Collective Argumentation: The Case of Aggregating Support-Relations of Bipolar Argumentation Frameworks

Weiwei Chen
Institute of Logic and Cognition and Department of Philosophy
Sun Yat-sen University
Guangzhou, China
chenww26@mail2.sysu.edu.cn

In many real-life situations that involve exchanges of arguments, individuals may differ on their assessment of which supports between the arguments are in fact justified, i.e., they put forward different support-relations. When confronted with such situations, we may wish to aggregate individuals' argumentation views on support-relations into a collective view, which is acceptable to the group. In this paper, we assume that under bipolar argumentation frameworks, individuals are equipped with a set of arguments and a set of attacks between arguments, but with possibly different support-relations. Using the methodology in social choice theory, we analyze what semantic properties of bipolar argumentation frameworks can be preserved by aggregation rules during the aggregation of support-relations.

1 Introduction

The attack relation has played a significant role in formal argumentation [2, 11, 23]. However, recent years have seen a revived interest in the support relation between arguments in argumentation systems [4, 5, 6, 7, 26]. In these systems, an argument can not only attack another argument, but it can also support another one. For example, an argument can support another argument by confirming its premise or undermining one of its attackers. The support relation between arguments is vital in modeling debates in real life. Due to the incompleteness of information, or different positions, agents may have different opinions regarding the support relation between arguments. To see this, consider the following example:

Example 1. Consider a debate regarding the possible influence of artificial intelligence (AI) to the job market. Suppose that there are two arguments in this debate:

A: Artificial intelligence improves the degree of work automation
B: More people will lose their jobs due to AI

Given the fact that AI is able to perform more of the tasks done by humans, some occupations will decrease. Therefore, some people hold that argument A supports argument B. On the other hand, given that AI will improve the quality of the work being done by humans, lower the prices of goods and services, create economic advantages, and allow for the creation of new jobs in new occupations, some people hold that argument A does not support argument B. $\triangle$

J. Y. Halpern and A. Perea (Eds.): Theoretical Aspects
of Rationality and Knowledge 2021 (TARK 2021)
EPTCS 335, 2021, pp. 87–102, doi:10.4204/EPTCS.335.8

In many scenarios, such as court debate, parliament debate, policy advisory committee decision-making, agents may have different opinions on which supports between arguments are acceptable, which form argumentative stances of them. When a group of agents are engaged in a debate, we may wish to aggregate stances possessed by agents to obtain a collective decision agreed on by the group. To model the support relation between arguments, we consider the *bipolar argumentation framework (BAF)* [5, 6, 7], a formalism of Dung's abstract argumentation framework [11]. Given that there is a broad discussion of the aggregation of argumentation systems with the attack relation [10, 25, 12, 9], it is far from being clear what consensuses can be achieved when the support relation is involved in this process. The goal of this paper is to investigate the aggregation of views of a group of agents in the context of bipolar argumentation. Given a set of arguments and a set of attack-relations between these arguments, agents might conflict with one another upon supports between arguments, i.e., for every pair of arguments that is being considered in a debate whether the first supports the second. In this scenario, we may wish to aggregate such support-relations.

In this paper, we use the method from *graph aggregation*, a recent discipline of social choice theory that deals with aggregating several graphs into a single output graph that constitutes a good compromise. Following the model introduced by Chen and Endriss [9], we consider the preservation of properties of bipolar argumentation frameworks, i.e., given a property that is satisfied by individual BAFs, we study whether it can be satisfied in the BAF returned by aggregation rules. For some properties, we show that there is an aggregation rule or a family of aggregation rules preserve them. For some others, we show that any aggregation rule that satisfies certain basic axioms and preserves them must be a dictatorship.

Paper overview The rest of the paper is organized as follows. In Section 2, we recapitulate the bipolar argumentation framework, along with its semantics. We introduce our model for the aggregation of support-relations of bipolar argumentation frameworks in Section 3, followed by our results of preservation in Section 4. In Section 5, we introduce some work related to our work. Finally, in Section 6, we conclude this work and point out some directions for future work.

2 Bipolar argumentation

An abstract bipolar argumentation framework [5, 6, 7] is an extension of Dung's abstract argumentation framework [11] in which a general support relation between arguments is added. Formally, an abstract bipolar argumentation framework is a triple $\langle Arg, \rightharpoonup, \rightsquigarrow \rangle$, where Arg is a set of arguments, $\rightharpoonup$ is a binary relation on Arg, which is called the attack relation, $\rightsquigarrow$ is a binary relation on Arg, which is called the support relation. Given two arguments $A, B \in Arg$, if $A \rightharpoonup B$ holds, then we say that A attacks B, if $A \rightsquigarrow B$, then we say that A supports B. The attack relation and the support relation must verify the following consistency constraint: $\rightharpoonup \cap \rightsquigarrow = \emptyset$, which is called *essential constraint*.

Definition 1. *Let $A, B \in Arg$, there is a sequence of supports for B by A iff there exists a sequence of elements $(A_1, \ldots, A_n)$ of Arg such that $n \geqslant 2, A = A_1, B = A_n, A_1 \rightsquigarrow A_2, \ldots, A_{n-1} \rightsquigarrow A_n$.*

Definition 2. *Let $A, B \in Arg$, a* supported attack *against B by A is a sequence of arguments $(A_1, \ldots, A_n)$ of Arg such that $A_1 \rightsquigarrow, \ldots, \rightsquigarrow A_{n-1}$, $A_{n-1} \rightharpoonup A_n$, $A = A_1$, $A_n = B$, and $n \geqslant 3$.*

Note that if $A \rightharpoonup B$ is the case, then we say that A directly attacks B.

Definition 3. *A* secondary attack *against an argument B by an argument A is a sequence $(A_1, \ldots, A_n)$ of arguments of Arg such that $A_1 \rightharpoonup A_2$, $A_2 \rightsquigarrow \ldots, \rightsquigarrow A_n$, $A = A_1$, $A_n = B$, and $n \geqslant 2$.*

For example, in Figure 1, A_1 supported attacks E_1, while A_2 secondary attacks E_2.

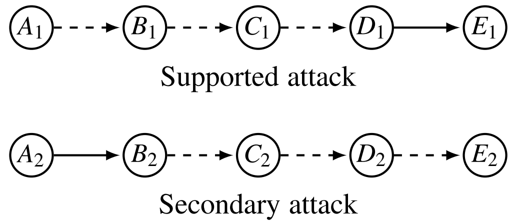

Figure 1: Illustration of supported attack and secondary attack

Definition 4. *Let $\Delta \subseteq Arg$ and $A \in Arg$. Δ set-attacks A iff there exists a supported attack or a secondary attack against A from an element of Δ. Δ set-supports A iff there exists a sequence of supports for A from an element of Δ.*

Definition 5. *Let $\Delta \subseteq Arg$ be a set of arguments, Δ is conflict-free iff $\nexists A, B \in \Delta$ such that $\{A\}$ set-attacks B.*

Definition 6. *Let $\Delta \subseteq Arg$ be a set of arguments, Δ is* safe *iff $\nexists B \in Arg$ such that Δ set-attacks B and either Δ set-supports B, or $B \in \Delta$.*

In the context of bipolar argumentation, admissibility can be translated into d-admissibility, s-admissibility and c-admissibility, based on different lines of coherence. In the following definition, the notion of *defense* is the same as classical defense, namely, we say $\Delta \subseteq Arg$ defends the argument $B \in Arg$, then, there is an argument $C \in \Delta$ with $C \rightharpoonup A$ for all arguments $A \in Arg$ such that $A \rightharpoonup B$.

Definition 7. *Let $\Delta \subseteq Arg$ be a set of arguments, Δ is called* d-admissible *iff Δ is conflict-free and defends all its elements; Δ is a d-preferred extension if it is maximal (w.r.t. set-inclusion) among all d-admissible sets.*

Definition 8. *Let $\Delta \subseteq Arg$ be a set of arguments, Δ is called* s-admissible *iff Δ is safe and defends all its elements; Δ is a s-preferred extension if it is maximal (w.r.t. set-inclusion) among all s-admissible sets.*

Let the *closure* of $\Delta \subseteq Arg$ be $\mathrm{CL}(\Delta) = \{A \in Arg \mid \text{there is a sequence of supports from } B \in \Delta \text{ to } A\}$, we say Δ is *closed* iff $\Delta = \mathrm{CL}(\Delta)$.

Definition 9. *Let $\Delta \subseteq Arg$ be a set of arguments, Δ is called* c-admissible *iff Δ is conflict-free, self-defending and closed; Δ is a c-preferred extension if it is maximal (w.r.t. set-inclusion) among all c-admissible sets.*

We restate a proposition in [5] that demonstrates the relation between safety and conflict-freeness.

Proposition 1. *Let $\Delta \subseteq Arg$ be a set of arguments, if Δ is safe, then Δ is conflict-free. If Δ is conflict-free and closed, then Δ is safe.*

Definition 10. *Let $\Delta \subseteq Arg$ be a set of arguments, Δ is stable if and only if Δ is conflict-free and for every argument $A \in Arg \backslash \Delta$, Δ set-attacks A.*

It is worth mentioning that in the original papers, [5, 6] consider a particular set of BAFs, namly acyclic BAFs, showing that such BAFs have some nice features. However, in this paper, we foucs on BAFs that are more general, i.e., we remove the restriction on BAFs and consider both acyclic and cyclic BAFs. From a technical point of view, the BAFs that are acyclic have only one stable extension, which is the only preferred extension as well, while the BAFs with cycles could have more than one stable extension and will be more general.

There are several interpretations of support in the literature, including the deductive support, the necessary support, and the evidential support (see an overview in [7]). The deductive support [4] is intended to capture the intuition that given two arguments A and B, if A supports B, then the acceptance of A implies the acceptance of B. The necessary support [19, 20] is intended to capture the intuition that if $A \rightsquigarrow B$ is the case, then the the acceptance of B implies the acceptance of A, i.e., the acceptance of A is necessary to obtain the acceptance of B. Finally, the evidential support [22, 21] proposes a new type of argument, namely *prima-facie* arguments. Every standard argument is supposed to be supported by at least one *prima-facie* argument, and every *prima-facie* argument does not require support from other arguments.

The supported attack is connected with deductive support. To see this, let us come back to Figure 1, according to the deductive support, the acceptance of A_1 implies the acceptance of B_1, and so the acceptance of C_1, the acceptance of D_1. In the meantime, the acceptance of D_1 implies the non-acceptance of E_1. Thus, the acceptance of A_1 implies the non-acceptance of E_1. The necessary support can be taken into account by considering secondary attack. We again consider Figure 1. First, the acceptance of A_2 implies the non-acceptance of B_2. Then, according to necessary support, the non-acceptance of B_2 implies the non-acceptance of C_2, and so the non-acceptance of D_2, the non-acceptance of E_2. Thus, the acceptance of A_2 implies the non-acceptance of E_2.

3 The model

Fix a finite set Arg of arguments, a set $(\rightharpoonup)$ of attacks between arguments, and a set $N = \{1, \ldots, n\}$ of n agents. Each agent $i \in N$ supplies us with a set of supports $\rightsquigarrow_i$, which together with Arg and $(\rightharpoonup)$ gives rise to a bipolar argumentation framework $\langle Arg, \rightharpoonup, \rightsquigarrow_i \rangle$, reflecting her individual views on which supports between arguments are acceptable. A *profile* of support-relations $\boldsymbol{\rightsquigarrow} = (\rightsquigarrow_1, \ldots, \rightsquigarrow_n)$ is a set of support-relations provided by agents. An aggregation rule $F : (2^{Arg \times Arg})^n \to 2^{Arg \times Arg}$ is a function that maps a given profile of support-relations into a single support-relation. We denote $N^{\boldsymbol{\rightsquigarrow}}_{sup}$ by the set of agents who accept sup under profile $\boldsymbol{\rightsquigarrow}$, i.e., $N^{\boldsymbol{\rightsquigarrow}}_{sup} = \{i \in N \mid sup \in \rightsquigarrow_i\}$, and $\#N^{\boldsymbol{\rightsquigarrow}}_{sup}$ denotes the number of such agents.

Here we define desirable properties of aggregation rules. These properties are referred as axioms in the social choice literature. We start with formal definitions, followed by informal descriptions.

Definition 11. *An aggregation rule F is* unanimous *if* $\rightsquigarrow_1 \cap \ldots \cap \rightsquigarrow_n \subseteq F(\boldsymbol{\rightsquigarrow})$.

Definition 12. *An aggregation rule F is* grounded *if* $F(\boldsymbol{\rightsquigarrow}) \subseteq \rightsquigarrow_1 \cup \ldots \cup \rightsquigarrow_n$.

Definition 13. *An aggregation rule F is* neutral *if for any profile of support-relations $\boldsymbol{\rightsquigarrow}$, for any pair of supports sup_1, sup_2, $N^{\boldsymbol{\rightsquigarrow}}_{sup_1} = N^{\boldsymbol{\rightsquigarrow}}_{sup_2}$ then $sup_1 \in F(\boldsymbol{\rightsquigarrow})$ iff $sup_2 \in F(\boldsymbol{\rightsquigarrow})$.*

Definition 14. *An aggregation rule F is* independent *if for any pair of profiles of support-relations $\boldsymbol{\rightsquigarrow}_1$, $\boldsymbol{\rightsquigarrow}_2$, for any support sup, $N^{\boldsymbol{\rightsquigarrow}_1}_{sup} = N^{\boldsymbol{\rightsquigarrow}_2}_{sup}$ then $sup \in F(\boldsymbol{\rightsquigarrow}_1)$ iff $sup \in F(\boldsymbol{\rightsquigarrow}_2)$.*

Definition 15. *An aggregation rule F is* dictatorial *if there is an agent i such that for any profile of support-relations $\boldsymbol{\rightsquigarrow}, F(\boldsymbol{\rightsquigarrow}) = \rightsquigarrow_i$.*

The *unanimity* axiom states that the support agreed by all agents should be included in the collective BAF. The *groundedness* axiom expresses that all supports in the collective BAF should be supported by at least one agent. The *neutrality* axiom requires that given a profile of support-relations, any pair of supports should be treated equally in this profile. The *independent* axiom states that all support-relations should be treated equally in any profile of support-relations. The *dictatorship* axiom indicates that there is an agent who is dictatorial.

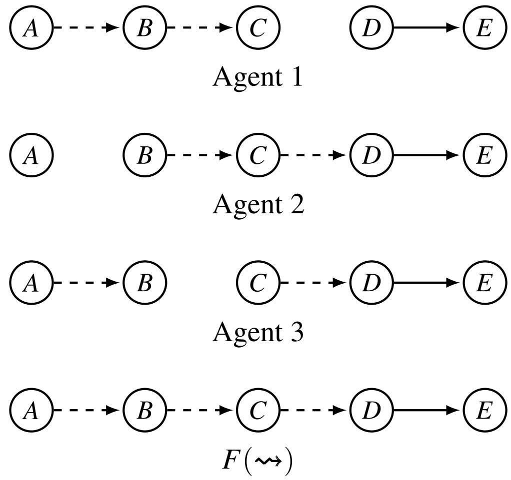

Figure 2: Example for a profile with $Arg = \{A,B,C,D,E\}$

Definition 16. *The unanimity rule is an aggregation rule F with $F(\rightsquigarrow) = \{sup \in Arg \times Arg \mid sup \in \rightsquigarrow_1 \cap \ldots \cap \rightsquigarrow_n\}$.*

Definition 17. *Let $i \in N$ be an agent, the dictatorship rule of individual i is the aggregation rule with $F(\rightsquigarrow) = \rightsquigarrow_i$.*

The unanimity rule only accepts those supports approved by all agents: it is a demanding aggregation rule. The dictatorships always return the supports submitted by dictators.

Example 2. Suppose that there are three agents have to decide on the acceptance of supports between four arguments. Agent 1 supports $A \rightsquigarrow B$ and $B \rightsquigarrow C$, agent 2 supports $B \rightsquigarrow C$ and $C \rightsquigarrow D$, agent 3 supports $A \rightsquigarrow B$ and $C \rightsquigarrow D$. We assume that the attack relation from D to E is accepted by all agents. The scenario is illustrated in Figure 2. If we apply the majority rule, then we obtain a bipolar argumentation framework consisting of the three supports $A \rightsquigarrow B$, $B \rightsquigarrow C$, and $C \rightsquigarrow D$. We observe that the set $\{A,E\}$ is conflict-free for all agents. However, it is not conflict-free in the outcome of the majority rule (which returns a set containing only the majoritarian supports) since A supported attacks E. So conflict-freeness as a semantic property is not preserved by the majority rule in this specific example. $\triangle$

But what about the preservation results of other semantic properties? Can they be preserved in general? Before going any further, we introduce more semantic properties of particular interest.

The problem we are considering in this paper is the preservation of semantic properties in the context of bipolar argumentation. Given a property $P \subseteq 2^{Arg \times Arg}$ that is a set of supports on Arg, and P is satisfied by all agents, whether the output of the aggregation rule satisfies P? A formal definition is as follows.

Definition 18. *An aggregation rule F preserves a property P if whenever for every profile $\rightsquigarrow$ we have that $P(\rightsquigarrow_i)$ for all $i \in N$, then we have $P(F(\rightsquigarrow))$.*

The problem of preservation is a special problem of *collective rationality* which has been discussed extensively in other parts of social choice, such as preference aggregation [1], judgment aggregation [16], graph aggregation [13], as well as attack aggregation in the context of abstract argumentation [3, 24, 9].

In the scenario where each agent possesses a BAF, agents might disagree on some details, such as whether a support between two arguments can be justified. Nevertheless, they may agree on some high-level features of BAFs. The *essential constraint* is an example of a high-level feature that requires no agent accepts both the attack relation and the support relation between a pair of arguments. When we observe that all agents verify such semantic feature, we would like to see what aggregation rule preserves this basic constraint under aggregation.

Given a set of arguments $\Delta \in Arg$ that is conflict-free in every agent's bipolar argumentation framework, we may wish to preserve its conflict-freeness in the outcome. Therefore, conflict-freeness as a semantic property is of particular interest. Similar definition can be posed to the *preservation of safety and admissibility*. Recall that if a set of arguments Δ is conflict-free and closed, then Δ is safe (Proposition 1). Thus, the *closedness* is of interest to us as well. Finally, we are also interested in the preservation of semantic extensions. Given a set of arguments $\Delta \subseteq Arg$ that is an extension of a specific semantics of $\langle Arg, \rightharpoonup, \rightsquigarrow_i \rangle$ for all $i \in N$, we are interested in under what circumstances Δ is an extension of such semantics of $F(\boldsymbol{\rightsquigarrow})$ as well. Finally, given an argument that is acceptable under a specific semantics for all agents, we would like to see whether such argument is acceptable in the collective outcome.

4 Preservation results

In this section, we present the preservation results for semantic properties. We start with *essential constraint* and *closedness*, two basic requirements of bipolar argumentation frameworks. Then, we turn to consider the preservation of *conflict-freeness*, followed by considering *safety*, followed by considering *d-admissibility*, *s-admissibility*, and *c-admissibility*. Then, we proceed with the study the properties of being an extension, including the property of *being a d-preferred extension*, *being a s-preferred extension*, *being a c-preferred extension* and *being a stable extension*. Finally, we study the preservation of *acceptability of arguments*. Proofs of results in this section can be found in the appendix.

4.1 Preservation results for essential constraint, closedness, conflict-freeness, safety and admissibility

Recall that a bipolar AF satisfies essential constraint if it does not contain two arguments for which the first one simultaneously attacks and supports the second one.

Proposition 2. *Every aggregation rule F that is grounded preserves essential constraint.*

The closedness is also an important property. Our result demonstrates that every reasonable rule preserves it.

Lemma 3. *Every aggregation rule F that is grounded preserves closedness.*

For conflict-freeness, we obtain that the unanimity rule, a demanding rule preserves the conflict-freeness of arbitrary sets of arguments.

Proposition 4. *The unanimity rule preserves conflict-freeness.*

The preservation of the safety of arbitrary sets of arguments can be accomplished by the unanimity rule.

Proposition 5. *The unanimity rule preserves safety.*

Proof. This proposition is a consequence of Proposition 4, Lemma 3, and Proposition 1. □

The concepts of d-admissibility and s-admissibility are based on different coherences, but the preservation results for them are similar, as the following proposition demonstrates.

Proposition 6. *The unanimity rule preserves either d-admissibility or s-admissibility.*

4.2 Preservation results for properties of being an extension

We are going to present preservation results for more demanding properties. Before proceeding, we introduce some necessary terminology and a simple result, as well as a technique developed by Endriss and Grandi for the more general framework of graph aggregation [13]. Let $sup \in \rightsquigarrow$ be a support, let $N = \{1, \ldots, n\}$ be a finite set of individuals (or agents, we assume that there are two or more agents), and let $\boldsymbol{\rightsquigarrow}$ be a profile of support-relations. Recall that $N^{\boldsymbol{\rightsquigarrow}}_{sup}$ is the set of agents who accept *sup* under profile $\boldsymbol{\rightsquigarrow}$. A *winning coalition* $\mathscr{W} \subseteq N$ is a set of agents who can decide whether to accept or reject a given support *sup*. Given an aggregation rule F, if F is *neutral* and *independent*, then F can be fully determined by a single set $\mathscr{W}$ of winning coalitions, i.e., for every profile $\boldsymbol{\rightsquigarrow}$ and every support *sup* it is the case that $sup \in F(\boldsymbol{\rightsquigarrow}) \Leftrightarrow N^{\boldsymbol{\rightsquigarrow}}_{sup} \in \mathscr{W}$.

In our proofs, we will rely on the concept of *ultrafilter* familiar from set theory [17]. An *ultrafilter* is a collection of subsets of N satisfying *closure under intersection*, *maximality*, and $\emptyset \notin \mathscr{W}$.

Definition 19. *An* ultrafilter $\mathscr{W}$ *on a set N is a collection of subsets of N satisfying the following conditions:*

(1) $\emptyset \notin \mathscr{W}$

(2) *for any pair of sets* $C_1, C_2 \subseteq N$, $C_1, C_2 \in \mathscr{W}$ *implies* $C_1 \cap C_2 \in \mathscr{W}$ *(closure under intersection)*

(3) *for any set C, one of C and* $N \backslash C$ *is in* $\mathscr{W}$ *(maximality)*

We restate a simple result, which interprets a well-known fact of ultrafilter in our context:

> *Let F be an independent and neutral aggregation rule and let* $\mathscr{W}$ *be the corresponding set of winning coalitions for supports, i.e.,* $sup \in F(\boldsymbol{\rightsquigarrow}) \Leftrightarrow N^{\boldsymbol{\rightsquigarrow}}_{sup} \in \mathscr{W}$ *for all* $sup \in \rightsquigarrow$*. Then, F is dictatorial if and only if* $\mathscr{W}$ *is an ultrafilter.*

Besides the properties identified in Section 3, we introduce two meta-properties:

Definition 20. *A property P is called* ***non-simple*** *if there exists a set* $Sup \subseteq Arg \times Arg$ *of supports and three individual supports* $sup_1, sup_2, sup_3 \in Arg \times Arg \setminus Sup$ *such that* $\langle Arg, \rightharpoonup, Sup \cup S \rangle$ *with* $S \subseteq \{sup_1, sup_2, sup_3\}$ *satisfies P if and only if* $S \neq \{sup_1, sup_2, sup_3\}$.

Definition 21. *A property P is called* ***disjunctive*** *if there exists a set* $Sup \subseteq Arg \times Arg$ *of supports and two individual supports* $sup_1, sup_2 \in Arg \times Arg \setminus Sup$ *such that* $\langle Arg, \rightharpoonup, Sup \cup S \rangle$ *with* $S \subseteq \{sup_1, sup_2\}$ *satisfies P if and only if* $S \neq \emptyset$.

Non-simplicity requires that, in the context of *Sup*, accepting any proper subset of $\{sup_1, sup_2, sup_3\}$ is possible, while accepting $\{sup_1, sup_2, sup_3\}$ is not. Disjunctiveness requires that, in the context of *Sup*, we should accept at least one of sup_1 and sup_2. The term of *simplicity* was introduced by Nehring and Puppe [18] as the *median property*; the term of *disjunctiveness* was introduced by Endriss and Grandi [13] as a graph meta-property. It is worth noting that a meta-property is a class of properties, a property satisfies or does not satisfy a specific meta-property. Even though the definitions of meta-properties are not complicated, deciding whether a given property belongs to a meta-property is still not straightforward.

Meta-properties have connections with properties of BAFs and aggregation rules: on the one hand, meta-properties outline high-level features of properties of BAFs, with which we are able to systematically study the preservation of semantic properties of BAFs with simple proofs; on the other hand, as we will see in the following lemmas, meta-properties allow us to generalize specific result for specific properties by instantiating the general results. To be more specific, if an aggregation rule preserves a property that belongs to a meta-property, then it is a dictatorship.

Lemma 7. *Let P be a property that is non-simple and disjunctive. Then, for $|Arg| \geqslant 3$, any unanimous, grounded, neutral, and independent aggregation rule F that preserves P must be a dictatorship.*

If a property we are interested in is non-simple and disjunctive, then we can apply Lemma 7 to obtain an axiomatic result for it.

Theorem 8. *For $|Arg| \geqslant 5$, any unanimous, grounded, neutral, and independent aggregation rule F that preserves either d-preferred, s-preferred, or c-preferred extensions must be a dictatorship.*

For the scenarios when $|Arg| = 4$, or even $|Arg| = 3$, we are not able to verify whether the theorem can still apply, we conjecture that the bound on the cardinality of Arg is sharp and we believe that the theorem has covered all cases of practical interest. By comparison, we recall that for the property of being a preferred extension of Dung's argumentation framework, Chen and Endriss have shown that only dictatorships preserve it [9]. They have made the assumption that every agent is equipped with a different set of attack relations while they hold the same set of arguments.

Note that Theorem 8 is an impossibility result that indicates the preservation of specific properties is impossible, unless the aggregation rule under consideration is dictatorial. They relate to generalisations of Arrow's Impossibility Theorem [1] to graph aggregation and attack-relation aggregation. One of the features of the aggregation rules we used in this section is that they accept the axiom of independence. Even though this axiom is attractive in some sense, to escape the impossibilities, a prime direction is to relax it. For example, we can consider distance-based rules and investigate whether we are able to obtain some positive results.

Theorem 8 has assumed that the interpretation of support is deductive support. Even though it is enough for our purposes, namely it is enough to show that only dictatorships preserve d-preferred (s-preferred, c-preferred) extensions, we are still interested in what happens when the interpretation is restricted to necessary support. The bad news is, we still cannot overcome impossibility results.

Theorem 9. *If the interpretation of support is necessary support, for $|Arg| \geqslant 5$, any unanimous, grounded, neutral, and independent aggregation rule F that preserves either d-preferred, s-preferred, or c-preferred extensions must be a dictatorship.*

4.2.1 Preservation result for stable extensions

For stable extensions, by using the same techniques, we obtain a similar impossibility result.

Theorem 10. *For $|Arg| \geqslant 5$, any unanimous, grounded, neutral, and independent aggregation rule F that preserves stable extensions must be a dictatorship.*

By comparison, we recall that the nomination rule preserves stable extensions of Dung's argumentation frameworks [9].

4.3 Preservation of argument acceptability

Now, we move to study the preservation of acceptability of arguments. Before proceeding further, it is thus important to keep in mind that our model has assumed that each agent $i \in N$ reports a set of supports $\rightsquigarrow_i$ on the same set of arguments and attack relations. Given two BAFs $\rightsquigarrow_1$ and $\rightsquigarrow_2$, if $\rightsquigarrow_1 \supseteq \rightsquigarrow_2$, namely the supports of $\rightsquigarrow_1$ is a superset of $\rightsquigarrow_2$, then a d-admissible (s-admissible, c-admissible, respectively) set of $\rightsquigarrow_1$ remains d-admissible (s-admissible, c-admissible, respectively) in $\rightsquigarrow_2$.

Lemma 11. *Given two BAFs $\rightsquigarrow_1$ and $\rightsquigarrow_2$, if $\rightsquigarrow_1 \supseteq \rightsquigarrow_2$, then every d-admissible set of arguments of $\rightsquigarrow_1$ is a d-admissible set of $\rightsquigarrow_2$.*

Lemma 12. *Given two BAFs $\rightsquigarrow_1$ and $\rightsquigarrow_2$, if $\rightsquigarrow_1 \supseteq \rightsquigarrow_2$, then every s-admissible set of arguments of $\rightsquigarrow_1$ is a s-admissible set of $\rightsquigarrow_2$.*

Lemma 13. *Given two BAFs $\rightsquigarrow_1$ and $\rightsquigarrow_2$, if $\rightsquigarrow_1 \supseteq \rightsquigarrow_2$, then every c-admissible set of arguments of $\rightsquigarrow_1$ is a c-admissible set of $\rightsquigarrow_2$.*

Fact 14. *Given a BAF $\rightsquigarrow$, if $\Delta \subseteq Arg$ is a d-preferred (s-preferred, c-preferred, respectively) extension of $\rightsquigarrow$, then Δ is a d-admissible (s-admissible, c-admissible, respectively) set of arguments of $\rightsquigarrow$.*

Thus, every d-admissible (s-admissible, c-admissible, respectively) set of arguments is included in a d-preferred (s-preferred, c-preferred, respectively) extension. With this, we are ready to present the preservation results for argument acceptability under preferred semantics, including d-preferred semantics, s-preferred semantics, and c-preferred semantics. Note that in the following we say that an argument under a d-preferred extension, we mean that such argument is a member of a d-preferred extension.

Proposition 15. *The unanimity rule preserves the property of argument acceptability under d-preferred semantics.*

Proposition 16. *The unanimity rule preserves the property of argument acceptability under either s-preferred or c-preferred semantics.*

The proof is similar to the proof for Theorem 15. The only difference is that every s-admissible (c-admissible) set of $\rightsquigarrow_i$ is a s-admissible set of $F(\boldsymbol{\rightsquigarrow})$ is because of Lemma 12 (Lemma 13).

5 Related work

Previous work on obtaining argumentative consensus among a group of agents are mainly focus on abstract argumentation frameworks [10, 25, 9]. Among them, Chen and Endriss [9] study of the preservation of semantic properties during the aggregation of attack-relations of abstract argumentation frameworks. As a potential domain of application for the model they develop, they do not make explicit reference to bipolar argumentation frameworks. In addition, similar to us, they have made use of meta-properties proposed by Endriss and Grandi for graph aggregation [13], which serve as technical devices to obtain preservation results for semantic properties.

The problem of aggregating bipolar opinions has received interests in the literature. The idea of aggregating support-relations of bipolar argumentation frameworks has been outlined in a preliminary version of this paper [8]. Lauren *et al*. [15] consider aggregating bipolar assumption-based argumentation frameworks under the assumption that agents propose the same set of arguments, but propose different sets of attacks and supports. Their focus is quota rules and oligarchic rules. Kontarinis *et al*. [14] study designing mechanisms for "regulating" debates under the setting of each agent in the debate equipped with a bipolar argumentation framework. We note that in their settings, agents report the same set of arguments, but with possibly different attack- and support-relations.

6 Conclusion

In this paper, we have studied the aggregation of agents' view in the context of bipolar argumentation. To be more specific, we have explored the problem of aggregating support-relations of bipolar argumentation frameworks by making use of the methodology of social choice theory. To achieve this, we have designed a model, in which agents are equipped with a set of arguments and a set of attacks, but with possibly different support-relations. We have shown which semantic properties of BAFs can be preserved by aggregation rules. We have proposed two BAF meta-properties, namely the property of "non-simplicity" and "disjunctiveness", both of which are high level features of BAFs. We show that such meta-properties can be used to obtain impossibility results, namely for quickly proving what kind of aggregation rules (or no desirable aggregation rule) is collectively rational with respect to BAF properties.

For future work, it is worth having an investigation of further meta-properties. We point out that Lemma 7 is a variant of Theorem 16 in graph aggregation [13]. This indicates that there is space for other variants of impossibility results for different assumptions. The preservation of some desirable properties (for example, the property of being a d-preferred extension) during the aggregation of support-relations is difficult. Thus, it is worth studying whether such properties can be preserved in different settings. For instance, agents might be equipped with the same set of arguments, but with possibly different attack- and support-relations, and we aggregate attack- and support-relations by making use of different quota rules. Finally, recall that there are at least three interpretations of support, in this paper, we focus on the deductive support and necessary support. More interpretations of support, for example, evidential support, should be investigated.

Acknowledgements. I would like to thank Ulle Endriss for his generous suggestions on an earlier version and his guidance at the beginning of this work. I also thank three anonymous reviewers for their constructive feedback. This work was supported by the China Postdoctoral Science Foundation (grant no. 2019M663352).

References

[1] Kenneth J. Arrow (1963): *Social Choice and Individual Values*, 2nd edition. John Wiley and Sons. First edition published in 1951.

[2] Philippe Besnard & Anthony Hunter (2008): *Elements of Argumentation*. MIT Press, doi:10.7551/mitpress/9780262026437.001.0001.

[3] Gustavo M. Bodanza, Fernando A. Tohmé & Marcelo R. Auday (2017): *Collective Argumentation: A Survey of Aggregation Issues around Argumentation Frameworks*. Argument & Computation 8(1), pp. 1–34, doi:10.3233/AAC-160014.

[4] Guido Boella, Dov Gabbay, Leon van der Torre & Serena Villata (2010): *Support in abstract argumentation*. In: Proceedings of the Third International Conference on Computational Models of Argument (COMMA-10), Frontiers in Artificial Intelligence and Applications, IOS Press, pp. 40–51, doi:10.3233/978-1-60750-619-5-111.

[5] Claudette Cayrol & Marie-Christine Lagasquie-Schiex (2005): *Gradual valuation for bipolar argumentation frameworks*. In: European Conference on Symbolic and Quantitative Approaches to Reasoning and Uncertainty, Springer, pp. 366–377, doi:10.1007/11518655_32.

[6] Claudette Cayrol & Marie-Christine Lagasquie-Schiex (2005): *On the acceptability of arguments in bipolar argumentation frameworks*. In: European Conference on Symbolic and Quantitative Approaches to Reasoning and Uncertainty, Springer, pp. 378–389, doi:10.1007/11518655_33.

[7] Claudette Cayrol & Marie-Christine Lagasquie-Schiex (2013): *Bipolarity in argumentation graphs: Towards a better understanding*. International Journal of Approximate Reasoning 54(7), pp. 876–899, doi:10.1016/j.ijar.2013.03.001.

[8] Weiwei Chen (2020): *Aggregation of Support-Relations of Bipolar Argumentation Frameworks (Extended Abstract)*. In: Proceedings of the 19th International Conference on Autonomous Agents and Multiagent Systems (AAMAS-2020), IFAAMAS.

[9] Weiwei Chen & Ulle Endriss (2019): *Preservation of Semantic Properties in Collective Argumentation: The Case of Aggregating Abstract Argumentation Frameworks*. Artificial Intelligence 269, pp. 27–48, doi:10.1016/j.artint.2018.10.003.

[10] Sylvie Coste-Marquis, Caroline Devred, Sébastien Konieczny, Marie-Christine Lagasquie-Schiex & Pierre Marquis (2007): *On the Merging of Dung's Argumentation Systems*. Artificial Intelligence 171(10–15), pp. 730–753, doi:10.1016/j.artint.2007.04.012.

[11] Phan Minh Dung (1995): *On the Acceptability of Arguments and its Fundamental Role in Nonmonotonic Reasoning, Logic Programming and n-Person Games*. Artificial Intelligence 77(2), pp. 321–358, doi:10.1016/0004-3702(94)00041-X.

[12] Paul E. Dunne, Pierre Marquis & Michael Wooldridge (2012): *Argument Aggregation: Basic Axioms and Complexity Results*. In: Proceedings of the 4th International Conference on Computational Models of Argument (COMMA), IOS Press, doi:10.3233/978-1-61499-111-3-129.

[13] Ulle Endriss & Umberto Grandi (2017): *Graph Aggregation*. Artificial Intelligence 245, pp. 86–114, doi:10.1016/j.artint.2017.01.001.

[14] Dionysios Kontarinis, Elise Bonzon, Nicolas Maudet & Pavlos Moraitis (2011): *Regulating Multiparty Persuasion with Bipolar Arguments: Discussion and Examples*. In: Modèles Formels de l'Interaction, pp. 119–129.

[15] Stefan Lauren, Francesco Belardinelli & Francesca Toni (2021): *Aggregating Bipolar Opinions*. In: Proceedings of the 20th International Conference on Autonomous Agents and Multiagent Systems (AAMAS).

[16] Christian List & Philip Pettit (2002): *Aggregating Sets of Judgments: An Impossibility Result*. Economics and Philosophy 18(1), pp. 89–110, doi:10.1017/S0266267102001098.

[17] Bernard Monjardet (1983): *On the use of ultrafilters in social choice theory*. Social choice and welfare 5, pp. 73–78, doi:10.1016/B978-0-444-86487-1.50012-5.

[18] Klaus Nehring & Clemens Puppe (2007): *The structure of strategy-proof social choice—Part I: General characterization and possibility results on median spaces*. Journal of Economic Theory 135(1), pp. 269–305, doi:10.1016/j.jet.2006.04.008.

[19] Farid Nouioua & Vincent Risch (2010): *Bipolar argumentation frameworks with specialized supports*. In: 2010 22nd IEEE International Conference on Tools with Artificial Intelligence, 1, IEEE, pp. 215–218, doi:10.1109/ICTAI.2010.37.

[20] Farid Nouioua & Vincent Risch (2011): *Argumentation frameworks with necessities*. In: Proceedings of the International Conference on Scalable Uncertainty Management, Springer, pp. 163–176, doi:10.1007/978-3-642-23963-2_14.

[21] Nir Oren, Michael Luck & Chris Reed (2010): *Moving between argumentation frameworks*. In: Proceedings of the 2010 International Conference on Computational Models of Argument (COMMA-2010), IOS Press, doi:10.3233/978-1-60750-619-5-379.

[22] Nir Oren & Timothy J. Norman (2008): *Semantics for evidence-based argumentation*. Proceedings of the 2008 International Conference on Computational Models of Argument (COMMA-2008) 172, p. 276.

[23] Iyad Rahwan & Guillermo R. Simari (2009): *Argumentation in Artificial Intelligence*. Springer-Verlag, doi:10.1007/978-0-387-98197-0.

[24] Iyad Rahwan & Fernando A. Tohmé (2010): *Collective Argument Evaluation as Judgement Aggregation*. In: Proceedings of the 9th International Conference on Autonomous Agents and Multiagent Systems (AAMAS).

[25] Fernando A. Tohmé, Gustavo A. Bodanza & Guillermo R. Simari (2008): *Aggregation of Attack Relations: A Social-choice Theoretical Analysis of Defeasibility Criteria*. In: Proceedings of the 5th International Symposium on Foundations of Information and Knowledge Systems (FoIKS), Springer-Verlag, doi:10.1007/978-3-540-77684-0_4.

[26] Kristijonas Čyras, Claudia Schulz & Francesca Toni (2017): *Capturing Bipolar Argumentation in Non-flat Assumption-Based Argumentation*. In: International Conference on Principles and Practice of Multi-Agent Systems, pp. 386–402, doi:10.1007/978-3-319-69131-2_23.

Appendix: Remaining Proofs

In this appendix we present the proofs omitted from the body of the paper.

Proof of Proposition 2

Let $\boldsymbol{\rightsquigarrow} = (\rightsquigarrow_1, \ldots, \rightsquigarrow_n)$ be a profile of BAFs, in which $\rightsquigarrow_i$ satisfies the essential constraint for all $i \in N$. Let F be an aggregation rule that is grounded. For the sake of contradiction, we suppose that the essential constraint is violated in $F(\boldsymbol{\rightsquigarrow})$. Without loss of generality, we suppose that both $A \rightharpoonup B$ and $A \rightsquigarrow B$ get accepted in $F(\boldsymbol{\rightsquigarrow})$. From the assumption, we know that every agent agrees $A \rightharpoonup B$. In the meantime, at least one agent accepts $A \rightsquigarrow B$ under grounded aggregation rules, which cannot be the case. Thus, we have the proposition. □

Proof of Lemma 3

We again let $\boldsymbol{\rightsquigarrow} = (\rightsquigarrow_1, \ldots, \rightsquigarrow_n)$ be a profile of BAFs, let $\Delta \subseteq Arg$ be the set of arguments under consideration, and let F be an aggregation rule that is grounded. For the sake of contradiction, we suppose that Δ is closed in $\rightsquigarrow_i$ for all $i \in N$ and not closed in $F(\boldsymbol{\rightsquigarrow})$, i.e., there is an argument $A \in \Delta$ and an argument $B \in Arg \backslash \Delta$ such that $A \rightsquigarrow B$ in $F(\boldsymbol{\rightsquigarrow})$. As rules under considering are grounded, there is at least one agent $\rightsquigarrow_i$ for which $A \rightsquigarrow_i B$ is the case. This will lead to the situation that Δ not being closed in $\rightsquigarrow_i$, contradicting our assumption. □

Proof of Proposition 4

Recall that the unanimity rule is the quota rule F with the quota of n. Let $\boldsymbol{\rightsquigarrow} = (\rightsquigarrow_1, \ldots, \rightsquigarrow_n)$ be a profile of BAFs. Let $\Delta \subseteq Arg$ be the set of arguments under consideration. For the sake of contradiction, we suppose that Δ is conflict-free in $\rightsquigarrow_i$ for all $i \in N$, and is not conflict-free in $F(\boldsymbol{\rightsquigarrow})$. This means that there are two arguments $A, B \in \Delta$ such that A supported or secondary attacks B in $F(\boldsymbol{\rightsquigarrow})$.

We now show that in the scenario where A is supported attacking B, i.e., there is a sequence of arguments in $F(\boldsymbol{\rightsquigarrow})$ such that $A_1 \rightsquigarrow, \ldots, \rightsquigarrow A_m, A_m \rightharpoonup B$ in which $A_1 = A$, our proposition holds. From the assumption we know that $(A_m \rightharpoonup B) \in \rightharpoonup_i$ for all $i \in N$. In addition, every agent agrees $A_1 \rightsquigarrow, \ldots, \rightsquigarrow A_m$.

Thus, every agent agrees $A_1 \rightsquigarrow, \ldots, \rightsquigarrow A_m$, $A_m \rightharpoonup B$, i.e., Δ is not conflict-free in $\rightsquigarrow_i$ for all $i \in N$, in contradiction to our earlier assumption.

For the scenario where A is secondary attacking B, we note that the proof is similar to the proof of the one where A is supported attacking B. Thus, we have the proposition. □

Proof of Proposition 6

Let F be the unanimity rule. Let $\boldsymbol{\rightsquigarrow} = (\rightsquigarrow_1, \ldots, \rightsquigarrow_n)$ be a profile of bipolar argumentation frameworks. Let $\Delta \subseteq Arg$ be the set of arguments under consideration.

We suppose that Δ is d-admissible in $\rightsquigarrow_i$ for all $i \in N$. Then, Δ is conflict-free in $\rightsquigarrow_i$ for all $i \in N$ as well. By Proposition 4, Δ is conflict-free in $F(\boldsymbol{\rightsquigarrow})$. By the assumption that all agents report the same set of attacks, we get that for each argument $A \in \Delta$, A is defended by Δ in $\rightsquigarrow_i$ for all $i \in N$. It follows that A is defended by Δ in $F(\boldsymbol{\rightsquigarrow})$ as well. Thus, Δ is d-admissible in $F(\boldsymbol{\rightsquigarrow})$.

We omit the relative easy proof for s-admissibility. □

Proof of Lemma 7

Take any property P that is non-simple and disjunctive. Take any aggregation rule F that is unanimous, grounded, neutral, independent and preserves P. By the assumption that F is neutral and independent, F is determined by a set of winning coalitions $\mathscr{W} \subseteq 2^N$. What we need to do is proving that $\mathscr{W}$ is an ultrafilter, i.e., to show that $\mathscr{W}$ is closed under intersection, $\mathscr{W}$ satisfies maximality, and $\emptyset \notin \mathscr{W}$.

$\emptyset \notin \mathscr{W}$ This is a direct consequence of the assumption that F is grounded.

Maximality Take any set of agents $C \subseteq N$. Consider a profile in which exactly the individuals in C propose sup_1 and exactly those in $N \backslash C$ propose sup_2. Since P is disjunctive, we know that one of sup_1 and sup_2 must be part of $F(\boldsymbol{\rightsquigarrow})$. Hence $C \in \mathscr{W}$ or $N \backslash C \in \mathscr{W}$.

Closure under intersection Take any two winning coalitions $C_1, C_2 \in \mathscr{W}$. Assume toward a contradiction that $C_1 \cap C_2 \notin \mathscr{W}$. Consider a profile in which exactly the individuals in C_1 propose sup_1, exactly the individuals in C_2 propose sup_2, and exactly the individuals in $N \backslash (C_1 \cap C_2)$ propose sup_3. Now, since C_1 and C_2 are winning coalitions, sup_1 and sup_2 must be part of $F(\boldsymbol{\rightsquigarrow})$. Hence, due to $C_1 \cap C_2 \notin \mathscr{W}$ and $\mathscr{W}$ satisfying maximality, we have $N \backslash (C_1 \cap C_2) \in \mathscr{W}$. Since the individuals in $N \backslash (C_1 \cap C_2)$ propose sup_3, we have $sup_3 \in F(\boldsymbol{\rightsquigarrow})$. But we have assumed that F preserves non-simplicity, i.e., sup_1, sup_2, sup_3 cannot be accepted together in $F(\boldsymbol{\rightsquigarrow})$. Thus, $C_1 \cap C_2 \in \mathscr{W}$. □

Proof of Theorem 8

Suppose $|Arg| \geqslant 5$. Let P be the properties representing a given set of arguments being either a d-preferred, a s-preferred or a c-preferred extension. Thus, by Lemma 7, we need to show that P is non-simple and disjunctive in this case.

Non-simplicity Let $Arg = \{A, B, C, D, E, \ldots\}$, let $\rightharpoonup = \{D \rightharpoonup E, E \rightharpoonup D, B \rightharpoonup B, C \rightharpoonup C\}$. Now we focus on $Arg \setminus \{B, C, D\}$ as the subset of arguments that may (or may not) form either a d-preferred, a s-preferred or a c-preferred extension. We define $Sup = \emptyset$, $sup_1 = (A \rightsquigarrow B)$, $sup_2 = (B \rightsquigarrow C)$, and $sup_3 = (C \rightsquigarrow D)$. This scenario is depicted in the top part of Figure 3. Consider all BAFs of the form

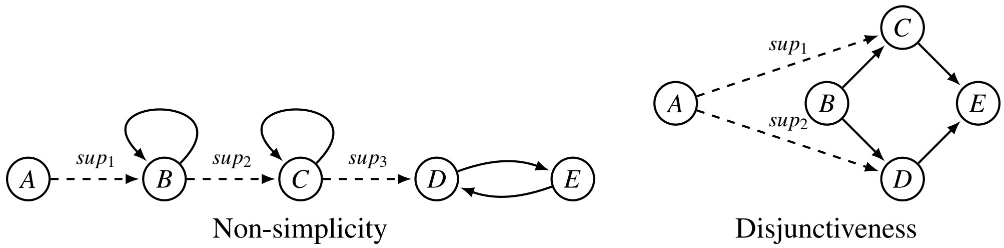

Figure 3: Scenarios used in the proof of Theorem 8

$BAF = \langle Arg, \rightharpoonup, Sup \cup S\rangle$ with $S \subseteq \{sup_1, sup_2, sup_3\}$. It is not difficult for the reader to verify that, for $S \neq \{sup_1, sup_2, sup_3\}$, B and C are self-attacking, D is attacked by E. Thus, they are unacceptable with respect to $\{A, E\}$, i.e., $Arg \setminus \{B, C, D\}$ is a d-preferred, a s-preferred and a c-preferred extension. On the other hand, for $S = \{sup_1, sup_2, sup_3\}$, $Arg \setminus \{B, C, D\}$ is not a d-preferred nor a s-preferred or c-preferred extension as E is set-attacked by A. Thus, P is non-simple.

Disjunctiveness Let $Arg = \{A, B, C, D, E \ldots\}$, let $\rightharpoonup = \{B \rightharpoonup C, B \rightharpoonup D, C \rightharpoonup E, D \rightharpoonup E\}$. We focus on $Arg \setminus \{C, D, E\}$ as the subset of arguments that may (or may not) form a s-preferred extension. We define $Sup = \emptyset$, $sup_1 = (A \rightsquigarrow C)$, $sup_2 = (A \rightsquigarrow D)$. This scenario is depicted in the bottom part of Figure 3. Consider all BAFs of the form $BAF = \langle Arg, \rightharpoonup, Sup \cup S\rangle$ with $S \subseteq \{sup_1, sup_2\}$. For $S \neq \emptyset$, $Arg \setminus \{C, D, E\}$ is a d-preferred, a s-preferred, and a c-preferred extension. On the other hand, for $S = \emptyset$, $Arg \setminus \{C, D, E\}$ is not a s-preferred nor a d-preferred or a c-preferred extension as E is defended by B. Thus, P is disjunctive. $\square$

Proof of Theorem 9

Similar to Theorem 8, we still need to show that the property of being a d-preferred, a s-preferred, or a c-preferred extension is non-simple and disjunctive when the interpretation of support is necessary support.

Non-simplicity Let $Arg = \{A, B, C, D, E, \ldots\}$, let $\rightharpoonup = \{D \rightharpoonup E, E \rightharpoonup D, B \rightharpoonup B, C \rightharpoonup C\}$. Now we focus on $Arg \setminus \{B, C, D\}$ as the subset of arguments that may (or may not) form either a d-preferred, a s-preferred or a c-preferred extension. We define $Sup = \emptyset$, $sup_1 = (B \rightsquigarrow A)$, $sup_2 = (C \rightsquigarrow B)$, and $sup_3 = (D \rightsquigarrow C)$, as illustrated in the top part of Figure 4. Consider all BAFs of the form $BAF = \langle Arg, \rightharpoonup, Sup \cup S\rangle$ with $S \subseteq \{sup_1, sup_2, sup_3\}$. It is not difficult for the reader to verify that, for $S \neq \{sup_1, sup_2, sup_3\}$, B and C are self-attacking, D is attacked by E. Thus, they are unacceptable with respect to $\{A, E\}$. In the meantime, A is not attacked by any other argument, E defends itself, $\{A, E\}$ is conflict-free, i.e., $Arg \setminus \{B, C, D\}$ is a d-preferred, a s-preferred, and a c-preferred extension. On the other hand, for $S = \{sup_1, sup_2, sup_3\}$, $Arg \setminus \{B, C, D\}$ is neither a d-preferred nor a s-preferred nor c-preferred extension as E secondary attacks A. Thus, P is non-simple.

Disjunctiveness Let $Arg = \{A, B, C, D \ldots\}$, let $\rightharpoonup = \{B \rightharpoonup C, B \rightharpoonup D\}$. We focus on $Arg \setminus \{A, C, D\}$ as the subset of arguments that may (or may not) form a s-preferred extension. We define $Sup = \emptyset$, $sup_1 = (C \rightsquigarrow A)$, $sup_2 = (D \rightsquigarrow A)$, as illustrated in the bottom part of Figure 4. Consider all BAFs of the form $BAF = \langle Arg, \rightharpoonup, Sup \cup S\rangle$ with $S \subseteq \{sup_1, sup_2\}$. For $S \neq \emptyset$, B secondary attacks A and directly

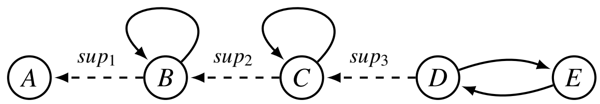

Non-simplicity

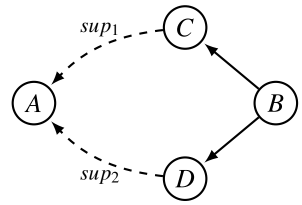

Disjunctiveness

Figure 4: Scenarios used in Theorem 9.

attacks C and D. Thus, $Arg \setminus \{A,C,D\}$ is a d-preferred, a s-preferred, and a c-preferred extension. On the other hand, for $S = \emptyset$, $Arg \setminus \{A,C,D\}$ is neither a s-preferred nor a d-preferred nor a c-preferred extension as A is not attacked by any other argument and thus should be included in every d-preferred (s-preferred, c-preferred) extension. Thus, P is disjunctive. □

Proof of Theorem 10

Suppose $|Arg| \geqslant 5$. Let P be the BAF-properties representing a given set of arguments being a stable extension. We need to demonstrate that P is non-simple and disjunctive in this case.

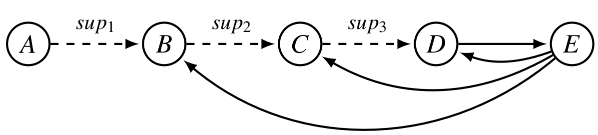

Non-simplicity

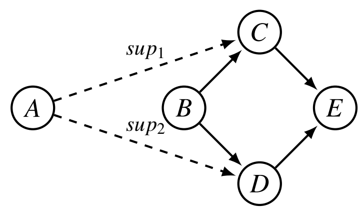

Disjunctiveness

Figure 5: Scenarios used in the proof of Theorem 10

Non-simplicity Let $Arg = \{A,B,C,D,E,\ldots\}$, let $\rightharpoonup = \{D \rightharpoonup E, E \rightharpoonup B, E \rightharpoonup C, E \rightharpoonup D\}$. We focus on $Arg \setminus \{B,C,D\}$ as the subset of arguments that may (or may not) form a stable extension. We define $Sup = \emptyset$. $sup_1 = (A \rightsquigarrow B)$, $sup_2 = (B \rightsquigarrow C)$, and $sup_3 = (C \rightsquigarrow D)$. This scenario is depicted in the top part of Figure 5. Consider all BAFs of the form $BAF = \langle Arg, \rightharpoonup, Sup \cup S \rangle$ with $S \subseteq \{sup_1, sup_2, sup_3\}$. The reader should be able to verify that, indeed, for $S \neq \{sup_1, sup_2, sup_3\}$, $Arg \setminus \{B,C,D\}$ is a stable extension. For example, for $S = \{sup_1, sup_2\}$, B, C, and D are attacked by E. In the meantime, $\{A,E\}$ is conflict-free, i.e., $Arg \setminus \{B,C,D\}$ is a stable extension. On the other hand, for $S = \{sup_1, sup_2, sup_3\}$, $Arg \setminus \{B,C,D\}$ is not a stable extension as E is set-attacked by A. Thus, P is non-simple.

Disjunctiveness Let $Arg = \{A,B,C,D,E\ldots\}$, let $\rightharpoonup = \{B \rightharpoonup C, B \rightharpoonup D, C \rightharpoonup E, D \rightharpoonup E\}$. We focus on $Arg \setminus \{C,D,E\}$ as the subset of arguments that may (or may not) form a stable extension. We define $Sup = \emptyset$, $sup_1 = (A \rightsquigarrow C)$, $sup_2 = (A \rightsquigarrow D)$. This scenario is depicted in the bottom part of Figure 5. Consider all BAFs of the form $BAF = \langle Arg, \rightharpoonup, Sup \cup S \rangle$ with $S \subseteq \{sup_1, sup_2\}$. The reader should be able to verify that, indeed, for $S \neq \emptyset$, $Arg \setminus \{C,D,E\}$ is a stable extension. On the other hand, for $S = \emptyset$, $Arg \setminus \{C,D,E\}$ is not a stable extension as E is defended by B. Thus, P is disjunctive. □

Proof of Lemma 11

Let $\Delta \subseteq Arg$ be a d-admissible set of $\rightsquigarrow_1$. We need to show that Δ is a d-admissible set of $\rightsquigarrow_2$. To achieve this, we need to demonstrate that in $\rightsquigarrow_2$, (*i*) Δ is conflict-free, and (*ii*) Δ defends all of its members.

For (*i*), we need to show that Δ is conflict-free in $\rightsquigarrow_2$. If not, then there are two arguments $A, B \in \Delta$ such A directly, supported, or secondary attacks B. If A directly attacks B in $\rightsquigarrow_2$, then A directly attacks B in $\rightsquigarrow_1$ as the pair of BAFs report the same set of attacks, which contradicts the assumption that Δ is conflict-free in $\rightsquigarrow_1$. If A supported attacks B in $\rightsquigarrow_2$, then there is a sequence of argument $(A_1, \ldots, A_n)$ such that $A_1 \rightsquigarrow A_2, \ldots, A_{n-1} \rightharpoonup A_n$, $A = A_1$, and $A_n = B$. As $\rightsquigarrow_1 \supseteq \rightsquigarrow_2$ and $A_{n-1} \rightharpoonup A_n$ is the case, we know that $A_1 \rightsquigarrow A_2, \ldots, A_{n-1} \rightharpoonup A_n$ in $\rightsquigarrow_2$ as well, which means that there are two arguments $A, B \in \Delta$ such that A supported attacks B, contradicting the fact that Δ is conflict-free in $\rightsquigarrow_1$. If A supported attacks B in $\rightsquigarrow_2$, this case is similar to the case that A supported attacks B, which will lead to that Δ failing to satisfy conflict-freeness in $\rightsquigarrow_1$. Thus, Δ is conflict-free in $\rightsquigarrow_2$.

For (*ii*), we need to show that for every argument $A \in \Delta$, if $B \rightharpoonup A$, then there is a $C \in \Delta$ such that $C \rightharpoonup B$, i.e, Δ defends all its members in $\rightsquigarrow_2$. Clearly, this is true as $\rightsquigarrow_1$ and $\rightsquigarrow_2$ report the same set of attacks, and Δ defends all its members in $\rightsquigarrow_1$. $\square$

Proof of Lemma 12

We need to show that a s-admissible set of arguments $\Delta \subseteq Arg$ of $\rightsquigarrow_1$ is a s-admissible set of $\rightsquigarrow_2$. To arrive at this goal, we need to demonstrate that in $\rightsquigarrow_2$, (*i*) Δ is conflict-free, (*ii*) Δ defends all of its members, and (*iii*) Δ is safe. For (*i*) and (*ii*), the proofs are the same as the ones in Lemma 11. It remains to show that Δ is safe in $\rightsquigarrow_2$. If not, then there is an argument $B \in Arg$ such that Δ set-attacks B and Δ set-supports B, or $B \in \Delta$. If Δ set-supports B. there are two arguments $A \in \Delta$ such A directly, supported, or secondary attacks B. Using the construction similar to Lemma 11, it is easy to verify that under this assumption, Δ set-attacks B in $\rightsquigarrow_1$. According to the assumption that $\rightsquigarrow_1 \supseteq \rightsquigarrow_2$, Δ set-supports B in $\rightsquigarrow_1$. Then, Δ is not safe in $\rightsquigarrow_1$, contradiction. $\square$

Proof of Lemma 13

Once again, we need to show that every c-admissible set Δ of arguments of $\rightsquigarrow_1$ is a c-admissible set of $\rightsquigarrow_2$. To arrive at this goal, we need to show that in $\rightsquigarrow_2$ (*i*), Δ is conflict-free, (*ii*) Δ defends all of its members, and (*iii*) Δ is closed. For (*i*) and (*ii*), the proofs are the same as the ones in Lemma 11. It remains to show that Δ is closed in $\rightsquigarrow_2$. If not, then there is an argument $A \in \Delta$ and an argument $B \in Arg$ such that A supports B, and $B \notin \Delta$. According to the assumption that $\rightsquigarrow_1 \supseteq \rightsquigarrow_2$, A supports B, and $B \notin \Delta$ in $\rightsquigarrow_1$, we get that Δ is not closed in $\rightsquigarrow_1$, contradiction. $\square$

Proof of Proposition 15

Let $A \in Arg$ be the argument under consideration, and we suppose that A is acceptable under a d-preferred extension of $\rightsquigarrow_i$ for all $i \in N$. Let F be the unanimity rule. Clearly, $F(\boldsymbol{\rightsquigarrow})$ is a subset of $\rightsquigarrow_i$ for all $i \in N$. Without loss of generality, we take $\rightsquigarrow_i$ to be the BAF under consideration. Then, $\rightsquigarrow_i \supseteq F(\boldsymbol{\rightsquigarrow})$. Furthermore, A is acceptable under a d-preferred extension $\Delta_1 \subseteq Arg$ of $\rightsquigarrow_i$ i.e., $A \in \Delta_1$. Note that Δ_1 is a d-admissible set as well. According to Lemma 11, Δ_1 is a d-admissible set of $F(\boldsymbol{\rightsquigarrow})$ as well. By Fact 14, we know that there is a d-preferred extension $\Delta_2 \subseteq Arg$ of $F(\boldsymbol{\rightsquigarrow})$ such that $\Delta_2 \supseteq \Delta_1$, and A is a member of Δ_2. That is to say, A is acceptable under a d-preferred extension of $F(\boldsymbol{\rightsquigarrow})$. We are done. $\square$