Aggregating Bipolar Opinions

Stefan Lauren
Imperial College London, UK
s18@imperial.ac.uk

Francesco Belardinelli
Imperial College London, UK
Universite d'Evry, France
francesco.belardinelli@imperial.ac.uk

Francesca Toni
Imperial College London,UK
f.toni@imperial.ac.uk

ABSTRACT

We introduce a novel method to aggregate Bipolar Argumentation (BA) Frameworks expressing opinions by different parties in debates. We use Bipolar Assumption-based Argumentation (ABA) as an all-encompassing formalism for BA under different semantics. By leveraging on recent results on judgement aggregation in Social Choice Theory, we prove several preservation results, both positive and negative, for relevant properties of Bipolar ABA.

KEYWORDS

Bipolar Argumentation; Judgement Aggregation; Social Choice

ACM Reference Format:
Stefan Lauren, Francesco Belardinelli, and Francesca Toni. 2021. Aggregating Bipolar Opinions. In *Proc. of the 20th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2021), Online, May 3–7, 2021*, IFAAMAS, 9 pages.

1 INTRODUCTION

There is a long and well-established tradition in knowledge representation and reasoning to formally describe debates as exchanges of opinions by the parties involved through attacks [9] and supports [21] between arguments understood abstractly as in *Abstract Argumentation (AA)* [13] or as in *Bipolar Argumentation* (BA) [7, 8]. When these debates emerge in multi-agent systems [15, 21] a key question concerns opinion aggregation, namely how we can obtain a collective consensus from several opinions expressed as argumentation frameworks, in such a way that the agents' opinions are well-portrayed in the collective outcome [4]. Recently, Chen and Endriss [9, 10] applied aggregation procedures from Social Choice Theory [18] to AA Frameworks (AAFs), by leveraging especially on judgement aggregation [17], and proving the preservation of interesting properties of AAFs, such as conflict-freeness, acyclicity and extensions according to several semantics. Some efforts had been made to provide procedures for the aggregation of *Quantitative Argumentation Debate Frameworks* [1, 21, 22] (a form of *BA Frameworks (BAFs)* incorporating both attacks and supports but equipped with gradual, rather than extension-based, semantics), but the aggregation of BAFs is an open problem, and a challenging one mainly because there are different semantic interpretations of support in BAFs (e.g., as deductive or necessary [5, 8, 20]), which might generate "inconsistencies" in the aggregated framework.

Moving from the considerations above, we advance the state of the art on the application of Social Choice Theory to opinion aggregation in computational argumentation by focusing on BA. To address the problem that different parties may adopt different interpretations of support in their opinions (BAFs), we use *Bipolar Assumption-based Argumentation (ABA) frameworks* [12] for representing opinions. Bipolar ABA is a restricted (but "non-flat") form of ABA providing a unified formalism to accommodate different interpretations of support [12]. Thus, by adopting Bipolar ABA, we let parties choose their interpretation of support before aggregation takes place. Then, our contribution is twofold. Firstly, we define aggregation procedure for Bipolar ABA frameworks based on Social Choice Theory. Our investigations mainly focus on quota and oligarchic rules [17]. Secondly, we conduct a study on the preservation of properties using the defined aggregation procedures. In some cases, restrictions need to be placed, for example, towards specific aggregation rules.

Structure of the Paper. In Section 2 we provide the necessary background on Bipolar ABA. Section 3 sets the ground, by formulating the aggregation problem, aggregation rules and preservation properties in our Bipolar ABA setting. Section 4 gives the main theoretical contribution of the paper, providing preservation results for various properties of Bipolar ABA frameworks. Finally, Section 5 concludes and elaborates on several promising directions for future works. Because of space limit, we omit some proofs: these can be found in http://arxiv.org/abs/2102.02881.

2 BIPOLAR ABA

Bipolar Assumption-based Argumentation [12] (Bipolar ABA) is a form of structured argumentation, where arguments and attacks are derived from assumptions, rules, and a contrary map from assumptions. Note that contrary should not be confused with negation, which may or may not occur in the underlying language [23].

Definition 1. *[12] A* Bipolar ABA framework *is a quadruple* $\langle \mathcal{L}, \mathcal{R}, \mathcal{A}, \overline{} \rangle$*, where*

- *$(\mathcal{L}, \mathcal{R})$ is a* deductive system *with $\mathcal{L}$ a* language *(i.e. a set of sentences) and $\mathcal{R}$ a set of* rules *of the form $\phi \leftarrow \alpha$ where $\alpha \in \mathcal{A}$ and either $\phi \in \mathcal{A}$ or $\phi = \bar{\beta}$ for some $\beta \in \mathcal{A}$; ϕ is the* head *and α the* body *of rule $\phi \leftarrow \alpha$;*
- *$\mathcal{A} \subseteq \mathcal{L}$ is a non-empty set of* assumptions*;*
- *$\overline{} : \mathcal{A} \rightarrow \mathcal{L}$ is a total map; for $\alpha \in \mathcal{A}$, $\bar{\alpha}$ is the* contrary *of α.*

Then, a *deduction* for $\phi \in \mathcal{L}$ supported by $A \subseteq \mathcal{A}$ and $R \subseteq \mathcal{R}$, denoted $A \vdash^R \phi$, is a finite tree with the root labelled by ϕ; leaves labelled by assumptions, with A the set of all such assumptions; and each non-leaf node ψ has a single child labelled by the body of some ψ-headed rule in $\mathcal{R}$, with R the set of all such rules.

Note that in Bipolar ABA rules are of a restricted kind, in comparison with generic ABA [6]: their bodies amount to a single assumption, and thus, in particular, there are no rules with an empty body; also, their heads are either assumptions or contraries thereof. Because assumptions may be "deducible" from rules in Bipolar ABA

Proc. of the 20th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2021), U. Endriss, A. Nowé, F. Dignum, A. Lomuscio (eds.), May 3–7, 2021, Online.

Table 1: Bipolar ABA Semantics (for extension $A \subseteq \mathcal{A}$).

Semantics	Conditions
Admissible	A is closed, conflict-free and for every $B \subseteq \mathcal{A}$, if B is closed and attacks A, then A attacks B.
Preferred	A is $\subseteq$-maximally admissible.
Complete	A is admissible and $A = \{\alpha \in \mathcal{A} : A \text{ defends } \alpha\}$ where A *defends* $\alpha \in \mathcal{A}$ iff for all closed $B \subseteq \mathcal{A}$: if B attacks α then A attacks B.
Set-stable	A is closed, conflict-free, and attacks $Cl(\beta)$ for each $\beta \in \mathcal{A} \setminus A$.
Well-founded	A is the intersection of all complete extensions.
Ideal	A is $\subseteq$-maximal such that it is admissible and $A \subseteq B$ for all preferred extensions $B \subseteq \mathcal{A}$.

frameworks, though, these frameworks may be *non-flat* in general, thus lacking some of the properties that *flat* ABA frameworks (where assumptions are not "deducible" from rules) exhibit [11].

In Bipolar ABA, $A \subseteq \mathcal{A}$ *attacks* $\beta \in \mathcal{A}$ iff $\exists A' \vdash^R \bar{\beta}$, such that $A' \subseteq A$; $\alpha \in \mathcal{A}$ *attacks* $\beta \in \mathcal{A}$ iff $\{\alpha\}$ attacks β; $A \subseteq \mathcal{A}$ *attacks* $B \subseteq \mathcal{A}$ iff $\exists \beta \in B$ such that A attacks β. Then, A is *conflict-free* iff A does not attacks A. Let the *closure of* $A \subseteq \mathcal{A}$ be $Cl(A) = \{\alpha \in \mathcal{A} : \exists A' \vdash^R \alpha, A' \subseteq A, R \subseteq \mathcal{R}\}$. Then, A is *closed* iff $A = Cl(A)$.

Several conditions can be imposed on Bipolar ABA frameworks, which characterise different sets of assumptions (also called *extensions*) according to as many *semantics*. Table 1 gives the semantics for Bipolar ABA frameworks we will analyse in the paper, in addition to properties of conflict-freeness and closedness. As a simple illustration of these semantics, consider a Bipolar ABA framework with $\mathcal{L} = \{\alpha, \beta, \gamma, \bar{\alpha}, \bar{\beta}, \bar{\gamma}\}$, $\mathcal{A} = \{\alpha, \beta, \gamma\}$ and $\mathcal{R} = \{\bar{\beta} \leftarrow \gamma, \gamma \leftarrow \alpha\}$.[1] Then, $\{\alpha\}$ and $\{\alpha, \beta\}$ are not closed (and thus not admissible etc.), $\{\beta\}$ is closed and conflict-free but not admissible etc., and $\{\alpha, \gamma\}$ is closed, conflict-free, admissible (etc.). Note that, in Bipolar ABA (as in flat ABA, but not in general ABA [11]), admissible, preferred and ideal extensions are guaranteed to exist: in particular, since rules cannot have an empty body, the empty set of assumption is closed, and thus admissible. Instead, complete, well-founded and set-stable extensions may not exist [12].

Bipolar ABA provides an all-encompassing framework for capturing different interpretations of support (under admissible, preferred and set-stable semantics) [12], as illustrated in the following example, adapted from [21].[2]

Example 1. *The UK public may hold a range of views on Brexit:*
A: The UK should leave the EU.
B: The UK staying in the EU is good for its economy.
C: The EU's immigration policies are bad for the UK's economy.
D: EU membership fees are too high.
E: The UK staying in the EU is good for world peace.
Here A may be deemed to be attacked by B and E, but supported by C and D. A Bipolar ABA representation for the deductive interpretation *of support [8] is* $\langle \mathcal{L}, \mathcal{R}, \mathcal{A}, \overline{} \rangle$*, where*

- $\mathcal{L} = \{A, B, C, D, E, \bar{A}, \bar{B}, \bar{C}, \bar{D}, \bar{E}\}$;
- $\mathcal{R} = \{\bar{A} \leftarrow B, \quad \bar{A} \leftarrow E, \quad A \leftarrow C, \quad A \leftarrow D\}$;
- $\mathcal{A} = \{A, B, C, D, E\}$.

Instead, a Bipolar ABA representation for the necessary interpretation *of support [20] is* $\langle \mathcal{L}, \mathcal{R}', \mathcal{A}, \overline{} \rangle$ *with* $\mathcal{R}' = \{\bar{A} \leftarrow B, \quad \bar{A} \leftarrow E, \quad C \leftarrow A, \quad D \leftarrow A\}$. *(Other interpretations of support can also be represented but are not illustrated here for lack of space).*

Note that Bipolar ABA can naturally capture *supported attacks* under the deductive interpretation of support [8], for example the supported attack from α to β in a (standard) BAF where α supports γ and γ attacks β is matched by α attacking β in a Bipolar ABA framework where $\{\bar{\beta} \leftarrow \gamma, \gamma \leftarrow \alpha\} \subseteq \mathcal{R}$.

3 SOCIAL CHOICE FOR BIPOLAR ABA

Social choice theory mainly focuses on how to aggregate (people's or agents') opinions into a single collective decision. Broadly speaking, there are mainly two types of aggregation: preference aggregation [2, 18] and judgement aggregation [18] (other aggregation types exist but are omitted here). Here we focus on the latter, and adapt notions given by [9] to accommodate opinions drawn from Bipolar ABA frameworks.

Hereafter we assume a set of agents $N = \{1, \ldots, n\}$ $(n > 1)$, with agents' opinions represented as Bipolar ABA frameworks. We assume that the (Bipolar ABA frameworks of) agents have the same language, assumptions, and contraries. Thus, the aggregation combines the rules of the frameworks. This is intuitive because aggregating the Bipolar ABA frameworks means aggregating the attacks and supports in the original BAFs, which correspond to the rules in Bipolar ABA frameworks. We also assume that agents behave "rationally", for example that their opinions are not "self-attacking" , e.g., rules $\bar{\alpha} \leftarrow \alpha$ (for $\alpha \in \mathcal{A}$) will never belong to $\mathcal{R}$; further elaboration on agents' rationality, when they are defined argumentatively, can be found in [21].

To collectively combine the opinions of all agents, i.e., the agents' Bipolar ABA frameworks, we need aggregation rules.

Definition 2. *Let* $\boldsymbol{\mathcal{F}}$ *be the set of all Bipolar ABA frameworks with the same language* $\mathcal{L}$*, set of assumptions* $\mathcal{A}$ *and contrary mapping* $\overline{}$*. A* Bipolar ABA aggregation rule *is a mapping* $F : \boldsymbol{\mathcal{F}}^n \rightarrow \boldsymbol{\mathcal{F}}$ *from n Bipolar ABA frameworks into a single Bipolar ABA framework. Given n (as opinions of agents in* $N = \{1, \ldots, n\}$*) Bipolar ABA frameworks* $\langle \mathcal{L}, \mathcal{R}_1, \mathcal{A}, \overline{} \rangle, \ldots, \langle \mathcal{L}, \mathcal{R}_n, \mathcal{A}, \overline{} \rangle$*, F returns a single* aggregated Bipolar ABA framework $\langle \mathcal{L}, \mathcal{R}_{agg}, \mathcal{A}, \overline{} \rangle$.

Inspired by graph aggregation [14], we restrict attention in this paper to aggregation rules in the form of either quota rules or oligarchic rules, defined below in our setting.

Quota Rules. These set a quota $q \in N$ as a threshold to accept some set of rules $R \subseteq \boldsymbol{\mathcal{R}} = \bigcup_{i \in N} \mathcal{R}_i$, i.e., there should be at least q agents that accept the rules in R.

Definition 3. *The* quota rule F_q*, for* $q \in N$*, is a Bipolar ABA aggregation rule such that* $F_q(\langle \mathcal{L}, \mathcal{R}_1, \mathcal{A}, \overline{} \rangle, \ldots, \langle \mathcal{L}, \mathcal{R}_n, \mathcal{A}, \overline{} \rangle) = \{r \in \boldsymbol{\mathcal{R}} : r \in \bigcap_{j \in N'} \mathcal{R}_j \text{ for } N' \subseteq N, \ |N'| \geq q\}$.

The quota q can be any number, but there are several special quotas that are commonly used: *weak majority* has a quota $q = \lfloor \frac{n}{2} \rfloor$;

[1]With an abuse of notation, we use $\bar{x}$ to denote the $\mathcal{L}$-sentence amounting to the contrary of $x \in \mathcal{A}$ and omit to specify the contrary map explicitly.
[2]The formal definitions of how different interpretations of support are captured in Bipolar ABA are in [12], and outside the scope of this paper.

strict majority has a quota $q = \lceil \frac{n}{2} \rceil$; *nomination* accepts all rules accepted by at least 1 agent, i.e., $q = 1$; and *unanimity* requires all agents to accept the same rules, i.e., $q = n$.

For example, assume that $n = 3$ and agents accept sets of rules $\mathcal{R}_1 = \{\bar{A} \leftarrow B\}$, $\mathcal{R}_2 = \{A \leftarrow C\}$, and $\mathcal{R}_3 = \{\bar{A} \leftarrow B, \quad A \leftarrow D\}$. Using weak majority, the aggregated Bipolar ABA framework has set of rules $\mathcal{R}_{agg} = \{\bar{A} \leftarrow B, \quad A \leftarrow C, \quad A \leftarrow D\}$; with strict majority, $\mathcal{R}_{agg} = \{\bar{A} \leftarrow B\}$; nomination gives $\mathcal{R}_{agg} = \{\bar{A} \leftarrow B, \quad A \leftarrow C, \quad A \leftarrow D\}$, while unanimity returns $\mathcal{R}_{agg} = \{\}$.

Oligarchic Rules. These give agents the power to veto the accepted rule sets. Clearly, oligarchic rules are not fair in that the opinions of agents without veto power are disregarded. However, in some cases they are necessary to avoid conflicts among the agents.

Definition 4. *Let $N_v \subseteq N$ be the* agents with veto power. *The* oligarchic rule *F_o is a Bipolar ABA aggregation rule such that* $F_o(\langle \mathcal{L}, \mathcal{R}_1, \mathcal{A}, \overline{} \rangle, \ldots, \langle \mathcal{L}, \mathcal{R}_n, \mathcal{A}, \overline{} \rangle) = \{r \in \boldsymbol{\mathcal{R}} : r \in \bigcap_{j \in N_v} \mathcal{R}_j\}$. *If $|N_v| = 1$ then the oligarchic rule is called* dictatorship.

If all agents have veto powers, then the oligarchic rule coincides with unanimity. As an example, assume, as above, that $n = 3$ and agents accept sets of rules $\mathcal{R}_1 = \{\bar{A} \leftarrow B\}$, $\mathcal{R}_2 = \{A \leftarrow C\}$, and $\mathcal{R}_3 = \{\bar{A} \leftarrow B, \quad A \leftarrow D\}$. If agents 1 and 3 are given veto powers, then the aggregated Bipolar ABA framework has set of rules $\mathcal{R}_{agg} = \{\bar{A} \leftarrow B\}$. On the other hand, if all agents have veto powers, then $\mathcal{R}_{agg} = \{\}$, as with unanimity.

We will study whether the properties of the agents' Bipolar ABA frameworks, including semantics, conflict-freeness and closedness, and other "graph" properties, are preserved in the aggregated Bipolar ABA framework. To produce stronger preservation results, we assume that a property under consideration for an aggregated Bipolar ABA framework (obtained by applying some Bipolar ABA aggregation rule) needs to be satisfied by all agents.

Definition 5. *Let P be a property of Bipolar ABA frameworks. If $\Delta \subseteq \mathcal{A}$ is P in each agents' Bipolar ABA framework $\langle \mathcal{L}, \mathcal{R}_i, \mathcal{A}, \overline{} \rangle$ (with $i \in N$), then P is* preserved *in the aggregated Bipolar ABA framework $\mathcal{F} = \langle \mathcal{L}, \mathcal{R}_{agg}, \mathcal{A}, \overline{} \rangle$ if and only if Δ is P in $\mathcal{F}$.*

4 PRESERVATION RESULTS

In this section, we present preservation results for several properties P (see Definition 5), specifically for P equal to conflict-freeness and closedness of sets of assumptions, P any of the semantics in Table 1, P amounting to assumption acceptability under these semantics, and P amounting to (implicative and disjunctive) "graph" properties adapted from [14]. Throughout, we will assume that $\mathcal{F} = \langle \mathcal{L}, \mathcal{R}_{agg}, \mathcal{A}, \overline{} \rangle$ is the aggregated Bipolar ABA framework resulting from the Bipolar ABA aggregation rules as considered, and, for any such rule, we will say that the rule preserves a property P to mean that P is preserved in $\mathcal{F}$, as specified in Definition 5.

4.1 Conflict-freeness

Conflict-freeness is a basic property in argumentation, always preserved in our setting.

Theorem 4.1. *Every quota rule and oligarchic rule preserves conflict-freeness.*

Proof. Assume that $\Delta \subseteq \mathcal{A}$ is conflict-free in $\langle \mathcal{L}, \mathcal{R}_i, \mathcal{A}, \overline{} \rangle$ for all $i \in N$. By contradiction, assume that Δ is not conflict-free in $\mathcal{F}$. Then $\exists \alpha, \beta \in \Delta$ such that α attacks β, i.e., $\exists R = \{\bar{\beta} \leftarrow \gamma_1, \ \ldots, \ \gamma_{m-1} \leftarrow \gamma_m\} \subseteq \mathcal{R}_{agg}$ for $m \geq 1$, $\{\gamma_1, \ldots, \gamma_m\} \subseteq \mathcal{A}$ and $\gamma_m = \alpha$. By definition of quota and oligarchic rules, there has to be at least one agent $i \in N$ such that $R \subseteq \mathcal{R}_i$ and thus α attacks β in $\langle \mathcal{L}, \mathcal{R}_i, \mathcal{A}, \overline{} \rangle$, thus contradicting our assumption that Δ is conflict-free in $\langle \mathcal{L}, \mathcal{R}_i, \mathcal{A}, \overline{} \rangle$ for all $i \in N$. □

Note that Theorem 4.1 is a direct extension of Theorem 2 in [9] because conflict-freeness in Bipolar ABA frameworks holds under the same conditions as in AAFs.

4.2 Closedness

Closedness is an important property in Bipolar ABA frameworks, made so by the presence of "support" between assumptions (in the form of rules whose head and body are assumptions). Instead, it is not meaningful in AA, and indeed it is not studied in [9].

Theorem 4.2. *Every quota and oligarchic rule preserves closedness.*

Proof. Assume that $\Delta \subseteq \mathcal{A}$ is closed in $\langle \mathcal{L}, \mathcal{R}_i, \mathcal{A}, \overline{} \rangle$ for all $i \in N$. By contradiction, assume that Δ is not closed in $\mathcal{F}$. Then $\exists \alpha \in \Delta$ and $\beta \notin \Delta$ such that $\beta \in Cl(\{\alpha\})$, i.e. there exists $R = \{\beta \leftarrow \gamma_1, \ \ldots, \ \gamma_{m-1} \leftarrow \gamma_m\} \subseteq \mathcal{R}_{agg}$ for $m \geq 1$, $\{\gamma_1, \ldots, \gamma_m\} \subseteq \mathcal{A}$ and $\gamma_m = \alpha$. By definition of quota and oligarchic rules, there has to be at least one agent $i \in N$ such that $R \subseteq \mathcal{R}_i$ and thus $\beta \in Cl(\{\alpha\})$ in $\langle \mathcal{L}, \mathcal{R}_i, \mathcal{A}, \overline{} \rangle$, thus contradicting our assumption that Δ is closed in $\langle \mathcal{L}, \mathcal{R}_i, \mathcal{A}, \overline{} \rangle$ for all $i \in N$.

□

4.3 Admissible Extensions

The preservation result below (Theorem 4.3) for admissibility extends Theorem 3 in [9] by considering also support. It assumes constraints on the number of assumptions. Theorem 4.4 below instead analyses the preservation of admissibility for corner cases.

Theorem 4.3. *For $|\mathcal{A}| \geq 4$, nomination is the only quota rule that preserves admissibility.*

Proof. First we prove that nomination preserves admissibility. Assume that $\Delta \subseteq \mathcal{A}$ is admissible in $\langle \mathcal{L}, \mathcal{R}_i, \mathcal{A}, \overline{} \rangle$ for all agents $i \in N$. By contradiction, assume that Δ is not admissible in $\mathcal{F}$. Then $\exists \alpha \in \Delta$ that is attacked by $\beta \in \mathcal{A} \setminus \Delta$, i.e., $R = \{\bar{\alpha} \leftarrow \gamma_1, \ \ldots, \ \gamma_{m-1} \leftarrow \gamma_m\} \subseteq \mathcal{R}_{agg}$, for $m \geq 1$ and $\gamma_m = \beta$, and $\nexists \gamma \in \Delta$ such that γ attacks β in $\mathcal{F}$. By definition of nomination rule, $R \subseteq \mathcal{R}_i$ for some $i \in N$ and β attacks α in the Bipolar ABA framework $\mathcal{F}_i$ of agent i. Then, given that Δ is admissible in $\mathcal{F}_i$, $\exists \gamma \in \Delta$ such that γ attacks β in $\mathcal{F}_i$, i.e., $\exists R' = \{\bar{\beta} \leftarrow \delta_1, \ \ldots, \ \delta_{l-1} \leftarrow \delta_l\} \subseteq \mathcal{R}_i$, for $l \geq 1$ and $\delta_l = \gamma$. But, by definition of nomination rule, $R' \subseteq \mathcal{R}_{agg}$, and thus γ attacks β in $\mathcal{F}$: contradiction. To complete the proof, we need to show that for $|\mathcal{A}| \geq 4$, other quota rules except for nomination do not preserve admissibility. If $N - q$ agents choose rule $\mathcal{R} = \{\}$, $q - 1$ agents choose rule $\mathcal{R} = \{\bar{D} \leftarrow B, \quad \bar{C} \leftarrow D\}$, and 1 agent chooses rule $\mathcal{R} = \{\bar{D} \leftarrow A, \quad \bar{C} \leftarrow D, \quad A \leftarrow B\}$, as illustrated in Figure 1, then $\Delta = \{A, B, C\}$ is admissible in all frameworks. Using quota rules with $q > 1$, $\mathcal{R}_{agg} = \{\bar{C} \leftarrow D\}$ and Δ is not admissible anymore as assumption C is attacked by D and it is undefended by other assumptions in Δ. □

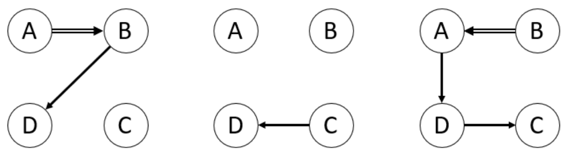

Figure 2: Counter Example of the Preservation of Set-stable Extensions (for Theorem 4.5). Here, we use BAFs as a graphical representation of Bipolar ABA frameworks (assuming the deductive interpretation of support). Single-lined edges are attacks, double-lined edges are supports.

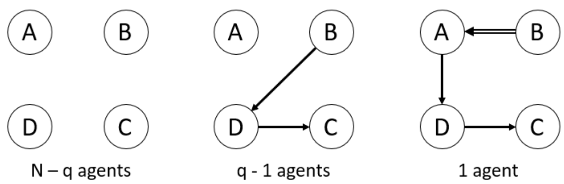

Figure 1: Counter Example for the Preservation of Admissibility (for Theorem 4.3). Here, we use BAFs as a graphical representation of Bipolar ABA frameworks (assuming the deductive interpretation of support). Single-lined edges are attacks, double-lined edges are supports.

We remark that for at most three assumptions, quota and oligarchic rules are guaranteed to preserve admissibility.

Theorem 4.4. *For $|\mathcal{A}| \leq 3$, every quota rule and oligarchic rule preserves admissibility.*

Proof. If $|\mathcal{A}| = 1$, the result holds vacuously. If $|\mathcal{A}| = 2$, assume that $\mathcal{A} = \{\alpha, \beta\}$ and $\Delta = \{\alpha\}$ is admissible in $\left\langle \mathcal{L}, \mathcal{R}_i, \mathcal{A}, \overline{} \right\rangle$ for all $i \in N$. Then $r = \bar{\alpha} \leftarrow \beta \notin \mathcal{R}_i$ for all $i \in N$, and thus $r \notin \mathcal{R}_{agg}$. Then, every quota and oligarchic rule yields $\{\alpha\}$ as an admissible extension in $\mathcal{F}$. The other cases ($\Delta = \{\beta\}$, $\Delta = \{\}$ or $\Delta = \mathcal{A}$) can be proven similarly. If $|\mathcal{A}| = 3$, assume that $\mathcal{A} = \{\alpha, \beta, \gamma\}$. Consider the case where $\Delta = \{\alpha\}$ is admissible in $\left\langle \mathcal{L}, \mathcal{R}_i, \mathcal{A}, \overline{} \right\rangle$ for all $i \in N$. By contradiction, assume that Δ is not admissible in $\mathcal{F}$. Then, given that Δ is conflict-free and closed in $\mathcal{F}$ no matter which aggregation rule, by Theorems 4.1 and 4.2, there are $R \subseteq \mathcal{R}_{agg}$ and $A \subseteq \{\beta, \gamma\}$ such that $A \vdash^R \bar{\alpha}$ in $\mathcal{F}$. By quota and oligarchic rules, there must be $i \in N$ such that $R \subseteq \mathcal{R}_i$; thus Δ is not admissible for agent i: contradiction. The other cases can be proven similarly. □

Note that the restriction to consider at most three assumptions may be useful in some settings, e.g., when at most three options are up for debate.

4.4 Set-stable Extensions

The set-stable semantics for Bipolar ABA frameworks generalises the stable semantic for AAFs to accommodate supports (see [12]). This generalisation though does not affect preservation. Thus, Theorem 4.5 extends Proposition 5 in [9].

Theorem 4.5. *Nomination is the only quota rule that preserves set-stable extensions.*

Proof. Assume that $\Delta \subseteq \mathcal{A}$ is set-stable in $\left\langle \mathcal{L}, \mathcal{R}_i, \mathcal{A}, \overline{} \right\rangle$ for all $i \in N$. By Theorem 4.1 and 4.2, nomination preserves closedness and conflict-freeness. Therefore, Δ is both closed and conflict-free in $\mathcal{F}$. To be set-stable, Δ has to attack the closure of every assumption β not Δ, i.e., $\exists R \subseteq \mathcal{R}_{agg}$ such that $\{\alpha\} \vdash^R \bar{\gamma}$ for some $\gamma \in Cl(\{\beta\})$. This is trivially the case if $\Delta = \mathcal{A}$. Otherwise, as Δ is set-stable in all agents' frameworks, then $R_i \subseteq \mathcal{R}_i$ must exist, for all $i \in N$, such that $\{\alpha\} \vdash^{R_i} \bar{\gamma}$ for some $\gamma \in Cl(\{\beta\})$ in $\left\langle \mathcal{L}, \mathcal{R}_i, \mathcal{A}, \overline{} \right\rangle$. Thus, by using nomination, $Cl(\{\beta\})$, for $\beta \in \mathcal{A} \setminus \Delta$, is attacked also in $\mathcal{F}$.

Other quota rules do not preserve set-stable extensions because, while preserving conflict-freeness and closedness, they do not guarantee that the closure of every assumption not in the extension is attacked. A counter example follows: assume three Bipolar ABA frameworks with rules $\mathcal{R}_1 = \{\bar{D} \leftarrow B, \quad B \leftarrow A\}$, $\mathcal{R}_2 = \{\bar{D} \leftarrow C\}$, and $\mathcal{R}_3 = \{\bar{D} \leftarrow A, \quad \bar{C} \leftarrow D, \quad A \leftarrow B\}$, as illustrated in Figure 2. In each framework, the set of assumptions $\{A, B, C\}$ is set-stable. Using other quota rules with $q > 1$, $\mathcal{R}_{agg} = \{\}$, and $\{A, B, C\}$ is not set-stable anymore as the assumption D is not included in it and it is not attacked either ($Cl(\{D\}) = \{D\}$ in this example). □

4.5 Assumption Acceptability

Assumption acceptability concerns the preferred, complete, set-stable, well-founded, and ideal semantics but at the level of single assumption rather than full extensions. If an assumption is acceptable in one of those semantics (by belonging to a set of assumptions accepted by the semantics) in each of the agents' framework, then preservation amounts to that assumption being still acceptable under the same semantics in the aggregated Bipolar ABA framework.

Definition 6. *(Acceptability of Assumptions) An assumption $\alpha \in \mathcal{A}$ is* acceptable *(in a Bipolar ABA framework) under preferred, complete, set-stable, well-founded, or ideal semantic iff there is $\Delta \subseteq \mathcal{A}$ with $\alpha \in \Delta$ such that Δ is (respectively) a preferred, complete, set-stable, well-founded, or ideal extension (in the Bipolar ABA framework).*

The proof of preservation results regarding acceptability use adaptations of results from [14] on *implicative* and *disjunctive* properties to represent impossibility results with dictatorship. We cast these properties for Bipolar ABA frameworks as follows:

Definition 7. *(Implicative Properties). A Bipolar ABA framework property P is* implicative *in $\left\langle \mathcal{L}, \mathcal{R}, \mathcal{A}, \overline{} \right\rangle$ iff there exist three rules $R_1, R_2, R_3 \notin \mathcal{R}$ such that P holds in $\left\langle \mathcal{L}, \mathcal{R}_{agg}, \mathcal{A}, \overline{} \right\rangle$ for $\mathcal{R}_{agg} = \mathcal{R} \cup \mathcal{S}$, for all $\mathcal{S} \subseteq \{R_1, R_2, R_3\}$, except for $\mathcal{S} = \{R_1, R_2\}$.*

Intuitively, if the aggregated Bipolar ABA framework includes R_1 and R_2 as additional rules in $\mathcal{S}$, then it should adopt R_3 as well to preserve property P.

Definition 8. *(Disjunctive Properties). A Bipolar ABA framework property P is* disjunctive *in $\left\langle \mathcal{L}, \mathcal{R}, \mathcal{A}, \overline{} \right\rangle$ iff there exist two rules $R_1, R_2 \notin \mathcal{R}$, such that P holds in $\left\langle \mathcal{L}, \mathcal{R}_{agg}, \mathcal{A}, \overline{} \right\rangle$ for $\mathcal{R}_{agg} = \mathcal{R} \cup \mathcal{S}$, for all $\mathcal{S} \subseteq \{R_1, R_2\}$, except for $\mathcal{S} = \{\}$.*

Intuitively, the aggregated Bipolar ABA framework has to include at least one of R_1 or R_2 to preserve the property P. Definitions of implicative and disjunctive properties lead us to prove two lemmas on preservation.

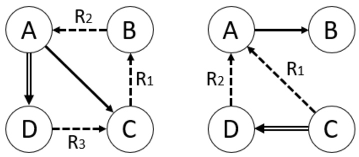

Figure 3: Acceptability of an Assumption: Implicative case (Left) and Disjunctive case (Right) (for the proof of Theorem 4.6). Here, we use BAFs as a graphical representation of Bipolar ABA frameworks (assuming the deductive interpretation of support). Single-lined and hyphenated edges are attacks, double-lined edges are supports.

Lemma 1. *Let a Bipolar ABA framework property P be implicative in $\left\langle \mathcal{L}, \mathcal{R}_i, \mathcal{A}, \overline{} \right\rangle$, for each $i \in N$. Then, unanimity preserves P.*

Proof. Let $\mathcal{R}_1 \supseteq \mathcal{S}_1, \ldots, \mathcal{R}_n \supseteq \mathcal{S}_n$ and let $\mathcal{S}_i \subseteq \{R_1, R_2, R_3\}$ for all $i \in N$. Let $\mathcal{S}_i \neq \{R_1, R_2\}$ for all $i \in N$. Then, unanimity preserves P because it is impossible to get $\mathcal{S}_{agg} = \{R_1, R_2\}$, with $\mathcal{R}_{agg} \supseteq \mathcal{S}_{agg}$. □

Note that, even if P is implicative, nomination and majority do not preserve P in general, as it is possible to get $\mathcal{S}_{agg} = \{R_1, R_2\}$. For example, let $\mathcal{S}_1 = \{R_1\}$, and $\mathcal{S}_2 = \{R_2\}$. Using nomination and majority, $\mathcal{R}_{agg} \supseteq \mathcal{S}_{agg} = \{R_1, R_2\}$; hence, P is not guaranteed to be preserved.

Lemma 2. *Let a Bipolar ABA framework property P be implicative and disjunctive in $\left\langle \mathcal{L}, \mathcal{R}_i, \mathcal{A}, \overline{} \right\rangle$, for each $i \in N$. Then, the only Bipolar ABA aggregation rule that preserves P is dictatorship.*

Proof. The proof for the implicativeness can be found in Lemma 1. For the disjunctiveness, let $\mathcal{R}_1 \supseteq \mathcal{S}_1, \ldots, \mathcal{R}_n \supseteq \mathcal{S}_n$ and let $\mathcal{S}_i \subseteq \{R_1, R_2\}$ for all $i \in N$. Let $\mathcal{S}_i \neq \{\}$ for all $i \in N$. Then, nomination and majority preserve P because it is impossible to get $\mathcal{S}_{agg} = \{\}$, with $\mathcal{R}_{agg} \supseteq \mathcal{S}_{agg}$. As P is implicative and disjunctive, P is preserved only with dictatorship. None of the quota rules preserve P as using nomination or majority rule, it is possible to get $\mathcal{S}_{agg} = \{R_1, R_2\}$ and violating the implicativeness; and using unanimity rule, it is possible to get $\mathcal{S}_{agg} = \{\}$ and violating the disjunctiveness. □

Notice that, in the definition of implicative and disjunctive properties, the rules R_1, R_2, and (if applicable) R_3 can only be in the form of $\bar{\alpha} \leftarrow \beta$ for some $\alpha, \beta \in \mathcal{A}$, thus bringing attacks between assumptions. They cannot be in the form of $\alpha \leftarrow \beta$ that denote supports between assumptions because then some agents may have different closures of assumptions from the other agents. As a consequence, some agents' Bipolar ABA frameworks may satisfy P, while some others may not because of closedness.

The preservation result on the acceptability of an assumption in Theorem 4.6 below is an extension, within our more general setting, of Theorem 1 in [9]. This result is true for all five semantics: preferred, complete, set-stable, well-founded, or ideal. If the agents' Bipolar ABA frameworks have no (rules for) support, then this Theorem 4.6 is the same as Theorem 1 in [9].

Theorem 4.6. *For $|\mathcal{A}| \geq 4$, the only Bipolar ABA aggregation rule that preserves the acceptability of an assumption under preferred, complete, set-stable, well-founded, or ideal semantic is dictatorship.*

Proof. Let P be acceptability of an assumption under preferred, complete, set-stable, well-founded, or ideal semantics. We need to prove that for $|\mathcal{A}| \geq 4$, P is implicative and disjunctive. Then, by Lemma 2, the theorem holds. The proof has the same structure for each of the five semantics. Consider a set of at least four assumptions $\mathcal{A} = \{A, B, C, D, \ldots\}$.

To show that P is implicative, let B be the accepted assumption. Let $\mathcal{R} = \{\bar{C} \leftarrow A, \quad D \leftarrow A\}$, $R_1 = \{\bar{B} \leftarrow C\}$, $R_2 = \{\bar{A} \leftarrow B\}$, and $R_3 = \{\bar{C} \leftarrow D\}$ (see the left graph of Figure 3). Consider an aggregated framework with $\mathcal{R}_{agg} = \mathcal{R} \cup \mathcal{S}$ with $\mathcal{S} \subseteq \{R_1, R_2, R_3\}$. If $\mathcal{S} = \{\}$, $\{R_2\}$, $\{R_3\}$, or $\{R_2, R_3\}$ then B is unattacked. If $\mathcal{S} = \{R_1\}$, $\{R_1, R_3\}$, or $\{R_1, R_2, R_3\}$ then B is defended by other assumptions. Therefore, B is either unattacked or defended in all seven cases, and B is acceptable under preferred, complete, set-stable, well-founded, and ideal semantics. However, if $\mathcal{S} = \{R_1, R_2\}$, $\{A, B, C\}$ forms cyclic attacks so that the assumptions A, B, and C are not acceptable under preferred, complete, set-stable, well-founded, and ideal semantics. Thus, we have identified a set of rules $\mathcal{R}$ and three rules R_1, R_2, R_3 such that P holds in $\left\langle \mathcal{L}, \mathcal{R} \cup \mathcal{S}, \mathcal{A}, \overline{} \right\rangle$ iff $\mathcal{S} \neq \{R_1, R_2\}$. Accordingly, P is implicative.

To show that P is disjunctive, let B be the accepted assumption. Let $\mathcal{R} = \{\bar{B} \leftarrow A, \quad D \leftarrow C\}$, $R_1 = \{\bar{A} \leftarrow C\}$, and $R_2 = \{\bar{A} \leftarrow D\}$ (see the right graph of Figure 3). Consider $\mathcal{R}_{agg} = \mathcal{R} \cup \mathcal{S}$ with $\mathcal{S} \subseteq \{R_1, R_2\}$. If $\mathcal{S} = \{R_1\}$, $\{R_2\}$, or $\{R_1, R_2\}$ then the assumption B is defended. Therefore, B is acceptable under the five semantics. However, if $\mathcal{S} = \{\}$, the assumption B is attacked by A and is not defended, thus B is unacceptable under preferred, complete, set-stable, well-founded, and ideal semantics. Thus, we have identified a set of rules $\mathcal{R}$ and two rules R_1, R_2 such that P holds in $\left\langle \mathcal{L}, \mathcal{R} \cup \mathcal{S}, \mathcal{A}, \overline{} \right\rangle$ iff $\mathcal{S} \neq \{\}$. Therefore, P is disjunctive. □

Theorem 4.6 shows that it is not easy to even preserve the acceptability of one assumption, as dictatorship is needed. We will see, in Section 4.6 below, that it is even more difficult to preserve whole extensions. On the other hand, for $|\mathcal{A}| \leq 3$, acceptability of an assumption can be preserved with both quota and oligarchic rules.

Theorem 4.7. *For $|\mathcal{A}| = 3$, majority, unanimity, and oligarchic rules preserve assumption acceptability under preferred, complete, set-stable, well-founded, and ideal semantics.*

Proof. Let a Bipolar ABA framework property P be the acceptability of an assumption under preferred, complete, set-stable, well-founded, or ideal semantics. Assume that P holds in $\left\langle \mathcal{L}, \mathcal{R}_i, \mathcal{A}, \overline{} \right\rangle$ for all $i \in N$, where $\mathcal{A} = \{\alpha, \beta, \gamma\}$ and assume that α is acceptable under preferred, complete, set-stable, well-founded, and ideal semantics in all frameworks.

By contradiction, assume P does not hold in $\mathcal{F}$. In other words, $\exists R \subseteq \mathcal{R}_{agg}$ such that $\{\delta\} \vdash^R \bar{\alpha}$ and not $\exists R \subseteq \mathcal{R}_{agg}$ such that $\{\theta\} \vdash^R \bar{\delta}$ for some $\theta \in \{\beta, \gamma\}$ and $\delta \in \{\beta, \gamma\}$, $\theta \neq \delta$. As a result, α is not acceptable under preferred, complete, set-stable, well-founded, and ideal semantics. By definition of of majority rule, unanimity rule, and oligarchic rules, the deduction from rules $\{\delta\} \vdash^R \bar{\alpha}$ must exist in the majority (majority rule), all (unanimity rule), or veto powered (oligarchic rules) agents' frameworks, but there is at least one framework $\left\langle \mathcal{L}, \mathcal{R}_i, \mathcal{A}, \overline{} \right\rangle$ for some $i \in N$ without the

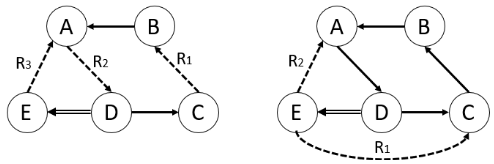

Figure 4: Preferred, Complete, Well-founded, and Ideal extensions: Implicative case (Left) and Disjunctive case (Right) (for the proof of Theorem 4.9). Here, we use BAFs as a graphical representation of Bipolar ABA frameworks (assuming the deductive interpretation of support). Single-lined and hyphenated edges are attacks, double-lined edges are supports.

deduction $\{\theta\} \vdash^R \bar{\delta}$ for $\theta \neq \delta$, thus α is not acceptable under preferred, complete, set-stable, well-founded, and ideal semantics in the agents' frameworks as well. It contradicts the initial assumption that α is acceptable in $\langle \mathcal{L}, \mathcal{R}_i, \mathcal{A}, \overline{} \rangle$ for all $i \in N$.

To show that for $|\mathcal{A}| = 3$, nomination rule do not preserve the assumption acceptability under preferred, complete, set-stable, well-founded, or ideal semantics, a counter example is given. Take three Bipolar ABA frameworks with rules $\mathcal{R}_1 = \{\bar{A} \leftarrow C, \quad \bar{B} \leftarrow C\}$, $\mathcal{R}_2 = \{\bar{B} \leftarrow A, \quad \bar{C} \leftarrow B\}$, and $\mathcal{R}_3 = \{\bar{C} \leftarrow A, \quad \bar{A} \leftarrow B\}$. Let assumption C be the accepted assumption in check. From the first framework, a set of assumptions $\{C\}$ is preferred, complete, set-stable, well-founded, and ideal. On the second framework is $\{A, C\}$ and third framework is $\{B, C\}$, both extensions are preferred, complete, set-stable, well-founded, and ideal as well. In all three frameworks, the assumption C is acceptable. It is still acceptable using unanimity rule as the aggregated rule is $\mathcal{R} = \{\}$. However, using nomination rule the preferred, complete, and set-stable extensions are $\{A\}$, $\{B\}$, and $\{C\}$; while the well-founded and ideal extensions are $\{\}$. Hence, $\{C\}$ is not acceptable in those five semantics. □

Theorem 4.8. *For $|\mathcal{A}| \leq 2$, every quota and oligarchic rule preserves the acceptability of an assumption under preferred, complete, set-stable, well-founded, and ideal semantics.*

4.6 Preferred, Complete, Well-founded, and Ideal Extensions

The proof of preservation for the preferred, complete, well-founded, and ideal semantics uses the concept of implicative and disjunctive properties from Section 4.5. The preservation result is an extension of Theorem 4 in [9] with the addition of ideal semantics.

Theorem 4.9. *For $|\mathcal{A}| \geq 5$, the only Bipolar ABA aggregation rule that preserves preferred, complete, well-founded, and ideal semantics is dictatorship.*

Proof. Let P be that an extension $\Delta \subseteq \mathcal{A}$ is preferred, complete, well-founded, or ideal. We need to prove that for $|\mathcal{A}| \geq 5$, P is implicative and disjunctive. Then by Lemma 2, the theorem holds. The proof has the same structure for all four semantics. It uses a generic $\mathcal{A} = \{A, B, C, D, E, \ldots\}$ with at least five assumptions.

To show that P is implicative, let $\Delta = \{B, D, E\}$. Let $\mathcal{R} = \{\bar{C} \leftarrow D, \quad \bar{A} \leftarrow B, \quad E \leftarrow D\}$, $R_1 = \{\bar{B} \leftarrow C\}$, $R_2 = \{\bar{D} \leftarrow A\}$, and $R_3 = \{\bar{A} \leftarrow E\}$ (see the left graph of Figure 4). Consider $\mathcal{R}_{agg} = \mathcal{R} \cup \mathcal{S}$ with $\mathcal{S} \subseteq \{R_1, R_2, R_3\}$. If $\mathcal{S} = \{\}, \{R_1\}, \{R_2\}, \{R_3\}, \{R_1, R_3\}, \{R_2, R_3\}$, or $\{R_1, R_2, R_3\}$. Then Δ is preferred, complete, well-founded, and ideal. However, if $\mathcal{S} = \{R_1, R_2\}$, $\{A, B, C, D\}$ forms cyclic attacks such that Δ is not preferred, complete, well-founded, or ideal. Thus, we have identified a set of rules $\mathcal{R}$ and three rules R_1, R_2, R_3 such that P holds in $\langle \mathcal{L}, \mathcal{R} \cup \mathcal{S}, \mathcal{A}, \overline{} \rangle$ iff $\mathcal{S} \neq \{R_1, R_2\}$. Hence, P is implicative.

To show that P is disjunctive, let $\Delta = \{B, D, E\}$. Let $\mathcal{R} = \{\bar{C} \leftarrow D, \quad \bar{B} \leftarrow C, \quad \bar{A} \leftarrow B, \quad \bar{D} \leftarrow A, \quad D \leftarrow E\}$, $R_1 = \{\bar{C} \leftarrow E\}$, and $R_2 = \{\bar{A} \leftarrow E\}$ (see the right graph of Figure 4). Consider $\mathcal{R}_{agg} = \mathcal{R} \cup \mathcal{S}$ with $\mathcal{S} \subseteq \{R_1, R_2\}$. If $\mathcal{S} = \{R_1\}, \{R_2\}$, or $\{R_1, R_2\}$, then Δ is preferred, complete, well-founded, and ideal. However, if $\mathcal{S} = \{\}$, $\{A, B, C, D\}$ forms cyclic attacks such that Δ is not preferred, complete, well-founded, or ideal. Therefore, we have identified a set of rules $\mathcal{R}$ and two rules R_1, R_2 such that P holds in $\langle \mathcal{L}, \mathcal{R} \cup \mathcal{S}, \mathcal{A}, \overline{} \rangle$ iff $\mathcal{S} \neq \{\}$. Thus, P is disjunctive. □

Although preferred and complete semantics may accept multiple extensions, as long as all agents agree on the extensions, then Theorem 4.9 still holds. The only restriction is the presence of supports in the agents' frameworks. If all agents agree on the supports, i.e., the supports are included in $\mathcal{R}_i$ for all $i \in N$, then support does not affect the preservation. Otherwise, some agents will have different sets of closed assumptions from the other agents. This may lead into different preferred, complete, well-founded, and ideal extensions.

The corner cases show an impossibility result for $|\mathcal{A}| = 3$ and $|\mathcal{A}| = 4$. Thus, preserving whole extensions is more difficult than preserving the acceptability of single assumptions as in Section 4.5.

Theorem 4.10. *For $|\mathcal{A}| = 3$ and $|\mathcal{A}| = 4$, quota and oligarchic rules do not preserve preferred, complete, well-founded, and ideal semantics.*

Theorem 4.11. *For $|\mathcal{A}| \leq 2$, every quota and oligarchic rule preserves preferred, complete, well-founded, and ideal semantics.*

4.7 Non-emptiness of Well-founded Extensions

The well-founded extension is guaranteed to exist in a Bipolar ABA framework. However, to make sure that the well-founded extension is not empty, then the framework must have at least one unattacked assumption. This way, the unattacked assumptions are included in all complete extensions, and the intersection always has the unattacked assumptions in it. The preservation of non-emptiness of the well-founded extension guarantees the existence of unattacked assumption with a concept called *k-exclusivity* [9].

Definition 9. *(k-exclusivity). Let P be a property of Bipolar ABA framework. P is k-exclusive if there exist rules $\mathcal{S} = \{R_1, \ldots, R_k\}$ such that if $\mathcal{R} \supseteq \mathcal{S}$ then P does not hold, but if $\mathcal{R} \subset \mathcal{S}$ then P holds.*

Thus, to preserve P, the rules $\mathcal{S}$ cannot be adopted together, but only a subset of them. It leads us to the lemma for the preservation.

Lemma 3. *Let P be a k-exclusive property of Bipolar ABA framework. For $k \geq N$, where N is the number of agents, P is preserved if at least one of the N agents has veto power.*

Proof. It needs to be showed that if an aggregation rule preserves P, then it has to give at least one agent with veto powers.

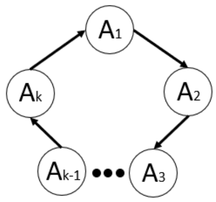

Figure 5: Graphical illustration of the *k-exclusivity* property (for the proof of Theorem 4.12)

Notice that if all agents accept a rule r, then it must be accepted in the aggregated rules, i.e., $r \in \mathcal{R}_{agg}$ iff $r \in \mathcal{R}_i$ for all $i \in N$.

For some agents $i \in N$ to have veto powers means that $\mathcal{R}_{agg} = (\bigcap \mathcal{R}_i)$. In other words, some agents have veto power, if the intersection of the agents' rules in $\bigcap \mathcal{R}_i$ are all accepted in $\mathcal{R}_{agg}$. Then, take any rule $r \in \mathcal{R}_{agg}$; as r is accepted in the aggregated framework, then all agents with veto powers must accept r as well such that the intersection of the set of rules $\bigcap \mathcal{R}_i$ is not empty.

Thus, the next step is to show that if an aggregation rule preserves P, then the intersection of k set of rules must be non-empty, i.e., $\mathcal{R}_1 \cap \ldots \cap \mathcal{R}_k \neq \{\}$. To prove by contradiction, assume there exist a profile of set of rules $\{\mathcal{R}_1 \cup \ldots \cup \mathcal{R}_k\} \subseteq \mathcal{R}_{agg}$ such that $\mathcal{R}_1 \cap \ldots \cap \mathcal{R}_k = \{\}$. Then, it means that for every $j \in \{1, \ldots, k\}$, exactly (the agent with rule set) $\mathcal{R}_j$ accepts a rule r_j. As no rule exist in all $\mathcal{R}_i$ for $i \in N$, no agents accept all k rules. However, as each of the k rules is accepted by an agent and $\{\mathcal{R}_1 \cup \ldots \cup \mathcal{R}_k\} \subseteq \mathcal{R}_{agg}$, they are all accepted in the aggregated framework, i.e., $\{r_1, \ldots, r_k\} \subseteq \mathcal{R}_{agg}$, such that P does not hold due to it being an *k-exclusive* property. This contradicts the initial assumption that the aggregation rule preserves P.

Therefore, as it can be showed that the intersection of the agents' rules is not empty, then some agents must have veto powers. □

Theorem 4.12. *For $|\mathcal{A}| \geq N$, at least one agent must have veto power to preserve the non-emptiness of the well-founded extension.*

Proof. Let a Bipolar ABA framework property P be the non-emptiness of the well-founded extension. We need to show that P is *k-exclusive*. Let $k = |\mathcal{A}|$ and $\{A_1, \ldots, A_k\} \subseteq \mathcal{A}$. Assume that $\mathcal{S}$ consists of all rules $\overline{A_{i+1}} \leftarrow A_i$ for $i < |\mathcal{A}|$ as well as $\overline{A_1} \leftarrow A_k$, illustrated in Figure 5. This $\mathcal{S}$ fits the definition of *k-exclusive*. Indeed, if $\mathcal{S} \subseteq \mathcal{R}$, then in the case of $\mathcal{S} = \mathcal{R}$, the well-founded extension is empty due to the cyclic attacks. However, if only a subset of it is adopted, $\mathcal{R} \subset \mathcal{S}$, the well-founded extension is not empty as at least one assumption is not attacked. Thus, P is preserved when at least one agent has veto power to prevent cyclic attacks. □

Supports in Bipolar ABA framework do not affect the preservation of non-emptiness of the well-founded extension because supports between assumptions do not affect the unattacked assumption: if there is a rule $\overline{\alpha} \leftarrow \beta$ for $\alpha, \beta \in \mathcal{A}$ and β is unattacked, then supports from and into β do not change the fact that β is unattacked; and supports from and into α also leave β unattacked.

4.8 Acyclicity

It is clear that *k-exclusivity* deals with cyclic attacks. A Bipolar ABA framework is acyclic if there does not exist any cyclic attacks among the assumptions. Corollary 1 is extended from Theorem 8 in [9] in the way that supports are also considered.

Definition 10. *(Cyclic Attacks) The rule set $\mathcal{R}$ in $\langle \mathcal{L}, \mathcal{R}, \mathcal{A}, \overline{} \rangle$ contain cyclic attacks if there exist a chained connection between some of the assumptions in $\mathcal{A}$, such that $\mathcal{R} \supseteq \{\overline{\alpha_1} \leftarrow \alpha_2, \overline{\alpha_2} \leftarrow \alpha_3, \ldots, \overline{\alpha_k} \leftarrow \alpha_1\}$ for $\alpha_i \in \mathcal{A}$ and $k \geq 2$.*

The preservation result for acyclicity has a similar proof structure as the preservation of the non-emptiness of the well-founded extension in Theorem 4.12. Thus, it is presented as a corollary.

Corollary 1. *For $|\mathcal{A}| \geq N$, at least one agent must have veto power to preserve acyclicity.*

Proof. Let P be acyclicity. We need to show that P is *k-exclusive*. To get a cycle, a minimum number of two assumptions are required. Thus, let $k = |\mathcal{A}| \geq 2$ and $\{A_1, \ldots, A_k\} \subseteq \mathcal{A}$. Assume that the rule set $\mathcal{S}$ consists of $\overline{A_{i+1}} \leftarrow A_i$ for $i < |\mathcal{A}|$ as well as $\overline{A_1} \leftarrow A_k$, illustrated in Figure 5. This $\mathcal{S}$ fits the definition of *k-exclusivity*. Indeed, if $\mathcal{S} \subseteq \mathcal{R}$, then in the case of $\mathcal{S} = \mathcal{R}$, the cyclic attacks remain in the framework. However, if only a subset of $\mathcal{S}$ is adopted ($\mathcal{R} \subset \mathcal{S}$), then the cyclic attacks are broken because at least one rule that connects the cycle disappears. Therefore, P is preserved when at least one agent has veto power. □

The presence of supports does not make an acyclic framework to become cyclic, but instead may break any existing cycle. Let $k = |\mathcal{A}|$ with $|\mathcal{A}| \geq 2$, $\{A_1, \ldots, A_k\} \subseteq \mathcal{A}$, and $\mathcal{S} = \{\overline{A_{i+1}} \leftarrow A_i\colon i < |\mathcal{A}|\}$. The rules in $\mathcal{S}$ are acyclic and if a support $A_1 \leftarrow A_k$ or $A_k \leftarrow A_1$ is added, for example, then they will remain acyclic. On the contrary, if there exist cyclic attacks, then support may break the cycle due to closedness.

4.9 Coherence

Coherence amounts to two or more semantics coinciding (in other words, given a Bipolar ABA framework, two or more semantics give identical extensions thereof). For example, if a set of assumptions is set-stable, then it is preferred as well. Our next preservation result extends Theorem 9 in [9] and shows that, in order to preserve coherence, the aggregation rule must be dictatorial. The proof for the result uses the concept of implicativeness and disjunctiveness introduced in Section 4.5.

Theorem 4.13. *For $|\mathcal{A}| \geq 4$, the only aggregation rule preserving coherence is dictatorship.*

Proof. Let P be coherence. We need to prove that, for $|\mathcal{A}| \geq 4$, P is implicative and disjunctive. Take a Bipolar ABA framework with at least four assumptions $\mathcal{A} = \{A, B, C, D, \ldots\}$.

To show that P is implicative, let $\mathcal{R} = \{\bar{C} \leftarrow A, \quad D \leftarrow A\}$, $R_1 = \{\bar{B} \leftarrow C\}$, $R_2 = \{\bar{A} \leftarrow B\}$, and $R_3 = \{\bar{C} \leftarrow D\}$, as illustrated in the left graph of Figure 6. Consider an aggregated framework with $\mathcal{S} \subseteq \{R_1, R_2, R_3\}$. If $\mathcal{S} = \{\}$, $\{R_1\}$, $\{R_3\}$, or $\{R_1, R_3\}$, the only preferred extension is $\{A, B, D\}$, which is set-stable as well. If $\mathcal{S} = \{R_2\}$, the set of assumptions $\{B, C, D\}$ is both preferred and set-stable. If $\mathcal{S} = \{R_2, R_3\}$ or $\{R_1, R_2, R_3\}$; then the set of assumptions $\{B, D\}$ is both preferred and set-stable as well. However, if $\mathcal{S} = \{R_1, R_2\}$, the only preferred extension is $\{D\}$ and it is not set-stable as the other assumptions are not attacked. Thus, there exists a set of rules $\mathcal{R}$

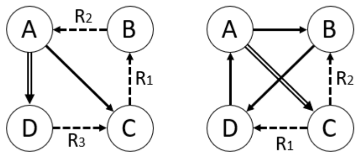

Figure 6: Coherence: Implicative case (Left) and Disjunctive case (Right) (for the proof of Theorem 4.13). Here, we use BAFs as a graphical representation of Bipolar ABA frameworks (assuming the deductive interpretation of support). Single-lined and hyphenated edges are attacks, double-lined edges are supports.

and three rules R_1, R_2, R_3 such that P holds in $\langle \mathcal{L}, \mathcal{R} \cup \mathcal{S}, \mathcal{A}, \overline{} \rangle$ iff $\mathcal{S} \neq \{R_1, R_2\}$. Accordingly, P is an implicative property.

To show that P is disjunctive, let $\mathcal{R} = \{\bar{A} \leftarrow D, \quad \bar{B} \leftarrow A, \quad \bar{D} \leftarrow B, \quad C \leftarrow A\}$, $R_1 = \{\bar{D} \leftarrow C\}$, and $R_2 = \{\bar{B} \leftarrow C\}$, as illustrated in the right graph of Figure 6. Consider an aggregated framework $\langle \mathcal{L}, \mathcal{R}_{agg}, \mathcal{A}, \overline{} \rangle$, where $\mathcal{R}_{agg} = \mathcal{R} \cup \mathcal{S}$ with $\mathcal{S} \subseteq \{R_1, R_2\}$. If $\mathcal{S} = \{R_1\}$ or $\{R_1, R_2\}$, the set of assumptions $\{A, C\}$ is both preferred and set-stable. If $\mathcal{S} = \{R_2\}$, the set of assumptions $\{C, D\}$ is also preferred and set-stable. However, if $\mathcal{S} = \{\}$, the preferred extension is $\{C\}$ and it is not set-stable because the other assumptions are not attacked. Therefore, there exists a set of rules $\mathcal{R}$ and two rules R_1, R_2 such that P holds in $\langle \mathcal{L}, \mathcal{R} \cup \mathcal{S}, \mathcal{A}, \overline{} \rangle$ iff $\mathcal{S} \neq \{\}$. Hence, P is a disjunctive property.

As P is proven to be both implicative and disjunctive, then by Lemma 2, for P to be preserved, the aggregation rule must be dictatorial. □

Note that Theorem 4.13 also works for other semantics (indeed, in the proof, the accepted sets of assumptions may be complete and well-founded, rather than just preferred and set-stable).

The presence of supports is acceptable in the preservation of coherence only if the supports are adopted by each agent, such that all agents have the same closure of assumptions. If supports join the additional rules in $\mathcal{S}$ as either R_1, R_2, or R_3, then coherence is not preserved in the aggregated framework as some agents have different set of closures than the other agents.

For corner cases, it is easier to preserve coherence, as indeed unanimity preserves it for $|\mathcal{A}| \leq 3$. Moreover, admittedly less interestingly, both quota and oligarchic rules preserve coherence when there is one assumption.

Theorem 4.14. *For $|\mathcal{A}| = 2$ or $|\mathcal{A}| = 3$, unanimity rule is the only quota rule that preserves coherence.*

Theorem 4.15. *For $|\mathcal{A}| = 1$, every quota rule and oligarchic rule preserve coherence.*

5 CONCLUSION

We have considered Bipolar ABA Frameworks [12] to account for both attack and support relationships between arguments in Bipolar Argumentation, as it allows to capture uniformly different interpretations of support. The aggregation of Bipolar ABA Frameworks combines the rules of all agents into a collective set of rules. We made use of the aggregation rules from judgement aggregation [17, 18], specifically, quota and oligarchic rules, to combine these agents' rules and extended results (on Abstract Argumentation) from [9]. Generally, the preservation results show that most properties can be preserved, but with significant restrictions sometimes. The results assume agreement among the agents on language, assumptions, contraries, and assume that agents accept the same properties (Definitions 2 and 5). We observe that, when the notion of agreement comes into play, the presence of supports does not greatly affect the performance of the aggregation rules towards preservation.

As regards positive results, conflict-freeness and closedness are preserved by any quota and oligarchic rule (Theorems 4.1 and 4.2). We proved positive results for admissibility and set-stability as well, albeit with some restrictions, such as limiting the number of assumptions or the choice of aggregation rules. Admissibility is preserved by nomination for at least four assumptions, else it is preserved by every quota and oligarchic rule; while the set-stable semantics is preserved by nomination (Theorems 4.3, 4.4, and 4.5).

We show that some properties can only be preserved by oligarchic rules or dictatorship. These particular aggregation rules are actually not ideal, as they ignore most opinions. However, we still deem this better than not being able to preserve the properties at all. For the properties of acceptability of an assumption, and coherence when the number of assumptions is at least four, dictatorship is the only preserving rule (Theorems 4.6 and 4.13). The same holds for preferred, complete, well-founded, and ideal semantics, but by assuming at least five assumptions (Theorem 4.9). Unsurprisingly, in the corner cases, these properties can be preserved with other quota rules (Theorems 4.7, 4.8, 4.10, 4.11 , 4.14, and 4.15).

Preservation results also involve the non-emptiness of the well-founded extension and acyclicity. We proved that both properties are preserved when at least one agent has veto power and the number of assumptions is greater or equal than the number of agents (Theorem 4.12 and Corollary 1). This unique constraint is meant to avoid cyclic relationships.

To conclude, our preservation study produces stronger results to fill the gaps in [9] since we consider more properties, some relevant to Bipolar Argumentation only (closedness) others also relevant to Abstract Argumentation (ideal semantics); we also provide preservation results for corner cases.

There are several possible directions for future work. First of all, here the preservation of properties relies on the agreement of all agents. However, in real applications it is likely that some agents have different opinions, i.e., some of them might disagree on the properties. Thus, it would be interesting to study preservation when a number of agents disagree. Another path to work on in the future is to have agents with different knowledge about the environment, meaning that they might have different languages, assumptions, or contraries. A further promising direction for future work is to expand the choice of aggregation rules with a more complex formalisation. Finally, it would be worth to generalise this study for the more general ABA Frameworks of [6, 11], as well as for other forms of structured argumentation, such as ASPIC [19], DeLP [16] or logic-based argumentation [3]. In particular, the possibility of having rules with empty body might need specific attention when it comes to aggregation.

REFERENCES

[1] Pietro Baroni, Marco Romano, Francesca Toni, Marco Aurisicchio, and Giorgio Bertanza. 2015. Automatic evaluation of design alternatives with quantitative argumentation. *Argument & Computation* 6, 1 (2015), 24–49.

[2] Dorothea Baumeister and Jörg Rothe. 2016. Preference Aggregation by Voting. In *Economics and Computation*. Springer, Berlin, Heidelberg, 197–325.

[3] Philippe Besnard and Anthony Hunter. 2014. Constructing Argument Graphs with Deductive Arguments: A Tutorial. *Argument & Computation* 5, 1 (2014), 5–30. https://doi.org/10.1080/19462166.2013.869765

[4] Gustavo Adrian Bodanza, Fernando Tohmé, and Marcelo Auday. 2017. Collective argumentation: A survey of aggregation issues around argumentation frameworks. *Argument & Computation* 8, 1 (2017), 1–34.

[5] Guido Boella, Dov M. Gabbay, Leendert W. N. van der Torre, and Serena Villata. 2010. Support in Abstract Argumentation. In *Computational Models of Argument: Proceedings of COMMA 2010, Desenzano del Garda, Italy, September 8-10, 2010*. IOS Press, Desenzano del Garda, Italy, 111–122.

[6] Andrei Bondarenko, Phan Minh Dung, Robert A. Kowalski, and Francesca Toni. 1997. An Abstract, Argumentation-Theoretic Approach to Default Reasoning. *Artif. Intell.* 93 (1997), 63–101.

[7] Claudette Cayrol and Marie-Christine Lagasquie-Schiex. 2005. On the Acceptability of Arguments in Bipolar Argumentation Frameworks. In *Symbolic and Quantitative Approaches to Reasoning with Uncertainty, 8th European Conference, ECSQARU 2005, Barcelona, Spain, July 6-8, 2005, Proceedings*. Springer, Berlin, Heidelberg, 378–389.

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

[9] Weiwei Chen and Ulle Endriss. 2017. Preservation of Semantic Properties during the Aggregation of Abstract Argumentation Frameworks. In *Proceedings Sixteenth Conference on Theoretical Aspects of Rationality and Knowledge, TARK 2017, Liverpool, UK, 24-26 July 2017*. OPA, Liverpool, UK, 118–133.

[10] Weiwei Chen and Ulle Endriss. 2018. Aggregating Alternative Extensions of Abstract Argumentation Frameworks: Preservation Results for Quota Rules. In *Computational Models of Argument - Proceedings of COMMA 2018, Warsaw, Poland, 12-14 September 2018*. IOS Press, Warsaw, Poland, 425–436.

[11] Kristijonas Cyras, Xiuyi Fan, Claudia Schulz, and Francesca Toni. 2017. Assumption-based Argumentation: Disputes, Explanations, Preferences. *FLAP* 4, 8 (2017), 2407–2455.

[12] Kristijonas Cyras, Claudia Schulz, and Francesca Toni. 2017. Capturing Bipolar Argumentation in Non-flat Assumption-Based Argumentation. In *PRIMA 2017: Principles and Practice of Multi-Agent Systems - 20th International Conference, Nice, France, October 30 - November 3, 2017, Proceedings*. Springer, Cham, 386–402.

[13] Phan Minh Dung. 1995. On the Acceptability of Arguments and its Fundamental Role in Nonmonotonic Reasoning, Logic Programming and n-Person Games. *Artif. Intell.* 77, 2 (1995), 321–358.

[14] Ulle Endriss and Umberto Grandi. 2017. Graph aggregation. *Artif. Intell.* 245 (2017), 86–114.

[15] Jordi Ganzer-Ripoll, Natalia Criado, Maite Lopez-Sanchez, Simon Parsons, and Juan A. Rodriguez-Aguilar. 2019. Combining Social Choice Theory and Argumentation: Enabling Collective Decision Making. *Group Decision and Negotiation* 28, 1 (2019), 127–173.

[16] Alejandro Javier García and Guillermo Ricardo Simari. 2014. Defeasible Logic Programming: DeLP-servers, Contextual Queries, and Explanations for Answers. *Argument & Computation* 5, 1 (2014), 63–88. https://doi.org/10.1080/19462166.2013.869767

[17] Davide Grossi and Gabriella Pigozzi. 2014. *Judgment Aggregation: A Primer*. Morgan & Claypool Publishers, San Rafael.

[18] Christian List. 2013. Social Choice Theory. In *The Stanford Encyclopedia of Philosophy* (winter 2013 ed.), Edward N. Zalta (Ed.). Metaphysics Research Lab, Stanford University, Stanford.

[19] Sanjay Modgil and Henry Prakken. 2014. The ASPIC+ Framework for Structured Argumentation: A Tutorial. *Argument & Computation* 5, 1 (2014), 31–62. https://doi.org/10.1080/19462166.2013.869766

[20] Farid Nouioua and Vincent Risch. 2010. Bipolar Argumentation Frameworks with Specialized Supports. In *22nd IEEE International Conference on Tools with Artificial Intelligence, ICTAI 2010, Arras, France, 27-29 October 2010 - Volume 1*. IEEE Computer Society, Arras, 215–218.

[21] Antonio Rago and Francesca Toni. 2017. Quantitative Argumentation Debates with Votes for Opinion Polling. In *PRIMA 2017: Principles and Practice of Multi-Agent Systems - 20th International Conference, Nice, France, October 30 - November 3, 2017, Proceedings*. Springer, Cham, 369–385.

[22] Antonio Rago, Francesca Toni, Marco Aurisicchio, and Pietro Baroni. 2016. Discontinuity-Free Decision Support with Quantitative Argumentation Debates. In *Principles of Knowledge Representation and Reasoning: Proceedings of the Fifteenth International Conference, KR 2016, Cape Town, South Africa, April 25-29, 2016*. AAAI Press, Cape Town, South Africa, 63–73.

[23] Francesca Toni. 2014. A tutorial on assumption-based argumentation. *Argument & Computation* 5, 1 (2014), 89–117.