-based evaluation of extension-based argumentation semantics. *Artificial Intelligence*, 171(10-15):675 700, 2007.

[Baroni and Giacomin, 2008] Pietro Baroni and Massimiliano Giacomin. A systematic classi cation of argumentation frameworks where semantics agree. In *Proc. of the 2nd Int. Conf. on Computational Models of Argument (COMMA 2008)*, pages 37 48, 2008.

[Baroni and Giacomin, 2009a] Pietro Baroni and Massimiliano Giacomin. Semantics of abstract argument systems. In I. Rahwan and G.R. Simari, editors, *Argumentation in Artificial Intelligence*, pages 25 44. Springer, 2009.

[Baroni and Giacomin, 2009b] Pietro Baroni and Massimiliano Giacomin. Skepticism relations for comparing argumentation semantics. *Int. J. of Approximate Reasoning*, 50(6):854 866, 2009.

[Baroni *et al.*, 2005a] Pietro Baroni, Massimiliano Giacomin, and Giovanni Guida. Self-stabilizing defeat status computation: dealing with con ict management in multi-agent systems. *Artificial Intelligence*, 165(2):187 259, 2005.

[Baroni *et al.*, 2005b] Pietro Baroni, Massimiliano Giacomin, and Giovanni Guida. SCC-recursiveness: a general schema for argumentation semantics. *Artificial Intelligence*, 168(1-2):165 210, 2005.

[Baroni *et al.*, 2009] Pietro Baroni, Paul E. Dunne, and Massimiliano Giacomin. Computational properties of resolution-based grounded semantics. In *Proc. of the 21st International Joint Conference on Artificial Intelligence (IJCAI 2009)*, pages 683 689, 2009.

[Baroni *et al.*, 2011a] Pietro Baroni, Martin W. A. Caminada, and Massimiliano Giacomin. An introduction to argumentation semantics. *Knowledge Engineering Review*, 26(4):365 410, 2011.

[Baroni *et al.*, 2011b] Pietro Baroni, Federico Cerutti, Massimiliano Giacomin, and Giovanni Guida. AFRA: argumentation framework with recursive attacks. *Int. J. of Approximate Reasoning*, 52(1):19 37, 2011.

[Baroni *et al.*, 2011c] Pietro Baroni, Paul E. Dunne, and Massimiliano Giacomin. On the resolution-based family of abstract argumentation semantics and its grounded instance. *Artificial Intelligence*, 175(3-4):791 813, 2011.

[Baumann *et al.*, 2016] Ringo Baumann, Thomas Linsbichler, and Stefan Woltran. Veri a-bility of argumentation semantics. In P. Baroni, T. F. Gordon, T. Sche er, and M. Stede, editors, *Proc. of the 6th Int. Conf. on Computational Models of Argument (COMMA 2016)*, pages 83 94. IOS Press, 2016.

[Bench-Capon, 2003] Trevor J.M. Bench-Capon. Persuasion in practical argument using value-based argumentation frameworks. *J. of Logic and Computation*, 13(3):429 448, 2003.

[Caminada and Gabbay, 2009] Martin W. A. Caminada and Dov M. Gabbay. A logical account of formal argumentation. *Studia Logica*, 93(2-3):109 145, 2009. Special issue: new ideas in argumentation theory.

[Caminada and Pigozzi, 2011] Martin W. A. Caminada and Gabriella Pigozzi. On judgment aggregation in abstract argumentation. *Autonomous Agents and Multi-Agent Systems*, 22(1):64 102, 2011.

[Caminada and Wu, 2009] Martin W. A. Caminada and Yining Wu. An argument game of stable semantics. *Logic Journal of IGPL*, 17(1):77 90, 2009.

[Caminada and Wu, 2010] Martin W. A. Caminada and Yining Wu. On the justi cation status of arguments. In *Proc. of the 22nd Benelux Conference on Artificial Intelligence*, 2010.

[Caminada *et al.*, 2012] Martin W. A. Caminada, Walter A. Carnielli, and Paul E. Dunne. Semi-stable semantics. *J. of Logic and Computation*, 22(5):1207 1254, 2012.

[Caminada *et al.*, 2015] Martin W. A. Caminada, Samy Sa, Joao Alcantara, and Wolfgang Dvorak. On the equivalence between logic programming semantics and argumentation semantics. *Int. J. of Approximate Reasoning*, 58:87 111, 2015.

[Caminada *et al.*, 2016] Martin W. A. Caminada, Wolfgang Dvorak, and Srdjan Vesic. Preferred semantics as socratic discussion. *J. of Logic and Computation*, 26(4):1257 1292, 2016.

[Caminada, 2006a] Martin W. A. Caminada. On the issue of reinstatement in argumentation. In M. Fischer, W. van der Hoek, B. Konev, and A. Lisitsa, editors, *Proc. of the 10th European Conf. on Logics in Artificial Intelligence (JELIA 2006)*, number 4160 in Lecture Notes in Computer Science, pages 111 123. Springer, 2006.

[Caminada, 2006b] Martin W. A. Caminada. Semi-stable semantics. In P. E. Dunne and T.J.M. Bench-Capon, editors, *Proc. of the 1st Int. Conf. on Computational Models of Argument (COMMA 2006)*, pages 121 130. IOS Press, 2006.

[Caminada, 2007a] Martin W. A. Caminada. An algorithm for computing semi-stable semantics. In *Proc. of the 9th European Conference on Symbolic and Quantitalive Approaches to Reasoning with Uncertainty (ECSQARU 2007)*, number 4724 in Springer Lecture Notes in AI, pages 222 234, Berlin, 2007. Springer Verlag.

[Caminada, 2007b] Martin W. A. Caminada. Comparing two unique extension semantics for formal argumentation: ideal and eager. In M. M. Dastani and E. de Jong, editors, *Proc. of the 19th Belgian-Dutch Conference on Artificial Intelligence (BNAIC 2007)*, pages 81 87, 2007.

[Caminada, 2010] Martin W. A. Caminada. Preferred semantics as socratic discussion. In A. E. Gerevini and A. Saetti, editors, *Proceedings of the 11th AI*IA Symposium on Artificial Intelligence*, pages 209 216, 2010.

[Caminada, 2011] M. W. A. Caminada. A labelling approach for ideal and stage semantics. *Argument and Computation*, 2(1):1 21, 2011.

[Caminada, 2014] Martin W. A. Caminada. Strong admissibility revisited. In S. Parsons, N. Oren, C. Reed, and F. Cerutti, editors, *Proc. of the 5th Int. Conf. on Computational Models of Argument (COMMA 2014)*, pages 197 208. IOS Press, 2014.

[Caminada, 2015] Martin W. A. Caminada. A discussion game for grounded semantics. In E. Black, S. Modgil, and N. Oren, editors, *Proc. of the 3rd Int. Workshop on Theory and Applications of Formal Argumentation (TAFA 2015)*, volume 9524 of *Lecture Notes in Computer Science*, pages 59 73. Springer, 2015.

[Cayrol and Lagasquie-Schiex, 2013] Claudette Cayrol and Marie-Christine Lagasquie-Schiex. Bipolarity in argumentation graphs: Towards a better understanding. *Int. J. of Approximate Reasoning*, 54(7):876 899, 2013.

[Cayrol *et al.*, 2003] Claudette Cayrol, Sylvie Doutre, and Jerome Mengin. On decision problems related to the preferred semantics for argumentation frameworks. *J. of Logic and Computation*, 13(3):377 403, 2003.

[Cohen *et al.*, 2014] Andrea Cohen, Sebastian Gottifredi, Alejandro J. Garc a, and Guillermo R. Simari. A survey of di erent approaches to support in argumentation systems. *Knowledge Engineering Review*, 29(5):513 550, 2014.

[Dung *et al.*, 2007] Phan Minh Dung, Paolo Mancarella, and Francesca Toni. Computing ideal sceptical argumentation. *Artificial Intelligence*, 171(10-15):642 674, 2007.

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

[Dunne *et al.*, 2011] Paul E. Dunne, Anthony Hunter, Peter McBurney, Simon Parsons, and Michael Wooldridge. Weighted argument systems: Basic de nitions, algorithms, and complexity results. *Artificial Intelligence*, 175(2):457 486, 2011.

[Dunne *et al.*, 2013] Paul E. Dunne, Wolfgang Dvorak, and Stefan Woltran. Parametric properties of ideal semantics. *Artificial Intelligence*, 202:1 28, 2013.

[Dvorak and Gaggl, 2012a] Wolfgang Dvorak and Sarah A. Gaggl. Computational aspects of cf2 and stage2 argumentation semantics. In *Proc. of the 4th Int. Conf. on Computational Models of Argument (COMMA 2012)*, pages 273 284, 2012.

[Dvorak and Gaggl, 2012b] Wolfgang Dvorak and Sarah A. Gaggl. Incorporating stage semantics in the SCC-recursive schema for argumentation semantics. In *Proc. of the 14th Int. Workshop on Non-Monotonic Reasoning (NMR 2012)*, pages 273 284, 2012.

[Dvorak and Gaggl, 2016] Wolfgang Dvorak and Sarah A. Gaggl. Stage semantics and the SCC-recursive schema for argumentation semantics. *J. of Logic and Computation*, 26(4):1149 1202, 2016.

[Dvorak, 2012] Wolfgang Dvorak. On the complexity of computing the justi cation status of an argument. In S. Modgil, N. Oren, and F. Toni, editors, *Proc. of the 1st Int. Workshop on Theory and Applications of Formal Argumentation (TAFA 2011)*, number 7132 in Lecture Notes in Computer Science, pages 32 49. Springer, 2012.

[Gabbay, 2012] Dov M. Gabbay. Equational approach to argumentation networks. *Argument & Computation*, 3(2-3):87 142, 2012.

[Garc a and Simari, 2004] Alejandro J. Garc a and Guillermo R. Simari. Defeasible logic programming: an argumentative approach. *Theory and Practice of Logic Programming*, 4(1):95 138, 2004.

[Gelfond and Lifschitz, 1988] Michael Gelfond and Vladimir Lifschitz. The stable model semantics for logic programming. In R.A. Kowalski and K. Bowen, editors, *Proceedings of the 5th Int. Conference/Symposium on Logic Programming*, pages 1070 1080. MIT Press, 1988.

[Gelfond and Lifschitz, 1991] Michael Gelfond and Vladimir Lifschitz. Classical negation in logic programs and disjunctive databases. *New Generation Computing*, 9(3/4):365 385, 1991.

[Gorogiannis and Hunter, 2011] Nikos Gorogiannis and Anthony Hunter. Instantiating abstract argumentation with classical logic arguments: Postulates and properties. *Artificial Intelligence*, 175(9-10):1479 1497, 2011.

[Governatori *et al.*, 2004] Guido Governatori, Michael J. Maher, Grigoris Antoniou, and David Billington. Argumentation semantics for defeasible logic. *J. of Logic and Computation*, 14(5):675 702, 2004.

[Hunter, 2013] Anthony Hunter. A probabilistic approach to modelling uncertain logical arguments. *Int. J. of Approximate Reasoning*, 54(1):47 81, 2013.

[Hunter, 2014] Anthony Hunter. Probabilistic quali cation of attack in abstract argumentation. *Int. J. of Approximate Reasoning*, 55(2):607 638, 2014.

[Jakobovits and Vermeir, 1999] Hadassa Jakobovits and Dirk Vermeir. Robust semantics for argumentation frameworks. *J. of Logic and Computation*, 9(2):215 261, 1999.

[Kaci *et al.*, 2006] Souhila Kaci, Leendert W. N. van der Torre, and Emil Weydert. Acyclic argumentation: Attack = con ict + preference. In *Proc. of the 17th European Conference on Artificial Intelligence (ECAI 2006)*, pages 725 726, Riva del Garda, Italy, 2006. IOS Press.

[Modgil and Caminada, 2009] Sanjay Modgil and Martin W. A. Caminada. Proof theories and algorithms for abstract argumentation frameworks. In I. Rahwan and G.R. Simari, editors, *Argumentation in Artificial Intelligence*, pages 105 129. Springer, 2009.

[Modgil and Prakken, 2014] Sanjay Modgil and Henry Prakken. The ASPIC+ framework for structured argumentation: a tutorial. *Argument & Computation*, 5:31 62, 2014. Special Issue: Tutorials on Structured Argumentation.

[Modgil, 2009] Sanjay Modgil. Reasoning about preferences in argumentation frameworks. *Artificial Intelligence*, 173:901 1040, 2009.

[Nofal *et al.*, 2014] Samir Nofal, Katie Atkinson, and Paul E. Dunne. Algorithms for decision problems in argumentation systems under preferred semantics. *Artificial Intelligence*, 207:23 51, 2014.

[Pollock, 1992] John L. Pollock. How to reason defeasibly. *Artificial Intelligence*, 57(1):1 42, 1992.

[Pollock, 1995] John L. Pollock. *Cognitive Carpentry. A Blueprint for How to Build a Person.* MIT Press, Cambridge, MA, 1995.

[Pollock, 2001] John L. Pollock. Defeasible reasoning with variable degrees of justi cation. *Artificial Intelligence*, 133:233 282, 2001.

[Rahwan and Larson, 2008] Iyad Rahwan and Kate Larson. Pareto optimality in abstract argumentation. In D. Fox and C. P. Gomes, editors, *Proc. of 23rd AAAI Conference on Artificial Intelligence (AAAI 2008)*, pages 150 155. AAAI Press, 2008.

[Rahwan and Tohme, 2010] Iyad Rahwan and Fernando Tohme. Collective argument evaluation as judgement aggregation. In W. van der Hoek, G. A. Kaminka, Y. Lesperance, M. Luck, and S. Sen, editors, *Proc. of the 9th Int. Conf. on Autonomous Agents and Multiagent Systems (AAMAS 2010)*, pages 417 424, 2010.

[Reiter, 1980] Raymond Reiter. A logic for default reasoning. *Artificial Intelligence*, 13:81 132, 1980.

[Toni, 2014] Francesca Toni. A tutorial on assumption-based argumentation. *Argument & Computation*, 5:89 117, 2014. Special Issue: Tutorials on Structured Argumentation.

[Verheij, 1996] Bart Verheij. Two approaches to dialectical argumentation: admissible sets and argumentation stages. In J.-J.Ch. Meyer and L.C. van der Gaag, editors, *Proc. of the Eighth Dutch Conf. on Artificial Intelligence (NAIC'96)*, pages 357 368, 1996.

[Verheij, 2003] Bart Verheij. DEFLOG: on the logical interpretation of prima facie justi ed assumptions. *J. of Logic and Computation*, 13:319 346, 2003.

[Verheij, 2007] Bart Verheij. A labeling approach to the computation of credulous acceptance in argumentation. In M. M. Veloso, editor, *Proc. of the 20th Int. Joint Conference on Artificial Intelligence (IJCAI 2007)*, pages 623 628, 2007.

[Vreeswijk and Prakken, 2000] Gerard A. W. Vreeswijk and Henry Prakken. Credulous and sceptical argument games for preferred semantics. In *Proc. of the 7th European Workshop on Logic for Artificial Intelligence (JELIA-00)*, number 1919 in Lecture Notes in Computer Science, pages 239 253, Berlin, 2000. Springer Verlag.

[Vreeswijk, 1993] Gerard A. W. Vreeswijk. Studies in defeasible argumentation. *PhD thesis at Free University of Amsterdam*, 1993.

[Vreeswijk, 2006] Gerard A. W. Vreeswijk. An algorithm to compute minimally grounded and admissible defence sets in argument systems. In P. E. Dunne and T. J. M. Bench-Capon, editors, *Proc. of the 1st Int. Conf. on Computational Models of Argument (COMMA 2006)*, pages 109 120. IOS, 2006.

[Wu and Caminada, 2010] Yining Wu and Martin W. A. Caminada. A labelling-based justi cation status of arguments. *Studies in Logic*, 3(4):12 29, 2010.

[Wu *et al.*, 2009] Yining Wu, Martin W. A. Caminada, and Dov M. Gabbay. Complete extensions in argumentation coincide with 3-valued stable models in logic programming. *Studia Logica*, 93(1-2):383 403, 2009. Special issue: new ideas in argumentation theory.

Pietro Baroni
Dip. Ingegneria dell'Informazione
University of Brescia, Italy
Email: pietro.baroni@unibs.it

Martin Caminada
School of Computer Science & Informatics
Cardiff University, UK
Email: CaminadaM@cardiff.ac.uk

Massimiliano Giacomin
Dip. Ingegneria dell'Informazione
University of Brescia, Italy
Email: massimiliano.giacomin@unibs.it