Arguing with Doubt

An Approach to Structured Bipolar Argumentation

MSc Thesis *(Afstudeerscriptie)*

written by

Michael A. Müller

under the supervision of **Srdjan Vesic, Bruno Yun** and **Ulle Endriss**, and submitted to the Examinations Board in partial fulfillment of the requirements for the degree of

MSc in Logic

at the *Universiteit van Amsterdam.*

Date of the public defense:
11.10.2024

Members of the Thesis Committee:
Ulle Endriss
Davide Grossi
Aybüke Özgün
Sonja Smets
Srdjan Vesic
Bruno Yun

INSTITUTE FOR LOGIC, LANGUAGE AND COMPUTATION

Abstract

There is little agreement in formal argumentation on how to evaluate bipolar argumentation, where supporting arguments are considered in addition to attacking arguments. In contrast, informal argumentation has well established approaches to argumentation that deal with supports. In this thesis, we provide a new approach to bipolar argumentation that incorporates two features of informal argumentation: arguments have structure and one is allowed to doubt unattacked arguments. For that purpose, we define *structured bipolar argumentation frameworks* (SBAFs), which differ from bipolar argumentation frameworks by having structured arguments and from structured argumentation by having an explicit support relation. We provide two types of extension-based semantics for SBAFs: one that gives *argument extensions* and one that gives *language extensions*, where the latter are sets of sentences instead of arguments. We show that what we call *coherent* argument extensions and *adequate* language extensions correspond under certain assumptions. We further show how preferred semantics can be retrieved from *weakly* coherent argument extensions and how deductive support can be retrieved from *strongly* coherent argument extensions. We additionally provide a brief principled comparison of weakly and strongly coherent argument extensions and show that they are distinguished by the principle “directionality”. Finally, we indicate how a form of knowledge-based reasoning can be implemented in SBAFs by distinguishing between *doubtable* and *contestable* sentences, where the latter have to be accepted in absence of attacks.

Contents

Chapter 1

Introduction

Debates are ubiquitous in today's society. Be it in person amongst friends, in news media reporting about political topics, or online in social media, we are relentlessly exposed to controversial issues where multiple sides are trying to convince us of their viewpoints. As such debates can easily come to contain a vast number of arguments and claims, it can be exceedingly difficult to make sense of everything and arrive at a final verdict on what to believe. Perhaps we have some intuitions about a few arguments and claims that are made in a debate and end up accepting or rejecting them, but it is often unclear whether our position thus delineated is coherent. After all, by accepting some arguments and claims, it is easy to unwittingly incur inconsistent or indefensible commitments. What we would need, then, is a tool with which a user or agent could check whether their position in a debate is coherent or that could tell them how they would need to adapt their position in order to make it so.

A prominent approach that could lead to a tool for checking the coherence of a position in a debate has been introduced by Dung in form of *abstract argumentation* (Dung 1995). His approach models debates as directed graphs, where the nodes stand for arguments and the edges indicate attacks between them. Importantly, abstract argumentation forgets about any content arguments might have and does not consider any positive relation between them. While this is means that the model of a debate works with very limited information, we can nevertheless use it to evaluate a debate into rationally acceptable sets of arguments, called *extensions*. There might be some disagreement on what counts as a rationally acceptable set of arguments and, accordingly, various *semantics* have been proposed. Each semantics specifies which extensions map out the space of rationally acceptable positions in a debate and thus allows agents to check whether their positions can be seen as being coherent. Many different semantics for abstract argumentation have been proposed in the literature, but most of them share many basic assumptions, a crucial one being that unattacked arguments have to be accepted.

When it comes to the usefulness of abstract argumentation for evaluating actual debates, we should question whether forgetting all content of the arguments and only considering attacks between them really retains enough information. Indeed, if we look at relations between claims in a debate, i.e. propositional relations, then we find that attacks are the least common type of relation. Koszowy et al. analyse six debate corpora and distinguish three types of propositional relations: rephrase, inference, and conflict (Koszowy et al. 2022). Rephrase means that two claims have the same argumentative relations, but express their propositional content in different ways. Inference means that one claim supports an other and makes it more plausible, while two claims are in conflict if they cannot be true together. Interestingly, across all corpora, Koszowy et al. find that inference is the most frequent relation (62%), followed by rephrase (23%) and conflict coming in last (15%) (Koszowy et al. 2022, p. 64). Abstract argumentation can account for conflicts as they form the basis for attacks between arguments, and it seems reasonable to forget about rephrases when modelling a debate. However, this still leaves 62% of all propositional relations unaccounted for. Even though it can be argued that inferences occur mostly within arguments and are thus implicitly represented in abstract argumentation, empirical evidence suggests that this is not enough to capture the positive relations between arguments we find in practice (Polberg and Hunter 2018). These limitations of abstract argumentation have led to the development of *bipolar argumentation* (Cayrol and Lagasquie-Schiex 2005b), where a support relation is added to abstract argumentation, in order to explicitly take positive relations between arguments into account.

In contrast to purely attack based abstract argumentation, and as Yu et al. note, there is no commonly accepted semantics for bipolar argumentation that would tell us which sets of arguments are rationally acceptable once supports are taken into account (L. Yu et al. 2023). In that sense, the situation in the formal study of argumentation is markedly different compared to that of more philosophically and linguistically oriented approaches to argumentation (see e.g. Johnson 1996, van Eemeren and Grootendorst 2004). In the latter, informal approaches, inference and support play a much more central role. This difference is based on two features of informal argumentation: arguments are taken to be *structured*, that is, they are taken to consist of premises and a conclusion, and agents are allowed to *doubt* arguments and reject them even if they are not attacked. Both of these features have important implications for supports between arguments. Consider the following argument: "The certified meteorologist Alexa claims that it will rain next week, hence we can assume it will indeed rain next week." We can distinguish the two premises that (1) Alexa is a certified meteorologist and (2) she claims it will rain next week. This structure of the argument is important if, for instance, it is both attacked and supported. Without presupposing any specific semantics for bipolar argumentation, it is clear that the situation has to be evaluated very differently depending on whether, say, both attacks and supports concern premise (1) or whether (1) is attacked and (2) is supported. Doubt comes into play when there is no attack present. If unattacked arguments have to be accepted, as is the case in most abstract and bipolar semantics, there is no point in supporting, say, that Alexa really did claim that it will rain next week. In contrast, if an agent doubts whether Alexa made the claim, then it makes sense to cite some evidence and support that premise.

In this thesis, we import structured arguments and doubt into bipolar argumentation. We introduce the notion of structured bipolar argumentation frameworks that add structured arguments to bipolar frameworks with supports and attacks and provide a number of semantics for them. While there are approaches to structured argumentation in the literature (see Besnard et al. 2014), none of them make use of an explicit support relation and they employ a different notion of structured arguments as we understand it here. Further, our semantics come with two novelties: first, they allow agents to doubt arguments even if they are unattacked, thus implementing a notion of doubt, and, second, they evaluate argumentation both on the level of which arguments one should accept and on the level of which claims or statements one should accept. This duality of the argument and language levels gives two complementary perspectives that lets agents choose whether they want to think about whole arguments or just about individual claims in order to define their position in a debate.

In our approach, we focus on *extension-based* semantics, where a debate is evaluated into sets of rationally acceptable sets of arguments (see Baroni et al. 2018a for an overview). This contrasts with *gradual* semantics where each argument is given an acceptability degree, which is usually represented as a real number between 0 and 1 (see Baroni et al. 2019 for an overview). Our focus on extensions has two reasons: *First*, it makes it easier for our agents to interpret extensions and define their own positions, rather than them having to translate the real numbers given by gradual semantics into a position of accepted and unaccepted arguments. *Second*, extension-based semantics offer more conceptual clarity, as gradual approaches often suffer from technical difficulties in how the acceptability degrees are calculated.

The remainder of this chapter introduces all the necessary background in formal and informal argumentation theory as well a first account of how we can think of structured bipolar argumentation. Chapter 2 then examines existing semantics for bipolar and structured argumentation and details where they fail to take structured arguments or doubt into account. Our own proposal for semantics for structured bipolar argumentation is presented in Chapter 3. We introduce strongly and weakly coherent argument extensions as well as strongly and weakly adequate language extensions and show that the argument and language perspectives correspond. We also describe how support behaves according to our semantics and how they compare to standard semantics in abstract and structured argumentation. Finally, Chapter 4 concludes this thesis.

1.1 Argumentation Theory

Argumentation theory in its modern form as an academic field can be traced back to Toulmin's book *The Uses of Argument* (updated edition 2003) as well as Perelman and Olbrechts-Tyteca's *The New Rhetoric* (1969), both originally published in 1958. Each in their own way, the two books went against the then dominant logical view of argumentation (see van Eemeren et al. 2020, Chapter 3). There, assessment of an argument consists of two steps: First, it is checked whether the premises of an argument logically entail its conclusion, i.e. whether it is *valid*, and, second, it is checked whether the premises are actually true, i.e. whether the argument is also *sound*. In contrast, Toulmin focused on a richer account of argument structure, identifying several different components that culminated in the *Toulmin Scheme* (Figure 1.1). Perelman and Olbrechts-Tyteca took on a rhetorical perspective and put the audience, to which an argument should be convincing, in the centre. From these beginnings, various approaches have developed, the relevant ones here being *pragma-dialectics* and *informal logic*.

Pragma-dialectics has been developed by van Eemeren and Grootendorst (1984; 1992; 2004) in an effort to provide a comprehensive framework for studying argumentation. Its fundamental notion is that of a *difference of opinion* between, paradigmatically, two

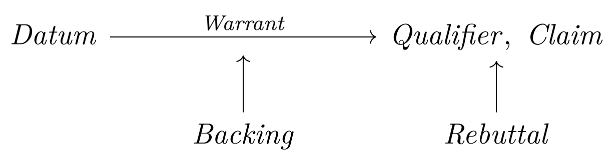

Figure 1.1: The Toulmin scheme, where a qualified claim is inferred from some datum via a warrant. Backing supports the warrant and the rebuttal lists exceptions for when the claim does not follow. The traditional example is the following: Because Harry was born in Bermuda (datum), he is presumably (qualifier) a British subject (claim), since a person born in Bermuda will generally be a British subject (warrant). This is because of legal provisions (backing) and holds unless, for instance, both parents are legal aliens (rebuttal) (Toulmin 2003, p. 92).

agents. Argumentation is then employed as a reasonable means of resolving this difference of opinion. What ensues is called a *critical discussion* and pragma-dialectics sets out the rules which such an ideal discussion should follow.

In brief, a critical discussion consists of four stages (van Eemeren and Grootendorst 2004, pp. 59-62): In the *confrontation stage*, the difference of opinion is established. Van Eemeren and Grootendorst distinguish between mixed and non-mixed differences of opinion. In a non-mixed difference, there is only one standpoint at play that is doubted by a participant. A mixed difference of opinion occurs if both a standpoint and its negation are in play.

After the confrontation, an *opening stage* follows. First, the participants determine their respective roles in the critical discussion. In a non-mixed difference, there is a *protagonist* that puts forth a *standpoint* (usually a statement they claim to be true) and an *antagonist* that has doubts about the standpoint of the protagonist. In a mixed difference of opinion, the antagonist can have a standpoint of their own, for instance if they not only doubt the standpoint of the protagonist but also assert its negation. In that sense, there is a significant difference between merely doubting a standpoint and claiming that a standpoint is false. There can also be more than one standpoint involved, which would lead to a "multiple" difference of opinion, but we will focus on the simplest case of a single non-mixed difference of opinion. The protagonist and antagonist also use the opening stage to determine which common ground they share. This can include a set of statements that both agents accept and an agreement on which types of inferences are allowed (Krabbe 2017).

The third and most extensive stage is the *argumentation stage*. Since the antagonist expressed doubt about the protagonist's standpoint in the confrontation stage, the protagonist has to produce arguments in its favour. The antagonist assesses the arguments and checks whether they are acceptable or require more justification.

Finally, a critical discussion ends with a *concluding stage*. The two agents establish the result of the discussion. The difference of opinion counts as resolved if both agents agree either that the protagonist's standpoint should be accepted or that the antagonist's doubt wins and the standpoint should be retracted. The four stages are summmarised in Figure 1.2.

Confrontation Stage	A difference of opinion is established.
Opening Stage	The roles of protagonist and antagonist as well as a common ground are determined.
Argumentation Stage	The protagonist defends their standpoint against the doubts of the antagonist.
Concluding Stage	Either the protagonist retracts their standpoint or the antagonist accepts it.

Figure 1.2: The four stages of a critical discussion with a single non-mixed difference of opinion.

The whole critical discussion is oriented around speech acts (Austin 1962). Pragma-dialectics sets our rules for which speech acts each agent is allowed to perform at each stage of the discussion and what consequences these have. For instance, a crucial rule for the argumentation stage reads as follows (van Eemeren and Grootendorst 2004, p. 144):

> Rule 6
>
> a. The protagonist may always defend the standpoint that he adopts in the initial difference of opinion or in a sub-difference of opinion by performing

a complex speech act of argumentation, which then counts as a provisional defense of this standpoint.

b. The antagonist may always attack a standpoint by calling into question the propositional content or the justificatory or refutatory force of the argumentation.

c. The protagonist and the antagonist may not defend or attack standpoints in any other way.

Pragma-dialectics is mainly interested in argumentation as a dialectical activity between two agents. In recent years, its model has also been extended to integrate rhetorical aspects of argumentation such as strategic manoeuvring within the allowed space of speech acts (van Eemeren 2010) and distinguishing different argumentative styles (van Eemeren et al. 2022). Further, there have been different formal dialogue frameworks modelled after pragma-dialectics, most notably by Krabbe (2017) and Visser (2013; 2015; 2017). However, this school of argumentation theory does not focus much on individual argumentative moves and their evaluation.

In contrast, the school of *informal logic* makes the evaluation of individual arguments its central goal. Originated by Johnson and Blair in 1977 (2006), it focused from the beginning on critical thinking and developing methods for dealing with argumentation in practice, such as as found in politics and newspapers (Johnson 1996). For our purposes, we are mostly interested in the development of argument schemes that can be used to categorise and evaluate arguments (Walton et al. 2008).

An argument scheme takes a typical pattern of reasoning found in practice and gives a general formulation of it. In doing so, schemes abstract away from some of the content of full arguments, but they go nowhere near the level of a logical reconstruction. A common example for an argument schemes is that from expert opinion (Walton and Koszowy 2017):

Argument from Expert Opinion

(1)	Source E is an expert in subject domain S containing proposition A.
(2)	E asserts that proposition A is true (false).
(3)	A is true (false).

This scheme gives the general form of arguments relying on someone's expertise to justify some claim. It consists of a first premise to establish the expertise of the cited source with respect to the proposition in question and a second premise stating that the proposition is indeed asserted (resp. rejected). It then concludes that the proposition is true (resp. false). Arguably, both premises are required to justify the conclusion. Without the source being an expert, we have little reason to trust their claim and without them saying anything about the proposition in question, we have no reason to conclude anything about it. But are the two premises themselves already enough to justify their conclusion?

Argument schemes not only allow us to identify arguments in practice and reconstruct them in a principled way, they also come with a list of critical questions that help us evaluate arguments that fit the schemes (S. Yu and Zenker 2020). The critical questions for the argument from expert opinion are the following (Walton and Koszowy 2017):

Expertise Question: How credible is E as an expert source?

Field Question: Is E an expert in the [subject domain S] that A is in?

Opinion Question: What did E assert that implies A?

Trustworthiness Question: Is E personally reliable as a source?

Consistency Question: Is A consistent with what other experts assert?

Backup Evidence Question: Is E's assertion based on evidence?

These questions point to potential problems with an argument from expert opinion and should any of these questions be answered negatively, the argument fails. Some of these questions point to potential problems with the premises (e.g. the field and opinion questions), others point to problems with the inference from an expert's claim to the truth of the claim (e.g. the trustworthiness and backup evidence questions), and finally some questions point to potential arguments for the opposite conclusion (e.g. the consistency question).

Importantly, the inferences in argument schemes are assumed to be generally warranted. If there is no reason to assume that some exceptional situation occurs, we are able to infer the truth or falsity of a claim from the assertion of an expert. That is, while most critical questions need to be answered once asked, questions that point to exceptions for the inference are only relevant if they themselves can be backed up with an argument (Walton et al. 2008, p. 386).

In terms of structured arguments and doubt, there are two things to note here. *First*, we can identify three types of components that constitute and argument: premises, an inference step, and a conclusion. Toulmin's scheme is more fine-grained in also recognising backing, qualifiers and rebuttals. However, the three part structure of arguments is the most widely shared in the literature (see also Brun and Betz 2016; Rigotti and Greco 2019 and with some caveats Johnson 2000). *Second*, both in pragma-dialectics and in argument schemes, whoever is confronted with an argument is allowed to have doubt about the components of the argument. In pragma-dialectics, this is reflected in Rule 6b that allows the antagonist to always express doubt which forces the protagonist to react. With argumentation schemes, it is the critical questions that allow expression of doubt. When it comes to questions concerning the premises of an argument, then uttering them already forces the arguer to react. In contrast, questions about the inference need to be substantiated. This is somewhat mirrored in pragma-dialectics where the inferences of arguments are supposed to be part of the common ground (Krabbe 2017). Thus we can speak of two types of doubt: *mere* doubt and *reasoned* doubt. The latter requires justification through an argument (usually doubt about the inference of an argument), whereas the former is effective simply by being expressed. Notably, allowing for mere doubt makes it non-trivial to evaluate argumentation where no attacking arguments are present. In particular, a non-mixed difference of opinion in pragma-dialectics proceeds with supports only. This is in stark contrast to formal argumentation. Of course, the latter recognises reasoned doubt in the sense that attacked arguments do not have to be accepted, but it lacks a notion of mere doubt.

1.2 Abstract and Bipolar Argumentation

Abstract argumentation takes a very different approach to studying arguments than the schools of the previous section. Initiated by Dung (1995), all internal content of arguments is abstracted away and only the relations between arguments are considered. More specifically, only attacks between arguments are recognised. As no content of arguments is considered, no one specific meaning is attached to the notion of attack, except that is expresses some general notion of incompatibility between arguments. Rather, the arguments and their attacks are assumed to be given. Abstract argumentation then uses this information to build an argument graph and aims to evaluate which sets of arguments are acceptable.

Accordingly, the basic notion is that of an *abstract argumentation framework*, which specifies which arguments there are and which attacks are present. Figure 1.3 gives an example. As we are interested in modelling debates, we can restrict our attention to finite frameworks.

Definition 1 (Abstract Argumentation Framework)**.** *An abstract argumentation framework (AF) is a tuple* $\mathcal{A} = \langle A, \rightarrow \rangle$ *where* A *is a finite set of arguments and* $\rightarrow \subseteq A \times A$ *an attack relation.*

We write $a \rightarrow b$ in case $(a, b) \in \rightarrow$ and generalise to sets of arguments, i.e. $E \rightarrow b$ in case $a \rightarrow b$ for some $a \in E$ and $a \rightarrow E$ in case $a \rightarrow b$ for some $b \in E$. We say that a set of arguments $E \subseteq A$ *defends* an argument $a \in A$ if $\forall b \in A : b \rightarrow a \implies E \rightarrow b$.

The next step is to define *semantics* that tell us for each AF which sets of arguments are acceptable. We can think of them as follows. If an agent is confronted with some arguments, they can go through each argument and decide whether they accept it. Given such a set of accepted arguments, we can then ask whether the agent is rational in accepting these arguments. Of course, without knowing the content of the arguments, we cannot say much about the rationality of accepting any individual argument (except, perhaps, if it is self-attacking). However, we can exclude some sets of arguments from being rational by considering the attacks between arguments. We will call a rationally acceptable set of arguments an *extension*. We will not explicitly model agents and their beliefs (see e.g. Sakama 2024; Sakama and Son 2020; Shi et al. 2018 for explicit representations of agents in abstract argumentation). Rather, we use agents informally as a background consideration and interpret extensions as possible positions an agent could have in a debate. Different semantics then implement different notions of what is a rational position, though almost all semantics share some basic assumptions. For instance, it seems uncontroversial enough to claim that arguments a_1 and a_2 in Figure 1.3 cannot rationally be accepted together, as they attack each other. This is encoded in the notion that extensions should be *conflict-free*, meaning that no accepted argument should attack another accepted argument.

Next to conflict-freeness, the other fundamental notion employed by Dung in defining his semantics is that of *defence* (Dung 1995). The basic idea is that if you want to accept an argument that is being attacked, then you also need to accept an argument that attacks the attacker. If there is a counter to an argument you accept, then you cannot simply ignore it, you need to react. In our example, this means that you should not accept only argument a_4, as it is being attacked by a_3. Rather, you should also react to that attack by accepting either a_1 or a_2. One way to encode this idea is by requiring that any acceptable extension E must be such that if there is an argument a such that $a \rightarrow E$, we also have $E \rightarrow a$. However, it is useful to go through the *characteristic function* of an AF, which calculates for each set of arguments the set of arguments that it defends (Dung 1995).

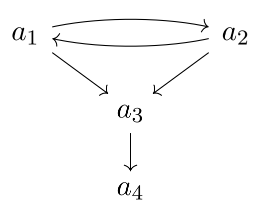

Figure 1.3: An example of an abstract argumentation framework.

Definition 2 (Characteristic Function)**.** *Let $\mathcal{A} = \langle A, \rightarrow \rangle$ be an AF. We define its characteristic function $F_{\mathcal{A}} : 2^A \rightarrow 2^A$ as*

$$F_{\mathcal{A}}(E) := \{a \in A \mid E \textit{ defends } a\}.$$

For instance, in Figure 1.3, we have $F_{\mathcal{A}}(\{a_1\}) = \{a_1, a_4\}$, as a_1 both defends itself from the attack by a_2 and a_4 from the attack by a_3. We can now formulate the defence requirement as $E \subseteq F_{\mathcal{A}}(E)$, meaning that all arguments in the set are defended.

If an extension is both conflict-free and defended, then we call it *admissible*, as it passes the most basic requirements for a rationally acceptable set of arguments. For instance, in Figure 1.3, $\{a_1\}$ is admissible, as is $\{a_1, a_4\}$, and also $\emptyset$. However, most semantics for abstract argumentation go a step further, requiring that a rational extension is not only admissible, but also contains all arguments it defends. This is called *completeness* and can be expressed by saying that an extension should be a fixpoint of the characteristic function. In our example, there are three complete extensions: $\emptyset, \{a_1, a_4\}, \{a_2, a_4\}$. Note that $\emptyset$ is complete, because it does not defend any argument, and, say, $\{a_1\}$ would not be complete because it defends a_4 without containing it. Since there can be multiple complete extensions, we can also talk about minimal and maximal (w.r.t. set-inclusion) complete extension, called *grounded* (in this case $\emptyset$) and *preferred* (in this case $\{a_1, a_4\}, \{a_2, a_4\}$), respectively. All of this is summarised in the following definition.

Definition 3 (Dung Semantics)**.** *Let $\mathcal{A} = \langle A, \rightarrow \rangle$ be an AF. An extension $E \subseteq A$ is called*

Conflict-Free: if $\forall a, b \in E : a \not\rightarrow b$,

Defended: if $E \subseteq F_{\mathcal{A}}(E)$,

Admissible: if it is conflict-free and defended,

Complete: if it is conflict-free and $E = F_{\mathcal{A}}(E)$,

Grounded: if it is $\subseteq$-minimal among complete extensions,

Preferred: if it is $\subseteq$-maximal among complete extensions.

Dung also defines an extension E to be *stable* if it is admissible and for each $a \in A \backslash E$, we have $E \rightarrow a$. However, this type of extension will not play much of a role in this thesis.

Now we examine some properties of these semantics. All the results listed in this subsection stem from Dung's original paper (1995), though the presentation also follows Baroni et al. (2018) and Grossi and Modgil (2019). Starting with the existence of extensions, we can first note that $\emptyset$ is always admissible. For complete extension, we require a few lemmas. The characteristic function is monotonic:

Lemma 1. *Let $\mathcal{A} = \langle A, \rightarrow \rangle$ be an AF and $F_{\mathcal{A}}$ its characteristic function. Then for any $E \subseteq E' \subseteq A$, we have $F_{\mathcal{A}}(E) \subseteq F_{\mathcal{A}}(E')$.*

We also observe that for admissible extensions, the characteristic function is increasing:

Lemma 2. *Let $\mathcal{A} = \langle A, \rightarrow \rangle$ be an AF and $F_{\mathcal{A}}$ its characteristic function. Then for any admissible extension $E \subseteq A$, we have $E \subseteq F_{\mathcal{A}}(E)$ and $F_{\mathcal{A}}(E)$ is also conflict-free.*

Since we only consider finite AFs, these two lemmas together give us our first existence result for complete semantics:

Proposition 1. *Let $\mathcal{A} = \langle A, \rightarrow \rangle$ be an AF. Then there exists a complete extension.*

Corollary 1. *Let $\mathcal{A} = \langle A, \rightarrow \rangle$ be an AF. Then there exists at least one grounded and at least one preferred extension.*

We can further note that preferred extensions are not only maximal among complete extensions, but also maximal among admissible extensions.

Proposition 2. *Let $\mathcal{A} = \langle A, \rightarrow \rangle$ be an AF. Then $E \subseteq A$ is a preferred extensions if and only if E is $\subseteq$-maximal amongst admissible extensions.*

This proposition gives us that every admissible extension is a subset of a preferred and thus also of a complete extension. Further, it can be shown that there exists a *unique* grounded extension for every AF, but the result requires some set-up. We use Davey and Priestly (2002) as a reference for the following definitions and fact. A partial order $\leq$ on a set L is a binary relation that is reflexive ($\forall x, y, z \in L$, we have $x \leq x$), anti-symmetric (if $x \leq y$ and $y \leq x$, then $x = x$), and transitive (if $x \leq y$ and $y \leq z$, then $x \leq z$). A greatest lower bound (resp. least upper bound) of a subset $S \subseteq L$ is the greatest (resp. least) element $x \in L$ such that $x \leq y$ (resp. $y \leq x$) for all $y \in S$.

Definition 4 (Complete Lattice)**.** *A complete lattice L is a non-empty, partially-ordered set, such that for any $S \subseteq L$, there exists a greatest lower bound $\bigwedge S$ and a least upper bound $\bigvee S$.*

Fact 1 (Knaster-Tarski Fixpoint Theorem)**.** *Let L be a complete Lattice and $F : L \rightarrow L$ a monotonic function. Then there exists a least fixpoint of F.*

Lemma 3. *Let $\mathcal{A} = \langle A, \rightarrow \rangle$ be an AF and E any admissible extension. Then $\{E' \subseteq A \mid E \subseteq E'\}$ is a complete lattice.*

It then follows from Lemmas, 1 and 3 and the Knaster-Tarski fixpoint theorem that the grounded extension is unique. To see that the least fixpoint is indeed conflict-free, we can note that it is contained in all fixpoints, hence we can infer conflict-freeness from Proposition 1.

Proposition 3. *Let $\mathcal{A} = \langle A, \rightarrow \rangle$ be an AF. Then there exists a unique grounded extension.*

We can also iteratively construct the grounded extension from the empty extension:

Proposition 4. *Let $\mathcal{A} = \langle A, \rightarrow \rangle$ be an AF and E its grounded extension.*
Then $E = \bigcup_{i \in \mathbb{N}} F^i_{\mathcal{A}}(\emptyset)$, where $F^0_{\mathcal{A}}(S) = S$ and $F^{i+1}_{\mathcal{A}}(S) = F_{\mathcal{A}}(F^i_{\mathcal{A}}(S))$.

We end by mentioning a result that shows how all completeness-based semantics coincide in a large number of AFs. We say an AF is *well-founded* if there exists no infinite path of attacking arguments $a_1 \rightarrow \cdots \rightarrow a_n \rightarrow \ldots$. Since we restrict our attention to finite AFs, this means that an AF does not contain any cycle.

Proposition 5. *Let $\mathcal{A} = \langle A, \rightarrow \rangle$ be a well-founded AF. Then there exists a unique complete extension.*

Recall that our main interest lies in also examining supports between arguments, rather than only attacks. For that purpose, we can add a support relation to AFs, leading to the following definition.

Definition 5 (Bipolar Argumentation Framework)**.** *A bipolar argumentation framework (BAF) is a triple $\mathcal{B} = \langle A, \rightarrow, \dashrightarrow \rangle$ where A is a finite set of arguments, $\rightarrow \subseteq A \times A$ an attack relation, and $\dashrightarrow \subseteq A \times A$ a support relation.*

Definitions of $a \dashrightarrow b, E \dashrightarrow b$, and $a \dashrightarrow E$ are as with attacks. Sometimes, it is assumed that the support and attack relations are disjoint (e.g. Oren and Norman 2008), but for now we leave open the possibility that an argument can at the same time support and attack another argument. We will discuss various semantics that have been proposed for BAFs in Chapter 2, but first we introduce some structure to the arguments.

1.3 Structured Bipolar Argumentation

Our goal is to provide a framework for bipolar argumentation that follows the insights of Section 1.1. Accordingly, it is useful to introduce here how we formally understand structured argumentation. Other frameworks such as ASPIC$^+$(to be discussed in detail in Section 2.2.2) approach structured argumentation as a form of *knowledge-based* reasoning. That is, arguments are constructed from a knowledge base by means of inference rules. Our approach is different in that it is somewhat closer to abstract argumentation. We assume the arguments to be given, for instance by a debate where various arguments are put forth. As in abstract argumentation, we simply start with a set of arguments, but we also add some content to the nodes of an AF.

As mentioned in Section 1.1, we assume arguments to consist of three types of components: premises, a conclusion, and an inference step from the premises to the conclusion. In order to represent these components, we need a *language*. The basic notion we will use is that of *sentences*. Sentences of our language represent the statements made in a debate and they can be used to express premises, conclusions, and inference steps. Arguments then consist of sentences and the relations between arguments will be based on relations between their sentences. What we need for that is an indication for which sentences are *incompatible* with each other. Take two sentences such as s: "This wall is red" and t: "This wall is blue". It cannot be the case that both of them are true at the same time, thus they are incompatible with each other. This notion will then underlie our definition of attacks. For instance, two arguments, one concluding with s and one with t would certainly attack each other.

Definition 6 (Language). *A language $\mathcal{L} = \langle L, \overline{}, n\rangle$ consists of a set of sentences L, an incompatibility function $\overline{} : L \to 2^L$, and a naming function $n : 2^L \times L \to L$.*

The naming function is how we can have sentences that express the inference steps of arguments. This will be useful to define undercutting attacks later on. In practice and in reconstructions using argument schemes, it is often only the premises and the conclusion of an argument that are made explicit. That there is an argument present can, for instance, be indicated through linguistic markers such as "hence" or "it follows that" (Johnson and Blair 2006, p. 13). We then infer that there is an implied inference step from the premises to the conclusion. For instance, when someone says "The wall looks red, hence it is red", there is no need to state the inference step explicitly. Adding something like "That the wall looks red implies that it is red" is somewhat redundant—it adds no new information that what was already present in the previous sentence. Nevertheless, on the formal level it is useful to talk about the inference claim and this is done using the naming function. For each argument, it specifies a sentence that represents its implicit inference claim.

An example for a language could be propositional logic with a defeasible conditional $\rightsquigarrow$. L would then be the set of well-formed formulas, incompatibility could be defined as $\overline{\varphi} := \{\psi \in L \mid \varphi \land \psi \models \bot\}$, and the naming function as $n(\langle\{\varphi_1, \dots, \varphi_n\}, \psi\rangle) := (\varphi_1 \land \dots \land \varphi_n) \rightsquigarrow \psi$. In the following, we will not rely on any specific language. The only assumption we will make is that our language contains a negation $\neg$, such that for any sentence s, we have $\neg s \in \overline{s}$. This is purely for ease of notation so that we have a canonical incompatible sentence to refer to.

Note that there is no claim towards the truth of the implicit inference claim of an argument. That is, we do not assume that arguments are *valid*. Rather, any combination of premises with a conclusion will count as an argument. Keeping with the propositional language, $\langle\{p\}, q\rangle$ would as much be an argument as $\langle\{p, q\}, p\rangle$ or $\langle\{p, p \rightsquigarrow q\}, q\rangle$. In the last case, it is useful to distinguish defeasible conditionals in general from implicit inference claims, as that of the latter argument would be $(p \land (p \rightsquigarrow q)) \rightsquigarrow q$. A further consequence of the lack of standards for inferences is that we assume all premises to be relevant for the argument. We simply assume that given an argument, the claim is that its premises justify the conclusion. If it contains irrelevant premises, then it might be a bad argument, but it is an argument nonetheless.

Definition 7 (Arguments)**.** *An argument in a language* $\mathcal{L} = \langle L, \overline{}, n \rangle$ *is a tuple* $a = \langle Prem(a), Conc(a) \rangle$ *where* $Prem(a) \subseteq L$ *(non-empty) and* $Conc(a) \in L$*. We also define* $Sent(a) := Prem(a) \cup \{Conc(a)\}$*.*

This view of arguments differs somewhat from that of other approaches to structured argumentation such as ASPIC$^+$. While arguments here are assumed to be built using only one inference step, ASPIC$^+$arguments can use an arbitrary number of inference rules to connect the premises to the conclusion. Every single inference step is then said to create a sub-argument of the original argument. This idea of sub-arguments is absent in our definition. Rather, in ASPIC$^+$'s terms, each inference rule creates its own argument. Nevertheless, our representation of arguments is closer to the view found in informal argumentation theory.

With this structure of arguments in mind, we can now think about how they can be related to each other. It seems clear that we should expect the relations to be based on the components of the arguments. Moreover, we assume that the relations are based on the conclusions of the arguments. That is, whether an argument supports or attacks another argument depends on how the conclusion of the former relates to the components of the latter. This is a fairly standard assumption, though it is not always made explicit (e.g. Betz 2010; Cohen et al. 2018; Modgil and Prakken 2014).

Definition 8 (Support and Attack)**.** *Let* a, b *be arguments in some language* $\mathcal{L}$*.*
We say that a *supports* b *if* $Conc(a) \in Prem(b)$*.*
We say that a *attacks* b *if* $Conc(a) \in \overline{s}$ *for some* $s \in Sent(b)$ *or if* $Conc(a) \in \overline{n(b)}$*.*

An argument a attacks an argument b if it is incompatible with any of its components. Since there are three types of components: premises, conclusion, and implicit inference claim, we can distinguish three types of attacks: undermines (contradicting a premise), rebuts (contradicting the conclusion), and undercuts (contradicting the implicit inference claim, represented using the naming function). This distinction can also be found in ASPIC$^+$(Modgil and Prakken 2014) and goes back to Pollock (1987).

According to our definition, supports are more limited than attacks. Namely, only confirming a premise counts as a support, while there is no notion of inference-support or conclusion-support. The absence of inference-support is a general phenomenon in the literature. ASPIC$^+$does not have it, Krabbe's formalisation of pragma-dialectics does not allow supporting an inference step through argument (Krabbe 2017), and even Cohen et al.'s survey of different types of support in structured argumentation does not consider it (Cohen et al. 2018). One notable exception of this is Toulmin's argument scheme, which explicitly includes the option of backing (supporting) a warrant (an inference) (Toulmin 2003). Nevertheless, we stay with the majority here and do not allow supports on inference claims. Conclusion-support can be found in Cohen et al. (2018). However, it is hard to see how two arguments with the same conclusion support each other. While they both support the same sentence and thus make their conclusion more plausible, the strength of an argument should not depend on the plausibility of its conclusion. Arguing that you will win the lottery because you saw it in your dreams is a bad argument, even if there is an independent argument that makes it very plausible that you will win the lottery (e.g. if the game is rigged). In sum, we only recognise premise-support as defined here.

It is important at this point not to confuse conditional sentences that occurs as premises with the implicit inference claims of arguments. The implicit inference claim is always

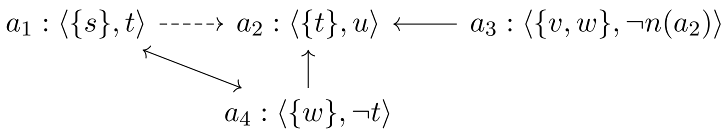

Figure 1.4: A first example of an SBAF.

that all premises together justify the conclusion, whereas the premises might well contain a conditional, for instance in a *modus ponens*. It is useful to distinguish between what we might informally think of general inference rules and the specific applications of these rules. An example for a general inference rule could be "If an expert claims something in their field of expertise, then we can assume their claim to be true". In the example from the introduction, this would be "If a certified meteorologist claims something about the weather, then we can assume their claim to be true". We can easily add that rule to the premise that the certified meteorologist Alexa claims that it will rain next week, to get a somewhat more complete argument in favour of the claim that it will rain next week. But note that the implicit inference claim is very different from the general inference rule. Namely, the implicit claim is that if the certified meteorologist Alexa claims that it will rain next week, then we can assume that it will rain next week. That is, it is the general inference rule applied to the specific case of Alexa and her claim about next week.

It is possible to support a general inference rule, since it will appear as a premise of the argument, but it is not possible to support the implicit inference claim. While one could take an argument a and add its inference claim $n(a)$ to its premises, this would create a new argument a' with a new implicit inference claim $n(a')$. It would also be somewhat redundant. Since the inference claim simply takes what is already implicit in the argument, it would not add anything to it. With these definitions in place, we can define a structured bipolar argumentation framework:

Definition 9 (Structured Bipolar Argumentation Framework)**.** *A structured bipolar argumentation framework (SBAF) is a tuple $\mathcal{SB} = \langle \mathcal{L}, A, \rightarrow, \dashrightarrow \rangle$ where A is a finite set of arguments in language $\mathcal{L}$ (Definition 7) and $\rightarrow, \dashrightarrow$ are the corresponding attack and support relations (Definition 8).*

As we only deal with finite SBAFs, it is worth noting that an SBAF is guaranteed to be finite if it is based on a finite language. Namely, the number of arguments in an SBAF has an upper bound of $(2^{|L|} - 1)|L|$, since any non-empty subset of sentences can be used as the premises and any sentence can be a conclusion.

Figure 1.4 presents a first example of an SBAF. a_1 supports a_2, a_3 undercuts a_2, and a_4 both rebuts a_1 and undermines a_2. Note that the rebut between a_4 and a_1 goes in both directions. This is a general feature of rebuts, as later on we will assume the incompatibility function to be symmetric.

The closest analogue of SBAFs in the literature can be found in logic-based argumentation (see e.g. Arieli et al. 2021; Besnard and Hunter 2001) and the statement graphs of Hecham et al. (2018; 2020). While arguments in SBAFs (as simple premise-conclusion structures) have very similar analogues in these approaches, there are two main differences: both logic-based argumentation and statement-graphs represent knowledge-based reasoning where the arguments are explicitly constructed, whereas SBAFs take arguments as given. Further, our notion of undercuts is not in the same way present in either of the other approaches.

Chapter 2

An Opinionated Survey on Bipolar Argumentation

In this chapter, we review a range of approaches to bipolar argumentation and evaluate them from the perspective of SBAFs and doubt. Early arguments in favour of adding a support relation to abstract argumentation go back to Amgoud, Cayrol, and Lagasquie-Schiex (Amgoud et al. 2008; Cayrol and Lagasquie-Schiex 2005b, 2009). From the beginning, bipolar argumentation has been approached from a wide variety of perspectives. Sometimes it is seen as a generalisation of Dung Semantics (Cayrol and Lagasquie-Schiex 2005b; Potyka 2021), other times it is used as a form of logic programming (Nouioua and Risch 2011), and in some cases it is taken as a form of legal reasoning (L. Yu et al. 2023). All these approaches are *extension-based* in that the bipolar semantics they propose define sets of acceptable arguments. During the same time, there has been a parallel development of *gradual* semantics in bipolar argumentation (e.g. Amgoud and Ben-Naim 2016, 2018; Amgoud et al. 2008; Baroni et al. 2015; Cayrol and Lagasquie-Schiex 2005a; Potyka 2018; Rago et al. 2016). There, the idea is that arguments should not only be evaluated into accepted and not accepted, as extensions do, but they should rather receive an acceptability degree, most often a value between 0 and 1. The acceptability degree of an argument then depends on the acceptability of its attackers and supporters.

All of the approaches mentioned above treat arguments as abstract entities and simply provide semantics for BAFs. There is an independent family of approaches for structured argumentation (Besnard et al. 2014). This tradition is most prominently manifested in ASPIC^+(Modgil and Prakken 2013, 2014) and assumption-based argumentation (ABA) (Toni 2014), so we will mostly discuss those.

When we survey all these approaches, we are not interested in whether they are adequate for their originally intended purposes. Rather, we investigate how they fare from the perspective of informal argumentation theory, namely from the perspective of structured arguments and mere doubt. Of course, as $\text{ASPIC}^+$and ABA are already structured, we will focus on doubt when discussing them.

2.1 Structure in Bipolar Argumentation

We have introduced SBAFs in Section 1.3, with the main departure from BAFs being that arguments now have premises and a conclusion. We argue that for capturing support between arguments, recognising their structure is crucial. In essence, we claim that the effect of supports and their interaction with attacks take on a different form, depending on how they relate to the components of the arguments. With abstract arguments as we have them in BAFs, we cannot distinguish between different cases and have to treat

$$a_1 \longrightarrow a_2$$

(a) A simple abstract attack.

$$a_1 : \langle \{s\}, \neg t \rangle \longrightarrow a_2 : \langle \{t, u\}, v \rangle$$

(b) A structured instance of an attack.

Figure 2.1: Comparing abstract and structured attacks.

all supports and interactions with attacks the same. Before seeing this problem at play in standard bipolar semantics, it is useful to consider why structured arguments are less important if we only have attacks to consider.

Take the following two arguments:

a_1:

s :	Alexa never finished her degree in meteorology.
$\neg t$:	Alexa is not a certified meteorologist.

a_2:

t :	Alexa is a certified meteorologist.
u :	Alexa claims that it will rain next week.
v :	It will rain next week.

Clearly, a_1 attacks a_2, since its conclusion, that Alexa is not a certified meteorologist, is incompatible with the premise of a_2, that Alexa is a certified meteorologist. The situation is depicted in Figure 2.1 in both abstract and structured terms. The effect of the attack is clear: as long as there is no response to a_1, a_2 cannot be accepted. Note that this effect does not depend the specifics of the attack. In order for a_2 to be acceptable, Alexa both needs to be a certified meteorologist and must have claimed that it will rain next week. Of course, we should also not have reason to doubt the inference of the argument. If it turns out that Alexa is not a certified meteorologist, we have not much reason to believe her claim about next week's weather. And if it turns out that Alexa never claimed that it will rain next week, there is again no reason to believe that it will rain. Similarly, if there would be evidence that Alexa was lying, then the inference of the argument fails and with it the whole argument. The only attack that would not directly result in the unacceptability of a_2 would be an argument for a contrary conclusion. If there would be another meteorologist claiming that it will be sunny next week, then we might still choose to accept a_2. However, this will also be clear from the abstract perspective, as there would be a mutual attack between the two arguments.

As this example illustrates, an argument fails as soon as one of its components is shown to be problematic. Which specific component it is or how many other components there are does not change the result. Thus, there seems to be no feature of the attack that we can recognise on the structured perspective that would lead to a different evaluation of the situation than the abstract perspective. For this reason, there is less need to use structured arguments when we only consider attacks. As we will see in the next two sections, the situation is different when it comes to support.

2.1.1 Deductive and Necessary Support

Some of the earliest proposed semantics for bipolar argumentation frameworks work with very strong notions of support. As we have seen that an attack can make an argument unacceptable, it might be reasonable to consider analogous notions of support, where supports can force acceptance of arguments. Suppose we have arguments a and b such that a supports b (i.e. $a \dashrightarrow b$). If we understand this support *deductively*, then accepting a will entail accepting b. In that sense, *a deductively entails b* (Boella et al. 2010; Cayrol

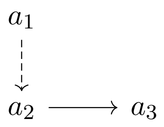

(a) A bipolar framework.

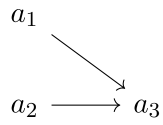

(b) An attack framework with a supported attack.

Figure 2.2: A supported attack.

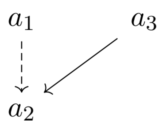

(a) A bipolar framework.

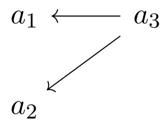

(b) An attack framework with a mediated attack.

Figure 2.3: A mediated attack.

and Lagasquie-Schiex 2013). Alternatively, we can understand support as *necessary*. For instance, we can say that we can only accept b if we can support it. Then accepting b would entail accepting a, i.e. a is *necessary* for b (Nouioua 2013; Nouioua and Risch 2011).

Interestingly, both ideas can be captured in pure attack frameworks. We start with deductive support. Given a bipolar framework, we can add suitable attacks and then use the usual Dung semantics. The idea can be illustrated the following way: Consider the bipolar framework in Figure 2.2a. There, a_1 deductively supports a_2, which in turn attacks a_3. But if you cannot accept a_1 without also accepting a_2, then in some sense a_1 also attacks a_2, namely through a *supported attack*. Similarly, in Figure 2.3a, a_1 again supports a_2, but this time a_3 attacks a_2. If we now think about what the support does, we can note that a_3 in some sense attacks a_1. This is because it attacks something that is directly implied by a_1. We call this a *mediated* attack.

Formally, we can define supported and mediated attacks as follows (Cayrol and Lagasquie-Schiex 2013; L. Yu et al. 2023):

Definition 10 (Supported and Mediated Attacks)**.** *Let $\mathcal{B} = \langle A, \to, \dashrightarrow \rangle$ be a bipolar argumentation framework and take $a, b \in A$.*

We say that a supported attacks b if there exists $c \in A$ such that $a \dashrightarrow c$ and $c \to b$.

We say that a mediated attacks b if there exists $c \in A$ such that $b \dashrightarrow c$ and $a \to c$.

Definition 11 (Deductive Support Semantics)**.** *Let $\mathcal{B} = \langle A, \to, \dashrightarrow \rangle$ be a bipolar argumentation framework. We define the set of complex attacks,*

$$\to^{ded} := \bigcup_{i \in \mathbb{N}} \to^i,$$

where $\to^0 := \to$ and $\to^{i+1} := \to^i \cup \{(a, b) \in A \times A \mid a$ supported or mediated attacks b w.r.t. $\to^i\}$.

An extension $E \subseteq A$ is called d-admissible, d-complete, d-grounded, or d-preferred if it is admissible and closed under $\dashrightarrow$, complete, grounded, or preferred in $\mathcal{A} = \langle A, \to^{ded} \rangle$.

The notion of necessary support can be developed analogously. The intuition remains the same, except that the direction changes. Now, if some argument a supports some other argument b, accepting b entails accepting a. Thus if we have a situation as in Figure 2.4a, where a_2 supports a_1 and is attacked by a_3, we can add a secondary attack from a_3 to

(a) A bipolar framework.

(b) An attack framework with a secondary attack.

Figure 2.4: A secondary attack.

(a) A bipolar framework.

(b) An attack framework with an extended attack.

Figure 2.5: An extended attack.

a_1. The intuition is that a_1 needs to be defended against a_3, since it attacks a supporter. Similarly, in Figure 2.5a, we can capture the support from a_2 to a_1 by adding an extended attack from a_1 to a_3. This is because accepting a_1 directly implies accepting an attacker of a_3. Thus in some sense, a_1 can be said to attack a_3.

We can then define the semantics for necessary supports as follows (Nouioua and Risch 2011; L. Yu et al. 2023):

Definition 12 (Secondary and Extended Attacks)**.** *Let $\mathcal{B} = \langle A, \rightarrow, \dashrightarrow \rangle$ be a bipolar argumentation framework and take $a, b \in A$.*

We say that a secondary attacks b if there exists $c \in A$ such that $c \dashrightarrow b$ and $a \rightarrow c$.

We say that a extended attacks b if there exists $c \in A$ such that $c \dashrightarrow a$ and $c \rightarrow b$.

Definition 13 (Necessary Support Semantics)**.** *Let $\mathcal{B} = \langle A, \rightarrow, \dashrightarrow \rangle$ be a bipolar argumentation framework. We define the set of complex attacks,*

$$\rightarrow^{nec} := \bigcup_{i \in \mathbb{N}} \rightarrow^{i},$$

where $\rightarrow^0 := \rightarrow$ and $\rightarrow^{i+1} := \rightarrow^i \cup \{(a,b) \in A \times A \mid a$ secondary or extended attacks b w.r.t. $\rightarrow^i\}$.

An extension $E \subseteq A$ is called n-admissible, n-complete, n-grounded, or n-preferred if it is admissible and closed under $\dashrightarrow^{-1}$, complete, grounded, or preferred in $\mathcal{A} = \langle A, \rightarrow^{nec} \rangle$, where $\dashrightarrow^{-1} = \{(a,b) \in A \times A \mid (b,a) \in \dashrightarrow\}$.

It is worth noting that deductive support can be translated to necessary support (and vice-versa) by inverting the support relation (Cayrol and Lagasquie-Schiex 2013).

Proposition 6. *Let $\mathcal{B} = \langle A, \rightarrow, \dashrightarrow \rangle$ and $\mathcal{B}' = \langle A, \rightarrow, \dashrightarrow^{-1} \rangle$ be bipolar argumentation frameworks. Then an extension $E \subseteq A$ is d-complete, d-grounded, or d-preferred in $\mathcal{B}$ iff it is n-complete, n-grounded, or n-preferred in $\mathcal{B}'$.*

There is a further notion of support that can be seen as a variant of necessary support, namely *evidential* support (Oren and Norman 2008). The basic idea here is that some arguments can only be accepted if they are supported by evidence. For that purpose, Oren and Norman introduce a special argument η, which stands for the environment (which we can understand as the evidence). Intuitively, for an argument a to be acceptable, there

(a) Simple abstract supports.

(b) Structured instances of supports.

Figure 2.6: Extended example from Figure 2.1

needs to be a support chain with η as its root and a as its head. This is then built into an new notion of admissibility, where each argument needs to be supported by evidence and defended only against attacks of arguments that are themselves supported by evidence. In a later paper, Polberg and Oren show that evidential support can be translated into necessary support (Polberg and Oren 2014, Theorem 5.8.). The main difference is that evidential argumentation systems allow for collective supports and attacks. That is, support and attack are not relations between arguments but are subsets of $2^A \times A$. While it is possible to define collective support and attack relations with structured arguments (see Yun et al. 2020 for structured collective attacks), SBAFs work with relations between arguments. Hence, we can consider evidential support a version of necessary support as defined above.

Let us now examine deductive and necessary support from the perspective of SBAFs and return to the example of Alexa the meteorologist (page 16). Previously, we only had an attack, so let us introduce the following two supporting arguments:

a_3: w : Alexa didn't pass her final exams.
———
s : Alexa never finished her degree in meteorology.

a_4: x : Alexa claimed it will rain next week in an interview with the local newspaper.
———
u : Alexa claims that it will rain next week.

The new SBAF and its corresponding AF are depicted in Figure 2.6. We now have argument a_3 supporting a_1 and argument a_4 supporting a_2 and we can ask: should these supports be deductive or necessary?

If we consider the structured arguments, the situation is clear. Accepting a_3 means that one accepts the claim that Alexa never finished her degree in meteorology. But note that argument a_1 infers directly from that claim that Alexa is not a certified meteorologist. Further, as far as the information in the SBAF is concerned, there is no reason to doubt the inference of a_1. All of this together strongly suggests that accepting a_3 should entail accepting a_1, just as deductive support would require. However, the situation is different when it comes to the support from a_4 to a_2. The difference is that a_2 relies on two premises: that Alexa is a certified meteorologist and that she claims it will rain next week. As before, accepting the supporting argument a_4 requires one to accept a premise of a_2, namely that Alexa claims it will rain next week. But now this is only one of two premises, whereas before the supported argument (a_1) relied on only one premise. It is very much reasonably possible to accept that Alexa's claim happened, while at the same time maintaining that she is not a certified meteorologist (e.g. by accepting a_1). Thus a_4 does not deductively support a_2.

We have two supports in our example, one that seems to be deductive and one that does not seem so. There is an important difference between the two supports and we should not treat them equally when evaluating the argumentation. Crucially, we cannot tell the supports apart in the abstract BAF version of the arguments, as seen in Figure 2.6a. The problem with bipolar semantics that only consider abstract arguments is that it has to treat all supports the same in all cases. While we have seen that this is plausible for attacks, it is less plausible for supports. We will make this point throughout this and the next section.

The situation is similar when it comes to necessary support. There are cases where it seems reasonable that accepting the supported argument should entail accepting the supporting argument. For instance, this is the case when somewhat minimal arguments are involved that use the same sentence as premise and conclusion. In essence, a minimal argument simply states the sentence and is an edge case of an argument. But if we have a situation such as this: $a_1 : \langle\{s\}, s\rangle \dashrightarrow a_2 : \langle\{s, t\}, u\rangle$, where the minimal argument a_1 supports a_2, this support seems to be necessary. If we accept a_2, then we should also accept sentence s. But a_1 is just that sentence restated, so a_1 should also be accepted, making its support necessary. But as we have seen before in Figure 2.6, most other supports are not necessary. Again, the point is that supports behave very differently in different situations, but we can only distinguish these situations once we recognise the structure of the arguments.

2.1.2 Gradual Support

Apart from deductive and necessary support, a large literature has developed around notions of gradual support. Gradual evaluation of BAFs differs from extension-based approaches in that instead of calculating sets of acceptable arguments, each argument is given an *acceptability degree*. Usually, this is a real number between 0 and 1, but in theory any real number could be seen as specifying the acceptability of an argument (Mossakowski and Neuhaus 2016). The idea then is that attacks and supports can influence the acceptability of arguments in degrees. Baroni et al. (2019) give a general overview on a range of graded approaches. Early approaches to gradual bioplar argumentation go back to Cayrol and Lagasquie-Schiex (2005) and Amgoud et al. (2008).

We have seen in the previous section that not all support is deductive. If a supporting argument only supports one of many premises of the supported argument, the support is not deductive, but might still be thought to increase the acceptability of the supported argument. While extension-based approaches have difficulties capturing support that is weaker than deductive (or necessary), gradual approaches are very well suited for that purpose. In these, supports can increase the acceptability of an argument to just *some degree*, without having to deductively entail it.

Most gradual approaches start by putting *initial weights* on arguments that represent their initial plausibility, independently of the other arguments in the framework. These weights will then be updated to take the whole framework into account and give each argument an acceptability degree.

Definition 14 (Weighted Bipolar Argumentation Framework)**.** *A weighted bipolar argumentation framework (WBAF)* $\mathcal{WB} = \langle A, w, \rightarrow, \dashrightarrow\rangle$ *is a BAF with a function* $w : A \rightarrow [0, 1]$*, specifying for each argument an initial weight.*

There are many different ways to update the initial weight of an argument (Amgoud and Ben-Naim 2016, 2018; Baroni et al. 2015; Mossakowski and Neuhaus 2016; Potyka 2018; Rago et al. 2016). Here is one such way (Amgoud and Ben-Naim 2018):

Definition 15 (Euler-based Semantics)**.** *Let* $\mathcal{WB} = \langle A, w, \rightarrow, \dashrightarrow\rangle$ *be a well-founded WBAF such that* $w(a) < 1$ *for all* $a \in A$*. The acceptability degree of an argument* $a \in A$ *is defined*

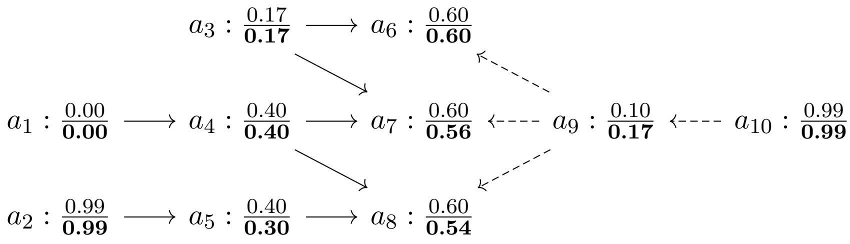

Figure 2.7: An example for the Euler-based semantics (Amgoud and Ben-Naim 2018, p. 51). The numbers are $\frac{\textit{initial weight}}{\textbf{\textit{acceptability degree}}}$.

recursively as

$$Deg_{\mathcal{WB}}(a) = 1 - \frac{1 - w(a)^2}{1 + w(a) \cdot e^E}, \text{ where } E = \sum_{b \dashrightarrow a} Deg_{\mathcal{WB}}(b) - \sum_{c \rightarrow a} Deg_{\mathcal{WB}}(c).$$

Figure 2.7 gives an example for the Euler-based semantics. Note the restriction to well-founded $\mathcal{WB}$s. This is a general restriction for other graded bipolar semantics as well (Baroni et al. 2015; Rago et al. 2016). We can attempt to use the Euler-based semantics on non-well-founded WBAFs by iteratively updating the acceptability degrees of each argument starting with the initial weight. Figure 2.8 gives an example of a non-well-founded WBAF where the iterative application of the Euler-based semantics does not converge on a unique acceptability degree. While there are two approaches in the literature that (potentially) give well-defined results under some assumptions in non-well-founded frameworks (Mossakowski and Neuhaus 2016; Potyka 2018), there is no approach in the literature that is known to give well-defined results in all frameworks.

Apart from the restriction to well-founded frameworks, gradual evaluations of bipolar argumentation again deal only with abstract arguments and thus have to treat all supports the same. It is interesting to consider the case where supports and attacks interact, that is, where one argument is both supported and attacked. This is the case, for instance, with argument a_6 in Figure 2.7, that is attacked by a_3 and supported by a_9. Notably, in this

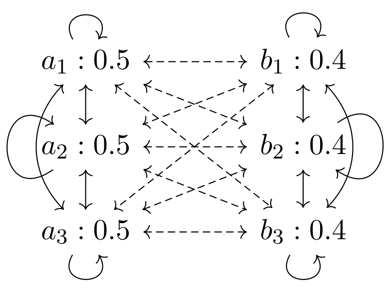

(a) A non-well-founded WBAF where the Euler-based semantics does not converge.

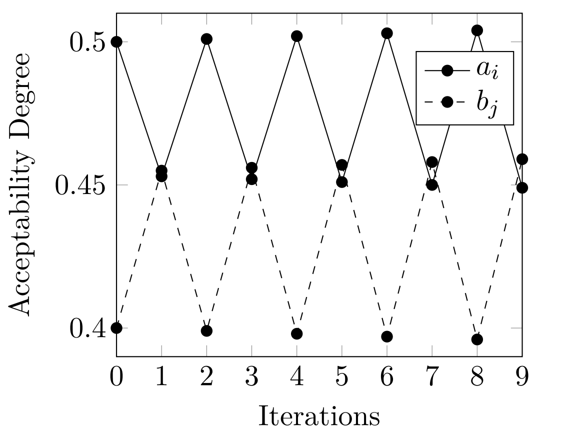

(b) The acceptability degrees of arguments a_i and b_j with multiple iterations of $Deg_{\mathcal{WB}}$.

Figure 2.8: The Euler-based semantics does not always converge on non-well-founded frameworks. The example stems from Mossakowski and Neuhaus (2018).

case the attacker and supporter have the same acceptability degree, so the acceptability of a_6 remains unchanged from its initial weight. In that sense, the attack and support cancel out. Amgoud and Ben-Naim call this the *strict Franklin* property (Amgoud and Ben-Naim 2018, p. 44).

In general, there are three possibilities when attacks and supports interact: (1) Attacks outweigh supports, (2) they cancel each other, or (3) Supports outweigh attacks. Amgoud and Ben-Naim dismiss option (3), at least in the case where the attacks and supports are equally strong, and in their Euler-based semantics, they choose option (2), while admitting that (1) would be reasonable as well (Amgoud and Ben-Naim 2018, p. 44). From the perspective of SBAFs, the answer depends on which components of the argument are attacked and supported. For instance, in the example of Figure 2.6, argument a_2, claiming that it will rain next week, since the certified meteorologist Alexa says so, is both attacked and supported. But note that the support concerns the premise that Alexa did indeed claim that it will rain next week, while the attack concerns the premise that Alexa is indeed a certified meteorologist. We have noted that for an argument to be acceptable, each premise has to be acceptable. That is, if one premise fails, the whole argument fails. Thus, the attacker a_1 renders a_2 unacceptable, irrespective of how its strength compares to that of the supporter a_4. In general, supports only have the potential to cancel or outweigh attacks if they concern the same component. The limitation of gradual semantics, even with weighted BAFs, is that they have to choose the same outcome of the interaction in all cases and cannot distinguish the different situations where supports either never outweigh attacks or where they sometimes might.

Part of this limitation could be overcome by putting weights on attacks as well as arguments (Yun and Vesic 2021). Then, if a support is turned ineffective because another component of the supported argument is attacked, this could be modelled by putting the weight of the support to 0. However, a good understanding of the role of structure in the interaction between supports and attacks is still needed in order to know how to put weights on attacks.

2.2 Doubt in Bipolar and Structured Argumentation

We have seen that many semantics for abstract bipolar argumentation cannot be directly adapted to structured bipolar argumentation. The reason is that they have to treat all supports and interactions between supports and attacks the same, whereas recognising the structure of arguments allows us to distinguish many situations that lead to different treatment of supports. But structure is not the only feature of informal argumentation that we discussed in Section 1.1. We also saw that there is a notion of *mere doubt* at play, meaning that one is at least sometimes allowed to reject an argument or sentence based on doubt, rather than only through counter-arguments. In this section, we review more bipolar semantics and some approaches to structured argumentation and discuss how a lack of mere doubt manifests.

In essence, the problem is that all the approaches we discuss here either use Dung's complete semantics or semantics based on it such as grounded or preferred. Recall that a complete extension has to contain all arguments that it defends. This means in particular that all unattacked arguments have to be accepted. In other words, you cannot doubt an argument without there being a counter-argument. This leads to a wide range of cases where supports between arguments do not have any or very little impact on the acceptable extensions. To see that mere doubt is needed even when we recognise the structure of arguments, we can consider what happens to support when we use complete semantics in SBAFs.

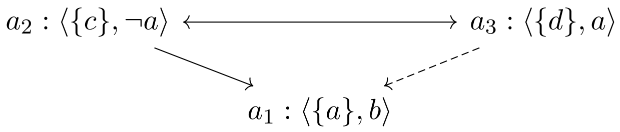

Figure 2.9: An attack and a support on the same component.

Going back to Figure 2.6, we can see that there is only one complete extensions, namely $\{a_1, a_3, a_4\}$, and everything we have said about the interpretation of support thus far agrees with that. Since a_1 is unattacked, it has to be accepted, making the support on it superfluous. As a_4 only supports one of the two premises of a_2, there is no issue in accepting the supporting argument without also accepting the supported argument. More generally, we have seen that attacks outweigh supports if they concern different components on an argument. Taken together with the feature of complete extensions that all defended arguments need to be accepted, it seems that support can at most play a role if it concerns an attacked component of an argument.

Figure 2.9 depicts the situation where an attack and a support go to the same component of an argument. Note that the attacker, a_2, and the supporter, a_3, rebut each other. This will in general be the case: the conclusion of the attacker is incompatible with the same premise of the attacked argument that forms the conclusion of the supporter. But now a_3 not only supports a_1, it also defends it. Thus, even in complete semantics, the support makes a_1 acceptable in that $\{a_1, a_3\}$ is a complete extension. This raises the question: is all impact from the support of a_3 to a_1 captured by the notion of defence? Indeed, it seems that complete semantics render support as a relation superfluous even in SBAFs. Implementing a notion of doubt, e.g. by using admissibility-based semantics, where one is not forced to accept defended arguments, can remedy this problem. We will present our semantics for SBAFs in Chapter 3. For now, we examine how other approaches to bipolar and structured argumentation fare with respect to doubt.

2.2.1 Selection and Defence-Based Support

Recall that Dung semantics sometimes give multiple extensions that are deemed acceptable. For instance, the BAF in Figure 2.10 has four preferred extensions: $\{a_1, a_3\}$, $\{a_1, a_4\}$, $\{a_2, a_3\}$, and $\{a_2, a_4\}$. If we want to use these semantics not as a way to check the reasonability of our own view, but to figure out what we should believe, this might be unsatisfying. Selection-based semantics reduce the number of acceptable extensions by using supports between arguments to make a selection of the extensions given by Dung semantics (Gargouri et al. 2021). Thus, supports are used to reduce the number of extensions that are deemed reasonable.

There are various ways to select extensions given by Dung semantics and indeed, the semantics we provide in Chapter 3 can be seen as versions of selection-based semantics. The point we focus on here, however, is that in the literature, all support-score based semantics select extensions that are complete. For the purpose of illustration, it then suffices to introduce a comparatively simple version of these semantics that selects preferred extensions (L. Yu et al. 2023):

Definition 16 (Internal Coherence Semantics)**.** *Let $\mathcal{B} = \langle A, \rightarrow, \dashrightarrow\rangle$ be a BAF. An extension $E \subseteq A$ is called internally coherent if*

(i) it is preferred and

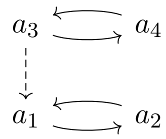

Figure 2.10: An example for support-score based semantics

(ii) for all preferred extensions E': $|\{a \dashrightarrow b \mid a, b \in E\}| \geq |\{a \dashrightarrow b \mid a, b \in E'\}|$.

In Figure 2.10, there is one internally coherent extension, namely $\{a_1, a_3\}$, as a_3 supports a_1. That way, support plays an active role in determining which extensions are acceptable in a BAF. Of course, we could also choose some other complete semantics than preferred semantics, and we could consider different scores. For instance, if we chose to count the supports coming from outside the extension, $\{a_1, a_4\}$ would be the chosen extension in Figure 2.10.

While this type of semantics does find a role for support in complete semantics, that role is fairly limited. As noted in Proposition 5 (p. 11), all frameworks with a well-founded attack relation have one unique complete extension. Thus, in these cases there is nothing to select and support cannot play a role. That is, no matter what the support relation looks like, selection-based semantics will always chose the same complete extension. A similar point goes for unattacked arguments. As they are always part of any complete extension, they will be accepted no matter which extension is selected. Hence, supporting them will not increase their acceptability in any way. This point can be extended by observing that, in attack-free frameworks, all arguments are accepted. Thus, support can at most play a role in relation to attacks. This is in stark contrast with the empirical prevalence of support (Koszowy et al. 2022) and the critical discussions of pragma-dialectics, where all non-mixed differences of opinion proceed without any attacks. Selection-based semantics hence have very limited space for support to play a role and the main reason for this is their reliance on complete semantics.

A different way of taking supports in BAFs into account is to strengthen the defence requirement in Dung semantics. Yu et al. propose three different ways of doing so: (i) defending arguments need to support the defended argument, (ii) defending arguments need to be supported, and (iii) supporters of attackers also need to be defended against (L. Yu et al. 2023).

Definition 17 (Strengthened Defence)**.** *Let $\mathcal{B} = \langle A, \rightarrow, \dashrightarrow \rangle$ be a BAF. An extension $E \subseteq A$ is said to defend an argument $a \in A$ if*

(i) for each $b \in A$ s.t. $b \rightarrow a$, there exists $c \in E$ s.t. $c \rightarrow b$ and $c \dashrightarrow a$.

(ii) for each $b \in A$ s.t. $b \rightarrow a$, there exists $c \in E$ s.t. $c \rightarrow b$ and $E \dashrightarrow c$.

(iii) for each $b \in A$ s.t. $b \rightarrow a$, $E \rightarrow b$, and for each $c \in A$ s.t. $c \dashrightarrow b$, $E \rightarrow c$.

Figure 2.11 illustrates the different types of defence. Based on these notions, complete, grounded, and preferred semantics are then defined as usual. Important for us is that semantics with strengthened defence are still completeness-based. As we have observed with selection-based semantics, the consequence of this is that support is only subsidiary to attacks. Generally with complete semantics this is implicit in that it turns out that supporting an argument in absence of attacks will not make a difference, but here it is explicit. Support is only taken into account in terms of defence. Hence, evaluating

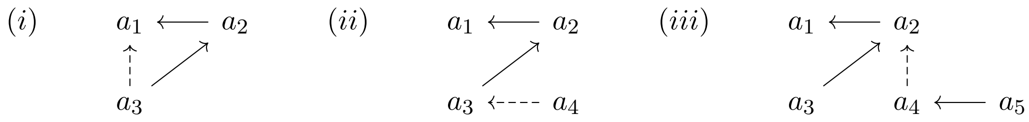

Figure 2.11: The three types of strengthened defence. In (i), a_3 defends a_1. In (ii), a_3 and a_4 together defend a_1. In (iii), a_3 and a_5 together defend a_1.

frameworks without attacks is trivial, as all arguments are simply accepted. Again, mere doubt is not allowed and thus support only plays a minor role.

In terms of structured arguments, defence (i) correlates with our observations about complete semantics in SBAFs. Namely, we saw that, assuming completeness, the only supports that seem to make a difference from the structured perspective are those that go to already attacked components of arguments (see Figure 2.9). But then the supporting argument will rebut the attacker and thus defend it. This suggest that defence (i) shares the following view: assuming completeness, the only supports that play a role are those of arguments that also defend the supported argument. However, as argued above, it is not satisfactory if the role of support is exhausted by that.

Of course, there are more proposals for extension based support semantics in the literature. For instance, Potyka (2021) provides a semantics that is suitable for argumentation where pro- and con-arguments are weighed against each other. A consequence of this perspective is that the resulting semantics does not need to be conflict-free, as it is very possible to accept both an attacker and a supporter in this context. Also, Cohen et al. (2012) propose a bipolar framework based on the Toulmin scheme. However, their semantics reduce to abstract argumentation frameworks in a similar way as deductive and necessary supports.

2.2.2 ASPIC$^+$and Assumption-Based Argumentation

All of the approaches we have discussed so far were abstract in the sense that arguments were abstract entities without structure. Now we discuss two notable approaches to structured argumentation: ASPIC$^+$(Modgil and Prakken 2013; Prakken 2010) and assumption-based argumentation (ABA) (first introcuded by Bondarenko et al. 1997, 1993). Interestingly, neither of them use an explicit support relation between arguments. They make up for this by giving arguments a richer structure than they have in SBAFs where supporting arguments are interpreted as sub-arguments. We first introduce ASPIC$^+$and then ABA.

Since we are interested in comparing ASPIC$^+$to SBAFs, we will only define its aspects relevant for our purposes. Thus, ASPIC$^+$as presented here is somewhat simplified compared to its full version. In general, ASPIC$^+$takes into account features of argumentation that are not present in SBAFs, such as strict inference rules that cannot be questioned and preference rankings between arguments to represent their relative strength. We will note the omissions as we go on. The presentation mainly follows Modgil and Prakken (2014).

As in SBAFs, we start with a language from which arguments are built. This is very similar to Definition 6, except that we add a set of inference rules R and the naming function n now gives names to rules instead of arguments. This stems from a difference in perspective between SBAFs and ASPIC$^+$. In SBAFs, we assume the arguments to be given, for instance through a reconstruction of a debate. We are then interested in evaluating the arguments as they are. In ASPIC$^+$, however, the arguments are constructed inside the framework. For that, we need to have inference rules that tell us which sentences can be inferred from which others.

Definition 18 (Argumentation System). *An argumentation system is a tuple $\mathcal{AS} = \langle L, \overline{}, R, n \rangle$ where L is a set of sentences, $\overline{} : L \to 2^L \backslash \{\emptyset\}$ maps sentences to sets of contrary sentences, R is a set of rules of the form $\varphi_1, \dots, \varphi_n \Rightarrow \psi$ with $\varphi_1, \dots, \varphi_n, \psi \in L$ ($\varphi_1, \dots, \varphi_n$ are called antecedents and ψ the consequent of the rule) and $n : R \to L$ gives names to rules.*

In the full version of ASPIC$^+$, the set of rules contains both defeasible and strict rules. Strict rules are intended to represent logically valid inferences and they cannot be undercut. However, as there is no notion of logically valid arguments in SBAFs, we omit these rules here.

In order to construct arguments, we not only need inference rules but also a knowledge base. It contains the starting assumptions that form the premises of arguments and to which rules can be applied. As with strict rules, the full version of ASPIC$^+$distinguishes between axioms and ordinary premises, where axioms cannot be attacked. As there is no notion of axioms present in SBAFs, we again omit them.

Definition 19 (Argumentation Theory). *An argumentation theory $\mathcal{AT} = \langle \mathcal{AS}, K \rangle$ is a tuple consisting of an argumentation system and a knowledge base $K \subseteq L$.*

Now we can construct arguments. As mentioned, they have a richer structure in ASPIC$^+$than in SBAFs, as they can contain sub-arguments. Intuitively, each SBAF argument contains one rule application, represented by what we call its implicit inference claim. ASPIC$^+$arguments, in contrast, can contain any number of rule applications, so long as the original premises are part of the knowledge base. Each rule application creates a sub-argument, whose conclusion can then be used together with further rules. In that sense, ASPIC$^+$arguments resemble *inference trees*, where the nodes are connected via rules and the leaves need to be in the knowledge base. Suppose, for instance, that we have the knowledge base $\{p\}$ with rules $\{p \Rightarrow q, q \Rightarrow r\}$. We get the following arguments: $a_1 : p$ itself, $a_2 : a_1 \Rightarrow q$, where the conclusion of a_1 is used together with the first inference rule to infer q, and $a_3 : a_2 \Rightarrow r$, where the conclusion of a_2 is used together with the second inference rule to infer r. We say that a_1 and a_2 are sub-arguments of a_3, as the inference tree for r traces back through these arguments to premise p in the knowledge base. In comparison, in an SBAF, we would say that a_1 supports a_2 and it in turn supports a_3.

Definition 20 (ASPIC$^+$Arguments). *Let $\mathcal{AT} = \langle \mathcal{AS}, K \rangle$ be an argumentation theory. An argument a is:*

(1) φ *if* $\varphi \in K$ *with:* $Prem(a) = \{\varphi\},\ Conc(a) = \varphi,\ Sub(a) = \{\varphi\},\ Rules(a) = \emptyset.$

(2) $a_1, \dots, a_n \Rightarrow \psi$

if $a_1, \dots, a_n$ *are arguments and there is a rule* $Conc(a_1), \dots, Conc(a_n) \Rightarrow \psi$ *in* R.

$Prem(a) = Prem(a_1) \cup \dots \cup Prem(a_n),$

$Conc(a) = \psi,$

$Sub(a) = Sub(a_1) \cup \dots \cup Sub(a_n) \cup \{a\},$

$Rules(a) = Rules(a_1) \cup \dots \cup Rules(a_n) \cup \{Conc(a_1), \dots, Conc(a_n) \Rightarrow \psi\}.$

$TopRule(a) = Conc(a_1), \dots, Conc(a_n) \Rightarrow \psi.$

Attacks between arguments (in absence of strict rules and axioms) are defined similarly as in SBAFs: ASPIC$^+$also distinguishes between undermines, undercuts, and rebuts. The main difference is that, since arguments need to draw their premises from the knowledge base, arguments can also be attacked on sub-arguments. That is, they can be attacked by showing that their conclusions cannot be traced back to the knowledge base.

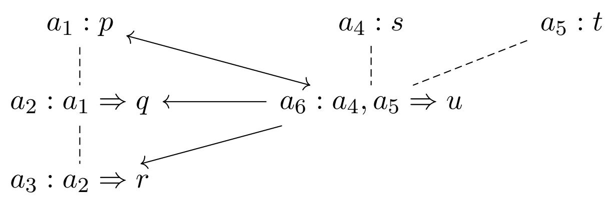

Figure 2.12: An ASPIC$^+$framework with $K = \{p, s, t\}$, $R = \{p \Rightarrow q; q \Rightarrow r; s, t \Rightarrow u\}$, and $u \in \overline{p}, p \in \overline{u}$. Dashed lines indicate sub-arguments.

Definition 21 (ASPIC$^+$Attacks)**.** *An argument a*

(i) undercuts argument b if $Conc(a) \in \overline{n(r)}$ *for some rule* $r \in Rules(b)$*,*

(ii) rebuts argument b if $Conc(a) \in \overline{\psi}$ *for some sub-argument* $b'_1, \ldots, b'_n \Rightarrow \psi$ *of b,*

(iii) undermines argument b if $Conc(a) \in \overline{\psi}$ *for some* $\psi \in Prem(b)$*.*

In its full version, ASPIC$^+$distinguishes between *attacks* and *defeats*, where the latter represents successful attacks. Here, ASPIC$^+$uses a preference ordering over the arguments. An attack is deemed successful if it is an undercut or the attacked argument is not preferred over the attacker. The preference ordering over arguments can be given directly or it can be aggregated from preferences over premises and inference rules. Since we omit preferences, we can deal directly with attacks and end up with a simplified definition of a structured argumentation framework.

Definition 22 (Structured Argumentation Framework)**.** *Given an argumentation theory* $\mathcal{AT} = \langle \mathcal{AS}, K \rangle$*, we can define a structured argumentation framework (SAF)* $\mathcal{SA} = \langle A, \rightarrow \rangle$*, where A is the set of arguments that can be created from* $\mathcal{AT}$ *according to definition 20 by using a finite number of inference rules, and* $\rightarrow$ *is the corresponding attack relation according to definition 21.*

Figure 2.12 presents an example for a *SAF*. Note that a_6 attacks both a_2 and a_3 by attacking their sub-argument a_1.

Even though ASPIC$^+$represents an approach to structured argumentation, SAFs are evaluated using standard Dung semantics, namely completeness-based semantics such as preferred or grounded. Thus, the structure of arguments is only used in order to compute their attacks. In Figure 2.12, there are the following complete extensions: $\{a_1, a_2, a_3, a_4, a_5\}$, $\{a_4, a_5\}$, and $\{a_4, a_5, a_6\}$. Note that since they are unattacked, a_4 and a_5 have to be accepted.

In spite of its abstract evaluation, ASPIC$^+$satisfies some rationality postulates that ensure that the extensions given by Dung semantics make sense also from the structured perspective. In the somewhat limited version of ASPIC$^+$presented here, there are two relevant rationality postulates: sub-argument closure and direct consistency (Modgil and Prakken 2013). For any extension E in some $\mathcal{SA} = \langle A, \rightarrow \rangle$, we can formulate them as follows:

Sub-argument Closure: For any $a \in E$, $Sub(a) \subseteq E$.

Direct Consistency: For all $a, b \in E$, $Conc(a) \notin \overline{Conc(b)}$.

Proposition 7. *Let E be an admissible extension in some* $\mathcal{SA} = \langle A, \rightarrow \rangle$*. Then it satisfies direct consistency. If it is also complete, is further satisfies sub-argument closure.*

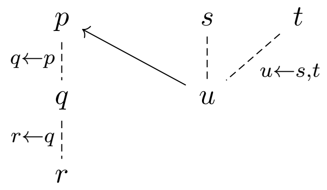

Figure 2.13: An ABA framework with $A = \{p, s, t\}$, $R = \{q \leftarrow p; r \leftarrow q; u \leftarrow s, t\}$, and $u = \overline{p}$.

As most of the discussion on ASPIC$^+$also applies to assumption-based argumentation, we now introduce ABA. The presentation here mainly follows Toni (2014) and Čyras et al. (2017). As ASPIC$^+$, ABA frameworks uses inference rules to construct arguments from a set of assumptions, which takes the place of the knowledge base. A notable difference is that ABA does not use a naming function, as rules cannot be attacked.

Definition 23 (ABA Framework)**.** *An ABA framework is a tuple $\mathcal{AB} = \langle L, R, A, \overline{} \rangle$ where L is a set of sentences and R a set of rules of the form $\sigma_0 \leftarrow \sigma_1, \ldots, \sigma_n$ with $\sigma_i \in L$ for all $i \leq n$ (σ_0 is called head and $\sigma_1, \ldots, \sigma_n$ are called body). $A \subseteq L$ is a (non-empty) set of assumptions and $\overline{} : A \to L$ maps any assumption to its contrary sentence.*

It is sometimes assumed that the ABA frameworks are *flat*, meaning that no rule concludes an assumption (Toni 2014).[1] In flat frameworks, assumptions are indeed only starting points for arguments and cannot be supported by further reasoning. However, for our purposes, we do not need this assumption.

ABA arguments are then constructed similarly to ASPIC$^+$arguments. They are full inference trees, starting with assumptions as premises and containing arbitrarily many rule applications. Rebutting and undercutting attacks are absent from ABA, as only the assumptions of an argument can be attacked. Nevertheless, it is possible to model more general defeasible reasoning in ABA, see Toni (2008).

Definition 24 (ABA Arguments)**.** *An argument is a deduction $S \vdash \sigma$ where $S \subseteq A$, $\sigma \in L$, and there is a finite tree with σ as the root, where each leave is in S or is empty in case of rules with empty bodies, and for each non-leave node there is a rule such that the node is the head and its children the body.*

A set of assumptions $S \subseteq A$ attacks $T \subseteq A$ ($S \to T$) if there is an argument $S' \vdash \overline{\sigma}$ such that $S' \subseteq S$ and $\sigma \in T$.

Figure 2.13 presents an example of an ABA framework. It is uses the same assumptions and rules as the example for ASPIC$^+$in Figure 2.12. Note that all arguments need to start from assumptions, thus instead of 6 arguments as in ASPIC$^+$, we only have 2 arguments here: $\{p\} \vdash r$ and $\{s, t\} \vdash u$. Further, only assumptions have contraries, thus we cannot express that p is contrary to u, as we specified in Figure 2.12. The attack on p is also not transferred to attacks on the sub-deductions for q and r as it would be in ASPIC$^+$.

ABA semantics come in two types: either they define acceptable sets of arguments or acceptable sets of assumptions. Acceptable sets of arguments are determined as in ASPIC$^+$according to Dung semantics (Toni 2014). In order to determine sets of acceptable

[1] Note that flat frameworks are defined differently in Čyras et al. 2018 using the notion of closed sets of assumptions, see Definition 25. According to them, an ABA framework is flat if all subsets of assumptions are closed. The two definitions are not quite equivalent in presence of rules that can never be reached in a deduction. However, for our purposes, it does not matter which definition we choose.

assumptions, Dung semantics are adapted in a fairly straightforward way (Čyras et al. 2018, 2017). The only complication is that since we do not restrict ourselves to flat ABA frameworks, acceptable sets of assumptions need to be closed under deductions of assumptions (not sentences in general).

Definition 25 (ABA Semantics)**.** *Let $\mathcal{AB} = \langle L, R, A, \overline{} \rangle$ be an ABA framework and $S \subseteq A$.*

Closure: S is closed if $S = \{\sigma \in A \mid \exists S' \subseteq S : S' \vdash \sigma\}$.

Conflict-free: S is conflict-free if $S \not\to S$

Defence: S is defended if $S \to T$ for each closed $T \subseteq A$ such that $T \to S$.

Admissible: S is admissible if it is closed, conflict-free, and defended.

Complete: S is complete if it is admissible and contains all assumptions it defends, where S defends an assumption σ if $S \to T$ for all closed $T \subseteq A$ such that $T \to \{\sigma\}$.

Grounded: S is grounded if it is $\subseteq$-minimal amongst complete sets.

Preferred: S is preferred if it is $\subseteq$-maximal amongst complete sets.

In Figure 2.13, there is only one complete set of assumptions, $\{s, t\}$, as anything containing p would not be defended.

As the sentences of the language are directly represented in frameworks for structured argumentation, it is very natural not only to ask about which arguments one should accept, but also about which sentences one should accept (Baroni et al. 2018b). Hence, focusing on sets of assumptions offers this additional perspective.

Neither ASPIC$^+$nor ABA rely on an explicit support relation. Nevertheless, they are able to capture various forms of support. ASPIC$^+$is mostly associated with necessary support. If we think of sub-arguments as supporting arguments, then their necessary character becomes clear. After all, sub-argument closure in ASPIC$^+$guarantees that if an argument is accepted, so are its sub-arguments, resp. its supporters. This is confirmed by Cohen et al. (2018). To make this observation formal, they first define simple SAFs by requiring that no argument contains an inference loop, where two sub-arguments share the same conclusion.

Definition 26 (Simple Structured Argumentation Framework)**.** *A simple structured argumentation framework (SSAF) is a SAF $\mathcal{SA} = \langle A, \to \rangle$ such that there exists no $a \in A$ such that for some $a', a'' \in Sub(a)$ we have $a' \neq a''$ and $Conc(a') = Conc(a'')$.*

For these SSAFs, Cohen et al. then prove that sub-arguments correspond to necessary support (Cohen et al. 2018, Proposition 7):

Proposition 8. *Let $\mathcal{SA} = \langle A, \to \rangle$ be a SSAF. We define its corresponding BAF as $\mathcal{B} = \langle A, \to, \dashrightarrow \rangle$ where $a \dashrightarrow b$ iff $a \in Sub(b)$ and $a \neq b$. Then an extension $E \subseteq A$ is complete, grounded, or preferred in $\mathcal{SA}$ iff it is n-complete, n-grounded, or n-preferred in $\mathcal{B}$.*

In addition, Prakken (2014) finds a natural correspondence between ASPIC$^+$and evidential support.

For ABA, Čyras et al. show that it can model both necessary and deductive support (Čyras et al. 2017). For that purpose, they introduce bipolar ABA frameworks:

Definition 27 (Bipolar ABA Frameworks)**.** *An ABA framework* $\mathcal{AB} = \langle L, R, A, \overline{} \rangle$ *is called bipolar if every rule in* R *is of the form* $\sigma_0 \leftarrow \sigma_1$ *where* $\sigma_1 \in A$ *and either* $\sigma_0 \in A$ *or* $\sigma_0 = \overline{\sigma_2}$ *for some* $\sigma_2 \in A$*.*

Essentially, bipolar ABA frameworks are such that all involved sentences are assumptions or contraries of assumptions and all rules have singleton bodies. This allows capturing necessary and deductive supports the following way: Given a BAF, it can be translated into a bipolar ABA framework such that admissible and preferred sets of assumptions correspond to admissible and preferred extensions either according to necessary or deductive support.

Definition 28 (From BAFs to ABA)**.** *Let* $\mathcal{B} = \langle A, \rightarrow, \dashrightarrow \rangle$ *be a BAF. We build a bipolar ABA framework as follows:*

$L \ = A \cup \{a^c \mid a \in A\}$

$R^N \ = \{b^c \leftarrow a \mid a \rightarrow b\} \cup \{a \leftarrow b \mid a \dashrightarrow b\}$

$R^D \ = \{b^c \leftarrow a \mid a \rightarrow b\} \cup \{b \leftarrow a \mid a \dashrightarrow b\}$

$\overline{} : \ \overline{a} = a^c$ *for all* $a \in A$

For necessary support, the corresponding ABA framework is $\mathcal{AB}^N = \langle L, R^N, A, \overline{} \rangle$ *and for deductive support it is* $\mathcal{AB}^D = \langle L, R^D, A, \overline{} \rangle$*.*

In accordance with the observation about the duality between necessary and deductive support (Proposition 6), the only difference between the translations for necessary and deductive support is the direction of the rules representing support. We then get the following correspondence (Čyras et al. 2017, Propositions 4, 5, 7, and 8):

Proposition 9. *Let* $\mathcal{B} = \langle A, \rightarrow, \dashrightarrow \rangle$ *be a BAF. Then* $E \subseteq A$ *is n-admissible or n-preferred in* $\mathcal{B}$ *iff it is admissible or preferred in* $\mathcal{AB}^N$ *and it is d-admissible or d-preferred iff it is admissible or preferred in* $\mathcal{AB}^D$*.*

It is notable that both ASPIC$^+$and ABA have natural connections with necessary and deductive support, since we have criticised those notions of support as being not well-suited for structured argumentation (Section 2.1.1). For ABA, though, capturing necessary and deductive support comes at the cost of severely limiting the structure of the arguments. Bipolar ABA frameworks only allow rules with singleton bodies, meaning that arguments with multiple premises are impossible to construct. Indeed, in such a setting, we will show a similar result for SBAFs (Proposition 15). However, we would expect for structured argumentation to also capture supports between arguments with multiple premises.

ASPIC$^+$fares better in that regard, as it implements a notion of necessary support with more complex arguments. However, support, here conceived as sub-argument, only plays a role in argument construction. The evaluation of the arguments follows Dung semantics and is accordingly completeness-based (as is ABA's for that matter). Thus, once an argument is constructed, it has to be accepted in absence of attacks (or presence of defenders). In that sense ASPIC$^+$(and ABA) do not leave room for doubt in their semantics. It is also worth mentioning that the rationality postulate of sub-argument closure only holds for complete extensions, hence the semantics are also not easily adapted to incorporate doubt.

As support only plays a role in the construction of ASPIC$^+$and ABA arguments, perhaps we could implement doubt there. Namely, all arguments need to be traced back to the knowledge base, resp. the set of assumptions. One way to incorporate doubt into

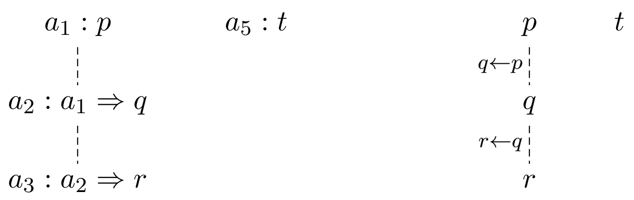

Figure 2.14: The frameworks of Figures 2.12 and 2.13 with $K = A = \{p, t\}$.

ASPIC$^+$and ABA would be to adapt the knowledge base and set of assumptions to the beliefs of an agent considering a debate (see Wallner et al. 2024 for such an approach). In the example above, the knowledge base and set of assumptions are the set $\{p, s, t\}$. But perhaps some agent doubts sentence s and would only accept it as the conclusion of an argument, but not as a premise or assumption. Hence, they might not want to have it as part of the knowledge base or set of assumptions. This would result in the frameworks depicted in Figure 2.14. Our agent would then end up with attack-free frameworks and could accept all arguments.

An issue with this way of incorporating doubt into ASPIC$^+$and ABA is that each agent considering a debate will end up with their own framework, rather than different extensions in the same framework. Thus it can be difficult to compare the views of different agents, as they would not even recognise the same set of arguments. In particular, this would open up doubt as a strategy for defence. Note, for instance, that argument a_3 and $\{p\} \vdash r$ are unattacked in Figure 2.14, but not in Figures 2.12 and 2.13. Hence, whether an agent accepting a_3 or $\{p\} \vdash r$ needs to defend it against a_4 or $\{s, t \vdash u\}$ depends on their knowledge base. In essence, if they do not want to (or cannot) defend the argument, they can simply doubt a premise of the attacker. Then, it will not even be constructed in their framework and there would be no need for defence. This strategy of ignoring attackers would work in all cases where the attacker is not solely based on the same premises as the accepted arguments. While ignoring opposing arguments might be an effective *rhetorical* strategy, it should hardly count as *reasonable* (see Betz (2016) for a similar argument).

We can also consider the question whether you need to defend an argument against an attacker of which you doubt a premise from the pragma-dialectical perspective. Recall the two roles of protagonist and antagonist in a critical discussion, where the protagonist tries to convince the antagonist from their standpoint. If you are in the role of the antagonist and you do not want to accept an argument of the protagonist, it is sufficient to express mere doubt about a premise. It is then the protagonists responsibility to provide further justification. Thus, if we ask the question "Does an agent have to accept argument a?", then it can be answered negatively in case the agent has doubts. However, the roles are different if the agent already accepts an argument which is then attacked. Now the question is "How do you defend your accepted argument a?", to which mere doubt is not an adequate response. In some sense, the agent is now in the role of the protagonist, defending their standpoint. Thus, the agent needs to actively defend their argument and show where the attacking argument goes wrong.

In sum, there might be ways of incorporating doubt into structured argumentation in the style of ASPIC$^+$or ABA, but it would require substantial work. The simple way of making the knowledge base agent-dependent would allow questionable dialectical strategies for defence that go against reasonable discourse.

Chapter 3

Structured Bipolar Argumentation

This chapter is dedicated to studying structured argumentation frameworks and their semantics. The first section introduces both coherent argument extensions and adequate language extensions and shows that they correspond. Afterwards, Section 3.2 takes the perspective of bipolar semantics and considers principles and support-related properties of semantics for SBAFs. The last two sections, 3.3 and 3.4, relate our semantics to Dung semantics and ASPIC$^+$.

3.1 A Framework for Doubt

We first reintroduce SBAFs with a bit more detail than before. Then, we introduce our semantics and add some discussion of the results we present.

3.1.1 Structured Bipolar Argumentation Frameworks

Now we can start to build a framework that takes all the discussion above into account. Let us first repeat some definitions for SBAFs, now with some technicalities added, that were not relevant before.

Definition 29 (Language (cf. Def. 6))**.** *A language $\mathcal{L} = \langle L, \overline{}, n\rangle$ consists of a set of sentences, an incompatibility function $\overline{} : L \to 2^L$, and a naming function $n : 2^L \times L \to L$.*

We assume $\overline{}$ to be symmetric, i.e. $\forall s, t \in L : s \in \overline{t} \iff t \in \overline{s}$. Additionally, we assume that $\overline{n(\langle\{t\}, t\rangle)} = \emptyset$.

We add the condition of symmetry to the incompatibility function, since it is intended to capture the notion of contrariness that says two sentences cannot be true together. This notion is clearly symmetric, but it is noteworthy that both ASPIC$^+$and ABA work with a more general, non-symmetric contrariness notion. It will further be useful to introduce the notion of a *minimal argument*, just consisting of one premise and concluding with the same sentence. Clearly, such an argument does not really contain an inference step and hence should not be able to be undercut, which is why its name should not have contraries. However, this condition can lead to some technical difficulties in defining the contrariness function with specific logical languages. For instance, the propositional contrariness mentioned in Section 1.3, $\overline{\varphi} = \{\psi \in L \mid \varphi \wedge \psi \models \bot\}$ together with the naming function $n(\langle\{\varphi_1, \dots, \varphi_n\}, \psi\rangle) = (\varphi_1 \wedge \dots \wedge \varphi_n) \rightsquigarrow \psi$ would give $\bot \in \overline{n(\{\varphi\}, \varphi)}$. Thus, extra care needs to be taken there.

Definition 30 (Arguments (cf. Def. 7))**.** *An argument in a language $\mathcal{L} = \langle L, \overline{}, n\rangle$ is a tuple $a =$ $Prem(a), Conc(a)\rangle$ where $Prem(a) \subseteq L$ (non-empty) and $Conc(a) \in L$.*

We also define the set of sentences of an argument a as $Sent(a) := Prem(a) \cup \{Conc(a)\}$. The set of sentences generalises for sets of arguments $Sent(A) = \bigcup_{a \in A} Sent(a)$.

We say a is a minimal argument for statement $s \in L$ if $a = \langle \{s\}, s \rangle$.

Definition 31 (Support and Attack (cf. Def. 8)). *Let a, b be arguments in language $\mathcal{L}$.*

We say that a supports b if $Conc(a) \in Prem(b)$.

We say that a attacks b if $Conc(a) \in \overline{s}$ for some $s \in state(b)$ or if $Conc(a) \in \overline{n(b)}$.

Note the asymmetries between supports and attacks. We have already discussed the more limited range of supports in Section 1.3. But it is also worth noting that while attacks are based on a contrariness function, which can be based on some notion of logical entailment as in the example of propositional logic, supports require the identity of the conclusion with a premise of the supported argument. Contrariness allows for sentences to conflict very indirectly, as it can sometimes be hard to tell whether two sentences are contraries or not. It would be possible to base support on a notion of logical entailment, where the conclusion of a supporting argument only needs to entail a premise of the supported argument. Of course, then it will also be very demanding to determine whether an argument supports another. However, it would also have beneficial effects. For instance, some very intuitive supports are not recognised according to the current definition. Take arguments $a_1 : \langle \{p\}, q \wedge r \rangle$ and $a_2 : \langle \{q, r\}, s \rangle$. It seems clear that a_1 should support a_2, as its conclusion is simply the conjunction of the premises of a_2. However, according to our definition, there is no support here, as the conclusion of a_1 is not identical to either premise of a_2. Defining support through entailment would solve this problem. Nevertheless, we keep to the more standard notion of support we have defined here. Thus, the problem of conjunctive conclusions needs to dealt with by splitting it into two arguments: $a_1' : \langle \{p\}, q \rangle$ and $a_1'' : \langle \{p\}, r \rangle$.

Definition 32 (Structured Bipolar Argumentation Framework (cf. Def. 9)). *A structured bipolar argumentation framework (SBAF) is a tuple $\mathcal{SB} = \langle \mathcal{L}, A, \rightarrow, \dashrightarrow \rangle$ where A is a finite set of arguments in language $\mathcal{L}$ and $\rightarrow, \dashrightarrow$ are the corresponding attack and support relations.*

$\mathcal{SB}$ is called saturated if $\forall s \in Sent(A)$ s.t. $\exists t \in Sent(A) \cap \overline{s}$, there is a minimal argument for s <u>or</u> for t in A, and $\forall u \in Sent(A)$ s.t. $u \in \overline{n(a)}$ for some $a \in A$, there is a minimal argument for u in A.

$\mathcal{SB}$ is called strongly saturated if $\forall s \in Sent(A)$ s.t. $\exists t \in Sent(A) \cap \overline{s}$, there is a minimal argument for s <u>and</u> for t in A, and $\forall u \in Sent(A)$ s.t. $u \in \overline{n(a)}$ for some $a \in A$, there is a minimal argument for u in A.

The notions of saturated and strongly saturated frameworks will be needed later on for some results. The idea is that if there are contrary sentences present in a framework, saturated frameworks should contain a minimal argument for at least one of them, while strongly saturated frameworks contain them for both. Also, both notions require minimal arguments for undercutting sentences. Clearly, strongly saturated frameworks are also saturated.

Figure 3.1 gives an example of an SBAF. Note that it is not saturated, since there is no minimal argument for t or $\neg t$, whereas there is a minimal argument for s. Further, it is useful to observe that minimal arguments always defend themselves against any attacks. This is because they cannot be undercut and since incompatibility is symmetric, any attack ends in mutual rebut.

Observation 1. *Let $\mathcal{SB} = \langle \mathcal{L}, A, \rightarrow, \dashrightarrow \rangle$ be an SBAF and $a \in A$ a minimal argument. Then $\{a\}$ is defended (and hence admissible).*

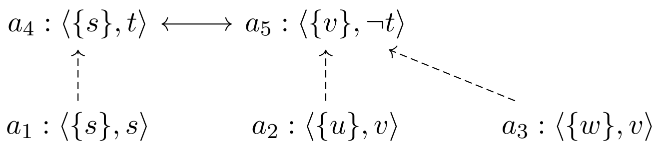

Figure 3.1: An example of an SBAF.

3.1.2 Semantics for SBAFs

What should a semantics for SBAFs look like? Since we deal with structured arguments, we can answer this question from two perspectives: we can talk about which arguments should be accepted (as in ASPIC$^+$) or we can talk about which sentences should be accepted (similar to ABA). Thus we are interested in both *argument extensions* and *language extensions*. Since argument extensions are more familiar, we start with those.

We argued above that it should be possible for agents to doubt arguments or sentences. In that sense, the semantics should be oriented around admissible semantics. But of course the requirements of admissible semantics are too weak, as it would allow to ignore support. Hence, we need to add some additional condition to admissibility that takes support into account. Note also that admissibility does not take the structure of the arguments into account. The basic idea is that we move somewhat towards the notion of complete extensions in the sense that what we call a *coherent* argument extensions should be admissible and include certain arguments—just not all arguments it defends. Essentially, if an argument extension commits an agent to accept all premises of an argument, then that argument should generally be accepted as well. We call this notion *sentence-respect*, since it requires argument extensions to take into account how its associated sentences relate to other arguments.

First, we need a notion of which sentences an (agent accepting an) argument extension is commited to. The easy answer is that the set of sentences corresponding to a set of arguments contains all and only the sentences that figure either as premises or as conclusion of accepted arguments. Second, we now need to specify under which circumstances accepting certain sentences commits an agent to accepting certain arguments. One condition is clear: all premises need to be accepted. There is also a clear condition for when one does not have to accept an argument: when there is reason to doubt the inference of the argument, i.e. if an undercut of the argument is accepted. These two conditions give us a notion of *strong sentence-respect*: If you accept all premises and no undercut of an argument, then you need to accept the argument as well.

Interestingly, strong sentence-respect might require an extension to contain conflicting arguments. In Figure 3.1, accepting arguments a_1 and a_2 would commit you to both accepting a_4 and a_5, even though these arguments rebut each other. Hence, one might want to weaken sentence-respect somewhat, so that it only applies to defended arguments. This gives us *weak sentence-respect*: If you accept all premises and no undercut of a defended argument, then you need to accept the argument as well. This would make accepting only a_1 and a_2 possible. Weak sentence-respect is somewhat closer to complete semantics. The two versions of sentence-respect then give us two notions of coherent argument extensions:

Definition 33 (Coherent Argument Extensions)**.** *A strongly coherent argument extension in an SBAF $\mathcal{SB} = \langle \mathcal{L}, A, \rightarrow, \dashrightarrow \rangle$ is an admissible extension $E \subseteq A$ that satisfies:*

Strong Sentence-Respect: $\forall a \in A$: if $Prem(a) \subseteq Sent(E)$ and $\overline{n(a)} \cap Sent(E) = \emptyset$, then $a \in E$.

A weakly coherent argument extension is an admissible extension E that satisfies:

Weak Sentence-Respect: $\forall a \in A$: *if* $Prem(a) \subseteq Sent(E)$, $\overline{n(a)} \cap Sent(E) = \emptyset$, *and* E *defends* a, *then* $a \in E$.

We will sometimes speak only of strongly or weakly coherent extensions and drop the mention of "argument" if it is clear that argument extensions are meant. Further, we use the shorthand of *coherent semantics* to express that strongly or weakly coherent extensions are seen as the rationally acceptable ones.

In Figure 3.1, we have the following strongly coherent argument extensions: $\emptyset$, $\{a_1, a_4\}$, $\{a_2, a_5\}$, $\{a_3, a_5\}$, $\{a_2, a_3, a_5\}$, and $\{a_5\}$. For instance, accepting a_1 means that strong sentence-respect forces you to also accept a_4, since you then accept all premises and no undercut. The situation is similar with a_2 or a_3, each of which would force acceptance of a_5 under strong sentence-respect. Thus, there can be no strongly coherent argument extension accepting a_1 together with either a_2 or a_3, even though these arguments are not directly in conflict with each other. In contrast, weak sentence-respect can only force acceptance of arguments that are already defended. Thus, it is allowed to accept, say, a_1 together with a_2, as neither a_4 nor a_5 are defended by these arguments. Weak sentence-respect then does not force acceptance of either a_4 or a_5. Thus, on the weakly coherent side, all admissible extensions are acceptable. That is, all extensions that do not contain both a_4 and a_5. This illustrates that weakly coherent argument extensions give somewhat less weight to supports between arguments.

It is worth noting that the empty extension is always strongly and weakly coherent, so we know that such argument extensions always exist. But for weakly coherent extensions, we can say more.

Observation 2. *In an SBAF, complete extensions are weakly coherent.*

Figure 3.1 contains an example of a complete (even preferred) extension, $\{a_1, a_2, a_3, a_5\}$, which is not strongly coherent, as strong sentence-respect would require to also contain a_4.

In terms of the relation between strongly and weakly coherent extensions, as their naming suggests, strongly coherent extensions are also weakly coherent, but not the other way around.

Observation 3. *In an SBAF, strongly coherent argument extensions are weakly coherent.*

We can also relate strongly coherent extensions to complete extensions as follows.

Proposition 10. *Let* $\mathcal{SB} = \langle \mathcal{L}, A, \rightarrow, \dashrightarrow \rangle$ *be an SBAF. Then for each strongly coherent argument extension* $E \subseteq A$*, there exists a weakly coherent argument extension* $E' \subseteq A$ *such that* $E \subseteq E'$ *that is also complete.*

Proof. We have seen that every admissible extension is a subset of a complete extension. Clearly, strongly coherent argument extensions are admissible, thus Observation 2 gives the result. □

We mentioned that we can both think about what arguments we should accept and about what sentences we should accept. Indeed, in some applications it might be more intuitive to think directly about sentences than about arguments. For instance, if you are listening to a debate, it can take a lot of mental effort to think about what exactly the arguments are and with which ones you agree. It is then easier to take each sentence as it comes and think about whether you agree with it or not. The latter approach gives you a language extension, i.e. a set of sentences that you accept. We can now examine what properties such an extension should have in order to count as rational.

The most basic condition we can require of a language extension is that it should not contain incompatible sentences.[2] This mirrors conflict-freeness of argument extensions in being a basic consistency requirement. It should at least be possible that all sentences you accept are true together.

Compatibility: A language extension $S \subseteq Sent(A)$ in an SBAF $\mathcal{SB} = \langle \mathcal{L}, A, \rightarrow, \dashrightarrow \rangle$ is compatible if $\forall s, t \in S$, we have $s \notin \overline{t}$.

Similarly as argument extensions should take the sentences of the arguments into account, we would also expect language extensions to take account of the arguments. After all, we are not presented with a mere set of sentences from which we can choose to accept some. Rather, the sentences are organised into arguments and they occur in a dialectical situation. Thus we need to know which arguments an agent should accept given their accepted sentences. As with sentence-respect above, we have two ways of determining that: either you should accept all arguments of which you accept all premises and no undercut (which gives you the *strong argument set*), or amongst those you only accept the ones you can also defend (which gives you the *weak argument set*). We then have strong and weak language extensions, depending on how we calculate their corresponding set of accepted arguments.

Using strong argument sets, we have a clear an easy criterion to determine whether an agent accepting some sentences is committed to accepting an argument. We only need to look at the accepted sentences and compare them with the sentences at play in a given argument. The situation for the weak argument set is different, as you only need to accept the arguments you can also defend. This opens up the possibility of there being multiple argument sets that could be a plausible interpretation of the commitments of an agent. To see this, consider the simple example in Figure 3.2 and suppose some agent accepts sentences s and t. To what argument are they committed to? It seems plausible for them to accept a_1, as they accept both involved sentences and a_1 defends itself. However, the agent could also claim to accept no argument at all, as they might accept s and t independently of the arguments. This would also give a set of arguments satisfying the conditions for a weak argument set, as the agent would not defend a_1 and thus would not have to accept it. This seems undesirable, as then we cannot tell by a set of accepted sentences which arguments an agent is required to accept. In general, it would dissociate the accepted sentences from the accepted arguments quite a lot and simply claiming to accept no arguments at all will be an allowed move in many cases. It is part of the point of putting forth arguments that they should have some bearing on agents that accept their premises. But if we are not careful with weak argument sets, this point gets somewhat lost. One way out of this is that we require an agent to at least accept the arguments of which they accept *all* sentences. This way, we have some basis from which we can determine which arguments should be accepted.

The problems with weak argument sets are not over yet, though. Figure 3.3 illustrates the next problem. Suppose some agent accepts again s and t. This seems plausible, as there are two minimal arguments for these sentences (a_3 and a_4). However, proceeding as above by requiring the agent to accept all arguments of which they accept all sentences would require them to accept also a_2, meaning they would end up with an undefended argument set. But our agent might accept these sentences based on the arguments a_1

[2]How basic this requirement is heavily depends on the specific incompatibility function we use. Using a logical notion of contrariness (i.e. two sentences together imply a contradiction) can also be extremely demanding on agents as it essentially requires logical omniscience.

$$a_1 : \langle \{s\}, t \rangle \longleftrightarrow a_2 : \langle \{\neg t\}, \neg t \rangle$$

Figure 3.2: In this SBAF, both $\emptyset$ and $\{a_1\}$ could be reasonable interpretations for the argument set of the language extension $\{s, t\}$.

$$\begin{array}{ccccc} & & a_1 : \langle \{\neg n(a_2)\}, \neg n(a_2) \rangle & & \\ & & \downarrow & & \\ a_3 : \langle \{s\}, s \rangle & \dashrightarrow & a_2 : \langle \{s\}, t \rangle & \dashrightarrow & a_4 : \langle \{t\}, t \rangle \end{array}$$

Figure 3.3: This example shows that not all argument of which all sentences and no undercuts are accepted belong to the argument set of a language extension, as otherwise $\{a_3, a_4\}$ would not have an acceptable corresponding language extension.

and a_4, without taking a_2 into account. Since that would be a weakly coherent argument extension, we should also allow that from the language perspective. The solution here is to allow an agent to reject an argument even though they accept all its sentences if they cannot defend it.

Formally defining the corresponding set of arguments to a language extension requires a bit of set-up. Calculating the set of arguments for strong language extensions is fairly straightforward, but the set for weak language extensions requires a fixpoint construction. Thus, we first define a respect function, which is both a version of the characteristic function in abstract argumentation and a formulation of weak sentence-respect.

Definition 34 (Respect Function)**.** *Given an SBAF $\mathcal{SB} = \langle \mathcal{L}, A, \rightarrow, \dashrightarrow \rangle$ and a set of sentences $S \subseteq Sent(A)$, we define its respect function $R^S_{\mathcal{SB}} : 2^A \rightarrow 2^A$ as*

$$R^S_{\mathcal{SB}}(E) := \{a \in A \mid Prem(a) \subseteq S, \overline{n(a)} \cap S = \emptyset \text{ and } E \text{ defends } a\}.$$

We drop the subscript $\mathcal{SB}$ if there reference is clear. It will prove useful to note that $R^S_{\mathcal{SB}}$, like the characteristic function, is monotonic and can preserve admissibility.

Lemma 4. *Let $\mathcal{SB} = \langle \mathcal{L}, A, \rightarrow, \dashrightarrow \rangle$ be an SBAF, $S \subseteq Sent(A)$, and $E, E' \subseteq A$. If $E \subseteq E'$, then $R^S(E) \subseteq R^S(E')$.*

Proof. Take any $E, E' \subseteq A$ such that $E \subseteq E'$. Let $a \in R^S(E)$. Then we know that $Prem(a) \subseteq S$, $\overline{n(a)} \cap S = \emptyset$, and E defends a. We immediately get $Prem(a) \subseteq S$, $\overline{n(a)} \cap S = \emptyset$, and E' defends a, which is sufficient for $a \in R^S(E')$ as desired. □

Lemma 5. *Let $\mathcal{SB} = \langle \mathcal{L}, A, \rightarrow, \dashrightarrow \rangle$ be an SBAF, $S \subseteq Sent(A)$. Then for any admissible extension $E \subseteq A$ such that for all $a \in E$: $Prem(a) \subseteq S$ and $\overline{n(a)} \cap S = \emptyset$, we have that $E \subseteq R^S(E)$ and $R^S(E)$ is admissible.*

Proof. Let E be admissible and $Prem(a) \subseteq S$ and $\overline{n(a)} \cap S = \emptyset$ for all $a \in E$. Take any $a \in E$, then it is clear that $Prem(a) \subseteq S$ and $\overline{n(a)} \cap S = \emptyset$. Also, since E is admissible, it defends a, thus $a \in R^S(E)$ as desired.

Note that $E \subseteq R^S(E)$ gives us that $R^S(E)$ is defended. Further, recall from Lemma 2 that $F_{\mathcal{SB}}(E)$ is conflict-free. Since we clearly have $R^S(E) \subseteq F_{\mathcal{SB}}(E)$, we can conclude that $R^S(E)$ is also conflict-free and thus admissible. □

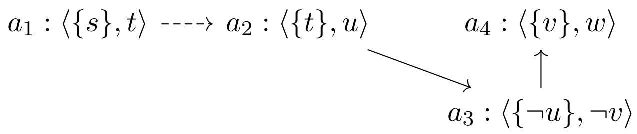

Figure 3.4: An example for argument sets.

Using this function, we can now define the strong and weak argument sets for a language extension.

Definition 35 (Argument Sets). *Given a set of sentences S in $\mathcal{SB} = \langle \mathcal{L}, A, \rightarrow, \dashrightarrow \rangle$, we define its strong argument set $Args_s(S) := \{a \in A \mid Prem(a) \subseteq S$ and $\overline{n(a)} \cap S = \emptyset\}$.*

If S is compatible, then its initial set, $Init(S)$ is defined as the largest admissible subset of $\{a \in A \mid Sent(a) \subseteq S$ and $\overline{n(a)} \cap S = \emptyset\}$.

The weak argument set of a compatible S, $Arg_w(S)$, then is the least fixpoint of $R^S_{\mathcal{SB}}$ that contains $Init(S)$.

Consider the example in Figure 3.4. Suppose we start with the language extension $S = \{s, t, v, w\}$. We can immediately see that its strong argument set is $Args_s(S) = \{a_1, a_2, a_4\}$. While its weak argument set arrives at the same result, it proceeds in steps. Based on the accepted sentences, we should in some sense directly include a_1 and a_4, but since a_4 is undefended, we start with $Init(S) = \{a_1\}$. The first application of the respect function yields $R^S(\{a_1\}) = \{a_1, a_2\}$ and the second then gives us the weak argument set $Arg_w(S) = R^S(\{a_1, a_2\}) = \{a_1, a_2, a_4\}$. To see the differences between the strong and weak argument sets, take the language extension $S' = \{t, \neg u\}$. The strong argument set would commit to the conflicting arguments $Args_s(S') = \{a_2, a_3\}$, since of both all premises and no undercuts are accepted. The weak argument set, however, starts out with $Init(S') = \emptyset$ and ends with $Arg_w(S') = R^{S'}(\emptyset) = \{a_2\}$. Finally, note that S' are somewhat unsatisfactory as they it does not accept all conclusion of the accepted arguments. This means we will have to require language extensions to *respect their arguments*.

Before defining adequate language extensions, note that the weak argument set is only well-defined if the underlying set of sentences is compatible. For instance, in Figure 3.2, there is no (unique) largest admissible subset of $\{a \in A \mid Prem(a) \subseteq \{s, t, \neg t\}$ and $\overline{n(a)} \cap \{s, t, \neg t\} = \emptyset\}$. However, if S is compatible, then $Arg_w(S)$ is well defined. We start by showing that $Init(S)$ is well-defined.

Lemma 6. *Let $\mathcal{SB} = \langle \mathcal{L}, A, \rightarrow, \dashrightarrow \rangle$ be an SBAF and $S \subseteq Sent(A)$ a compatible language extension. Then there exists a unique maximal admissible subset of $\{a \in A \mid Sent(a) \subseteq S$ and $\overline{n(a)} \cap S = \emptyset\}$.*

Proof. We first show that $\{a \in A \mid Sent(a) \subseteq S$ and $\overline{n(a)} \cap S = \emptyset\}$ is conflict-free. Suppose it contains argument a, b such that $a \rightarrow b$. Then either $Conc(a) \in \overline{s}$ for some $s \in Sent(b)$, which contradicts compatibility of S, or $Conc(a) \in \overline{n(b)}$, contradicting that $\overline{n(b)} \cap S = \emptyset$. Thus, the set is conflict-free.

Now we show that the union of all admissible subsets of the set is admissible, thus clearly being its unique maximal admissible subset. Since the whole set is conflict-free, so are all subsets and thus also their union. Further, the union of defended sets is defended, thus the union of all admissible subsets is also defended and hence admissible itself. □

In order to show that the weak argument set is well-defined, it remains to show that there indeed exists a least fixpoint of $R^S_{\mathcal{SB}}$ containing $Init(S)$. The proofs of Lemma 7 and Proposition 11 follow closely Grossi and Modgil (2019).

Lemma 7. *Let $\mathcal{SB} = \langle \mathcal{L}, A, \rightarrow, \dashrightarrow \rangle$ be an SBAF and $S \subseteq Sent(A)$ a compatible language extension. Then there exists a least fixpoint of R^S containing $Init(S)$.*

Proof. By Lemma 3, the set $\{E \subseteq A \mid Init(S) \subseteq E\}$ is a complete lattice. Further, Lemma 4 gives that R^S is a monotonic function. Thus the claim follows from the Knaster-Tarski fixpoint theorem (Fact 1). $\square$

This is enough to show that $Arg_w(S)$ is well-defined for compatible language extensions. However, we can go a step further and explicitly construct the set as $\bigcup_{i \in \mathbb{N}} R_i^S(Init(S))$, where $R_0^S(Init(S)) = Init(S)$ and $R_{i+1}^S(Init(S)) = R^S(R_i^S(Init(S))$. In order to show that, we first need another version of lemma 5.

Lemma 8. *Let $\mathcal{SB} = \langle \mathcal{L}, A, \rightarrow, \dashrightarrow \rangle$ be an SBAF and $S \subseteq Sent(A)$ a compatible language extension. Then for each $i \in \mathbb{N}$, $R_i^S(Init(S)) \subseteq R_{i+1}^S(Init(S))$, and $R_{i+1}^S(Init(S))$ is admissible.*

Proof. We proceed by induction on i.

For the base case, we know that $Init(S)$ is admissible and it is clear by its construction that for all $a \in Init(S)$, we have $Prem(a) \subseteq S$ and $\overline{n(a)} \cap S = \emptyset$. Thus Lemma 5 gives us $Init(S) \subseteq R^S(Init(S))$ and $R^S(Init(S))$ is admissible.

For the induction step, note that Lemma 5 applies to all $R_{i+1}^S(Init(S))$, as $R_i^S(Init(S))$ is admissible by the induction hypothesis and by definition of R^S, we have that for all $a \in R_i^S(Init(S))$, $Prem(a) \subseteq S$ and $\overline{n(a)} \cap S = \emptyset$. $\square$

This now lets us prove the following.

Proposition 11. *Let $\mathcal{SB} = \langle \mathcal{L}, A, \rightarrow, \dashrightarrow \rangle$ be an SBAF and $S \subseteq Sent(A)$ a compatible language extension. Then*

$$Arg_w(S) = \bigcup_{i \in \mathbb{N}} R_i^S(Init(S)).$$

Proof. Note that Lemma 8 gives us that $\bigcup_{i \in \mathbb{N}} R^S(R_i^S(Init(S))) = \bigcup_{i \in \mathbb{N}} R_i^S(Init(S))$. Thus we can prove that $\bigcup_{i \in \mathbb{N}} R_i^S(Init(S))$ is a fixpoint by showing that

$$R^S(\bigcup_{i \in \mathbb{N}} R_i^S(Init(S))) = \bigcup_{i \in \mathbb{N}} R^S(R_i^S(Init(S))).$$

$\subseteq$: Let $a \in R^S(\bigcup_{i \in \mathbb{N}} R_i^S(Init(S)))$. Then $Prem(a) \subseteq S$, $\overline{n(a)} \cap S = \emptyset$, and a is defended by $\bigcup_{i \in \mathbb{N}} R_i^S(Init(S))$. Since, by Lemma 8, $R_i^S(Init(S)) \subseteq R_{i+1}^S(Init(S))$, this gives us some i such that a is defended by $R_i^S(Init(S))$, meaning that $a \in R^S(R_i^S(Init(S)))$ and also $a \in \bigcup_{i \in \mathbb{N}} R^S(R_i^S(Init(S)))$ as desired.

$\supseteq$: Let $a \in \bigcup_{i \in \mathbb{N}} R^S(R_i^S(Init(S)))$. Again, by Lemma 8, this gives us some i such that $a \in R^S(R_i^S(Init(S)))$, meaning that $Prem(a) \subseteq S$, $\overline{n(a)} \cap S = \emptyset$, and a is defended by $R_i^S(Init(S))$. The latter gives us that a is defended by $\bigcup_{i \in \mathbb{N}} R_i^S(Init(S))$ and in sum we have $a \in R^S(\bigcup_{i \in \mathbb{N}} R_i^S(Init(S)))$.

It remains to show that $\bigcup_{i \in \mathbb{N}} R_i^S(Init(S))$ is the least fixpoint containing $Init(S)$. Assume for a contradiction that there is some fixpoint E such that $E \subsetneq \bigcup_{i \in \mathbb{N}} R_i^S(Init(S))$. By Lemma 8, this gives us some i such that $E \subsetneq R_i^S(Init(S))$. Note that we also have $Init(S) \subsetneq E$, since otherwise $Init(S)$ would itself be the least fixpoint and we would have $Init(S) = \bigcup_{i \in \mathbb{N}} R_i^S(Init(S))$. But then, by i applications of Lemma 4, $R_i^S(Init(S)) \subseteq R_i^S(E) = E \subsetneq R_i^S(Init(S))$, a contradiction.

Thus, we conclude that $\bigcup_{i \in \mathbb{N}} R_i^S(Init(S))$ is the least fixpoint containing $Init(S)$.

$\square$

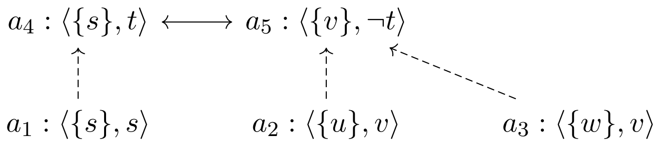

Figure 3.5: The SBAF from Figure 3.1

Corollary 2. *Let $\mathcal{SB} = \langle \mathcal{L}, A, \rightarrow, \dashrightarrow \rangle$ be an SBAF and $S \subseteq Sent(A)$ a compatible language extension. Then $Arg_w(S)$ is admissible.*

Given a language extension and its corresponding set of arguments, we can then add further conditions for it to be rationally acceptable. First, we should still require defence for the set of arguments. We argued above that defending your arguments is required even if we allow for doubt, thus we should also keep it here. Second, we should make sure that accepting an argument has consequences. Namely, we would not want someone to accept an argument, based on some sentences they accept, without concluding anything from it. Accordingly, we require that all sentences of an accepted argument are accepted.

Definition 36. *A strongly adequate language extension in an SBAF $\mathcal{SB} = \langle \mathcal{L}, A, \rightarrow, \dashrightarrow \rangle$ is a compatible set of sentences $S \subseteq Sent(A)$ such that $Arg_s(S)$ is defended and it satisfies:*

Argument-Respect: $\forall a \in Arg_s(S) : Sent(a) \subseteq S$.

A weakly adequate language extension $S \subseteq Sent(A)$ is a compatible language extension that satisfies argument-respect w.r.t. $Arg_w(S)$.

As with argument extensions, we will sometimes only talk about strongly or weakly adequate extensions if it is clear that language extensions are meant. We further use the shorthand *adequate semantics* to express the idea that strongly or weakly adequate extensions are the rationally acceptable ones. Also note that, by Corollary 2, Arg_w is always admissible, so we do not have to require defence explicitly as with strongly adequate extensions. Finally, Proposition 14 will show that all strongly adequate extensions are also weakly adequate.

Let us go back to the SBAF of Figure 3.1, seen here again as Figure 3.5, to illustrate language extensions. The extension $S = \{s, u, v, w\}$ is weakly adequate, but not strongly adequate. We have $Arg_w(S) = \{a_1, a_2, a_3\}$, while $Arg_s(S) = \{a_1, a_2, a_3, a_4, a_5\}$. With the latter, argument-respect is clearly not satisfied. In general, we have the following strongly adequate extensions: $\emptyset$, $\{t\}$, $\{\neg t\}$, $\{s, t\}$, $\{v, \neg t\}$, $\{u, v, \neg t\}$, $\{w, v, \neg t\}$, $\{u, v, w, \neg t\}$. Note that while some extensions have the same strong argument set (e.g. both $\{t\}$ and $\{\neg t\}$ accept no arguments), all argument sets are strongly coherent. For weakly adequate extensions, all compatible sets of sentences are weakly adequate, except if u or w occur without v. Again, all weak argument sets are weakly coherent.

Now we have two perspectives on the range of rationally acceptable positions. We can either view it through arguments or directly through sentences. The obvious question is how they relate, and indeed strongly coherent argument extensions correspond to strongly adequate language extensions and the same goes for weak extensions. However, as seen in the above example, the correspondence is not one-to-one. Nevertheless, this correspondence guarantees us that argument extensions are also reasonable when considered from the language perspective. This can be seen as satisfying rationality postulates such as *direct consistency*.

$$a_1 : \langle\{s\}, t\rangle \qquad a_2 : \langle\{\neg s\}, t\rangle$$

Figure 3.6: While $\{a_1, a_2\}$ is strongly coherent, its set of sentences $\{s, \neg s, t\}$ is not compatible and thus not strongly adequate.

Further, in general, it works only for saturated frameworks. Figure 3.6 shows a simple example where the correspondence fails in an unsaturated framework. The underlying reason that there is no correspondence is unsaturated SBAFs is that argumentation frameworks only recognise relations between arguments, but the compatibility condition for language extensions considers relations between sentences. Thus, to make them correspond, we need to import sufficient information on the relations between sentences into the argument level of SBAFs. In practice, saturated SBAFs contain that information.

Proposition 12. *Let $\mathcal{SB} = \langle \mathcal{L}, A, \rightarrow, \dashrightarrow \rangle$ be a saturated SBAF. Then for every strongly adequate language extension S, its strong argument set $Arg_s(S)$ is a strongly coherent argument extension.*

Also, for every strongly coherent argument extension E, its set of sentences $Sent(E)$ is a strongly adequate language extension.

Proof. Let S be a strongly adequate language extension. We check all conditions for strong coherence of $Arg_s(S)$.

Conflict-Free: Suppose there are arguments $a, b \in Arg_s(S)$ such that $a \rightarrow b$. Since $b \in Arg_s(S)$, we know that $\overline{n(b)} \cap S = \emptyset$, hence, by argument-respect, $Conc(a) \notin \overline{n(b)}$. Thus we know that $Conc(a) \in Sent(b)$, but then argument-respect leads to a violation of compatibility of S. Thus, $Arg_s(S)$ is conflict-free.

Defence: This is given by definition.

Strong Sentence-Respect: First, we show that $Sent(Arg_s(S)) \subseteq S$. So first take any $s \in Sent(Arg_s(S))$. Then there exists some $a \in Arg_s(S)$ such that $s \in Sent(a)$. By argument-respect, we have $Sent(a) \subseteq S$, meaning that $s \in S$ as desired. Now suppose that for some $a \in A$, we have $Prem(a) \subseteq Sent(Arg_s(S))$ and $\overline{n(a)} \cap Sent(Arg_s(S)) = \emptyset$. Since $Sent(Arg_s(S)) \subseteq S$, we directly have $Prem(a) \subseteq S$. It remains to show that $\overline{n(a)} \cap S = \emptyset$. Suppose for a contradiction that there is some $t \in \overline{n(a)} \cap S$. By saturatedness of $\mathcal{SB}$, there is a minimal argument $b \in A$ for t. Since $t \in S$, we also know that $b \in Arg_s(S)$. But then $t \in Sent(Arg_s(S))$, contradicting that $\overline{n(a)} \cap Sent(Arg_s(S)) = \emptyset$. Thus, we conclude that $\overline{n(a)} \cap S = \emptyset$ and hence that $a \in Arg_s(S)$ as desired.

Now let E be a strongly coherent argument extension. We check all conditions for strong adequacy of $Sent(E)$.

Compatibility: Suppose there are $s, t \in Sent(E)$ such that $s \in \bar{t}$. Then there are argument $a, b \in E$ such that $s \in Sent(a)$ and $t \in Sent(b)$. Further, by saturatedness of $\mathcal{SB}$, there is a minimal argument c for either s or t. Since we assume for minimal arguments that $\overline{n(c)} = \emptyset$, strong statement-respect gives us $c \in E$. But then either $c \rightarrow a$ or $c \rightarrow b$, contradicting conflict-freeness of E. Thus we conclude that $Sent(E)$ is compatible.

Defence: We show that $Arg_s(Sent(E)) = E$, from which defence follows directly. Thus take $a \in Arg_s(Sent(E))$. Then we know that $Prem(a) \subseteq Sent(E)$ and $\overline{n(a)} \cap$

$Sent(E) = \emptyset$. Strong sentence-respect then gives $a \in E$ as desired. Now take $a \in E$. Then we know that $Prem(a) \subseteq Sent(E)$. We need to show that $\overline{n(a)} \cap Sent(E) = \emptyset$ in order to get $a \in Arg_s(Sent(E))$. Thus assume there is some $t \in \overline{n(a)} \cap Sent(E)$. By saturatedness of $\mathcal{SB}$, we get a minimal argument b for t. Strong sentence-respect gives $b \in E$, but since $b \rightarrow a$, this contradicts conflict-freeness of E. We conclude that $\overline{n(a)} \cap Sent(E) = \emptyset$ and that $a \in Arg_s(Sent(E))$ as desired. In sum, $Arg_s(Sent(E)) = E$.

Argument-Respect: Recall that $Arg_s(Sent(E)) \subseteq E$. Thus for any $a \in Arg_s(Sent(E))$, we know that $a \in E$ and also that $Sent(a) \subseteq Sent(E)$, meaning that $Sent(E)$ satisfies argument-respect.

□

Note that we do not need saturatedness of $\mathcal{SB}$ for conflict-freeness of $Arg_s(S)$. We can also prove the same correspondence for weak extensions, but we need a few lemmas. First, we show that the conditions for weak sentence-respect can be reduced in saturated SBAFs if we know that the extension in question is admissible.

Lemma 9. *Let $\mathcal{SB} = \langle \mathcal{L}, A, \rightarrow, \dashrightarrow \rangle$ be a saturated SBAF, E an admissible extension, and $a \in A$ any argument. If $Prem(a) \subseteq Sent(E)$ and E defends a, then $\overline{n(a)} \cap Sent(E) = \emptyset$.*

Proof. Suppose that $Prem(a) \subseteq Sent(E)$ and E defends a. Further, suppose for a contradiction that there is some $t \in \overline{n(a)} \cap Sent(E)$. By saturatedness of $\mathcal{SB}$, there is a minimal argument b for t. Since E defends a, we have $E \rightarrow b$. This means there is some argument $c \in E$ with $Conc(c) \in \bar{t}$. But we also have that $t \in Sent(E)$, so there is an argument $d \in E$ with $t \in Sent(d)$. But then $c \rightarrow d$, contradicting conflict-freeness of E. □

Next, we show that in saturated SBAFs, weakly coherent argument extensions are fixpoints of the respect function. The restriction to saturated frameworks is required, since otherwise there could be undercutting sentences that only occur as premises and never as a conclusion. In such a case, the impact of such sentences on the acceptability of the arguments they undercut would not be visible on the argument level.

Lemma 10. *Let $\mathcal{SB} = \langle \mathcal{L}, A, \rightarrow, \dashrightarrow \rangle$ be a saturated SBAF. Then for any weakly coherent argument extension E, we have $R^{Sent(E)}(E) = E$.*

Proof. Weak sentence-respect of E gives us directly that $R^{Sent(E)}(E) \subseteq E$. For the other direction, take any $a \in E$. Then $Prem(a) \subseteq Sent(E)$ and E defends a, thus by Lemma 9, we have that $\overline{n(a)} \cap Sent(E) = \emptyset$, meaning that $a \in R^{Sent(E)}(E)$ as desired. □

Finally, we need a lemma showing that for weakly adequate language extensions, their initial set of arguments is already their weak set of arguments. This is essentially due to the condition of argument-respect.

Lemma 11. *Let $\mathcal{SB} = \langle \mathcal{L}, A, \rightarrow, \dashrightarrow \rangle$ be a SBAF and S a weakly adequate language extension. Then $Arg_w(S) = Init(S)$.*

Proof. We know by definition that $Init(S) \subseteq Arg_w(S)$. For the other direction, we show that $Arg_w(S) \subseteq \{a \in A \mid Sent(a) \subseteq S \text{ and } \overline{n(a)} \cap S = \emptyset\}$. Take any $a \in Arg_w(S)$. Then by argument-respect, $Sent(a) \subseteq S$. Now suppose there is some $t \in \overline{n(a)} \cap S$. Then certainly $a \notin Init(S)$. Recall that $Arg_w(S) = \bigcup_{i \in \mathbb{N}} R_i^S(Init(S))$. Thus there is some i such that $a \notin R_i^S(Init(S))$, but $a \in R_{i+1}^S(Init(S))$. However, $R_{i+1}^S = R^S(R_i^S(Init(S))) = \{a \in A \mid Prem(a) \subseteq S, \overline{n(a)} \cap S = \emptyset, \text{ and } R_i^S(Init(S)) \text{ defends } a\}$. Since the middle condition of the last set is not given for a, we have $a \notin R^S_{i+1}(Init(S))$, a contradiction. Thus we conclude that $\overline{n(a)} \cap S = \emptyset$, meaning that $a \in \{a \in A \mid Sent(a) \subseteq S \text{ and } \overline{n(a)} \cap S = \emptyset\}$. Thus also $Arg_w(S) \subseteq \{a \in A \mid Sent(a) \subseteq S \text{ and } \overline{n(a)} \cap S = \emptyset\}$, and since $Arg_w(S)$ is admissible (Lemma 8), we have $Arg_w(S) \subseteq Init(S)$ as desired. □

With all these lemmas in place, we are ready to prove the correspondence between weakly coherent argument extensions and weakly adequate language extensions.

Proposition 13. *Let $\mathcal{SB} = \langle \mathcal{L}, A, \rightarrow, \dashrightarrow \rangle$ be a saturated SBAF. Then for every weakly adequate language extension S, its weak argument set $Arg_w(S)$ is a weakly coherent argument extension.*

Also, for every weakly coherent argument extension E, its set of sentences $Sent(E)$ is a weakly adequate language extension.

Proof. Let S be a weakly adequate language extension. We check all conditions of $Arg_w(S)$.

Conflict-Free: Follows from Corollary 2.

Defence: Follows from Corollary 2.

Weak Sentence-Respect: Assume for some $a \in A$ that $Prem(a) \subseteq Sent(Arg_w(S))$, $\overline{n(a)} \cap Sent(Arg_w(S)) = \emptyset$, and $Arg_w(S)$ defends a. We need to show that $a \in Arg_w(S)$. First note that $Sent(Arg_w(S)) \subseteq S$, as Lemma 11 shows that $Arg_w(S) = Init(S)$ and all sentences occuring in arguments of $Init(S)$ are already contained in S. Thus we have $Prem(a) \subseteq S$. Further, we can show that $\overline{n(a)} \cap S = \emptyset$. Suppose for a contradiction that there is some $t \in \overline{n(a)} \cap S$. By saturatedness of $\mathcal{SB}$, there is a minimal argument b for t, for which we know that $b \rightarrow a$. Since $Arg_w(S)$ defends a, we know that there is some $c \in Arg_w(S)$ such that $c \rightarrow b$. By b being a minimal argument, we know that $Conc(c) \in \bar{t}$, and we further know by argument-respect that $Conc(c) \in S$. But this contradicts compatibility of S, since also $t \in S$. Thus we conclude that $\overline{n(a)} \cap S = \emptyset$. Now we can use that $Arg_w(S)$ is a fixpoint of R^S to conclude that, since $Arg_w(S)$ also defends a, we have $a \in Arg_w(S)$ as desired.

Now let E be a weakly coherent argument extension. We check all conditions for $Sent(E)$.

Compatibility: Suppose there are $s, t \in Sent(E)$ such that $s \in \bar{t}$. Then we have arguments $a, b \in E$ such that $s \in state(a)$ and $t \in state(b)$. Further, by saturatedness of $\mathcal{SB}$, there is w.l.o.g. a minimal argument c for s (note that $c \rightarrow b$). For Proposition 12, we could use the facts that $Prem(c) \subseteq Sent(E)$ and $\overline{n(c)} = \emptyset$ to conclude that $c \in E$. Here, we first need to additionally show that E defends c. Note that since $s \in Sent(a)$, any attack on c is also an attack on a. Since E defends a, E thus also defends c. Weak sentence-respect then gives $c \in E$. But since $c \rightarrow b$, this contradicts conflict-freeness of E. Thus we conclude that $Sent(E)$ is compatible.

Defence: We show that $E = Arg_w(Sent(E))$, from which defence follows directly. By Lemma 10, we know that $R^{Sent(E)} = E$, thus it suffices to show that $E = Init(S)$ (since then $Init(S)$ will itself be the smallest fixpoint containing it).

⊆: Take any $a \in E$. Then we have $Sent(a) \subseteq Sent(E)$. Further, we know that E defends a and $\mathcal{SB}$ is saturated, thus by Lemma 9, we know that $\overline{n(a)} \cap Sent(E) = \emptyset$. This gives us that $a \in \{a \in A \mid Sent(a) \subseteq Sent(E) \text{ and } \overline{n(a)} \cap Sent(E) = \emptyset\}$. But note that we have just now shown that $E \subseteq \{a \in A \mid Sent(a) \subseteq$

$Sent(E)$ and $\overline{n(a)} \cap Sent(E) = \emptyset\}$ and E being admissible, we can directly infer that $E \subseteq Init(Sent(E))$, as the latter is the largest admissible subset of $\{a \in A \mid Sent(a) \subseteq Sent(E)$ and $\overline{n(a)} \cap Sent(E) = \emptyset\}$.

$\supseteq$: Take any $a \in Init(Sent(E))$. Then we have that $Sent(a) \subseteq Sent(E)$ and $\overline{n(a)} \cap Sent(E) = \emptyset$. If we can show that E defends a, then weak sentence-respect gives us the desired $a \in E$. Thus take any attacker b of a. There are two cases. (1) $Conc(b) \in \overline{s}$ for some $s \in Sent(a)$. Since $Sent(a) \subseteq Sent(E)$, there is some argument $c \in E$ such that $s \in Sent(c)$. But then $b \rightarrow c$, and since E defends c, we also have $E \rightarrow b$. That is, E defends a against b. (2) $Conc(b) \in \overline{n(a)}$. By saturatedness of $\mathcal{SB}$, there is a minimal argument c for $Conc(b)$. Note that $c \rightarrow a$ and since a is defended by $Init(Sent(E))$, there is some argument $d \in Init(Sent(E))$ such that $d \rightarrow c$. This gives us in particular $Conc(d) \in Sent(E)$ and since c is a minimal argument, we also have $Conc(d) \in \overline{Conc(b)}$. Further, $Conc(d) \in Sent(E)$ gives us some argument $e \in E$ such that $Conc(d) \in Sent(e)$. Recall that incompatibility is symmetric, thus $b \rightarrow e$. Finally, since E defends e, we have $E \rightarrow b$, that is, E defends a against b.

In sum, E defends a and weak sentence-respect gives us $a \in E$ as desired.

Argument-Respect: Take any $a \in Arg_w(Sent(E))$. Since $Arg_w(Sent(E)) \subseteq E$, we have $a \in E$ and thus $Sent(a) \subseteq Sent(E))$ as desired.

□

Now we can show that, indeed, strongly adequate language extensions are also weakly adequate.

Proposition 14. *In an SBAF, strongly adequate language extensions are also weakly adequate.*

Proof. Let S be a strongly adequate language extension. We show that $Arg_s(S) = Arg_w(S)$. It is immediate that $Init(S) \subseteq Arg_s(S)$. Further, since R^S is monotonic (Lemma 4) and $R^S(Arg_s(S)) \subseteq Arg_s(S)$ (strong sentence-respect) we also have $Arg_w(S) \subseteq Arg_s(S)$. Now consider any $a \in Arg_s(S)$. Then by argument-respect, we have $Sent(a) \subseteq S$ and by the conditions of $Arg_s(S)$ also $\overline{n(a)} \cap S = \emptyset$. Thus $a \in \{a \in A \mid Sent(a) \subseteq S$ and $\overline{n(a)} \cap S = \emptyset\}$. This gives $Arg_s(S) \subseteq \{a \in A \mid Sent(a) \subseteq S$ and $\overline{n(a)} \cap S = \emptyset\}$ and since $Arg_s(S)$ is admissible (part of proof of Proposition 13 which does not depend on saturatedness), we further have $Arg_s(S) \subseteq Init(S) \subseteq Arg_w(S)$. □

Figure 3.7 gives an overview over the different semantics for SBAFs.

3.1.3 Discussion

In Section 1.1, we saw that doubt is a central idea that allows informal argumentation theory to capture a rich notion of support between arguments and in Section 2.2, we argued that existing formal approaches to argumentation have no straightforward way of incorporating doubt. Given its importance, it might be surprising that neither coherent nor adequate semantics explicitly employ a notion of doubt. This mainly due to doubt being an agent-centric notion, whereas coherence and adequacy just talk about sets of arguments and sentences respectively. But if we consider the situation from an agent's perspective, we can see that they are allowed to have *mere doubt*. Our semantics do not come with a completeness requirement on acceptable extensions, thus our agents are not required to accept any unattacked arguments or sentences. If they merely doubt something

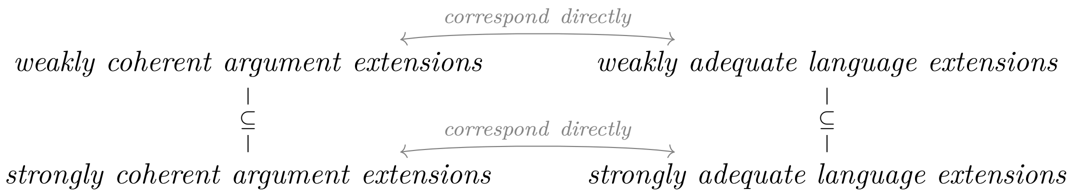

Figure 3.7: Overview on SBAF semantics. The set of strongly coherent argument extensions (resp. strongly adequate language extensions) is contained in the set of weakly coherent argument extensions (resp. weakly adequate language extensions). Argument extensions and language extensions correspond directly in that the set of sentences of a strongly (resp. weakly) coherent argument extension is a strongly (resp. weakly) coherent language extension. Further, the argument set of a strongly (resp. weakly) adequate language extension is a strongly (resp. weakly) coherent argument extension.

without having a direct counter-argument, they are generally allowed not to accept it. Of course, there are limits to doubt, otherwise there would not be much point in arguing at all. Hence, agents are not allowed to doubt conclusions of accepted arguments. This is the way coherent and adequate semantics incorporate doubt while retaining interesting argumentative constraints on acceptable extensions.

The requirement that conclusions of accepted arguments have to be accepted as well clearly relies on structured arguments. And with structured arguments there comes a language that allows us to specify the components of the arguments. This allows us to ask both "Which argument should one accept?" and "Which sentences should one accept?". That is, we can evaluate argumentation on both the perspective of the *argument level* and that of the *language level* and our approach to evaluating SBAFs accordingly uses both coherent argument extensions and adequate language extensions with equal importance. The correspondence results between coherent and adequate extensions then confirm that the resulting extensions make sense from both perspectives. While the two perspectives have some precedent in the literature around ABA (see also Baroni et al. 2018b), most approaches to structured argumentation almost exclusively deal with the argument level. This is somewhat surprising as, in some sense, argumentation is essentially about the language level: we argue in order to convince people of a specific claim or standpoint. In actual debates, we are also primarily exposed to sentences, whereas we often have to put extra effort into reconstructing arguments. It can thus be easier for agents to determine their accepted set of sentences and compare it to the language extensions of some semantics, than to determine the set of arguments they accept. In that sense, it is more natural to think about the beliefs or the position of an agent in terms of the sentences they accept.

Given a set of accepted sentences, it is not always easy to determine which arguments the agent should accept. We have seen how particularly weak argument sets are difficult to determine. In this respect, Corollary 2 is quite interesting. It says that for any not outright contradictory set of sentences, we can find a corresponding set of arguments that is admissible. This very much implements an attempt at making the agents out to be reasonable. But it is interesting that it is almost always possible to find a, at least somewhat, reasonable interpretation for a set of sentences. Of course, argument-respect might still fail, so agents that accept some sentences should still be careful, but it shows that they have quite a bit of leeway. This is in stark contrast with strong argument sets, where it is very clear for any set of sentences which arguments correspond to it.

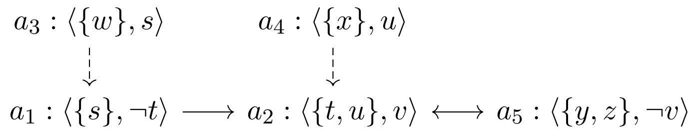

Figure 3.8: Extended example of Alexa, the meteorologist.

Strong argument sets will often be conflicting, even if the underlying set of sentences is compatible. This raises the question: How can we decide whether to use strong or weak argument sets (resp. strongly or weakly coherent argument extensions)?

Recall the example of Alexa, the meteorologist. She claimed that it will rain next week, but there was an objection claiming she is not in fact a certified meteorologist. In order to illustrate the differences between strong and weak argument sets, we can add another meteorologist, Susan, claiming that it will not rain next week. Figure 3.8 gives the example and here are the translations:

s: Alexa never finished her degree in meteorology.
t: Alexa is a certified meteorologist.
u: Alexa claims it will rain next week.
v: It will rain next week.
w: Alexa did not pass her final exams.
x: Alexa claimed it will rain next week in an interview with a local newspaper.
y: Susan is a certified meteorologist.
z: Susan claims it will not rain next week.

We have seen that weak argument sets are somewhat more generous to the agents than strong sets, as more effort is put in into interpreting the sentences they accept as an admissible argument extensions. Take, for instance, an agent that accepts t, u, y, and z while doubting everything else. That is, they accept that both Alexa and Susan are certified meteorologists and that they made contradictory claims about whether it will rain next week. If we take the strong argument set, then our agent is committed to both a_2 and a_5 and thus has to accept contradictory arguments. This is because without an explicit undercut, strong argument sets assume the arguments to work. Taking the weak argument set, in contrast, interprets the agent as accepting no argument at all, which makes the agent's position weakly adequate.

In case of two contradicting experts, it seems reasonable to accept neither argument and not take a stance as to whether it will rain next week or not. Taking weakly adequate language extensions allows for that possibility. However, there is still some tension there, as the set of accepted sentences strongly suggest contradictory claims. This tension can only be resolved by finding an undercut for one of the arguments or by rejecting a premise. Once this is done, we reached a strongly adequate language extension. In that sense, strongly adequate extensions are more settled and have less tension in them.[3]

A somewhat stranger case for weakly adequate extensions is when our agent still accepts t, u, y, and z, but additionally accepts $\neg v$ as well. The weak argument set now contains a_5, which makes the extension still weakly adequate. The inference of argument a_2 is assumed to fail without providing an explicit undercut. Rather, the reasoning allowed by the weak argument set is of the form "I accept all premises of this argument,

[3]The tension between rebutting arguments could also be resolved by introducing a preference ranking between arguments, representing their relative strengths. This could then be used to determine which of the rebutting arguments is successful (cf. Amgoud and Cayrol 1998; Amgoud and Vesic 2011; Kaci et al. 2018).

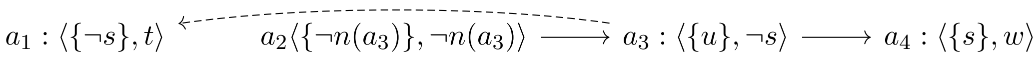

Figure 3.9: $\{a_1, a_2, a_4\}$ would be strongly and weakly coherent, but its set of sentences $\{s, \neg s, t, w, \neg n(a_3)\}$ is not compatible.

but I cannot defend it against attacks, hence I conclude that something in its inference must have gone wrong." This might work well as heuristic reasoning on how to continue the debate, but it is dialectically unsatisfactory. It already assumes that the agent has a reasonable position in the debate and uses that to conclude that some argument does not need to be accepted. This allows the protagonist of a standpoint to essentially ignore any direct rebut put forth by the antagonist, even if the rebutting argument is based on shared premises.

Interestingly, the weak argument set and weakly coherent argument extensions are closer to standard Dung semantics than their strong counterparts. We have already seen that complete extensions are weakly, but not strongly coherent (Observation 2). In Section 3.3, we will further see that a certain type of weakly coherent extensions correspond to preferred extensions.

There is much further work to be done in examining and understanding coherent argument extensions and adequate language extensions. One open question is whether the saturatedness requirement on SBAFs can be relaxed. It is clear that some form of saturatedness is required, as otherwise argument extensions can accept incompatible premises. Further, the following obvious relaxation does not work. It might be thought that it is enough to require that from a pair of incompatible sentences, at least one needs to occur as a conclusion of an argument. However, Figure 3.9 presents a counterexample to that claim.

Another open question is whether it is possible to formulate the defence requirement on adequate language extensions on the level of sentences. While coherent argument extensions can be mostly defined on the level of arguments, adequate language extensions import a lot of the argument level through the defence requirement. It would be interesting to see whether it is possible to formulate adequacy for language extensions more independently.

3.2 Support in Presence of Doubt

The previous section did not contain much discussion of the support relation. Indeed, while support never featured explicitly, it was there implicitly, namely in the notion of sentence-respect. The effect of this condition (in its weak and strong version) is that we get the following versions of support between (sets of) arguments:

Definition 37 (Collective Support)**.** *Let $\mathcal{SB} = \langle \mathcal{L}, A, \rightarrow, \dashrightarrow \rangle$ be an SBAF. A set of arguments $E \subseteq A$ is said to collectively support an argument $a \in A$ if $Prem(a) \subseteq Sent(E)$.*

Strong and weak sentence-respect now give us different conditions under which collective support is sufficient in the sense that any extension that collectively supports an argument also contains the argument. In essence, they are just a reformulation of sentence-respect in the language of support.

Observation 4. *Let $\mathcal{SB} = \langle \mathcal{L}, A, \rightarrow, \dashrightarrow \rangle$ be an SBAF and E a strongly coherent extension that collectively supports $a \in A$. If $\overline{n(a)} \cap Sent(E) = \emptyset$, then $a \in E$.*

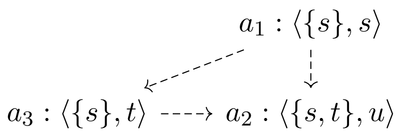

Figure 3.10: An argument (a_1) can sufficiently support another (a_2) without directly supporting all premises.

Observation 5. *Let $\mathcal{SB} = \langle \mathcal{L}, A, \rightarrow, \dashrightarrow \rangle$ be an SBAF and E a weakly coherent extension that collectively supports $a \in A$. If $\overline{n(a)} \cap Sent(E) = \emptyset$ and E defends a, then $a \in E$.*

It is worth noting that these observations do not capture all sufficient support there is. For instance, it only takes direct support into account, whereas indirect support can sometimes also be relevant. The latter would be a case shown in Figure 3.10, where supporting a supporting argument also leads to sufficient support. Further, somewhat counter-intuitively, collective support does not need to be based on the support relation. For instance, in the same figure, $\{a_2\}$ collectively supports a_1 and a_2.

Nevertheless, what these reformulations suggest, is that in some circumstances, strongly coherent extensions follow the notion of deductive support (see Section 2.1.1). While deductive support, contrary to collective support, is a binary relation, the two notions can coincide if all arguments have exactly one premise. Further, if we assume that there are no undercutting sentences in play, support in strongly coherent extensions starts to be come deductive in the following sense:

Proposition 15. *Let $\mathcal{SB} = \langle \mathcal{L}, A, \rightarrow, \dashrightarrow \rangle$ be an SBAF where $\forall a \in A : |Prem(a)| = 1$ and $\neg\exists t \in Sent(A), \exists a \in A : t \in \overline{n(a)}$. Then an extension $E \subseteq A$ that is strongly coherent is also d-admissible.*

If further $\forall a, b \in A$, we have $Prem(a) \neq Prem(b)$, then an extension $E \subseteq A$ that is d-admissible is also strongly coherent.

Proof. Let E be strongly coherent. We first check closure under $\dashrightarrow$. Suppose there is some $a \in E$ and $b \in A$ such that $a \dashrightarrow b$. Since $|Prem(a)| = 1$, we know that $Prem(a) \subseteq Sent(E)$ and since there are no undercuts, we know that $\overline{n(a)} \cap Sent(E) = \emptyset$. Thus, by strong sentence-respect, we have $b \in E$ and E is closed under $\dashrightarrow$.

Now we show inductively that E is admissible according to $\rightarrow^{ded}$. We know by definition that E is admissible w.r.t. $\rightarrow$, hence it remains to show that E is admissible w.r.t. $\rightarrow^{i+1}$, assuming it is admissible w.r.t. $\rightarrow^{i}$. We show defence first. Take any $a \in E$ such that there exists some $b \in A$ with $b \rightarrow^{i+1} a$. There are three cases: (i) if also $b \rightarrow^{i} a$, then we have $E \rightarrow^{i} b$ by assumption. (ii) if b supported attacks a, then there exists $c \in A$ such that $b \dashrightarrow c$ and $c \rightarrow^{i} a$. By assumption, we have $E \rightarrow^{i} c$ and thus E mediated attacks b, i.e. $E \rightarrow^{i+1} b$. (iii) if b mediated attacks a, then there exists $c \in A$ such that $a \dashrightarrow c$ and $b \rightarrow^{i} c$. By closure under $\dashrightarrow$, we have $c \in E$ and by assumption $E \rightarrow^{i} b$, thus also $E \rightarrow^{i+1} b$. In sum, E is defended. For conflict-freeness, take any $a, b \in E$. We know that $a \not\rightarrow^{i} b$ by assumption. Further, since both arguments are in E and E is closed under $\dashrightarrow$, both a supported or a mediated attack from a to b would contradict conflict-freeness of E under $\rightarrow^{i}$. In sum, E is admissible for each $\rightarrow^{i}$ and thus also for $\rightarrow^{ded}$. Together with closure under $\dashrightarrow$, this means that E is d-admissible.

Now assume that $\forall a, b \in A$, we have $Prem(a) \neq Prem(b)$ and let E be d-admissible. It is clear that E is admissible in $\mathcal{SB}$, as there are only fewer attacks to consider. It remains to check strong sentence-respect. Thus take any $a \in A$ such that $Prem(a) \in Sent(E)$

(recall that there are no undercuts). Our assumption guarantees that then there is some $b \in E$ such that $b \dashrightarrow a$, and by closure under $\dashrightarrow$, we get $a \in E$ as desired. □

The second direction indeed requires the extra condition. If we just have two arguments, $a_1 : \langle\{s\}, t\rangle$ and $a_2 : \langle\{s\}, u\rangle$, it would be d-admissible to just accept a_1. However, it would not be strongly coherent, because it would commit to accepting s and thus all premises (and no undercuts) of a_2. Hence, strongly coherent extensions would have to be either empty or contain both arguments. Interestingly, this shows that strongly (and weakly) coherent extensions differ from admissible extensions even in the absence of supports.

3.2.1 Support Principles

In comparison to other support semantics, it is instructive to consider principles that have been used to characterise different notions of support. Most of the presentation here follows a selection of the principles in Yu et al. (2023). The principles have been developed for abstract BAFs and consider changes in acceptable extensions when some supports are added or removed. For instance, there are *support removal robustness* principles, stating that acceptable extensions should not change if a support is removed. Such principles cannot be directly studied in SBAFs, since we are not free to change the support relation without changing the structure of the arguments. However, by reinterpreting the principles slightly, they can still give interesting characterisation of support in SBAFs.

Instead of changing the support relation, we will adapt the premises of an argument. In the process of arguing, it is common that not all premises are always made explicit immediately. Rather, we start with some first intuition that some sentences justify some other sentence. To use our original example, we might start to make an argument simply by saying that it will rain next week, because Alexa says so. Only later on we might notice that we relied on a further premise in our argument, namely that Alexa is a certified meteorologist. Thus, we add this premise to our argument. Similarly, we might realise at some point in a discussion that we used a redundant premise somewhere. Perhaps we got a bit carried away and claimed that Alexa is very popular amongst her colleagues in further support of her claim. But since this seems rather irrelevant as to whether she is correct in claiming that it will rain next week, we might retract that claim as a premise of our argument.

In that sense, changing the premises of arguments is a very common dynamic in practice (see also Pandžić 2022) and it is interesting to see how some semantics react to such changes. Further, comparing arguments with added or removed premises can give us a general insight into the relevance of the number of premises in an argument. We have seen above that it is at least somewhat relevant, as if there is always only one premise, we can end up with deductive support.

For the reformulation of the principles, we use the following notation for an argument a and a sentence s: $a\backslash\{s\} := \langle Prem(a)\backslash\{s\}, Conc(a)\rangle$ if $Prem(a) \neq \{s\}$ and $\langle\{Conc(a)\}, Conc(a)\rangle$ otherwise. $a \cup \{s\} := \langle Prem(a) \cup \{s\}, Conc(a)\rangle$. In the following, we assume a fixed SBAF $\mathcal{SB} = \langle\mathcal{L}, A, \rightarrow, \dashrightarrow\rangle$ and arguments $a, b, c \in A$.

Transitivity: Suppose $a \dashrightarrow b$ and $b \dashrightarrow c$. An argument extension E is strongly/weakly coherent in $\mathcal{SB}$ iff it is strongly/weakly coherent in $\mathcal{SB}' = \langle\mathcal{L}, A', \rightarrow', \dashrightarrow'\rangle$ where $A' = (A\backslash\{c\}) \cup \{c \cup \{Conc(a)\}\}$.

(Admissible) Extension Selection: If an argument extension E is strongly/weakly coherent in $\mathcal{SB}$, then it is also admissible.

Support Removal Robustness: Suppose $a \dashrightarrow b$. If an argument extension E is strongly/weakly coherent in $\mathcal{SB}$, then it remains so in $\mathcal{SB}' = \langle \mathcal{L}, A', \rightarrow', \dashrightarrow' \rangle$ where $A' = (A \backslash \{b\}) \cup \{b \backslash \{Conc(a)\}\}$.

Monotony of Status: If there exists a strongly/weakly coherent extension E containing b in $\mathcal{SB}$, then there is also a stongly/weakly coherent extension E in $\mathcal{SB}' = \langle \mathcal{L}, A', \rightarrow' , \dashrightarrow' \rangle$ where $A' = (A \backslash \{b\}) \cup \{b \cup \{Conc(a)\}\}$.

Directionality: Let $U \subseteq A$ be such that $A \backslash U \not\dashrightarrow U$ and $A \backslash U \not\rightarrow U$. We define $\mathcal{SB}_{|U} = \langle \mathcal{L}, U, \rightarrow \cap (U \times U), \dashrightarrow \cap (U \times U) \rangle$. Then an extension $E \subseteq U$ is strongly/weakly coherent in $\mathcal{SB}'$ iff there exists a strongly/weakly coherent extension $E' \subseteq A$ such that $E = E' \cap U$.

Briefly, *transitivity* states that the support is transitive, where a relation is transitive if $a \dashrightarrow b$ and $b \dashrightarrow c$ imply that $a \dashrightarrow c$. The claim then is that the set of acceptable extensions does not change if we make the support relation transitive. *Extension selection* says that strongly and weakly coherent extensions can be seen as using the support relation to select certain extensions given by some other semantics. In our case, it is admissible semantics. *Support removal robustness*, as mentioned above, says that removing supports does not change the acceptability of an extensions. Here this means that removing a supported premise from an argument will not impact the acceptable extensions. *Monotony of status* is intended to specify that adding a support to an argument should not decrease its acceptability. Here, specifically, this is interpreted as saying that adding new supported premise to an argument should not decrease its acceptability. *Directionality* expresses that the efffect of supports and attacks should only go in the direction of the relation. Thus, a subgraph of an SBAF with no ingoing edges should not have different acceptable extensions than the full graph, except, of course, that they only contain arguments from the subgraph. It says that the extensions of the subgraph are exactly the restrictions of the $\mathcal{SB}$ extensions to the subgraph.

Proposition 16. *Strongly coherent extensions satisfy extension selection and support removal robustness. They fail transitivity, monotony of status, and directionality.*

Weakly coherent extensions satisfy extension selection, support removal robustness, and directionality. They fail transitivity and monotony of status.

Proof. Figure 3.11 gives the counterexample for transitivity of both strongly and weakly coherent extensions. Extension selection is immediate as both strongly and weakly coherent extensions are defined to be admissible.

For support removal robustness, take any weakly coherent extension E. If $a \notin E$, then E will remain weakly coherent if the support to a is removed. If $a \in E$, we check its properties. It is clear that conflict-freeness remains. Similarly for defence, as there are only fewer ingoing attacks. Further, weak sentence-respect remains unaffected, as no new argument will fulfil the conditions. In sum, E remains weakly coherent and as strongly coherent extensions are also weakly coherent, this gives the result for those as well.

Figure 3.12 gives the counterexample for monotony of status of both strongly and weakly coherent extensions. Figure 3.13 gives the counterexample for directionality of strongly coherent extensions. For weakly coherent extensions, first let $E \subseteq U$ be such that $E = E' \cap U$ for some weakly coherent $E' \subseteq A$. Then E is clearly weakly sentence-respecting. Further, it is also admissible, since E is conflict-free and defended against any attacks within U. Thus E is weakly coherent. Now we assume that E is weakly coherent in $\mathcal{SB}_{|U}$. Note that the closure of E under $R_{\mathcal{SB}}^{Sent(E)}$ is weakly coherent, as shown by Lemma 5. Thus there exists a weakly coherent extension $E' \subseteq A$ such that $E = E' \cap U$. □

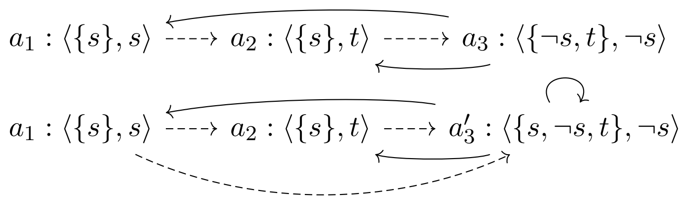

Figure 3.11: This example shows that support with weakly and strongly coherent extensions doesn't satisfy transitivity, as $\{a_3\}$ is strongly coherent, whereas $\{a'_3\}$ is not.

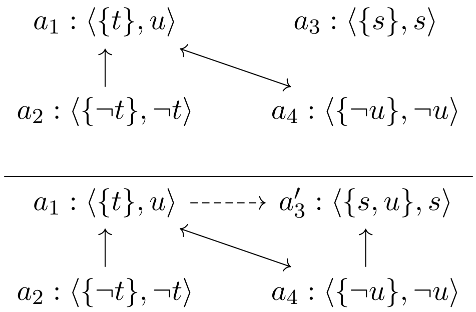

Figure 3.12: This example shows that monotony of status fails for strongly and weakly coherent extensions. $\{a_3\}$ is both strongly and weakly coherent, but there is no weakly coherent extension containing a'_3.

There are some insights to be drawn from these results. First, extension selection is fairly trivial, since strongly and weakly coherent extensions are defined to be admissible. Nevertheless, it is interesting that they can be seen as part of the same class of extensions as the support-score based semantics described in Section 2.2.1. Recall that these semantics rank, say, complete extensions according to the number of internal or external supports they receive and select those that maximise that number. Thus they have a very different approach to support than strongly and weakly coherent extensions.

The failure of transitivity and monotony of status as well as the satisfaction of support removal robustness all point to the same feature of coherent semantics: It is better for an argument to have as few premises as possible. Transitivity and monotony of status both fail because adding a supported premise to an argument also opens it up to new attacks. Support removal robustness, in contrast, is satisfied because removing a premise from an argument means that it is less susceptible to attacks. Thus extensions might lose their acceptability status when supported premises are added, but they will keep it if premises are removed. Whether this is a desirable result is debatable. On one hand, it incentivises arguments with no redundant premises. This can be useful, as we did not make any assumption on the relevance of the premises of an argument to its conclusion. Now we can see that while arguments can contain irrelevant premises, it would be strategically better to remove them. On the other hand, sometimes arguments get more plausible if more premises are added. Take Alexa, the meteorologist, again. We can give more credence to her claim that it will rain next week by saying that she is a certified meteorologist. If we further say that her claim is based on detailed calculations and she has an impeccable track record with these claims, we can be seen as adding premises that should make the argument as a whole more plausible. However, in our framework, this would be a

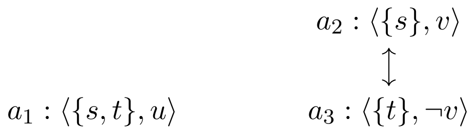

Figure 3.13: This example shows that directionality fails for strongly coherent extensions, as $\{a_1\}$ is strongly coherent if we disregard arguments a_2 and a_3. However, in the presence of those arguments, accepting a_1 would, by strong sentence-respect, require to accept both a_2 and a_3. But then the extension would not be conflict-free. Thus, there is no strongly coherent extension containing a_1 in the full framework.

strategically bad move, since it only opens up the argument to more potential attacks.

Directionality is interesting because it distinguishes between strongly and weakly coherent extensions. The reason it fails for strongly coherent extensions is that strong coherence has demanding requirements that go beyond the explicit attack and support relations. Since accepting an argument means you also accept its premises, this can commit you to accepting another argument as well if it shares the premises. In some cases, this can lead to accepting conflicting arguments. While weakly coherent extensions in general also have requirements that go beyond the explicit attack and support relations, the requirements are such that acceptability is not infringed. Since weak sentence-respect requires defended arguments, it will never force accepting conflicting arguments.

Table 3.1 compares strongly and weakly coherent extensions with the extension-based support semantics introduced in Chapter 2.

	Trans.	Ext. Select.	Supp. Rem. Rob.	Mon. of Stat.	Direct.
Strong Coherence	✗	✓	✓	✗	✗
Weak Coherence	✗	✓	✓	✗	✓
Deductive Support	✓	✗	✓	✓	✗
Necessary Support	✓	✗	✓	✗	✓
Internal Coherence	✗	✓	✗	✓	✗
Strength. Def. (i)	✗	✗	✗	✓	✓
Strength. Def. (ii)	✗	✗	✗	✓	✓
Strength. Def. (iii)	✗	✗	✓	✓	✓

Table 3.1: Principled comparison between support semantics. The values for BAF semantics come from L. Yu et al. (2023) and follow the original formulations of the principles.

3.2.2 Acceptability Degrees

We can further investigate the acceptability of arguments and how support impacts it. Note that the above principles mostly worked with adding or removing supports among the existing arguments in a framework. This cannot be directly translated to SBAFs, except for adding and removing premises of arguments. While this corresponds to some natural dynamics in argumentation, it is perhaps not the most obvious way of approaching the impact of support. Alternatively, we can investigate in what way the acceptable arguments change if new supporting arguments are added to an SBAF. This matches the situation where a debate continues over time and new arguments are added as time goes on.

Instead of going about this by means of principles as above, we will do it pseudo-empirically. That is, we will go through a number of examples to get an idea of how

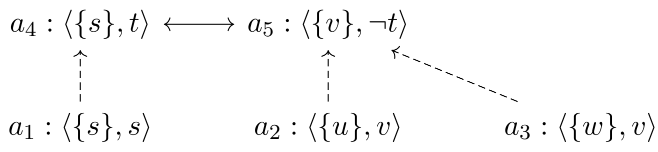

(a) The SBAF of Figure 3.1.

	ρ^A_s	ρ^A_w	ρ^L_s	ρ^L_s
a_1	0.167	0.6	0.125	0.5
a_2	0.334	0.5	0.25	0.4
a_3	0.334	0.5	0.25	0.4
a_4	0.167	0.2	0.125	0.167
a_5	0.667	0.4	0.5	0.267

(b) Acceptability degrees of Figure 3.1.

Figure 3.14: An example for acceptability degrees.

adding supporting arguments change the acceptability of the supported arguments. A very simple measure of the acceptability of an argument a is the following: we take the ratio of acceptable extensions that contain a over the acceptable extensions in general (cf. Betz 2012; Dondio 2018).

Definition 38 (Acceptability Degree). *Let $\mathcal{SB} = \langle \mathcal{L}, A, \rightarrow, \dashrightarrow \rangle$ be an SBAF and $a \in A$. We define its argument-based strong and weak acceptibility degree as*

$$\rho^A_{s/w}(a) := \frac{|\{E \subseteq A \mid a \in E \textit{ and } E \textit{ is strongly/weakly coherent}\}|}{|\{E \subseteq A \mid E \textit{ is strongly/weakly coherent}\}|}.$$

Its language-based strong and weak acceptability degree is

$$\rho^L_{s/w}(a) := \frac{|\{S \subseteq Sent(A) \mid a \in Arg_s(S)/Arg_w(S) \textit{ and } S \textit{ is strongly/weakly adequate}\}|}{|\{E \subseteq A \mid S \textit{ is strongly/weakly adequate}\}|}.$$

The table in Figure 3.14b gives the values for all four versions of the acceptability degree of an argument in the original SBAF example in Figure 3.1. First, we can observe that the values for strong acceptability are quite similar between the argument and language version. They give the same relative ranking and they both agree that a_1 and a_4 are equally strong. The values for weak acceptability are also quite similar between the argument and language version. The biggest difference between the strong and weak version is the value of a_1. Interestingly, according to strong coherence, a_1 is weaker than the other unattacked arguments a_2 and a_3, while weak coherence leads to a_1 being stronger.

The difference between strong and weak coherence we see at play here is that with strong sentence-respect, a_1 deductively entails a_4, whereas this is not the case with weak sentence-respect. Thus there is much more room for accepting a_1 in terms of weak coherence, which leads to its higher acceptability degree compared to strong coherence. That the acceptability of a_1 with weak coherence is higher than that of a_2 and a_3 come from it being a minimal argument. With weak sentence-respect, it is in some sense easier to accept a minimal argument than a non-minimal one, as the former is immediate as soon as it's one sentence is accepted. Accepting non-minimal arguments takes accepting more than one sentence and is thus generally a bit less likely. Interestingly, with strong coherence and a_1 being a minimal argument in favour of a_4, you can accept a_1 *if and only if* you accept a_4. This explains that the two arguments have the same acceptability with strong coherence.

Figure 3.15 gives an example that shows a very counterintuitive result for both strong and weak argument-based acceptability. Namely, if we just consider a_1 on its own, it has an acceptability of 0.5, which makes sense: either you accept it or not. But if we add the supporting argument a_2, the acceptability of a_1 decreases to 0.334. Thus, adding a support can decrease the acceptability of an argument even in complete absence of

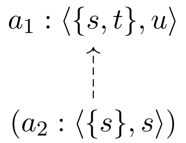

(a) An SBAF where a_2 gets added.

	ρ_s^A	ρ_w^A	ρ_s^L	ρ_s^L
a_1	0.5	0.5	0.143	0.143

	ρ_s^A	ρ_w^A	ρ_s^L	ρ_s^L
a_1	0.334	0.334	0.143	0.143
a_2	0.667	0.667	0.429	0.429

(b) Acceptability degrees without and with a_2.

Figure 3.15: A problematic case for argument-based acceptability.

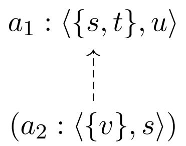

(a) An SBAF where a_2 gets added.

	ρ_s^A	ρ_w^A	ρ_s^L	ρ_s^L
a_1	0.5	0.5	0.143	0.143

	ρ_s^A	ρ_w^A	ρ_s^L	ρ_s^L
a_1	0.5	0.5	0.2	0.2
a_2	0.5	0.5	0.3	0.3

(b) Acceptability degrees without and with a_2.

Figure 3.16: Another problematic case for argument-based acceptability.

attacks! This happens because of a similar situation as in the previous example, where we also had a support by a minimal argument. If we add a_2, it is not acceptable to accept a_1 without a_2 (but it is allowed to accept a_2 without a_1). Thus amongst the three acceptable extensions, $\emptyset$, $\{a_2\}$, $\{a_1, a_2\}$, only one third contain a_1. This is clearly undesirable. Looking at the values of language-based acceptability, we get a much more reasonable result: the acceptability does not change. While a_2 does support a_1, note that a_2 is a minimal argument and thus does not provide any additional reason for premise s. Rather, it just restates it. Hence, it indeed should not increase the plausibility of a_1 and the language-based acceptability values get the right result here.

Even worse for argument-based acceptability, adding a supporting argument if the supported argument has more than one unsupported premises does not impact the acceptability at all. However, language-based acceptability is able to capture it, as seen in Figure 3.16b. Thus we conclude that language-based acceptability values should be used. Using them, we can confirm our observation earlier on that it is better for arguments to have fewer premises. Namely, arguments with fewer premises tend to have higher acceptability degrees. We can already see this in Figures 3.15 and 3.16. Figure 3.17 gives a direct comparison of unrelated arguments with a different number of premises. It is worth noting that the result stays the same when we change the example such that the two arguments in the figure rebut each other.

A final observation at this point is the following. We have seen in some examples that sentence-respect has effects that go beyond the explicit support relation. For instance, if arguments share premises, accepting one might also trigger sentence-respect, even though there is no support between the arguments (this leads to the failure of directionality in Figure 3.13). Thus one might wonder whether the support relation even matters. Perhaps sentence-respect is such that what actually matters is whether some argument share sentences, rather than whether one supports the other. However, we can now see that support matters for the acceptability degrees of arguments. Figure 3.18 compares two arguments when they simply share a premise and when one supports the other. We

$a_1 : \langle\{s,t\},u\rangle$

$a_2 : \langle\{v\},w\rangle$

(a) An SBAF with unrelated arguments.

	ρ_s^A	ρ_w^A	ρ_s^L	ρ_s^L
a_1	0.5	0.5	0.143	0.143
a_2	0.5	0.5	0.334	0.334

(b) Acceptability degrees without and with a_2.

Figure 3.17: Fewer premises seem to lead to higher acceptability.

$\mathcal{SB}$: $a_1 : \langle\{s\},t\rangle$ $a_2 : \langle\{s\},u\rangle$

$\mathcal{SB}'$: $a_1 : \langle\{s\},t\rangle \leftarrow\text{-----}\ a_2 : \langle\{u\},s\rangle$

(a) Two SBAFs, one with shared premises and one with support.

$\mathcal{SB}$	ρ_s^A	ρ_w^A	ρ_s^L	ρ_s^L
a_1	0.5	0.5	0.2	0.2
a_2	0.5	0.5	0.2	0.2

$\mathcal{SB}'$	ρ_s^A	ρ_w^A	ρ_s^L	ρ_s^L
a_1	0.334	0.334	0.5	0.5
a_2	0.667	0.667	0.25	0.25

(b) Acceptability degrees of $\mathcal{SB}$ and $\mathcal{SB}'$.

Figure 3.18: Support matters for acceptability degrees.

again see that the support decreases acceptability when we use argument-based values, but language-based values recognise the difference between sharing premises and support. Accordingly, the supported version of a_1 has higher acceptability than the unsupported one.

In this section, we examined the role of support in SBAFs, though everything we presented here are preliminary observations that should be investigated further. For instance, a full principled analysis of support is missing and it would be useful to better compare coherent semantics to other support semantics in the literature. Also, all examples and observations concerning acceptability degrees are very simple and while they might hint at interesting properties, there is much further work to be done in order to understand the values arguments receive. For instance, it is currently unclear whether acceptability degrees can be used as a full gradual or ranking semantics or not (for ranking semantics, see Amgoud and Ben-Naim 2013; Besnard and Hunter 2001; Pu et al. 2014, 2015). Nevertheless, we have characterised support to some extent, showing that it is worth investigating further. One important feature we have illustrated is that support has impact on acceptable extensions even in absence of any attacks.

Table 3.2 summarises distinguishing features for each SBAF semantics:

Strongly Coherent	Captures deductive support, does not satisfy directionality
Weakly Coherent	Satisfies directionality, closer to Dung semantics (e.g. includes all complete extensions)
Strongly Adequate	Useful for acceptability degrees, ranks minimal arguments low
Weakly Adequate	Useful for acceptability degrees, ranks minimal arguments high

Table 3.2: Summary of Properties for SBAF semantics.

$$a_1 : \langle \{s\}, s \rangle \longrightarrow a_2 : \langle \{t, \neg s\}, t \rangle$$

Figure 3.19: An example for maximising the set of accepted sentences.

3.3 Doubtless Extensions

We have argued that in order to find room for support in formal argumentation, we need a notion of *mere doubt*. We are now in a position to consider this claim formally. Namely, we can show that if we do not allow agents to merely doubt sentences, then weakly coherent argument extensions are just preferred extensions. As preferred extensions do not take support (or the structure of arguments) into account, we conclude that support plays no role in absence of mere doubt.

Before we start, it should be noted that there is no correspondence between strongly coherent argument extensions and preferred extensions, as we have seen back in Figure 3.1 (p. 34). There, the preferred extensions of the SBAF were not strongly coherent, but they were weakly coherent. Moreover, according to Observation 2, we know that complete extensions are weakly coherent. This, together with preferred extensions being $\subseteq$-maximal admissible, already gives the following further observation:

Observation 6. *Let $\mathcal{SB} = \langle \mathcal{L}, A, \rightarrow, \dashrightarrow \rangle$ be an SBAF. Then $E \subseteq A$ is $\subseteq$-maximal amongst weakly coherent extensions iff it is preferred.*

However, there is a more interesting perspective. Namely, we can model a lack of mere doubt about sentences by requiring language extensions to be $\subseteq$-maximal. That is, if you can accept a sentence, then you need to accept it. This way, there can be no mere doubt about sentences and we can show that we again end up with preferred extensions.

Definition 39 (Confident Extensions)**.** *Let $\mathcal{SB} = \langle \mathcal{L}, A, \rightarrow, \dashrightarrow \rangle$ be an SBAF. A language extension $S \subseteq Sent(A)$ is called a confident strongly (resp. weakly) adequate language extension if it is $\subseteq$-maximal amongst strongly (resp. weakly) adequate language extensions.*

An argument extensions $E \subseteq A$ is called a confident strongly (resp. weakly) coherent argument extension if there exists a confident strongly (resp. weakly) adequate language extension $S \subseteq Sent(A)$ such that $E = Arg_s(S)$ (resp. $E = Arg_w(S)$).

As before in Propositions 12 and 13, confident argument extensions and confident language extensions correspond. But now the set of sentences of a confident argument extension might not itself be confident. Take the example in Figure 3.19. There, $\{a_1\}$ is a confident strongly and weakly coherent argument extension, but its set of sentences only includes s, whereas $\{s, t\}$ would be the corresponding confident strongly and weakly adequate language extension. Nevertheless, we get the following correspondence.

Proposition 17. *Let $\mathcal{SB} = \langle \mathcal{L}, A, \rightarrow, \dashrightarrow \rangle$ be a saturated SBAF. Then for every confident strongly (resp. weakly) adequate language extension, its strong (resp. weak) argument set is confident strongly (resp. weakly) coherent.*

Also, for every confident strongly (resp. weakly) coherent argument extension E, there exists a confident strongly (resp. weakly) adequate language extension S such that $Arg_s(S) = E$ (resp. $Arg_w(S) = E$).

Proof. The first part of the proposition follows directly from Propositions 12 and 13.

The second part is immediate by definition. □

While it is obvious that any strongly or weakly adequate language extension can be extended to a confident one, it is less obvious for argument extensions. In fact, it will be a corollary of the following proposition. Since we are here interested in the interplay between the language and argument perspectives, we need to restrict ourselves to saturated SBAFs.

Proposition 18. *Let $\mathcal{SB} = \langle \mathcal{L}, A, \rightarrow, \dashrightarrow \rangle$ be a saturated SBAF. Then any preferred extension $E \subseteq A$ is confident weakly coherent.*

Proof. Let E be a preferred extension. Since it is then also complete, Observation 2 gives us that it is weakly coherent. It thus remains to show that there exists a confident weakly adequate language extension S such that $E = Arg_w(S)$.

Consider the set $\{S \subseteq Sent(A) \mid E = Arg_w(S)$ and S is weakly adequate$\}$. Note that it is non-empty, since $Sent(E)$ is weakly adequate and $Arg_w(Sent(E)) = E$ (Proposition 13). Recall that we only consider finite frameworks, thus there exists a $\subseteq$-maximal element S' that is weakly adequate and $Arg_w(S') = E$. It remains to show that S' is also $\subseteq$-maximal amongst weakly adequate language extensions.

Take any weakly adequate S'' such that $S' \subsetneq S''$. Since S' is $\subseteq$-maximal amongst weakly adequate extensions with $Arg_w(S) = E$, we know that $Arg_w(S'') \neq E$. We first show that $E \subsetneq \{a \in A \mid Sent(a) \subseteq S''$ and $\overline{n(a)} \cap S'' = \emptyset\}$. Take any $a \in E$. Then clearly, $Sent(a) \subseteq Sent(E) \subseteq S' \subsetneq S''$. Now suppose there is some $t \in \overline{n(a)} \cap S''$. Then, by saturatedness, there exists a minimal argument b for t. Note that $b \rightarrow a$ and since E is defended, there is some $c \in E$ such that $c \rightarrow b$, that is, $Conc(s) \in \bar{t}$. But since $Sent(c) \subseteq Sent(E) \subseteq S' \subseteq S''$, this contradicts compatibility of S''. Hence, $\overline{n(a)} \cap S'' = \emptyset$ and we can note that $E \subsetneq \{a \in A \mid Sent(a) \subseteq S''$ and $\overline{n(a)} \cap S'' = \emptyset\}$. Since E is admissible, this gives $E \subseteq Init(S'') \subseteq Arg_w(S'')$.

In sum, $E \subsetneq Arg_w(S'')$, but since $Arg_w(S'')$ is admissible (Proposition 13), this contradicts that E is preferred. Hence, there exists not weakly adequate S'' such that $S' \subsetneq S''$ and we conclude that E is confident weakly coherent. □

We get the following corollary, since any admissible extension can be extended to a preferred one.

Corollary 3. *Let $\mathcal{SB} = \langle \mathcal{L}, A, \rightarrow, \dashrightarrow \rangle$ be a saturated SBAF and $E \subseteq A$ a strongly (resp. weakly) coherent argument extension. Then there exists a confident strongly (resp. weakly) coherent argument extension E' such that $E \subseteq E'$.*

Interestingly, the other direction of Proposition 18 does not hold for saturated SBAFs, as can be seen again in Figure 3.19. There, $\emptyset$ is confident, because $\{\neg s, t\}$ is confident weakly adequate with an empty argument set. Nevertheless, we can show the other direction for *strongly saturated* SBAFs.

Proposition 19. *Let $\mathcal{SB} = \langle \mathcal{L}, A, \rightarrow, \dashrightarrow \rangle$ be a strongly saturated SBAF. Then any confident weakly coherent argument extension $E \subseteq A$ is preferred.*

Proof. Let E be weakly coherent and suppose it is not preferred. We show that E is not confident. We need to show that for any $S \subseteq Sent(A)$ such that $Arg_w(S) = E$, there exists an adequate language extension S' such that $S \subsetneq S'$.

Since E is admissible, there exists a preferred extension E' such that $E \subsetneq E'$. By Proposition 18, we know that E' is confident weakly coherent. Hence, there exists a confident weakly adequate language extension S' such that $Arg_w(S') = E'$.

Now take any $S \subseteq Sent(A)$ such that $Arg_w(S) = E$. We show that $S \subsetneq S'$. Take any $s \in Sent(A) \backslash S'$. Then $s \notin S'$ and since S' is confident, we know that $S' \cup \{s\}$ is not weakly adequate. We can show that $S' \cup \{s\}$ is not compatible. Suppose it is. Then

$Arg_w(S' \cup \{s\})$ is admissible (Lemma 8), and we know that $E' \not\subseteq Arg_w(S' \cup \{s\})$ (since otherwise either $S' \cup \{s\}$ would be weakly adequate, if $E' = Arg_w S' \cup \{s\}$, or E' not preferred, if $E \subsetneq Arg_w S' \cup \{s\}$). This lets us take an argument $a \in E' \backslash (S' \cup \{s\})$ and we know that $s \in \overline{n(a)}$ (since otherwise $a \in Init(S' \cup \{s\})$). By saturatedness of $\mathcal{SB}$, we then have a minimal argument b for s. Note that $b \rightarrow a$ and since E' is defended, there is some $c \in E'$ such that $Conc(c) \in \overline{s}$. But note that $Sent(c) \subseteq Sent(E) = Sent(Arg_w(S')) \subseteq S'$ (see proof of Proposition 13), contradicting compatibility of $S' \cup \{s\}$. Thus, we know that $S' \cup \{s\}$ cannot be compatible. But then there is some $t \in S'$ such that $t \in \overline{s}$. By strong saturatedness, there exists a minimal argument d for t. Since it is a minimal argument, $\{d\}$ is admissible (Observation 1) and hence $d \in Init(S') \subseteq E'$. Since $d \rightarrow b$, we also have $E' \rightarrow b$. But since $E \subseteq E'$, we know that $b \notin E$ (since otherwise E' would not be conflict-free). Finally, since b is a minimal argument for s, we know that $s \notin S$ (since otherwise $b \in Arg_w(S) = E$). In sum, $S \subseteq S'$.

It remains to note that $S \neq S'$, since otherwise $E = Arg_w(S) = warg(S') = E'$. Thus $S \subsetneq S'$ and E is not confident. This establishes the result, since we know that any preferred $E \subseteq A$ is weakly coherent. □

Finally, it is worth noting that we can also translate abstract argumentation frameworks into SBAFs and use confident weakly coherent extensions to calculate preferred extensions. Namely, we can get an SBAF with the same attack relations between arguments as an AF $\mathcal{A} = \langle A, \rightarrow \rangle$, by defining the structure of an argument $a \in A$ as $a_{\mathcal{A}} = \langle \{s_a\} \cup \{t_b \mid b \in A : b \rightarrow a\}, s_a \rangle$, where we assume that $t_b \in \overline{s_b}\ \forall b \in A$. It is then obvious that $\mathcal{SB}_{\mathcal{A}} = \langle \mathcal{L}, A_{\mathcal{A}}, \rightarrow, \emptyset \rangle$ has the same attack relation as $\mathcal{A}$. Note that the constructed SBAF might not be saturated. We then get the following result.

Proposition 20. *Let $\mathcal{A} = \langle A, \rightarrow \rangle$ be an AF and $\mathcal{SB}_{\mathcal{A}} = \langle \mathcal{L}, A_{\mathcal{A}}, \rightarrow, \emptyset \rangle$ be its corresponding SBAF. Then an argument extension $E \subseteq A_{\mathcal{A}}$ is confident weakly coherent if it is preferred.*

Proof. Let E be preferred. Then it is weakly coherent (Observation 2), so it remains to check whether it is confident. As in the proof of Proposition 18, we show that a $\subseteq$-maximal element S' of $\{S \subseteq Sent(A) \mid E = Arg_w(S)$ and S is weakly adequate$\}$ is confident. Thus let S'' be weakly adequate such that $S' \subsetneq S''$. Since there are no undercutting sentences in $\mathcal{SB}_{\mathcal{A}}$ and S' is $\subseteq$-maximal in the set above, we know that $Arg_w(S') \subsetneq Arg_w(S'')$. But this contradicts that E is preferred. Thus we conclude that E is confident weakly coherent. □

Interestingly, the other direction of this proposition does not hold. The counterexample is the AF $a \rightarrow b$, for which the corresponding SBAF is that of Figure 3.19. As we have seen, $\emptyset$ is confident weakly coherent, but not preferred.

While we know by definition that an agent choosing a preferred extension is one that wants to maximise the set of accepted arguments, we can now see that this does not always correspond to an agent that wants to maximise the set of accepted sentences. As Figure 3.19 shows, it is even possible that a maxiaml set of sentences corresponds to an empty set of arguments. Thus, from the language perspective, preferred extensions only seem natural in the specific class of strongly saturated SBAFs.

The translation from abstract AFs to SBAFs further illustrates the difficulties in characterising Dung semantics from the language perspective. Maximising accepted sentences does not correspond fully to preferred extension, though all preferred extensions maximise accepted sentences. For complete semantics in general, the situation is more unclear. We do not know of any natural condition for language extensions that would (at least partially) correspond to complete extensions. From the language perspective, there is nothing special about sentences that occur in defended arguments that would make it possible to single them out for maximisation. However, it seems that something like that would be needed in order to find a correspondence to complete extensions on the language level.

The results in this section also illustrate how a lack of doubt leads us to disregard support between arguments. Namely, if we have no mere doubt about any sentence, then we try to accept maximal sets of sentences, which leads us to accepting preferred extensions. As the latter does not take any kind of support into account, we conclude that support plays no role in the absence of doubt.

3.4 Knowledge-Based Reasoning

Most approaches to structured argumentation consider knowledge-based reasoning. That is, arguments are constructed from a knowledge base and the semantics essentially tell us what inference we can draw from a knowledge base. Our approach differs in that we do not construct arguments and there is no notion of a knowledge base. In this section, we examine to what extent knowledge-based reasoning can be recovered in our approach.

3.4.1 Frameworks with Contestable Sentences

As the various examples in the previous sections have illustrated, in any given SBAF, there tend to be many strongly or weakly coherent argument extensions. This is in stark contrast to Dung semantics, where there are usually only a few complete extensions and in the large class of well-founded AFs, there is even a unique complete extension. Having many rationally acceptable extensions is not necessarily a problem. For instance, if an agent just wants to check whether their position in a debate is reasonable, then it does not matter for them how many other such positions there are. In that sense, allowing for many acceptable extensions represents a form of pluralism with respect to what is a rational position. The situation is different if we want to know what we should believe, given the arguments of a debate. If we ask what *follows* from a debate, then it is not helpful to be given a dozen extensions for a rather small SBAF (e.g. as in Figure 3.1).

Perhaps, then, the notion of doubt we implemented in SBAFs is to sweeping. Allowing agents to doubt essentially all sentences naturally leads to a great number of acceptable extensions. Thus we might want to limit the extent to which sentences can be doubted. We have seen that disallowing mere doubt for all sentences leads us to preferred extensions, but we can still specify a subset of all sentences for which doubt needs to be justified. Let us call such sentences *contestable*, in contrast to the normal *doubtable* sentences. The idea then is that an agent needs to maximise the set of contestable sentences they accept, whereas they are free to accept or reject any doubtable sentences (within the scope of adequate extensions).

By adding contestable sentences, we also implement a notion of knowledge-based reasoning. It allows us to ask what follows from the contestable sentences given some arguments. A knowledge base of contestable sentences will be defeasible, but they give some starting point from which arguments can be constructed. Additionally, it allows us to distinguish sentences that come with different burdens of proof (Walton 2010).

Formally, we introduce contestable sentences by adding an indicator function to the language of a framework that tells us for each sentence whether it is contestable or doubtable.

Definition 40 (Language with Contestables)**.** *A language with contestables is a tuple* $\mathcal{L}^{\mathcal{C}} = \langle L, g, \overline{}, n\rangle$ *where* L, $\overline{}$, *and* n *are as in Definition 29 and* $g : L \to \{c, d\}$ *is an indicator function.*

For a sentence $s \in L$, $g(s) = c$ *indicates that sentence* s *is contestable, and* $g(s) = d$ *that it is doubtable. We use* s^c *and* s^d *as a shorthand. The set of all contestable sentences*

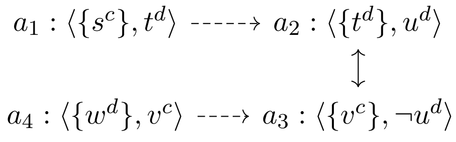

Figure 3.20: An example for an SBAF with contestables.

is denoted by L^c and that of doubtable sentences by L^d.

From this, we build SBAFs with contestables the same way as normal SBAFs.

Definition 41 (SBAFs with Contestables)**.** *An SBAF with contestables is a tuple $\mathcal{SBAF}^{\mathcal{C}} = \langle \mathcal{L}^{\mathcal{C}}, A, \rightarrow, \dashrightarrow \rangle$ where the arguments and relations are as in SBAFs.*

The addition of contestable sentences is recognised by requiring acceptable extension to maximise them. This mirrors the way we modelled a lack of mere doubt about all sentences in the previous section, just that this time we only have a lack of mere doubt about some sentences. This leads to the definition of *sensible* extensions.

Definition 42 (Sensible Extensions)**.** *Let $\mathcal{SBAF}^{\mathcal{C}} = \langle \mathcal{L}^{\mathcal{C}}, A, \rightarrow, \dashrightarrow \rangle$ be an SBAF with contestables. A strongly (weakly) adequate language extension $S \subseteq Sent(A)$ is called sensible if $S \cap L^c$ is $\subseteq$-maximal in $\{S' \cap L^c \mid S'$ is strongly (weakly) adequate and $S \subseteq S'\}$.*

A strongly (weakly) coherent argument extension $E \subseteq A$ is called sensible if there exists a cautious strongly (weakly) adequate language extension $S \subseteq Sent(A)$ such that $E = Arg_s(S)$ ($E = Arg_w(S)$).

Figure 3.20 gives an example for an SBAF with contestables. We have the following sensible strongly adequate language extensions: $\{s^c, t^d, u^d\}$, $\{v^c, \neg u^d\}$, and $\{w^d, v^c, \neg u^d\}$. Note that the latter two are not allowed to accept s^c, since then they would have to accept t^d and u^d as well, making them incompatible. Similarly, the first extension is not allowed to accept v^c. The corresponding sensible strongly coherent argument extensions are $\{a_1, a_2\}$, $\{a_3\}$, and $\{a_3, a_4\}$. Note that the option of accepting no sentence or argument is not allowed now. On the weak side, all compatible language extensions that contain s^c, v^c, and t^d are sensible weakly adequate. t^d has to be accepted because it is the conclusion of an unattacked argument with purely contestable premises. This gives the following corresponding sensible weakly coherent argument extensions: $\{a_1\}$, $\{a_1, a_2\}$, $\{a_1, a_2, a_4\}$, $\{a_1, a_3\}$, and $\{a_1, a_3, a_4\}$.

It is good to note that, since confident extensions are clearly sensible, we know that sensible extensions always exist in saturated frameworks. Also, as is routine by now, sensible adequate language extensions and sensible coherent argument extensions correspond in saturated frameworks.

Proposition 21. *Let $\mathcal{SB} = \langle \mathcal{L}, A, \rightarrow, \dashrightarrow \rangle$ be a saturated SBAF. Then for every sensible strongly (resp. weakly) adequate language extension, its strong (resp. weak) argument set is sensible strongly (resp. weakly) coherent.*

Also, for every sensible strongly (resp. weakly) coherent argument extension E, there exists a sensible strongly (resp. weakly) adequate language extension S such that $Arg_s(S) = E$ (resp. $Arg_w(S) = E$).

Proof. The first part of the proposition follows directly from Propositions 12 and 13. The second part is immediate by definition. □

The full picture of the different semantics for SBAFs is given in Figure 3.21. For computational interest, we provide an account of how some of these semantics can be calculated using SAT-solvers in the appendix.

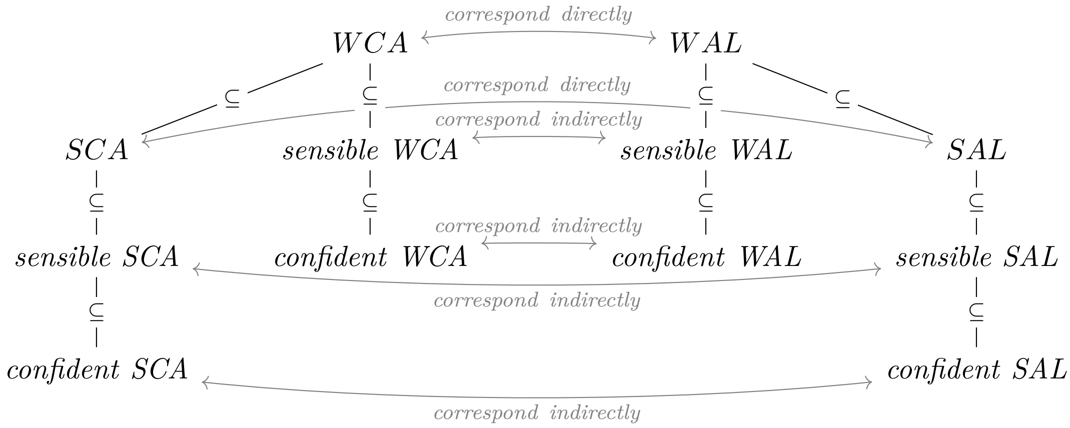

Figure 3.21: All semantics for SBAFs. SCA: strongly coherent argument extensions, WCA: weakly coherent argument extensions, WAL: weakly adequate language extensions, SAL: strongly adequate language extensions. We say that an argument-based semantics *corresponds directly* to a language-based semantics if the set of sentences of each argument extension belongs to the language-based semantics and the argument set of each language extension belongs to the argument-based semantics. An argument-based semantics *corresponds indirectly* to a language-based semantics if for each argument extension, there exists a language extension that accepts the same arguments and the argument set of each language extension belongs to the argument-based semantics.

The definition of sensible language extensions is a bit cumbersome, but it is easy to provoke undesired consequences if defined differently. For instance, requiring a sensible language extension to be maximal with respect to contestable sentences amongst *all* language extensions, rather than amongst those extending it, can make it almost impossible to reject a contestable sentence even if there are arguments against it (illustrated in Figure 3.22). Further, if we do not go through language extensions, but try to define sensible argument extensions directly, agents can be forced to accept unwanted doubtable sentences. For instance, we could define a sensible argument extension E to be such that $Sent(E) \cap L^c$ is $\subseteq$-maximal in $\{Sent(E)' \cap L^c \mid E'$ is strongly (resp. weakly) coherent and $E \subseteq E'\}$. But then an argument of the form $a_1 : \langle\{s^d, t^c\}, u^d\rangle$ would have to be accepted, effectively making s^d and u^d contestable sentences as well. Going through language extensions avoids that problem.

Some problems remain, but they seem to be confined to unsaturated SBAFs with contestables. Figure 3.23 gives an example of a framework where acceptance of a doubtable sentence is forced, even though the sentence is not implied by contestable sentences. The problem there can be avoided by making the framework saturated. This is not much of a limitation, as tying sensible argument extensions to sensible language extensions already relies on their correspondence in saturated frameworks. However, it is an open problem how to exactly formulate the issue with forced acceptance of unwanted doubtable sentences and, accordingly, whether it can always be avoided in saturated frameworks.

Another issue is that there is still a large number of sensible extensions in an SBAF, as illustrated by the example in Figure 3.20 above. Already in that small framework, there are too many sensible weakly adequate language extensions to easily write down. Also,

$$a_1 : \langle \{s^d, t^d\}, \neg u^d \rangle \longrightarrow a_2 : \langle \{u^c\}, v^d \rangle$$

Figure 3.22: If sensible language extensions had to be maximal with respect to contestable sentences amongst all language extensions, then it would not be possible to accept (all sentences of) a_1, as the extension $\{u^c\}$ would always be larger with respect to contestable sentences.

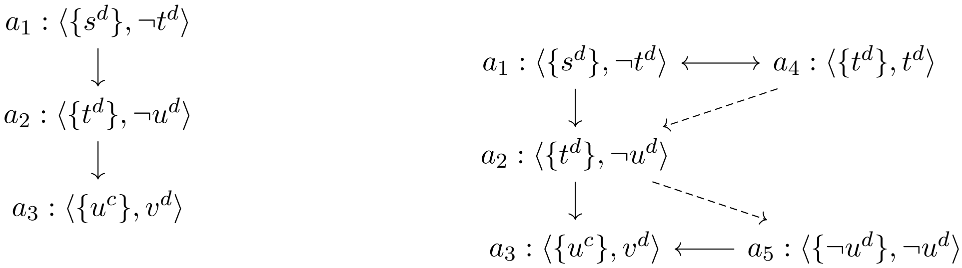

(a) In this SBAF with contestables, the only sensible strongly coherent extension is $\{a_1, a_3\}$, meaning that s^d has to be accepted, even though it is doubtable.

(b) The problem can be solved in this case by making the framework saturated, then e.g. $\{a_2, a_4, a_5\}$ is sensible strongly coherent.

Figure 3.23: An example showing a problem for sensible extensions in unsaturated SBAFs with contestables.

some of those are not especially close to the notion of knowledge-based reasoning, as they are free to contain doubtable sentences that are completely independent of the contestable ones. For instance, the sensible weakly adequate language extension $\{s^c, t^d, v^c, u^d, w^d\}$ accepts w^d, even though it is not in any way implied by the contestable sentences (our version of a knowledge base).

We can construct weak extensions that enforce the notion of reasoning only from the contestable sentences as follows. We start by using maximally compatible sets of contestable sentences as the initial set, then close under the respect function and argument-respect. Extensions constructed that way will only accept what follows from the contestable sentences they start with. For instance, in Figure 3.20, we would start with $\{s^c, v^c\}$ and end up with $\{s^c, t^d, v^c\}$. That way, we can drastically reduce the number of acceptable extensions. However, the construction requires that all contestable sentences come with a minimal argument.

Definition 43 (Canonical Extensions)**.** *Let $\mathcal{SBAF}^{\mathcal{C}} = \langle \mathcal{L}^{\mathcal{C}}, A, \rightarrow, \dashrightarrow \rangle$ be an SBAF with contestables such that for each contestable sentence s^c, we have $\langle \{s^c\}, s^c \rangle \in A$.*

A language extension $S \subseteq Sent(A)$ is called canonical if there exists a maximally compatible subset of contestable sentences $Cont \subseteq L^c$ such that $S = \bigcup_{i \in \mathbb{N}} S^i$, where $S^0 = Cont$ and $S^{i+1} = S^i \cup \{s \in Sent(A) \mid \exists a \in Arg_w(S^i) : s = Conc(a)\}$.

Figure 3.24 adapts the example of Figure 3.20 in order to illustrate the construction of canonical extensions. We start with $S^0 = \{s^c, v^c\}$, which gives $Arg_w(S^0) = \{a_5, a_6, a_1\}$. Adding the conclusions of accepted arguments, we get $S^1 = \{s^c, v^c, t^d\}$. As $Arg_w(S^1) = \{a_5, a_6, a_1\} = Arg_w(S^0)$. Thus we have $S = \{s^c, v^c, t^d\}$. The following proposition shows that the canonical extension is indeed sensible weakly adequate and that it is $\subseteq$-minimal amongst sensible weakly adequate extensions. This last property confirms that the canonical extension only accepts those doubtable sentences that follow from the contestables.

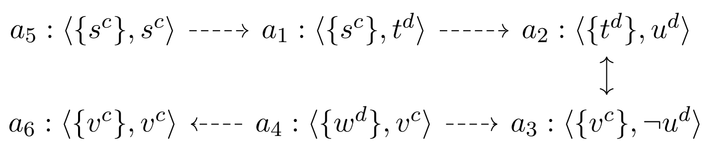

Figure 3.24: The SBAF of Figure 3.20 with minimal arguments for contestables.

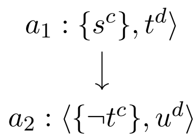

Figure 3.25: The construction of the canonical extension does not work if not all contestable sentences have a minimal argument. Start here with $S^0 = \{s^c, \neg t^c\}$, which has $Arg_w(S^0) = \{a_1\}$. Accordingly, $S^1 = \{s^c, \neg t^c, t^d\}$, which is not compatible.

Proposition 22. *Let* $\mathcal{SBAF}^{\mathcal{C}} = \langle \mathcal{L}^{\mathcal{C}}, A, \rightarrow, \dashrightarrow\rangle$ *be an SBAF with contestables such that for each contestable sentence* s^c*, we have* $\langle\{s^c\}, s^c\rangle \in A$*. Then any canonical extension is a* $\subseteq$*-minimal sensible weakly adequate language extension.*

Proof. We first show that each S^i is compatible. S^0 is compatible by definition. Thus assume S^i is compatible and we show that S^{i+1} is compatible as well. Since S^i is compatible, we know that $Arg_w(S^i)$ is admissible (Lemma 8). In particular, all its conclusions are compatible, meaning that S^{i+1} is compatible as well. Note that by definition, we have $S^i \subseteq S^{i+1}$, thus we conclude that S is compatible.

We further note that for each i, $Arg_w(S^i) \subseteq Arg_w(S^{i+1})$. Let $a \in Arg_w(S^i)$. Then clearly $Sent(a) \subseteq S^{i+1}$. Further, since $\overline{n(a)} \cap S^i = \emptyset$, the only way there can exists some $t \in \overline{n(a)} \cap S^{i+1}$ is if there is an argument $b \in Arg_w(S^i)$ such that $t = Conc(b)$. But then $b \rightarrow a$, contradicting conflict-freeness of $Arg_w(S^i)$. Thus $a \in \{a \in A \mid Sent(a) \subseteq S^{i+1}$ and $\overline{n(a)} \cap S^{i+1} =\}$. Indeed, $Arg_w(S^i) \subseteq \{a \in A \mid Sent(a) \subseteq S^{i+1}$ and $\overline{n(a)} \cap S^{i+1} =\}$, meaning that $Arg_w(S^i) \subseteq Init(S^{i+1}) \subseteq Arg_w(S^{i+1})$. This gives that $Arg_w(S) = \bigcup_{i \in \mathbb{N}} Arg_w(S^i)$.

For argument-respect, take any $a \in Arg_w(S)$. Then, since $Arg_w(S) = \bigcup_{i \in \mathbb{N}} Arg_w(S^i)$, for some i, we have $a \in Arg_w(S^i)$. This means that $Prem(a) \subseteq S^i$ and $Conc(a) \subseteq S^{i+1}$, meaning that $Sent(a) \subseteq S$. Thus S is weakly coherent. That S is sensible is immediate by definition, as S^0 is a maximally compatible set of contestables.

For $\subseteq$-minimality, take any sensible weakly coherent language extension $S' \subseteq S$. We show that for each i, we have $S^i \subseteq S'$, giving that $S' = S$. It is clear that $S^0 \subseteq S'$, otherwise it would not be sensible. Supposing that we have $S^i \subseteq S'$, we note that $Init(S^i) \subseteq \{a \in A \mid Sent(a) \subseteq S'$ and $\overline{n(a)} \cap S' = \emptyset\}$ (we get $\overline{n(a)} \cap S' = \emptyset$ from $S' \subseteq S$ and $Init(S^i) \subseteq Arg_w(S)$, as the latter gives $\overline{n(a)} \cap S = \emptyset$). Since $Init(S^i)$ is admissible, we get $Init(S^i) \subseteq Init(S')$. By the construction of the weak argument set (Proposition 11), we then know that $Arg_w(S^i) \subseteq Arg_w(S')$. Argument-respect then gives $S^{i+1} \subseteq S'$. This concludes the proof. □

Finally, it is worth noting that the construction of the canonical extension requires all contestable sentences to come with a minimal element. The problem that occurs if not is shown in Figure 3.25.

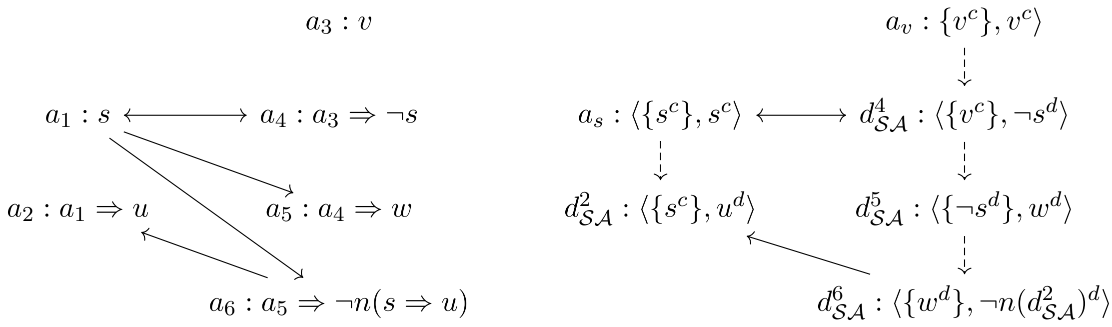

Figure 3.26: An example for a translation from a SAF (left) to an SBAF (right).

3.4.2 Revisiting ASPIC$^+$

We have seen in Section 2.2.2 that the straightforward ways of introducing doubt into ASPIC$^+$lead to a different notion of defence compared to standard formal argumentation and also structured bipolar argumentation. Namely, adapting the knowledge base according to which sentences some agent doubts allows them to ignore arguments if they doubt a premise. In contrast, we require defence even if a premise is doubted. Thus, there is no obvious way to translate an SBAF into an ASPIC$^+$framework that gives the same results. Nevertheless, it is possible to translate an ASPIC$^+$framework into an SBAF and compare the evaluations given by Dung semantics (used in ASPIC$^+$) and coherent extensions.

The translation goes as follows. Let $\mathcal{SA} = \langle A, \rightarrow\rangle$ be an ASPIC$^+$structured argumentation framework with knowledge base $\mathcal{KB}$ and argumentation system $\mathcal{AS} = \langle L, \overline{}, R, n\rangle$ with a symmetric contrariness function $\overline{}$. We construct an SBAF with contestables. Let each sentence in the knowledge base be contestable and let there be a minimal argument for each contestable sentence. That is, for each $s \in \mathcal{KB}$, we create an argument $a_s : \langle\{s^c\}, s^c\rangle$. These correspond to initial arguments in ASPIC$^+$. Then, for each inference rule $d : s_1, \ldots, s_n \Rightarrow t$ in R that occurs in some argument in $\mathcal{SA}$, we create an argument $d_{\mathcal{SA}} : \langle\{s_1^{d/c}, \ldots, s_n^{d/c}\}, t^{d/c}\rangle$, where the sentences are determined to be doubtable or contestable depending on whether they occur in the knowledge base. We need to limit ourselves to inference rules that occur somewhere in an argument, as there can be rules that never manifest in an argument as its premises cannot be deduced from the knowledge base. ASPIC$^+$'s naming function n for rules can then be directly taken over as a naming function for arguments. Also the contrariness function can directly be used as an incompatibility function. Attacks and supports are then determined as usual. Figure 3.26 gives an example.

We can now ask ourselves whether the acceptable argument extensions according to ASPIC$^+$and coherent argument extensions correspond. Figure 3.27 gives an example where a SAF has more arguments than its corresponding SBAF, since arguments a_3 and a_6 have the same top rule. Thus, we cannot directly compare extensions by the arguments they contain. Rather, we need to compare them on the sentences they accept. We have seen that extensions in a SAF are assumed to accept all conclusions of the arguments they containt. By sub-argument closure, this guarantees that all sentences involved of accepted arguments are accepted. Thus, if we want to know whether a an extension in a SAF should count as, say, weakly coherent, we need to find a corresponding extension in the translated SBAF with contestables, such that the two extensions accept the same sentences. In short, for an extension in the SAF, we construct its corresponding extension by adding all arguments that correspond to a rule used in some argument of the SAF

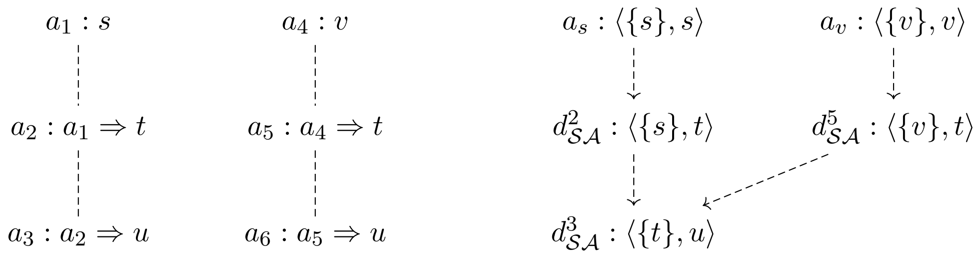

Figure 3.27: An example of a SAF, where the translation has fewer arguments.

extension.

Definition 44. *Let $\mathcal{SA} = \langle A, \rightarrow\rangle$ be a SAF with knowledge base $\mathcal{KB}$ and argumentation system $\mathcal{AS} = \langle L, \overline{\ \ }, R, n\rangle$. Further, let $\mathcal{SB}_{\mathcal{SA}} = \langle \mathcal{L}^{\mathcal{C}}, A_{\mathcal{SA}}, \rightarrow, \dashrightarrow\rangle$ be its corresponding SBAF with contestables. For any extension $E \subseteq A$, we define its corresponding extension $E_{\mathcal{SA}} = \{d_{\mathcal{SA}} \mid \exists a \in E : d \in Rules(a)\} \cup \{a_s \mid s \in E\}$.*

In the example of Figure 3.26, the extension $\{a_3, a_4, a_5, a_6\}$ in the SAF would be translated to the extension $\{a_v, d^4_{\mathcal{SA}}, d^5_{\mathcal{SA}}, d^6_{\mathcal{SA}}\}$. We can then note that extensions in SAFs only accept the conclusions of their arguments, while extensions in SBAFs accept all involved sentences. This discrepancy can be overcome by focusing on complete extensions in SAFs, as then sub-argument closure guarantees that they, too, accept all sentences involved in accepted argument. And indeed, in our example, both extensions accept the same sentences.

Lemma 12. *Let $\mathcal{SA} = \langle A, \rightarrow\rangle$ be a SAF with knowledge base $\mathcal{KB}$ and argumentation system $\mathcal{AS} = \langle L, \overline{\ \ }, R, n\rangle$. Further, let $\mathcal{SB}_{\mathcal{SA}} = \langle \mathcal{L}^{\mathcal{C}}, A_{\mathcal{SA}}, \rightarrow, \dashrightarrow\rangle$ be its corresponding SBAF with contestables.*

For any complete extension $E \subseteq A$, we have $\{Conc(a) \mid a \in E\} = Sent(E_{\mathcal{SA}})$.

Proof. We show set-inclusion in both directions. Take any $s \in \{Conc(a) \mid a \in E\}$. Then there exists an argument $a \in E$ such that $s = Conc(a)$. If $s \in \mathcal{KB}$, then it is immediate that $a_s \in E_{\mathcal{SA}}$. If $s \notin \mathcal{KB}$, We have that s is the conclusion of the top-rule d of a, which means that $s = Conc(d_{\mathcal{SA}})$ with $d_{\mathcal{SA}} \in E_{\mathcal{SA}}$. Hence, $s \in Sent(E_{\mathcal{SA}})$.

Now take any $s \in Sent(E_{\mathcal{SA}})$. Then $s \in \mathcal{KB}$ or there is a rule $d \in R$ such that $s \in Sent(d_{\mathcal{SA}})$ and $d \in Rules(a)$ for some $a \in E$. By sub-argument closure, there exists an argument $a' \in Sub(a)$ such that d is the top-rule of a'. If s is the conclusion of d, then we know that $s \in \{Conc(a) \mid a \in E\}$. If it is a premise, we can again use sub-argument closure to find a sub-argument $a'' \in Sub(a')$ such that $s = Conc(a'')$ and we again get $s \in \{Conc(a) \mid a \in E\}$. □

This confirms that the translation from a SAF to an SBAF with contestables makes sense in that we can find corresponding extensions. Of course, the question now is whether SAF extensions are rational from the SBAF perspective. One of the main differences between SAFs and SBAFs is the way attacks work. While each attack in an SBAF corresponds to an attack in the corresponding SAF, the other direction does not hold. In ASPIC$^+$, it is possible to attack an argument by attacking a sub-argument. As there is no notion of sub-arguments in SBAFs, such attacks disappear in the translation. This has consequences for which extensions count as defended. Thus, there are extensions that count as rational in ASPIC$^+$, but not when translated to SBAFs. Take again the example of Figure 3.26. In the SAF, there is a complete extension $\{a_1, a_2, a_3\}$, since a_1 defends a_2 by attacking a sub-argument of the attacker a_6. However, this defence gets lost in translation. Accordingly, the corresponding extension $\{a_s, a_v, d^2_{\mathcal{SA}}\}$ in the SBAF is not defended.

We have already seen in Section 2.2.2 that different notions of defence present a problem for translations from SBAFs to SAFs. Now defence is again an obstacle for translating between SAFs and SBAFs. The difference here is that ASPIC$^+$assumes that one can only accept an argument by accepting all its sub-arguments. Every argument needs to be linked back to the knowledge base and one can attack an argument by cutting that link. No similar notion is present in SBAFs. Rather, it is allowed to accept an argument without accepting all its supporting arguments. Perhaps there is a problem with the supporting argument, but one might accept the supported premise independently of that support. Hence, attacking the support does not impact the supported argument. This difference in interpretation leads to defended extensions in SAFs to be undefended in SBAFs.

Finally, we can observe that canonical extensions in SBAFs implement a different notion of knowledge-based reasoning than ASPIC$^+$. This can again be seen again in Figure 3.26. Namely, the canonical extension in the SBAF with contestables would be $\{s^c, v^c\}$, accepting arguments a_s and a_v. Without giving a fully formal translation for the general case, we can see that the corresponding extension in the SAF would be $\{a_1, a_3\}$. While this is an admissible extension, it is not complete, as it does not contain a_2, even though it is defended. Again, the problem lies in the different attacks that are present in the SAF but not in its corresponding SBAF. We can note, though, that the canonical extension in the SBAF is somewhat similar to the grounded extension in the SAF. But while the grounded extension never decides mutual conflicts (such as between a_1 and a_4), the canonical extension decides them if a minimal argument for a contestable sentence is involved. It does so by including the minimal argument.

Chapter 4

Conclusion

This thesis contributes to the growing literature on bipolar argumentation as well as to structured argumentation. We combine an explicit support relation with structured arguments and thus define the new notion of *structured bipolar argumentation frameworks.* When evaluating SBAFs, we define argument extensions as usual, but we also introduce the notion of language extensions, i.e. rationally acceptable sets of sentences that are used to build arguments. Language extensions might be easier for agents to interpret and give more direct access to what debates are really about: multiple sides trying to convince each other of some claim. We provide a range of semantics for both argument and language extensions, centred around the notions of *coherent argument extensions* and *adequate language extensions.* The main ideas incorporated into SBAFs and our semantics stem from pragma-dialectics and informal logic. This means that arguments have a simple premise-conclusion structure, in contrast to the full inference trees of ASPIC$^+$and ABA, and agents are allowed to doubt unattacked arguments and sentences.

We studied our semantics from a number of perspectives. *First*, we showed that coherent and adequate extensions correspond in saturated SBAFs, where contrary and undercutting sentences occur as minimal arguments. This condition is required, as otherwise not all relevant information about sentences is apparent in SBAFs, where only relations between arguments are considered. *Second*, we checked our semantics against common principles for bipolar semantics. We can distinguish strongly and weakly coherent extensions by the principle *directionality*, as only the latter satisfies it. In general, semantics for SBAFs would benefit from new principles, as many existing ones compare differences in extensions when manipulating the support relation. Given that support in SBAFs is based on the structure of the arguments, we cannot arbitrarily change the support relation. *Third*, we show that we can characterise preferred extensions as a certain type of weakly coherent argument extensions, called *confident*, where the set of accepted sentences is maximised. While confident weakly coherent extension, maximising sentences, and preferred extensions, maximising arguments, can differ in general, they are shown to coincide in strongly saturated SBAFs. *Fourth*, we investigated a form of knowledge-based reasoning in SBAF. Distinguishing sentences into two classes, *doubtable* and *contestable*, allows to define *sensible* extensions that maximise contestable sentences, while doubtable sentences can be doubted as usual. A particular version of sensible extensions, called the *canonical* extension implements directly the idea that we start by accepting all contestable sentences and only accept those doubtables that the arguments force to accept.

There is work to be done in checking whether some definitions can be refined. For instance, the asymmetry between support, which requires identity between sentences of two arguments, and attacks, which requires contrariness between sentences of two arguments, could be resolved by basing support on a notion of entailment. An argument could be seen as supporting another one if its conclusion entails the premise of the other. This would in some cases lead to a more natural notion of support. Further, it might be possible to weaken the conditions for saturated SBAFs and thus have our results for a larger class of frameworks. Further, it would be interesting to see whether the defence requirement on adequate language extensions can be formulated directly on the language level. In terms of studying the behaviour of coherent semantics, new principles need to be formulated. We have seen that principles that manipulate the support relation are difficult to apply to SBAFs, but adding and removing whole arguments might be a more promising way to study SBAF semantics. For instance, we would expect that adding new supporting arguments will not decrease the acceptability of the supported argument. It also remains to be seen whether the acceptability degrees introduced in Section 3.2.2 can be seen as a full gradual semantics.

With respect to knowledge-based reasoning, there might be a larger class of extensions to be investigated than canonical extensions. While we showed that canonical extensions are $\subseteq$-minimal amongst sensible extensions, we did not show that they are the only such extensions. Thus, it might be worthwhile to examine what we can call *cautious* extensions in general, defined as $\subseteq$-minimal sensible extensions. One benefit of that approach would be that cautious extensions are well-defined even if not all contestables have corresponding minimal arguments.

Finally, our approach provides opportunities for adding further features of argumentation. The use of preference rankings between arguments is widely established in the literature and it would be interesting to see how they could be incorporated into SBAFs. As with contestable sentences, we could add further types of sentences that come with their own dialectical obligations. For instance, we could add *controversial* sentences that can only be accepted if they are supported by an argument. This way, we could introduce a notion of necessary support into SBAFs.

Appendix A

CNFs for SBAFs

We detail how a solver for SBAF semantics can be set up using SAT methods. Some of the methods described here have been adapted from the pygarg module for Python, which is a SAT-based solver for Dung semantics (Mailly 2024).

Definition 45 (Conjunctive Normal Form)**.** *A propositional formula φ is said to be in conjunctive normal form (CNF) if it is of the form $\varphi = c_1 \wedge \cdots \wedge c_n$ with $c_i = l_1 \vee \cdots \vee l_m$ and each l_j is either a propositional variable or a negated propositional variable. We call l_j a literal and c_i a clause.*

A *SAT-model* of a CNF formula is a satisfying assignment. That is, an assignment of propositional variables to True and False such that each clause is made true.

The general idea is that for a given framework $\mathcal{SB} = \langle \mathcal{L}, A, \rightarrow, \dashrightarrow \rangle$, we construct a CNF such that models correspond to extensions. In order to construct such a CNF, we will often use the following equivalence: $(p_1 \wedge \cdots \wedge p_n) \rightarrow (q_1 \wedge \cdots \wedge q_m) \iff (\neg p_1 \vee \cdots \vee \neg p_n \vee q_1) \wedge \cdots \wedge (\neg p_1 \vee \cdots \vee \neg p_n \vee q_m)$.

We first define propositional variables for sentences and arguments. If the variable of a specific sentence or argument is true in a model of the CNF, we interpret the corresponding sentence or argument as being accepted.

The variable for a sentence s in $Sent(A)$ is denoted by p_s. The variable for an argument a in A is denoted by q_a. Following the approach in pygarg, we also introduce helper variables that indicate whether an attacker of an argument a in A is accepted. They are denoted by d_a.

We first define a CNF to encode the interpretation of the d-variables. It specifies that if a variable d_a is true, then for at least one attacker of b of a, the variable q_b must be true, and if for any attacker b of a the variable q_b is true, so is d_a.

$$CNF_d := \bigwedge_{a \in A} [(\neg d_a \vee \bigvee_{b \rightarrow a} q_b) \wedge \bigvee_{b \rightarrow a} (\neg q_b \vee d_a)].$$

Next, we make sure that the sets of accepted arguments and sentences correspond properly. Given a set of arguments, we want that all accepted sentence are components of accepted arguments. For that, it is useful to first encode argument-respect. It simply says that for each accepted argument, all its sentences must be accepted. For instance, for an argument $a = \langle \{s, t\}, u \rangle$, we use the formula $q_a \rightarrow (p_s \wedge p_t \wedge p_u)$.

$$CNF_{Argument-Respect} := \bigwedge_{a \in A} \bigwedge_{s \in Sent(a)} (\neg q_a \vee p_s).$$

The CNF for the set of sentences of an argument extension adds to argument-respect the other direction, meaning that for each accepted sentence, there must be an accepted argument containing that sentence.

$$CNF_{Language\ Set} := CNF_{Argument-Respect} \wedge \bigwedge_{s \in Sent(A)} (\neg p_s \vee \bigvee_{a \in A : s \in Sent(a)} q_a).$$

Next, we define the strong argument set, given a language extension. For that, it is useful to first define strong sentence-respect.

$$CNF_{Strong\ Sentence-Respect} := \bigwedge_{a \in A} [(\bigvee_{s \in Prem(a)} \neg p_s) \vee (\bigvee_{t \in \overline{n(a)}} p_t) \vee q_a].$$

Now we can define the strong argument set. We only need to add that no undercut of an accepted argument can be accepted. That all sentences of the argument are accepted is taken care of by the CNF for the language set.

$$CNF_{Strong\ Argument\ Set} := CNF_{Strong\ Sentence-Respect} \wedge \bigwedge_{a \in A} (\neg q_a \vee \bigvee_{t \in \overline{n(a)}} \neg p_t).$$

As the weak argument set is more difficult to define through CNFs, we will discuss it later and continue now with the CNFs for the remaining properties of extensions. Compatibility and conflict-freeness are straightforward.

$$CNF_{Compatibility} := \bigwedge_{s \in Sent(A)} \bigwedge_{t \in \overline{s}} (\neg p_s \vee \neg p_t).$$

$$CNF_{Conflict-Free} := \bigwedge_{a \in A} \bigwedge_{b \in A : b \to a} (\neg q_a \vee \neg q_b).$$

For defence, we make use of the additional d-variables. They allow us to define defence by specifying that if an argument is accepted, an attacker for each attacker also needs to be accepted.

$$CNF_{Defence} := \bigwedge_{a \in A} \bigwedge_{b \in A : b \to a} (\neg q_a \vee d_b).$$

The last property at this point is weak sentence-respect, which also makes use of d-variables the same way as in the CNF for defence.

$$CNF_{Weak\ Sentence-Respect} := \bigwedge_{a \in A} [(\bigvee_{s \in Prem(a)} \neg p_s) \vee (\bigvee_{t \in \overline{n(a)}} p_t) \vee (\bigvee_{b \in A : b \to a} \neg d_b) \vee q_a].$$

This allows us to calculate strongly and weakly coherent argument extensions as well as strongly adequate language extensions. Namely, these are models of the following CNFs:

$$CNF_{Strongly\ Coherent} :=$$
$$CNF_{Language\ Set} \wedge CNF_{Conflict-Free} \wedge CNF_d \wedge CNF_{Defence} \wedge CNF_{Strong\ Sentence-Respect}.$$

$$CNF_{Weakly\ Coherent} :=$$
$$CNF_{Language\ Set} \wedge CNF_{Conflict-Free} \wedge CNF_d \wedge CNF_{Defence} \wedge CNF_{Weak\ Sentence-Respect}.$$

$$CNF_{Strongly\ Adequate} :=$$
$$CNF_{Strong\ Argument\ Set} \wedge CNF_{Compatibility} \wedge CNF_d \wedge CNF_{Defence} \wedge CNF_{Argument-Respect}.$$

Calculating weakly adequate language extensions is more difficult. The difficulties mirror the more complex construction of the weak argument set. For instance, in order to get the initial set, we need to find a maximally admissible subset of arguments. For a given language extension, we can construct it using the following algorithm.

Algorithm 1 Finding $Init(S)$

Set $i = 1$
Define $Arg_i = \{a \in A \mid Sent(a) \subseteq S \text{ and } \overline{n(a)} \cap S = \emptyset\}$.
while Arg_i is not defended **do**
    Set i = i+1
    Define $Arg_i = \emptyset$
    for $a \in Arg_{i-1}$ **do**
        if a is defended by Arg_{i-1} **then**
            Add a to Arg_i
Return Arg_i

The next step then to use the fact that $Arg_w(S) = \bigcup_{i \in \mathbb{N}} R^S(Init(S))$ in order to construct the weak argument set by iteratively applying the respect function. Once we have the weak argument set, checking whether the language extension is weakly adequate is straightforward. However, this procedure is not directly adaptable to SAT-based methods for finding weakly adequate language extensions.

We will revisit weakly adequate language extensions, but first we show how we can calculate confident and sensible extensions. At this point, it is useful to introduce the notion of a *weighted CNF formula.*

Definition 46 (Weighted CNF)**.** *A formula in weighted CNF is a CNF formula where some clauses are associated with natural numbers representing their weight.*

A clause that is given weight is called a *soft clause* and a clause without weight is called a *hard clause.* For our purposes, it is only important whether a clause is hard or soft, thus we will denote soft-clauses with *. It is useful to consider weighted CNF formulas when unsatisfiable formulas are used. We are then interested in maximal satisfiable subsets (MSS) of the weighted CNF formula. An important constraint is that hard clauses *have to be* satisfied, while soft clauses can be removed in order to make the formula satisfiable.

Using weighted CNF formulas, it is then easy to encode confident and sensible strongly adequate language extensions. Confident and sensible strongly coherent argument extensions can then be calculated from them, at least in saturated SBAFs.

$$CNF_{Confident\ Strongly\ Adequate} := CNF_{Strongly\ Adequate} \wedge \bigwedge_{s \in Sent(A)} p_s^*.$$

$$CNF_{Sensible\ Strongly\ Adequate} := CNF_{Strongly\ Adequate} \wedge \bigwedge_{s \in L^c \cap Sent(A)} p_s^*.$$

Models of MSSs of these CNF formulas correspond to confident and sensible extensions respectively. The hard clauses make sure the extensions are strongly adequate, while the soft clauses specify that we either want to accept a maximal set of sentences overall or a maximal set of contestable sentences.

It seems clear that if we want to find weakly adequate extensions through CNF formulas, we also need them to be weighted. This is mainly because the initial argument set of a language extension is a maximally admissible subset. More precisely, given a language extension S, we can encode its initial set as follows, where the last clauses specify that the initial set should be maximal w.r.t. the arguments of which all sentences and no undercut are accepted.

$$CNF_{Initial\ Set} :=$$
$$CNF_{Conflict-Free} \wedge CNF_d \wedge CNF_{Defence} \wedge CNF_{Argument-Respect} \wedge$$
$$\bigwedge_{s \in S} p_s \wedge \bigwedge_{s \notin S} \neg p_s \wedge \bigwedge_{a \in A} [(\bigvee_{s \in Sent(a)} \neg p_s) \vee (\bigvee_{t \in \overline{n(a)}} p_t) \vee q_a]^*.$$

The problem with this approach is that we need to put in hard clauses that specify the language extension for which we want to find the initial set. Removing these clauses make the formula easily satisfiable if no sentences are accepted. Thus, we do not know how to efficiently find weakly adequate language extensions, but, given a language extension, we can efficiently test whether it is weakly adequate.

Appendix B

Index for Argumentation Frameworks and Semantics

Bibliography

Amgoud, Leila and Jonathan Ben-Naim (2013), “Ranking-based semantics for argumentation frameworks”, in *Scalable Uncertainty Management (SUM13)*, ed. by Weiru Liu, V. S. Subrahmanian, and Jef Wijsen, Springer, pp. 134-147.

— (2016), “Evaluation of arguments from support relations: Axioms and semantics”, in *International Joint Conference on Artificial Intelligence (IJCAI16)*, ed. by Subbarao Kambhampati, International Joint Conferences on Artificial Intelligence, pp. 900-906.

— (2018), “Evaluation of arguments in weighted bipolar graphs”, *International Journal of Approximate Reasoning*, 99, pp. 39-55.

Amgoud, Leila and Claudette Cayrol (1998), “On the Acceptability of Arguments in Preference-based Argumentation”, in *Uncertainty in Artificial Intelligence (UAI98)*, ed. by Gregory F. Cooper and Serafín Moral, Morgan Kaufmann Publishers, pp. 1-7.

Amgoud, Leila, Claudette Cayrol, and Marie-Christine Lagasquie-Schiex (2008), “On Bipolarity in Argumentation Frameworks”, *International Journal of intelligent Systems*, 23, pp. 1062-1093.

Amgoud, Leila and Srdjan Vesic (2011), “A new approach for preference-based argumentation frameworks”, *Annals of Mathematics and Artificial Intelligence*, 63, pp. 149-183.

Arieli, Ofer, AnneMarie Borg, Jesse Heyninck, and Christian Straßer (2021), “Logic-Based Approaches to Formal Argumentation”, *Journal of Applied Logics*, 8, 6, pp. 1793-1898.

Austin, John L. (1962), *How to Do Things With Words*, Oxford University Press.

Baroni, Pietro, Martin Caminada, and Massimiliano Giacomin (2018a), “Abstract Argumentation Frameworks and Their Semantics”, in *Handbook of Formal Argumentation. Volume 1*, ed. by Pietro Baroni, Dov Gabbay, Massimiliano Giacomin, and Leendert Van Der Torre, College Publications.

Baroni, Pietro, Massimiliano Giacomin, and Beishui Liao (2018b), “A general semi-structured formalism for computational argumentation: Definition, properties, and examples of application”, *Artificial Intelligence*, 257, pp. 158-207.

Baroni, Pietro, Antonio Rago, and Francesca Toni (2019), “From fine-grained properties to broad principles for gradual argumentation: A principled spectrum”, *International Journal of Approximate Reasoning*, 105, pp. 252-286.

Baroni, Pietro, Marco Romano, Francesca Toni, Marco Aurisicchio, and Giorgio Bertanza (2015), “Automatic evaluation of design alternatives with quantitative argumentation”, *Argument and Computation*, 6, 1, pp. 24-49.

Besnard, Philippe, Alejandro J. García, Anthony Hunter, Sanjay Modgil, Henry Prakken, Guillermo R. Simari, and Francesca Toni (2014), “Introduction to structured argumentation”, *Argument and Computation*, 5, 1, pp. 1-4.

Besnard, Philippe and Anthony Hunter (2001), “A logic-based theory of deductive arguments”, *Artificial Intelligence*, 128, pp. 203-235.

Betz, Gregor (2010), *Theorie dialektischer Strukturen*, Klostermann.
— (2012), “On Degrees of Justification”, *Erkenntnis*, 77, pp. 237-272.
— (2016), “Assessing the epistemological relevance of Dung-style argumentation theories”, *Annals of Mathematics and Artificial Intelligence*, 78, pp. 303-321.

Boella, Guido, Dov Gabbay, Leendert van der Torre, and Serena Villata (2010), “Support in Abstract Argumentation”, in *Computational Models of Argument (COMMA10)*, ed. by Pietro Baroni, Federico Cerutti, Massimiliano Giacomin, and Guillermo R. Simari, IOS Press, pp. 111-122.

Bondarenko, Andrei, Phan Minh Dung, Robert Anthony Kowalski, and Francesca Toni (1997), “An abstract, argumentation-theoretic approach to default reasoning”, *Artificial Intelligence*, 93, pp. 63-101.

Bondarenko, Andrei, Francesca Toni, and Robert Anthony Kowalski (1993), “An assumption-based framework for non-monotonic reasoning”, in *Proceedings of the second international workshop on Logic programming and non-monotonic reasoning*, ed. by Luís Moniz Pereira and Anil Nerode, MIT Press.

Brun, Georg and Gregor Betz (2016), “Analysing Practical Argumentation”, in *The Argumentative Turn in Policy Analysis*, ed. by Sven Ove Hansson and Gertrude Hirsch Hadorn, Springer International Publishing, pp. 39-77.

Cayrol, Claudette and Marie-Christine Lagasquie-Schiex (2005a), “Gradual Valuation for Bipolar Argumentation Frameworks”, in *Symbolic and Quantitative Approaches to Reasoning with Uncertainty (ECSQARU05)*, ed. by Lluís Godo, Springer, pp. 366-377.
— (2005b), “On the acceptability of arguments in bipolar argumentation frameworks”, in *Symbolic and Quantitative Approaches to Reasoning with Uncertainty (ECSQARU05)*, ed. by Lluis Godo, Springer, pp. 378-389.
— (2009), “Bipolar abstract argumentation systems”, in *Argumentation in Artificial Intelligence*, ed. by Iyad Rahwan and Guillermo R. Simari, Springer, pp. 65-84.
— (2013), “Bipolarity in argumentation graphs: Towards a better understanding”, *International Journal of Approximate Reasoning*, 54, pp. 876-899.

Cohen, Andrea, Alejandro J. García, and Guillermo R. Simari (2012), “Backing and undercutting in abstract argumentation frameworks”, in *Foundations of Information and Knowledge Systems (FoIKS12)*, ed. by Thomas Lukasiewicz and Attila Sali, Springer, pp. 107-123.

Cohen, Andrea, Simon Parsons, Elizabeth I. Sklar, and Peter McBurney (2018), “A characterization of types of support between structured arguments and their relationship with support in abstract argumentation”, *International Journal of Approximate Reasoning*, 94, pp. 76-104.

Čyras, Kristijonas, Xiuyi Fan, Claudia Schulz, and Francesca Toni (2018), “Assumption-Based Argumentation: Disputes, Explanations, Preferences”, in *Handbook of Formal Argumentation. Volume 1*, ed. by Pietro Baroni, Dov Gabbay, Massimiliano Giacomin, and Leendert van der Torre, College Publications, pp. 365-408.

Čyras, Kristijonas, Claudia Schulz, and Francesca Toni (2017), “Capturing Bipolar Argumentation in Non-flat Assumption-Based Argumentation”, in *Principles and Practice of Multi-Agent Systems (PRIMA17)*, ed. by Bo An, Ana Bazzan, João Leite, Serena Villata, and Leendert van der Torre, Springer International Publishing, pp. 386-402.

Davey, Brian A. and Hilary A. Priestly (2002), *Introduction to Lattices and Order. Second Edition*, Cambridge University Press.

Dondio, Pierpaolo (2018), "Ranking semantics based on subgraphs analysis", in *Autonomous Agents and Multiagent Systems (AAMAS18)*, ed. by Mehdi Dastani, Gita Sukthankar, Elisabeth Andre, and Sven Koenig, International Foundation for Autonomous Agents and Multiagent Systems, pp. 1132-1140.

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

Gargouri, Anis, Sébastien Konieczny, Pierre Marquis, and Srdjan Vesic (2021), "On a notion of monotonic support for bipolar argumentation frameworks", in *Autonomous Agents and Multiagent Systems (AAMAS21)*, ed. by Ulle Endriss, Ann Nowé, Frank Dignum, and Alessio Lomuscio, International Foundation for Autonomous Agents and Multiagent Systems, pp. 546-554.

Grossi, Davide and Sanjay Modgil (2019), "On the graded acceptability of arguments in abstract and instantiated argumentation", *Artificial Intelligence*, 275, pp. 138-173.

Hecham, Abdelraouf, Pierre Bisquert, and Madalina Croitoru (2018), "On a flexible representation for defeasible reasoning variants", in *Autonomous Agents and Multiagent Systems (AAMAS18)*, ed. by Mehdi Dastani, Gita Sukthankar, Elisabeth Andre, and Sven Koenig, International Foundation for Autonomous Agents and Multiagent Systems, pp. 1123-1131.

— (2020), "A Formalism Unifying Defeasible Logics and Repair Semantics for Existential Rules", in *Ontologies and Concepts in Mind and Machine (ICCS20)*, ed. by Mehwish Alam, Tanya Braun, and Bruno Yun, Springer Nature Switzerland, pp. 3-17.

Johnson, Ralph H. (1996), *The Rise of Informal Logic: Essays on Argumentation, Critical Thinking, Reasoning and Politics*, Vale Press.

— (2000), *Manifest rationality: A pragmatic theory of argument*, Routledge.

Johnson, Ralph H. and J. Anthony Blair (2006), *Logical Self-Defense*, International Debate Education Association.

Kaci, Souhila, Leendert van der Torre, Srdjan Vesic, and Serena Villata (2018), "Preference in abstract argumentation", *Frontiers in Artificial Intelligence and Applications*, 305, pp. 405-412.

Koszowy, Marcin, Katarzyna Budzynska, Barbara Konat, Steve Oswald, and Pascal Gygax (2022), "A Pragmatic Account of Rephrase in Argumentation: Linguistic and Cognitive Evidence", *Informal Logic*, 42, 1, pp. 49-82.

Krabbe, Erik C. W. (2017), "The Formalization of Critical Discussion", *Argumentation*, 31, pp. 101-119.

Mailly, Jean-Guy (2024), "pygarg : A Python Engine for Argumentation", *Technical report No. IRIT/RR–2024–02–FR (version 1.1)*.

Modgil, Sanjay and Henry Prakken (2013), "A general account of argumentation with preferences", *Artificial Intelligence*, 195, pp. 361-397.

— (2014), "The ASPIC+ framework for structured argumentation: A tutorial", *Argument and Computation*, 5, 1, pp. 31-62.

Mossakowski, Till and Fabian Neuhaus (2016), "Bipolar Weighted Argumentation Graphs", *arXiv*.
— (2018), "Modular Semantics and Characteristics for Bipolar Weighted Argumentation Graphs", *arXiv*.

Nouioua, Farid (2013), "AFs with Necessities: Further Semantics and Labelling Characterization", in *Scalable Uncertainty Management (SUM13)*, ed. by Weiru Liu, V. S. Subrahmanian, and Jef Wijsen, Springer, pp. 120-133.

Nouioua, Farid and Vincent Risch (2011), "Argumentation frameworks with necessities", in *Scalable Uncertainty Management (SUM11)*, ed. by Salem Benferhat and John Grant, Springer, pp. 163-176.

Oren, Nir and Timothy J. Norman (2008), "Semantics for Evidence-Based Argumentation", in *Computational Models of Argument (COMMA08)*, ed. by Philippe Besnard, Sylvie Doutre, and Anthony Hunter, IOS Press, pp. 276-284.

Pandžić, Stipe (2022), "Structured argumentation dynamics", *Annals of Mathematics and Artificial Intelligence*, 90, pp. 297-337.

Perelman, Chaïm and Lucie Olbrechts-Tyteca (1969), *The New Rhetoric. A Treatise on Argumentation*, University of Notre Dame Press.

Polberg, Sylwia and Anthony Hunter (2018), "Empirical evaluation of abstract argumentation: Supporting the need for bipolar and probabilistic approaches", *International Journal of Approximate Reasoning*, 93, pp. 487-543.

Polberg, Sylwia and Nir Oren (2014), "Revisiting Support in Abstract Argumentation Systems", *dbai Technical Report*.

Pollock, John L. (1987), "Defeasible Reasoning", *Cognitive Science*, 11, pp. 481-518.

Potyka, Nico (2018), "Continuous Dynamical Systems for Weighted Bipolar Argumentation", in *Principles of Knowledge Representation and Reasoning (KR18)*, Association for the Advancement of Artificial Intelligence, pp. 148-157.
— (2021), "Generalizing Complete Semantics to Bipolar Argumentation Frameworks", in *Symbolic and Quantitative Approaches to Reasoning with Uncertainty (ECSQARU21)*, ed. by Jiřina Vejnarová and Nic Wilson, Springer, pp. 130-143.

Prakken, Henry (2010), "An abstract framework for argumentation with structured arguments", *Argument and Computation*, 1, 2, pp. 93-124.
— (2014), "On support relations in abstract argumentation as abstractions of inferential relations", *Frontiers in Artificial Intelligence and Applications*, 263, pp. 735-740.

Pu, Fuan, Jian Luo, Yulai Zhang, and Guiming Luo (2014), "Argument ranking with categoriser function", in *Knowledge Science, Engineering and Management (KSEM14)*, ed. by Robert Buchmann, Claudiu Vasile Kifor, and Jian Yu, Springer, pp. 290-301.
— (2015), "Attacker and Defender Counting Approach for Abstract Argumentation", in *Proceedings of the 37th Annual Meeting of the Cognitive Science Society (CogSci15)*, ed. by David C. Noelle, Rick Dale, Anne S. Warlaumont, Jeff Yoshimi, Teenie Matlock, Carolyn D. Jennings, and Paul P. Maglio, pp. 1913-1918.

Rago, Antonio, Francesca Toni, Marco Aurisicchio, and Pietro Baroni (2016), "Discontinuity-free decision support with quantitative argumentation debates", in *Principles Of Knowledge Representation And Reasoning (KR16)*, Association for the Advancement of Artificial Intelligence, pp. 63-72.

Rigotti, Eddo and Sara Greco (2019), *Inference in Argumentation. A Topics-Based Approach to Argument Schemes*, Springer Nature Switzerland.

Sakama, Chiaki (2024), "Argument and Belief", in *Computational Models of Argument (COMMA24)*, ed. by Chris Reed, Matthias Thimm, and Tjitze Rienstra, IOS Press, pp. 241-252.

Sakama, Chiaki and Tran Cao Son (2020), "Epistemic argumentation framework: Theory and computation", *Journal of Artificial Intelligence Research*, 69, pp. 1103-1126.

Shi, Chenwei, Sonja Smets, and Fernando R. Velázquez-Quesada (2018), "Beliefs supported by binary arguments", *Journal of Applied Non-Classical Logics*, 28, 2-3, pp. 165-188.

Toni, Francesca (2008), "Assumption-Based Argumentation for Closed and Consistent Defeasible Reasoning", in *New Frontiers in Artificial Intelligence*, ed. by Ken Satoh, Akihiro Inokuchi, Katashi Nagao, and Takahiro Kawamura, Springer, pp. 390-402.
— (2014), "A tutorial on assumption-based argumentation", *Argument and Computation*, 5, 1, pp. 89-117.

Toulmin, Stephen E. (2003), *The Uses of Argument: Updated Edition*, Cambridge University Press.

Van Eemeren, Frans H. (2010), *Strategic Maneuvering in Argumentative Discourse*, John Benjamins.

Van Eemeren, Frans H., Bart Garssen, Sara Greco, Ton van Haaften, Nanon Labrie, Fernando Leal, and Peng Wu (2022), *Argumentative Style. A pragma-dialectical study of functional variety in argumentative discourse*, John Benjamins.

Van Eemeren, Frans H., Bart Garssen, Bart Verheij, Erik C. W. Krabbe, A. Francisca Snoeck Henkemans, and Jean H. M. Wagemans (2020), *Handbook of Argumentation Theory*, Springer.

Van Eemeren, Frans H. and Rob Grootendorst (1984), *Speech acts in argumentative discussions. A theoretical model for analysis of discussions directed towards solving conflicts of opinion*, Foris Publications.
— (1992), *Argumentation, Communication, and Fallacies: A Pragma-dialectical Perspective*, Routledge.
— (2004), *A Systematic Theory of Argumentation. The pragma-dialectical approach*, Cambridge University Press.

Visser, Jacky (2013), "A formal account of complex argumentation in critical discussion", in *Virtues of argumentation: proceedings of the 10th International Conference of the Ontario Society for the Study of Argumentation (OSSA13)*, ed. by Dima Mohammed and Marcin Lewiński, OSSA, pp. 1-14.
— (2015), "A Formal Perspective On The Pragma-Dialectical Discussion Model", in *International Society for the Study of Argumentation (ISSA15)*, ed. by Bart Garssen, David Godden, Mitchell Green, and A. Francisca Snoeck Henkemans, University of Amsterdam, pp. 1471-1482.
— (2017), "Speech Acts in a Dialogue Game Formalisation of Critical Discussion", *Argumentation*, 31, pp. 245-266.

Wallner, Johannes P., Adam Wyner, and Tomasz Zurek (2024), "Value-Based Reasoning in ASPIC^{+}", in *Computational Models of Argument (COMMA24)*, ed. by Chris Reed, Matthias Thimm, and Tjitze Rienstra, IOS Press, pp. 325-336.

Walton, Douglas (2010), “Types of dialogue and burdens of proof”, *Frontiers in Artificial Intelligence and Applications*, 216, pp. 13-24.

Walton, Douglas and Marcin Koszowy (2017), “Arguments from authority and expert opinion in computational argumentation systems”, *AI and Society*, 32, pp. 483-496.

Walton, Douglas, Chris Reed, and Fabrizio Macagno (2008), *Argumentation Schemes*, Cambridge University Press.

Yu, Liuwen, Caren Al Anaissy, Srdjan Vesic, Xu Li, and Leendert van der Torre (2023), “A Principle-Based Analysis of Bipolar Argumentation Semantics”, in *Logics in Artificial Intelligence (JELIA23)*, ed. by Sarah Gaggl, Maria Vanina Martinez, and Magdalena Ortiz, Springer, pp. 209-224.

Yu, Shiyang and Frank Zenker (2020), “Schemes, Critical Questions, and Complete Argument Evaluation”, *Argumentation*, 34, pp. 469-498.

Yun, Bruno and Srdjan Vesic (2021), “Gradual Semantics for Weighted Bipolar SETAFs”, in *Symbolic and Quantitative Approaches to Reasoning with Uncertainty (ECSQARU21)*, ed. by Jiřina Vejnarová and Nic Wilson, Springer, pp. 201-214.

Yun, Bruno, Srdjan Vesic, and Madalina Croitoru (2020), “Sets of attacking arguments for inconsistent datalog knowledge bases”, *Frontiers in Artificial Intelligence and Applications*, 326, pp. 419-430.