

Preservation of Semantic Properties during the Aggregation of Abstract Argumentation Frameworks

Weiwei Chen

Institute of Logic and Cognition and Department of Philosophy
Sun Yat-sen University, P.R. China
ILLC, University of Amsterdam, The Netherlands
chenww26@mail2.sysu.edu.cn

Ulle Endriss

Institute for Logic, Language and Computation
University of Amsterdam, The Netherlands
ulle.endriss@uva.nl

An abstract argumentation framework can be used to model the argumentative stance of an agent at a high level of abstraction, by indicating for every pair of arguments that is being considered in a debate whether the first attacks the second. When modelling a group of agents engaged in a debate, we may wish to aggregate their individual argumentation frameworks to obtain a single such framework that reflects the consensus of the group. Even when agents disagree on many details, there may well be high-level agreement on important semantic properties, such as the acceptability of a given argument. Using techniques from social choice theory, we analyse under what circumstances such semantic properties agreed upon by the individual agents can be preserved under aggregation.

1 Introduction

Formal argumentation theory provides tools for modelling both the arguments an agent may wish to employ in a debate and the relationships that hold between such arguments [4, 22, 23]. In the widely used system of *abstract argumentation*, introduced in the seminal work of Dung [11], we abstract away from the internal structure of arguments and only model whether or not one argument *attacks* another argument. This is a useful perspective when we require a high-level understanding of how different arguments relate to each other. But when several agents engage in a debate, they may differ on their assessment of some of the arguments and their relationships. How best to model such scenarios of *collective argumentation* is a question of considerable interest. Over the past decade or so, several authors have started to contribute to its resolution (see, e.g., [1, 5, 7, 8, 9, 13, 24, 25]).

Specifically, when agents differ on their assessment of which attacks between the arguments are in fact justified, i.e., when they put forward different *attack-relations*, we may wish to *aggregate* these individual pieces of information to obtain a global view. In this paper, we analyse under what circumstances a given *aggregation rule* will *preserve* relevant properties of the individual attack-relations, particularly properties that relate to the various *semantics* that have been proposed for abstract argumentation. For example, if all agents agree that argument *A* is *acceptable*, either because it is not attacked by any other argument or because it can be successfully defended against any such attack, then we would like *A* to also be considered acceptable relative to the attack-relation returned by our aggregation rule. Thus, *argument acceptability* is an example for a property that, ideally, should be preserved under aggregation.

Our approach uses techniques originating in *social choice theory*, the study of collective decision making [2, 16], the relevance of which to collective argumentation has previously been noted by several of the aforementioned authors, starting with Tohmé et al. [25]. In particular, we make use of recent results on *graph aggregation* [14]. Besides the formulation of a clear and simple model for the axiomatic study of the preservation of semantic properties during aggregation, our contribution consists in delineating

how fundamental axiomatic properties of aggregation rules interact with such preservation requirements. Our technical results range from characterisation results that indicate what kind of aggregation rule can satisfy certain combinations of desiderata, to impossibility results that show that only aggregation rules that are clearly unacceptable from an axiomatic point of view (namely, so-called dictatorships) can preserve the most demanding semantic properties.

This paper is organised as follows. Section 2 is a brief review of relevant concepts from the theory of abstract argumentation. Section 3 introduces our model and Section 4 present our technical results on the preservation of semantic properties under aggregation. Finally, we discuss related work in Section 5 and conclude, in Section 6, with suggestions for possible directions for future work.

2 Abstract Argumentation

In this section, we recall some of the fundamentals of abstract argumentation as originally introduced by Dung [11]. An *argumentation framework* is a pair $AF = \langle Arg, \rightarrow \rangle$, where Arg is a finite set of *arguments* and \rightarrow is an irreflexive binary relation on Arg .¹ If $A \rightarrow B$ holds for two arguments $A, B \in Arg$, we say that A *attacks* B . An *attack-relation* $(\rightarrow) \subseteq Arg \times Arg$ is made up of a set of *attacks* $att \in Arg \times Arg$. For a set of arguments $\Delta \subseteq Arg$ and an argument $B \in Arg$, we say that Δ *attacks* B , denoted as $\Delta \rightarrow B$, if $A \rightarrow B$ holds for some argument $A \in \Delta$. Given an argumentation framework AF , the question arises which arguments to accept. For example, we may not want to accept two arguments that attack each other. A *semantics* for abstract argumentation specifies which sets of arguments can be accepted together. Any such set of arguments is called an *extension* of AF under the semantics in question.

For all the definitions of specific choices of semantics that follow, consider an arbitrary but fixed argumentation framework $AF = \langle Arg, \rightarrow \rangle$ and a set of arguments $\Delta \subseteq Arg$. We say that Δ is *conflict-free*, if there exist no arguments $A, B \in \Delta$ such that $A \rightarrow B$. We further say that Δ *defends* the argument $B \in Arg$, if $\Delta \rightarrow A$ for all arguments $A \in Arg$ such that $A \rightarrow B$. Finally, Δ is called *admissible* if it is conflict-free and defends every single one of its members.

Definition 1. A *stable extension* of AF is a conflict-free set Δ of arguments in Arg that attacks every other argument $B \in Arg \setminus \Delta$.

Definition 2. A *preferred extension* of AF is an admissible set of arguments in Arg that is maximal with respect to set inclusion.

Definition 3. A *complete extension* of AF is an admissible set of arguments in Arg that includes all of the arguments it defends.

The *characteristic function* of AF is the function $f_{AF} : 2^{Arg} \rightarrow 2^{Arg}$ with $f_{AF} : \Delta \mapsto \{A \in Arg \mid \Delta \text{ defends } A\}$, mapping any given set of arguments in Arg to the set of arguments it defends.

Definition 4. The *grounded extension* of AF is the least fixed point of its characteristic function f_{AF} .

We can compute the grounded extension Δ by initialising Δ with the empty set \emptyset and then repeatedly executing the program $\Delta := f_{AF}(\Delta)$, until no more changes occur. Unlike for the other three semantics, there always is exactly one grounded extension. However, that extension may be empty. It is nonempty if and only if there is at least one unattacked argument.

How do these semantics relate to each other? Every stable extension is also a preferred extension. The set of stable extensions may be empty, while there always is at least one preferred extension. Every

¹Neither the finiteness nor the irreflexivity assumption are crucial for our results, but they simplify exposition and clearly are natural for most applications.




Figure 1: Example for a profile with $Arg = \{A, B, C, D\}$.

preferred extension is also a complete extension. Finally, the grounded extension is always a complete extension as well.

An interesting question is under what circumstances two or more semantics coincide. Probably the clearest example is the case of an *acyclic* attack-relation: if \rightarrow does not include any cycles, then the grounded extension is also the only stable extension, the only preferred extension, and the only complete extension. A weaker condition is *coherence*: AF is called *coherent* if every preferred extension of AF is stable, i.e., if the two semantics coincide.

3 The Model

Fix a finite set of arguments Arg . Let $N = \{1, \dots, n\}$ be a finite set of n agents. Suppose each agent $i \in N$ supplies us with an argumentation framework $AF_i = \langle Arg, \rightarrow_i \rangle$, reflecting her individual views on the status of possible attacks between arguments. Thus, we are given a *profile* of attack-relations $\rightarrow = (\rightarrow_1, \dots, \rightarrow_n)$.² What would be a good method of aggregating these individual argumentation frameworks to arrive at a single argumentation framework that appropriately reflects the views of the group as a whole? This is the central question we address in this paper. An *aggregation rule* is a function $F : (Arg \times Arg)^n \rightarrow Arg \times Arg$ mapping any given profile of attack-relations into a single attack-relation.

Example 1. *The first aggregation rule that comes to mind is the majority rule: include attack $A \rightarrow B$ in the outcome if and only if a (weak) majority of the individual agents do. If we apply this rule to the profile shown in Figure 1, we obtain an argumentation framework with $A \rightarrow B$, $B \rightarrow C$, and $C \rightarrow A$. Δ*

In the remainder of this section, we first define a number of specific aggregation rules and review their properties. We focus on simple rules that are adaptations of well-known rules used in the social choice literature, particularly in judgment aggregation [17] and graph aggregation [14]. We then adapt several standard *properties of aggregation rules*, known as *axioms* in that same literature, to our setting. Finally, we review several semantic *properties of argumentation frameworks* and formulate the question of whether a given rule will *preserve* such a property.

3.1 Aggregation Rules

Recall that an aggregation rule is a function F , mapping any given profile $\rightarrow = (\rightarrow_1, \dots, \rightarrow_n) \in (Arg \times Arg)^n$ of attack-relations to a single attack-relation $F(\rightarrow) \subseteq Arg \times Arg$. We sometimes write $(A \rightarrow B) \in F(\rightarrow)$ for $(A, B) \in F(\rightarrow)$. We use $N_{att}^{\rightarrow} := \{i \in N \mid att \in (\rightarrow_i)\}$ to denote the set of *supporters* of the attack att in profile \rightarrow .

²Note that we assume that all agents report an attack-relation over *the same* set of arguments Arg . As argued by Coste-Marquis et al. [8], generalisations, where different agents may be aware of different subsets of Arg , are possible and interesting, but—in line with most existing work in the area—we shall not explore them here.

Definition 5. Let $q \in \{1, \dots, n\}$. The **quota rule** F_q with quota q accepts all those attacks that are supported by at least q agents:

$$F_q(\rightarrow) = \{att \in Arg \times Arg \mid \#N_{att}^{\rightarrow} \geq q\}$$

The *weak majority rule* is the quota rule F_q with $q = \lfloor \frac{n}{2} \rfloor$ and the *strict majority rule* is the quota rule F_q with $q = \lceil \frac{n}{2} \rceil$. Two further quota rules are also of special interest. The *unanimity rule* only accepts attacks that are supported by everyone, i.e., this is F_q with $q = n$. The *nomination rule* is the quota rule F_q with $q = 1$. Despite being a somewhat extreme choice, the nomination rule has some intuitive appeal in the context of argumentation, as it reflects the idea that we should take seriously any conflict between arguments raised by at least one member of the group.

Definition 6. Let $C \in 2^N \setminus \{\emptyset\}$ be a nonempty coalition of agents. The **oligarchic rule** F_C accepts all those attacks that are accepted by all members of C :

$$F_C(\rightarrow) = \{att \in Arg \times Arg \mid C \subseteq N_{att}^{\rightarrow}\}$$

Thus, any member of the oligarchy C can *veto* an attack from being accepted. Observe that the unanimity rule can also be characterised as the oligarchic rule F_C with $C = N$. A subclass of the oligarchic rules are the *dictatorships*. The dictatorship of dictator $i \in N$ is the oligarchic rule F_C with $C = \{i\}$. Thus, under a dictatorship, to compute the outcome, we simply copy the attack-relation of the dictator. Intuitively speaking, oligarchic rules, and dictatorships in particular, are unattractive rules, as they unfairly exclude everyone not in C from the decision process.

Some rules combine features of the quota rules and the oligarchic rules. For example, we may choose to accept an attack only if it is accepted by (i) a weak majority of all agents and (ii) a small number of distinguished agents to which we want to give the right to veto attacks. Such rules (sometimes called *qualified majority rules*) are certainly more attractive than the oligarchic rules, but they are still unfair in the sense of granting some agents more influence than others.

Definition 7. Agent $i \in N$ has **veto powers** under aggregation rule F , if $F(\rightarrow) \subseteq (\rightarrow_i)$ for every profile \rightarrow .

Thus, under an oligarchic rule F_C the agents in C , and only those, have veto powers. With the exception of the unanimity rule, a quota rule does not grant veto powers to any agent.

3.2 Axioms

Next, we introduce some basic *axioms*, encoding intuitively desirable properties of an aggregation rule F :

Definition 8. F is **anonymous**, if $F(\rightarrow) = F(\rightarrow_{\pi(1), \dots, \pi(n)})$ holds for all profiles \rightarrow and all permutations $\pi : N \rightarrow N$.

Definition 9. F is **neutral**, if $N_{att}^{\rightarrow} = N_{att'}^{\rightarrow}$ implies $att \in F(\rightarrow) \Leftrightarrow att' \in F(\rightarrow)$ for all profiles \rightarrow and all attacks att, att' .

Definition 10. F is **independent**, if $N_{att}^{\rightarrow} = N_{att'}^{\rightarrow'}$ implies $att \in F(\rightarrow) \Leftrightarrow att \in F(\rightarrow')$ for all profiles $\rightarrow, \rightarrow'$ and all attacks att .

Definition 11. F is **monotonic**, if $N_{att}^{\rightarrow} \subseteq N_{att'}^{\rightarrow'}$ (together with $N_{att'}^{\rightarrow} = N_{att'}^{\rightarrow'}$ for all attacks $att' \neq att$) implies $att \in F(\rightarrow) \Rightarrow att \in F(\rightarrow')$ for all profiles $\rightarrow, \rightarrow'$ and all attacks att .

Definition 12. F is **unanimous**, if $F(\rightarrow) \supseteq (\rightarrow_1) \cap \dots \cap (\rightarrow_n)$ holds for all profiles $\rightarrow = (\rightarrow_1, \dots, \rightarrow_n)$.

Definition 13. F is **grounded**, if $F(\rightarrow) \subseteq (\rightarrow_1) \cup \dots \cup (\rightarrow_n)$ holds for all profiles $\rightarrow = (\rightarrow_1, \dots, \rightarrow_n)$.

Anonymity is a symmetry (and thus fairness) requirement regarding agents, and neutrality is a symmetry requirement regarding attacks. Independence expresses that whether an attack is accepted should only depend on its supporters. Monotonicity says that additional support for an accepted attack should never cause it to be rejected. Unanimity postulates that an attack supported by everyone must be accepted, while groundedness means that only attacks with at least one supporter can be collectively accepted.³

We observe that all quota rules and all oligarchic rules are easily seen to be unanimous, grounded, neutral, independent, and monotonic. The quota rules furthermore are also anonymous. In fact, it is not difficult to adapt a well-known result from judgment aggregation to our setting [10], so as to see that the quota rules are *the only* aggregation rules that satisfy all of these six axioms (see also [14]).

3.3 Preservation of Semantic Properties

Typically, agents will disagree about which attacks are justified. But even then, there may be high-level agreement on certain features. For example, they may all agree that, under a particular semantics, argument A is acceptable. Whenever we observe such agreement on semantic features in a profile, we would like those features to be preserved under aggregation. Thus, for our example, under the same semantics, we would like A to be acceptable also in the argumentation framework computed by our aggregation rule. In other words, we are interested in the preservation of properties of argumentation frameworks under aggregation. Formally, an *AF-property* $P \subseteq 2^{Arg \times Arg}$ is simply the set of all attack-relations on Arg that satisfy P . But in the interest of readability, we write $P(\rightarrow)$ rather than $(\rightarrow) \in P$.

Definition 14. *Let F be an aggregation rule and let P be an AF-property. We say that F **preserves** P , if for every profile \rightarrow we have that $P(\rightarrow_i)$ being the case for all agents $i \in N$ implies $P(F(\rightarrow))$.*

This notion of preservation is known under the name of *collective rationality* in other parts of social choice theory [3, 14, 20].

We now review the specific AF-properties for which we study preservation in this paper. Two of them we have already introduced in Section 2, namely *acyclicity* and *coherence*. They are attractive properties, because—if satisfied by an argumentation framework—they reduce our dependence on a specific choice of semantics, thereby making decisions less controversial.

Recall that the grounded extension may be empty, i.e., this semantics may not suggest any arguments to be accepted. Thus, argumentation frameworks that satisfy the AF-property of *nonemptiness of the grounded extension* are of particular interest.

Let $A \in Arg$ be one of the arguments under consideration. Then, for any given argumentation framework, A may or may not belong to the grounded extension. Thus, every $A \in Arg$ defines an AF-property, namely the property of membership of A in the grounded extension, i.e., of acceptance of A under the grounded semantics. We say that F *preserves argument acceptability under the grounded semantics*, if it is the case that, for all arguments $A \in Arg$, whenever A belongs to the grounded extension of $\langle Arg, \rightarrow_i \rangle$ for all $i \in N$, then A also belongs to the grounded extension of $\langle Arg, F(\rightarrow) \rangle$. For the stable, preferred, and complete semantics, we require a more refined definition, given that extensions under these semantics need not be unique. We say that F *preserves argument acceptability under the stable semantics*, if it is the case that, for all arguments $A \in Arg$, whenever A belongs to *some* stable extension of $\langle Arg, \rightarrow_i \rangle$ for all agents $i \in N$, then A also belongs to *some* stable extension of $\langle Arg, F(\rightarrow) \rangle$. The corresponding concepts for the preferred and the complete semantics are defined accordingly.⁴

³Note that, in line with the existing literature in argumentation theory on the one hand and social choice theory on the other, we use the term “grounded” in two unrelated ways in this paper (grounded extensions *vs.* grounded aggregation rules).

⁴A further distinction between acceptability in *some* extension and acceptability in *all* extensions would be possible. We leave the investigation of this issue to future work.

Rather than just preserving the acceptability status of a single argument, we may also be interested in preserving entire extensions. For example, we say that F *preserves extensions under the stable semantics*, if it is the case that, for all sets $\Delta \subseteq \text{Arg}$, whenever Δ is a stable extension of $\langle \text{Arg}, \rightarrow_i \rangle$ for all agents $i \in N$, then Δ is also a stable extension of $\langle \text{Arg}, F(\rightarrow) \rangle$. Again, the corresponding concept can also be defined for the other three semantics. Similarly, we say that F *preserves conflict-freeness*, if it is the case that, for all sets $\Delta \subseteq \text{Arg}$, whenever Δ is conflict-free in $\langle \text{Arg}, \rightarrow_i \rangle$ for all agents $i \in N$, then Δ is also conflict-free $\langle \text{Arg}, F(\rightarrow) \rangle$. Finally, *preservation of admissibility* is defined accordingly.

To summarise, we have identified the following AF-properties that, in case all agents agree on one of them being satisfied, we would like to see preserved under aggregation:

- *acyclicity* and *coherence* (reducing semantic ambiguity),
- *nonemptiness of the grounded extension* (enabling a sceptical approach to argument evaluation),
- *argument acceptability* under different semantics (allowing for agreement on arguments even in the face of disagreement on the attacks between them), and
- *being an extension* under one of the four semantics or being either a *conflict-free* or an *admissible set* (also allowing for semantic agreement despite disagreement on attacks).

Example 2. Consider again the profile of Figure 1 and recall that the (weak or strict) majority rule will return the argumentation framework with $A \rightarrow B$, $B \rightarrow C$, and $C \rightarrow A$. Thus, the majority rule does not preserve acyclicity.⁵ What about some of the other AF-properties? The grounded extension of AF_1 is $\{A, C, D\}$, that of AF_2 is $\{B, D\}$, that of AF_3 is $\{A, D\}$, and that of the majority outcome is $\{D\}$. Thus, preservation of both nonemptiness of the grounded extension and argument acceptability under the grounded semantics are not violated by this particular example (which, of course, is not to say that they might not be violated for other profiles). \triangle

4 Preservation Results

In this section, we present our results on the preservation of semantic properties under aggregation. We begin with argument acceptability, and then turn to the various properties of sets of arguments, and eventually to acyclicity and coherence.

4.1 Acceptability of an Argument

Our first result is going to demonstrate that preserving acceptability of an argument when using a “simple” aggregation rule is impossible, unless we are willing to use a dictatorship. This is true under any of the four semantics. To prove this result—and some of those that follow—we are going to use a technique developed by Endriss and Grandi [14] for the more general framework of graph aggregation, which in turn has been inspired by the seminal work on preference aggregation of Arrow [3]. It amounts to showing that, under certain assumptions, the collection of coalitions of agents that are sufficiently powerful to force collective acceptance of an attack must form an ultrafilter.

Using our present terminology, Endriss and Grandi [14, Theorem 18] show that, for $|\text{Arg}| \geq 3$, any unanimous, grounded, neutral, and independent aggregation rule F that preserves some AF-property P must be a dictatorship whenever P belongs to what they call the family of *implicative* and *disjunctive* properties. P is *implicative* if there exist a set $\text{Att} \subseteq \text{Arg} \times \text{Arg}$ of attacks and three attacks $\text{att}_1, \text{att}_2, \text{att}_3 \in \text{Arg} \times \text{Arg} \setminus \text{Att}$ such that, for any $S \subseteq \{\text{att}_1, \text{att}_2, \text{att}_3\}$, we have $P(\text{Att} \cup S)$ if and only if $S \neq \{\text{att}_1, \text{att}_2\}$.

⁵This observation is closely related to the famous *Condorcet Paradox* in the theory of preference aggregation [21].

Thus, P should require that, in the context of Att , accepting att_1 and att_2 implies accepting att_3 (and all seven patterns of acceptance consistent with that requirement are possible). P is *disjunctive* if there exist $Att \subseteq Arg \times Arg$ and $att_1, att_2 \in Arg \times Arg \setminus Att$ such that, for any $S \subseteq \{att_1, att_2\}$, we have $P(Att \cup S)$ if and only if $S \neq \emptyset$. Thus, P should require that, given Att , we must accept at least one of att_1 and att_2 (and all three patterns of acceptance consistent with that requirement are possible).⁶

In the appendix, we show that, under any of the four semantics and for $|Arg| \geq 4$, the property of argument acceptability is both implicative and disjunctive. Thus, we obtain the following impossibility result, showing that there exists no aggregation rule that satisfies all of our requirements:

Theorem 1. *Let P be the property of argument acceptability under either the grounded, the stable, the preferred, or the complete semantics. For $|Arg| \geq 4$, any unanimous, grounded, neutral, and independent aggregation rule F that preserves P must be a dictatorship.*

This is bad news. For scenarios with $|Arg| \leq 3$, we can do better and show that also the nomination rule preserves argument acceptability. In the interest of space, we omit the details and also do not discuss such boundary conditions for our remaining results.

4.2 Conflict-Freeness and Admissibility

Argument acceptability is a property that relates to a *single* argument. Next, we turn to properties that relate to *sets* of arguments. Our most basic property of sets of arguments is preserved under essentially all reasonable aggregation rules:

Theorem 2. *Every aggregation rule F that is grounded preserves conflict-freeness.*

Proof. Let F be an aggregation rule that is *grounded*. Consider any set $\Delta \subseteq Arg$ and any profile $\rightarrow = (\rightarrow_1, \dots, \rightarrow_n)$ such that Δ is conflict-free in $\langle Arg, \rightarrow_i \rangle$ for all $i \in N$. For the sake of contradiction, assume Δ is *not* conflict-free in $\langle Arg, F(\rightarrow) \rangle$, i.e., there exist two arguments $A, B \in \Delta$ such that $(A \rightarrow B) \in F(\rightarrow)$. Due to the groundedness of F , there then must be at least one agent $i \in N$ such that also $A \rightarrow_i B$, i.e., Δ is not conflict-free in $\langle Arg, \rightarrow_i \rangle$ either, in contradiction to our original assumption. \square

For admissibility, we obtain a significantly less broad but still positive result. It shows that there exists a reasonable rule that preserves the admissibility of arbitrary sets of arguments:

Theorem 3. *For $|Arg| \geq 4$, the only unanimous, grounded, anonymous, neutral, independent, and monotonic aggregation rule F that preserves admissibility is the nomination rule.*

Proof. We first show that the *nomination rule* indeed preserves *admissibility*. So let F be the nomination rule. Consider any set $\Delta \subseteq Arg$ and any profile $\rightarrow = (\rightarrow_1, \dots, \rightarrow_n)$ such that Δ is admissible in $AF_i = \langle Arg, \rightarrow_i \rangle$ for all $i \in N$. For the sake of contradiction, assume Δ is *not* admissible in $\langle Arg, F(\rightarrow) \rangle$, i.e., there is an argument $A \in \Delta$ that, in $F(\rightarrow)$, is attacked by an argument $B \in Arg \setminus \Delta$ and there does not exist a $C \in \Delta$ such that $(C \rightarrow B) \in F(\rightarrow)$. As $(B \rightarrow A) \in F(\rightarrow)$ and as F is grounded, we must have $B \rightarrow_i A$ for some $i \in N$. And as there does not exist a $C \in \Delta$ such that $(C \rightarrow A) \in F(\rightarrow)$, given the definition of the nomination rule, there cannot exist an argument $C \in \Delta$ such that $C \rightarrow_i A$ for that same agent i . Hence, Δ is not admissible in AF_i , in contradiction to our original assumption.

We still need to show that there can be no other aggregation rule than the nomination rule that preserves admissibility and that satisfies all of the axioms mentioned in the statement of Theorem 3.

⁶Our definitions of implicativeness and disjunctiveness are special cases of the more general definitions given by Endriss and Grandi [14]. They simplify exposition and are sufficient for our purposes here.




Figure 2: Profile used in the proof of Theorem 3.

By the characterisation result for quota rules due to Dietrich and List [10] in the context of judgment aggregation, which has been adapted to graph aggregation by Endriss and Grandi [14] and which we have briefly recalled near the end of Section 3.1, this claim is equivalent to the claim that *no quota rule* F_q with a quota $q > 1$ always preserves admissibility. So let us prove this.

Consider the generic profile shown in Figure 2 (and note that $q > 1$ ensures $q - 1 > 0$, i.e., there is at least one agent of the first kind). The set $\{A, B, C\}$ is admissible in all argumentation frameworks in such a profile. But when we aggregate using a quota rule F_q with a quota $q > 1$, we obtain an argumentation framework with a single attack $D \rightarrow A$, which means that A cannot be part of any admissible set. Hence, no such rule can preserve admissibility. \square

4.3 The Property of Being an Extension

We continue our examination of properties relating to sets of arguments and turn to the property of being an extension of a given argumentation framework. The following result can again be proved by reference to the result of Endriss and Grandi [14]:

Theorem 4. *For $|Arg| \geq 5$, any unanimous, grounded, neutral, and independent aggregation rule F that preserves grounded extensions must be a dictatorship.*

The required proof of the fact that all three properties are both implicative and disjunctive can be found in the appendix. We conjecture that Theorem 4 can be extended to the preferred and the complete semantics. The added difficulty lies in the fact that these semantics admit multiple extensions. Interestingly, for stable extensions we obtain a much more positive result:

Proposition 5. *The nomination rule preserves stable extensions.*

Proof. Let F be the *nomination rule*. Consider any set $\Delta \subseteq Arg$ and any profile $\rightarrow = (\rightarrow_1, \dots, \rightarrow_n)$ such that Δ is stable in $\langle Arg, \rightarrow_i \rangle$ for all $i \in N$. According to Theorem 2, given that F is grounded, F preserves conflict-freeness. Thus, Δ is conflict-free in $\langle Arg, F(\rightarrow) \rangle$.

What remains to be shown is that Δ attacks every argument $B \in Arg \setminus \Delta$. In case $\Delta = Arg$, the claim holds vacuously. Otherwise, consider an arbitrary argument $B \in Arg \setminus \Delta$. We need to show that B is attacked by some argument in Δ in $F(\rightarrow)$. Take the argumentation framework $AF_i = \langle Arg, \rightarrow_i \rangle$ for some $i \in N$. As Δ is stable in AF_i by assumption, there exists an argument $A \in \Delta$ such that $A \rightarrow_i B$. As F is the nomination rule, we also get $(A \rightarrow B) \in F(\rightarrow)$ as claimed. \square

4.4 Nonemptiness of the Grounded Extension

We have seen that preserving grounded extensions is impossible for reasonable aggregation rules (see Theorem 4). What about the seemingly less demanding requirement of at least preserving nonemptiness

of the grounded extension? The good news is that we can do better for this property. For instance, it is easy to check that the unanimity rule preserves nonemptiness of the grounded extension. Still, as we shall see next, we cannot do *much* better: only rules that grant veto powers to some agents will work.

Recall that the grounded extension is nonempty if and only if at least one argument is not attacked by any other argument. Thus, this AF-property is about the *absence* of attacks, while the technique we employed to prove Theorem 4 (and Theorem 1) exploits the *presence* of certain attacks. We are now going to present our preservation result regarding the nonemptiness of the grounded extension as a corollary to a more general theorem about the preservation of AF-properties that require the absence of certain attacks. Let $k \in \mathbb{N}$. Let us call an AF-property P *k-exclusive* if there exist k distinct attacks $att_1, \dots, att_k \in Arg \times Arg$ such that (i) $\{att_1, \dots, att_k\} \subseteq (\rightarrow)$ for no attack-relation \rightarrow with $P(\rightarrow)$, and (ii) for every $S \subsetneq \{att_1, \dots, att_k\}$ there exists an attack-relation \rightarrow such that $S \subseteq (\rightarrow)$ and $P(\rightarrow)$. Thus, you cannot accept all k attacks, but you should be able to accept any proper subset of them. We are able to prove the following powerful theorem (recall that n is the number of agents in N):

Theorem 6. *Let $k \geq n$ and let P be an AF-property that is k -exclusive. Then under any neutral and independent aggregation rule F that preserves P at least one agent must have veto powers.*

Proof. Let $k \geq n$, let P be an AF-property that is k -exclusive, and let F be an aggregation rule that is neutral and independent. We need to show that, if F preserves P , then F must give some agents the power to veto the collective acceptance of attacks.

First, observe that, if an aggregation rule F is neutral and independent, then we can describe F by listing all the coalitions $C \subseteq N$ for which it is the case that, whenever exactly the agents in C support an attack att , then att must be collectively accepted. Indeed, independence says that acceptance of an attack should only depend on its supporters, and neutrality adds that this dependence must be the same for all attacks. More formally, there exists a family of *winning coalitions* $\mathcal{W} \subseteq 2^N$ such that, for all profiles \rightarrow and all potential attacks $att \in Arg \times Arg$, the following relationship holds:

$$att \in F(\rightarrow) \quad \text{if and only if} \quad N_{att}^{\rightarrow} \in \mathcal{W}$$

Recall that $i \in N$ having veto powers under F means that $F(\rightarrow) \subseteq (\rightarrow_i)$ for every profile \rightarrow . We can now show that an agent $i \in N$ has veto powers, if she is a member of all winning coalitions:

$$i \in \bigcap_{C \in \mathcal{W}} C \quad \text{implies} \quad F(\rightarrow) \subseteq (\rightarrow_i) \quad \text{for all profiles } \rightarrow$$

If $\bigcap_{C \in \mathcal{W}} C = \emptyset$, then the above claim holds vacuously. Otherwise, take any attack $att \in F(\rightarrow)$. As att got accepted, N_{att}^{\rightarrow} must be a winning coalition, i.e., $N_{att}^{\rightarrow} \in \mathcal{W}$ and therefore $i \in N_{att}^{\rightarrow}$. But this is just another way of saying $att \in (\rightarrow_i)$, so we are done.

Next, we are going to show that the fact that F preserves the k -exclusive AF-property P implies that the intersection of any k winning coalitions must be nonempty:

$$C_1 \cap \dots \cap C_k \neq \emptyset \quad \text{for all } C_1, \dots, C_k \in \mathcal{W}$$

For the sake of contradiction, assume there do exist winning coalitions $C_1, \dots, C_k \in \mathcal{W}$ such that $C_1 \cap \dots \cap C_k = \emptyset$. We construct a profile $\rightarrow = (\rightarrow_1, \dots, \rightarrow_n)$ with $P(\rightarrow_i)$ for all $i \in N$ as follows: for every $j \in \{1, \dots, k\}$, exactly the agents in C_j accept attack att_j (for all other attacks, it is irrelevant which agents accept them). As no agent is a member of all k winning coalitions, no agent accepts all k attacks, so this construction indeed is possible. However, as each of the k attacks is supported by a winning coalition, they all get accepted, i.e., $\{att_1, \dots, att_k\} \subseteq F(\rightarrow)$, meaning that the outcome does *not* satisfy P . Thus, we have found a contradiction to our assumption of F preserving P and are done.

Let us briefly recap where we are at this point. We know that F is characterised by a family of winning coalitions \mathscr{W} . We also know that $C_1 \cap \dots \cap C_k \neq \emptyset$ for all $C_1, \dots, C_k \in \mathscr{W}$. We need to show that some agents have veto powers, and we know that this is the case if we can prove that $C^{(1)} \cap \dots \cap C^{(\ell)} \neq \emptyset$, where $\{C^{(1)}, \dots, C^{(\ell)}\}$ is some enumeration of the coalitions in \mathscr{W} . Thus, we are done, if we can show that $C_1 \cap \dots \cap C_k \neq \emptyset$ for *all* $C_1, \dots, C_k \in \mathscr{W}$ implies $C^{(1)} \cap \dots \cap C^{(\ell)} \neq \emptyset$. We are going to prove the contrapositive, namely that the following holds for *some* $C_1, \dots, C_k \in \mathscr{W}$:

$$C^{(1)} \cap \dots \cap C^{(\ell)} = \emptyset \quad \text{implies} \quad C_1 \cap \dots \cap C_k = \emptyset$$

In words, we need to show that in case the intersection of *all* winning coalitions is empty, then so is at least one intersection of just k winning coalitions.

Recall that we have assumed $k \geq n$. We construct a set $\mathscr{W}' \subseteq \mathscr{W}$ of k (or fewer) winning coalitions as follows. Initially, set $\mathscr{W}' := \emptyset$. Then, for every j from 1 to ℓ in turn, add $C^{(j)}$ to \mathscr{W}' if and only if the following condition is satisfied:⁷

$$\left(C^{(j)} \cap \bigcap_{C \in \mathscr{W}'} C \right) \not\subseteq \left(\bigcap_{C \in \mathscr{W}'} C \right)$$

Thus, every additional $C^{(j)}$ is selected only if it causes the removal of at least one further agent from the intersection. As there are only n agents, we therefore will pick at most n coalitions. Hence, we will indeed arrive at a family \mathscr{W}' of n or fewer—and thus certainly at most k —winning coalitions, the intersection of which is empty. This completes the proof. \square

It now suffices to show that the property of having a nonempty grounded extension is an $|\text{Arg}|$ -exclusive property to obtain the following result (the proof of this fact can be found in the appendix):

Theorem 7. *If $|\text{Arg}| \geq n$, then under any neutral and independent aggregation rule F that preserves nonemptiness of the grounded extension at least one agent must have veto powers.*

It is not difficult to prove that the converse holds as well: all rules that grant veto powers to at least one agent preserve nonemptiness of the grounded extension.

4.5 Acyclicity and Coherence

Finally, if we apply our techniques to the properties of acyclicity and coherence, we obtain the following results (the proofs of which can be found in the appendix):⁸

Theorem 8. *If $|\text{Arg}| \geq n$, then under any neutral and independent aggregation rule F that preserves acyclicity at least one agent must have veto powers.*

Theorem 9. *For $|\text{Arg}| \geq 4$, any unanimous, grounded, neutral, and independent aggregation rule F that preserves coherence must be a dictatorship.*

Thus, surprisingly, even though acyclicity is a stronger property than coherence, it is easier to preserve under aggregation.

⁷By convention, let $\bigcap_{C \in \emptyset} C = N$, i.e., the intersection of *no* winning coalitions is defined as the universe N of *all* agents.

⁸Theorem 8 was anticipated in the work of Tohmé et al. [25], who make a similar claim, but without appealing to the neutrality axiom. We stress that Theorem 8 cannot be strengthened by dropping neutrality. Indeed, there are rules that preserve acyclicity, that are independent (but not neutral), and that do not give veto powers to any of the agents. An example, for $N = \{1, 2\}$ and $\text{Arg} = \{A, B\}$, is the rule that accepts $A \rightarrow B$ if at least one agent does and that accepts $B \rightarrow A$ if both agents do.

5 Related Work

In this section, we briefly review related work on the aggregation of abstract argumentation frameworks. Coste-Marquis et al. [8] were the first to address this problem, but without making explicit reference to social choice theory. Instead, they focus on aggregation rules that minimise the *distance* between the input argumentation frameworks and the output argumentation framework.

Tohmé et al. [25] were the first to explicitly use social choice theory to analyse the aggregation of argumentation frameworks. Their focus is on the preservation of acyclicity and they show that qualified majority rules (which grant veto powers to some agents) will always preserve this property.

Dunne et al. [13] define several preservation requirements on aggregation rules that directly refer to the semantics of the argumentation frameworks concerned. This includes variants of what we call preservation of extensions (“ σ -unanimity”) and preservation of argument acceptability (“credulous acceptance unanimity”). Their focus is on analysing the computational complexity of deciding whether a given aggregation rule has a given property, rather than on the axiomatic method. In follow-up work, Delobelle et al. [9] establish for several concrete rules whether or not they satisfy the preservation requirements introduced by Dunne et al. [13].

While Endriss and Grandi [14] explicitly mention abstract argumentation as a possible domain of application for the model of graph aggregation they develop, they do not present any technical results related to argumentation.

Airiau et al. [1] introduce the concept of the *rationalisability* of a profile of argumentation frameworks. A profile is rationalisable if the diversity of views it contains can be explained in terms of (i) an underlying factual argumentation framework shared by all agents and (ii) everyone’s individual preferences. Thus, their work is concerned with understanding what kind of profiles a good aggregation rule should be able to deal with, rather than with aggregation itself.

Finally, social choice theory has also been used to analyse the aggregation of different *extensions* for *one* argumentation framework see, (e.g., [7, 24]). We note that this problem is different from the one studied here and refer to the survey by Bodanza et al. [5] for a comparison.

6 Conclusion

Using a variety of techniques, we have attempted to paint a clear picture of the capabilities and limitations of simple aggregation rules regarding the preservation of properties related to the semantics of abstract argumentation frameworks. While the significance of this issue and the promise of social choice theory for its resolution have previously been emphasised in the work of several authors [5, 9, 13, 25], this is the first systematic analysis of its kind. Our results show that only the most basic of properties, namely conflict-freeness, is preserved by essentially all rules. More demanding properties require either the nomination rule, a rule granting some agents veto powers, or a rule that is dictatorial.

We stress that these results only apply to simple rules, in particular, to rules that satisfy the axiom of independence. An alternative route, the one chosen by Coste-Marquis et al. [8], is to use distance-based rules (which violate independence). Such rules can be designed so as to guarantee specific properties of the outcome, so the question of preservation does not arise. On the downside, distance-based rules are computationally intractable [15, 18, 19]. We also stress that our results are based on the assumption that all agents report attack-relations over a single *common* set of arguments. Richer models, where different agents may be aware of different sets of arguments, are clearly of great interest as well.

There are multiple directions in which to extend this work. First, our conjecture regarding the preser-

vation of preferred and complete extensions should get settled. Second, one could study further properties of argumentation frameworks. We have already hinted at the possibility of distinguishing between argument acceptability in *some* extensions (the property studied in this paper) and argument acceptability in *all* extensions. Third, we should eventually go beyond the four classical semantics introduced by Dung [11] and also consider others, such as the *semi-stable semantics* [6] or the *ideal semantics* [12]. Fourth, one could vary the axioms imposed on aggregation rules. The most immediately promising direction here would be to investigate whether neutrality can be replaced by additional preservation requirements, in analogy to results in preference and graph aggregation [3, 14]. Finally, it would be interesting to investigate the strategic incentives of agents who are reporting an argumentation framework to an aggregation rule and whose objective might be to get a certain argument accepted.

Acknowledgements. We would like to thank Sirin Botan, Umberto Grandi, Ronald de Haan, and Zoi Terzopoulou for numerous enlightening discussions on the material presented in this paper and three anonymous reviewers for their constructive feedback.

References

- [1] S. Airiau, E. Bonzon, U. Endriss, N. Maudet & J. Rossit (2016): *Rationalisation of Profiles of Abstract Argumentation Frameworks*. In: *Proceedings of the 15th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*.
- [2] K.J. Arrow, A.K. Sen & K. Suzumura, editors (2002): *Handbook of Social Choice and Welfare*. North-Holland.
- [3] K.J. Arrow (1963): *Social Choice and Individual Values*, 2nd edition. John Wiley and Sons. First edition published in 1951.
- [4] P. Besnard & A. Hunter (2008): *Elements of Argumentation*. MIT Press, doi:10.7551/mitpress/9780262026437.001.0001.
- [5] G.M. Bodanza, F.A. Tohmé & M.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.
- [6] M. Caminada (2006): *Semi-stable Semantics*. In: *Proceedings of the 1st International Conference on Computational Model of Argument (COMMA)*, IOS Press.
- [7] M. Caminada & G. Pigozzi (2011): *On Judgment Aggregation in Abstract Argumentation*. *Journal of Autonomous Agents and Multiagent Systems* 22(1), pp. 64–102, doi:10.1007/s10458-009-9116-7.
- [8] S. Coste-Marquis, C. Devred, S. Konieczny, M.-C. Lagasque-Schiex & P. 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.
- [9] J. Delobelle, S. Konieczny & S. Vesic (2015): *On the Aggregation of Argumentation Frameworks*. In: *Proceedings of the 24th International Joint Conference on Artificial Intelligence (IJCAI)*.
- [10] F. Dietrich & C. List (2007): *Judgment Aggregation by Quota Rules: Majority Voting Generalized*. *Journal of Theoretical Politics* 19(4), pp. 391–424, doi:10.1177/0951629807080775.
- [11] P.M. 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] P.M. Dung, P. Mancarella & F. Toni (2007): *Computing Ideal Sceptical Argumentation*. *Artificial Intelligence* 171(10), pp. 642–674, doi:10.1016/j.artint.2007.05.003.

- [13] P.E. Dunne, P. Marquis & M. 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.
- [14] U. Endriss & U. Grandi (2017): *Graph Aggregation*. *Artificial Intelligence* 245, pp. 86–114, doi:10.1016/j.artint.2017.01.001.
- [15] U. Endriss, U. Grandi & D. Porello (2012): *Complexity of Judgment Aggregation*. *Journal of Artificial Intelligence Research (JAIR)* 45, pp. 481–514, doi:10.1613/jair.3708.
- [16] W. Gaertner (2006): *A Primer in Social Choice Theory*. LSE Perspectives in Economic Analysis, Oxford University Press, doi:10.1007/s00355-007-0273-9.
- [17] D. Grossi & G. Pigozzi (2014): *Judgment Aggregation: A Primer*. Synthesis Lectures on Artificial Intelligence and Machine Learning, Morgan & Claypool Publishers, doi:10.2200/S00559ED1V01Y201312AIM027.
- [18] E. Hemaspaandra, H. Spakowski & J. Vogel (2005): *The Complexity of Kemeny Elections*. *Theoretical Computer Science* 349(3), pp. 382–391, doi:10.1016/j.tcs.2005.08.031.
- [19] S. Konieczny, J. Lang & P. Marquis (2004): *DA² Merging Operators*. *Artificial Intelligence* 157(1-2), pp. 49–79, doi:10.1016/j.artint.2004.04.008.
- [20] C. List & P. Pettit (2002): *Aggregating Sets of Judgments: An Impossibility Result*. *Economics and Philosophy* 18(1), pp. 89–110, doi:10.1017/S0266267102001098.
- [21] I. McLean & A.B. Urken, editors (1995): *Classics of Social Choice*. University of Michigan Press, doi:10.3998/mpub.12736.
- [22] S. Modgil, N. Oren & F. Toni, editors (2012): *Proceedings of the 1st International Workshop on Theory and Applications of Formal Argumentation (TAFa)*. Springer-Verlag, doi:10.1007/978-3-642-29184-5.
- [23] I. Rahwan & G.R. Simari (2009): *Argumentation in Artificial Intelligence*. Springer-Verlag, doi:10.1007/978-0-387-98197-0.
- [24] I. Rahwan & F.A. Tohmé (2010): *Collective Argument Evaluation as Judgement Aggregation*. In: *Proceedings of the 9th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*.
- [25] F.A. Tohmé, G.A. Bodanza & G.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.

Appendix: Remaining Proofs

In this appendix, we present the proofs omitted from the body of the paper. All of the proofs in this appendix have the same structure: they show that a given semantic AF-property of interest has certain meta-properties, for which a general (impossibility) result is available.

Proof of Theorem 1

Recall that, for each of the four semantics, we have to show that, for $|Arg| \geq 4$, the AF-property of *argument acceptability* is both *implicative* and *disjunctive*. We are able to use the same construction in all four cases. Let P be the AF-property of argument acceptability under either the grounded, the stable, the preferred, or the complete semantics. Furthermore, let $Arg = \{A, B, C, D, \dots\}$ be a set of at least four arguments.

Let us first show that P is implicative. Let C be the argument under consideration. Let $Att = \{D \rightarrow B\}$, $att_1 = B \rightarrow C$, $att_2 = C \rightarrow D$, and $att_3 = A \rightarrow B$. This scenario is sketched in the lefthand part of Figure 3. Now consider the argumentation frameworks of the form $\langle Arg, Att \cup S \rangle$ with $S \subseteq \{att_1, att_2, att_3\}$. If




Figure 3: Scenarios used in the proof of Theorem 1.

$S \subseteq \{att_2, att_3\}$, then C is not attacked by any other argument. If $S = \{att_1\}$ or $S = \{att_1, att_3\}$, then C is defended by D , which is not attacked by any other argument. If $S = \{att_1, att_2, att_3\}$, then C is defended by A , which is not attacked by any other argument. Thus, in all of these seven cases, either C is not attacked by any other argument or it is defended by an argument that is not attacked by any other argument, i.e., in all cases C is acceptable under the grounded, the stable, the preferred, and the complete semantics. On the other hand, if $S = \{att_1, att_2\}$, then $\{B, C, D\}$ forms an isolated odd-length cycle. This means that all of B , C , and D will be unacceptable under the grounded, the stable, the preferred, and the complete semantics. We have thus found a set of attacks Att and three individual attacks att_1 , att_2 , att_3 such that $P(Att \cup S)$ if and only if $S \neq \{att_1, att_2\}$. Hence, P is an implicative AF-property.

Next, we show that P is also disjunctive. Let D be the argument under consideration. Let $Att = \{C \rightarrow D\}$, $att_1 = A \rightarrow C$, and $att_2 = B \rightarrow C$. This scenario is depicted on the righthand side of Figure 3. Consider all argumentation frameworks $\langle Arg, Att \cup S \rangle$ with $S \subseteq \{att_1, att_2\}$. If $S = \{att_1\}$, then D is defended by A . If $S = \{att_2\}$, then D is defended by B . If $S = \{att_1, att_2\}$, then D is defended by both A and B . In all three cases, D is defended by some argument that is not attacked by any other argument. Thus, D is acceptable under each of our four semantics. However, if $S = \emptyset$, then D is attacked by C and not defended by any other argument, which means that D is unacceptable under all four semantics. To summarise, we have seen that $P(Att \cup S)$ if and only if $S \neq \emptyset$. Hence, P is a disjunctive AF-property. \square

Proof of Theorem 4

Recall that we need to show that, for $|Arg| \geq 5$, the property P of being a grounded extension is an AF-property that is both *implicative* and *disjunctive*. In fact, as we shall see, we can show disjunctiveness even in case $|Arg| \geq 4$. Let $Arg = \{A, B, C, D, E, \dots\}$. For the proofs of both properties, we focus on $\Delta = \{A, C, E\}$ as the subset of arguments that may (or may not) form the grounded extension.

We first show that P is implicative. Let $Att = \{A \rightarrow B, C \rightarrow D\}$, $att_1 = B \rightarrow C$, $att_2 = D \rightarrow A$, and $att_3 = E \rightarrow D$. This scenario is depicted on the lefthand side of Figure 4. Consider all argumentation frameworks of the form $AF = \langle Arg, Att \cup S \rangle$ with $S \subseteq \{att_1, att_2, att_3\}$, and for each of them the corresponding characteristic function f_{AF} . If $S = \{att_1\}$, then $f_{AF}(\emptyset) = \{A, E\}$, $f_{AF}^2(\emptyset) = \{A, C, E\} = \Delta$, and $f_{AF}^3(\emptyset) = f_{AF}^2(\emptyset)$. Thus, the grounded extension is Δ in this case. Using the same kind of reasoning, it is easy to verify that the grounded extension is Δ whenever $S \neq \{att_1, att_2\}$. On the other hand, if $S = \{att_1, att_2\}$, then we get $f_{AF}(\emptyset) = \{E\}$ and $f_{AF}^2(\emptyset) = f_{AF}(\emptyset)$, i.e., the grounded extension now is $\{E\}$. Thus, for our argumentation framework to have Δ as its grounded extension, we must insist that, if both att_1 and att_2 are accepted, then also att_3 is accepted. Hence, P is implicative.

Next, we show that P is also a disjunctive AF-property. Let $Att = \{A \rightarrow B, B \rightarrow C, C \rightarrow D, D \rightarrow A\}$, $att_1 = E \rightarrow B$, and $att_2 = E \rightarrow D$. This is shown on the righthand side of Figure 4. Consider $AF = \langle Arg, Att \cup S \rangle$ with $S \subseteq \{att_1, att_2\}$. If $S = \{att_1\}$, then $f_{AF}(\emptyset) = \{E\}$, $f_{AF}^2(\emptyset) = \{C, E\}$, $f_{AF}^3(\emptyset) =$




Figure 4: Scenarios used in the proof of Theorem 4.

$\{A, C, E\} = \Delta$, and $f_{AF}^4(\emptyset) = f_{AF}^3(\emptyset)$, i.e., Δ is the grounded extension of AF . By analogous reasoning, Δ is the grounded extension also for $S = \{att_2\}$ and for $S = \{att_1, att_2\}$. However, for $S = \emptyset$, we get $f_{AF}(\emptyset) = \{E\}$ and $f_{AF}^2(\emptyset) = f_{AF}(\emptyset)$, i.e., the grounded extension now is just $\{E\}$. Thus, Δ is the grounded extension if and only if $S \neq \emptyset$, meaning that att_1 or att_2 are accepted. Hence, P is disjunctive. \square

Proof of Theorem 7

To obtain the claim as a corollary to Theorem 6, we need to show that the property of an argumentation framework *having a nonempty grounded extension* is a k -exclusive AF-property for $k = |Arg|$. Recall that having a nonempty grounded extension is equivalent to the property of *having at least one argument that is not attacked* by any other argument. We are going to show that the latter property is k -exclusive for $k = |Arg|$.

So let $k = |Arg|$. Take an arbitrary enumeration $\{A^{(1)}, \dots, A^{(k)}\}$ of Arg and consider the set of attacks $\{att_1, \dots, att_k\}$ with $att_i := A^{(i)} \rightarrow A^{(i+1)}$ for $i < k$ and $att_k := A^{(k)} \rightarrow A^{(1)}$. Clearly, this set of attacks meets our requirements: (i) if $\{att_1, \dots, att_k\} \subseteq (\rightarrow)$, then \rightarrow does not have the property of leaving at least one argument without an attacker and (ii) for every $S \subsetneq \{att_1, \dots, att_k\}$ there exists an attack-relation \rightarrow with $S \subseteq (\rightarrow)$, namely S itself, that does leave one argument without an attacker. \square

Proof of Theorem 8

Recall that, if we can show that *acyclicity* is an $|Arg|$ -exclusive AF-property, then the claim follows from Theorem 6. In fact, it is straightforward to show that acyclicity is a k -exclusive property for every $k \in \{2, \dots, |Arg|\}$. To see this, consider the case where the attacks $\{att_1, \dots, att_k\}$ form a cycle, and observe that the shortest (proper) cycle has length 2, while the longest cycle visits every argument exactly once and thus has length $|Arg|$. \square

Proof of Theorem 9

Recall that we need to show that, for $|Arg| \geq 4$, *coherence* is an AF-property that is both *implicative* and *disjunctive*. Let P represent coherence and let $Arg = \{A, B, C, D, \dots\}$.

Let us first show that P is an implicative AF-property. Let $Att = \{D \rightarrow B\}$, $att_1 = B \rightarrow C$, $att_2 = C \rightarrow D$, and $att_3 = A \rightarrow B$. This scenario is shown on the lefthand side of Figure 5 and is identical to the scenario used in the proof of Theorem 1. Now consider argumentation frameworks $\langle Arg, Att \cup S \rangle$ with $S \subseteq \{att_1, att_2, att_3\}$. If either $S = \{att_1\}$, $S = \{att_3\}$, $S = \{att_1, att_3\}$ or $S = \emptyset$, the only preferred extension is $\{A, C, D\}$, which is also stable. If $S = \{att_2\}$, the only preferred extension is $\{A, B, C\}$, which is also stable. If $S = \{att_2, att_3\}$ or $S = \{att_1, att_2, att_3\}$, the only preferred extension is $\{A, C\}$




Figure 5: Scenarios used in the proof of Theorem 9.

which once again also is stable. Thus, in all seven cases, AF is coherent. However, if $S = \{att_1, att_2\}$, the only preferred extension is $\{A\}$, which is not stable, as B , C , and D are not attacked by A . We have thus found a set of attacks Att and three individual attacks att_1 , att_2 , att_3 such that $P(Att \cup S)$ if and only if $S \neq \{att_1, att_2\}$. In other words, P is an implicative AF-property.

Next, we show that P is also a disjunctive AF-property. Let $Att = \{B \rhd C, C \rhd D, D \rhd B\}$, $att_1 = A \rhd B$, and $att_2 = A \rhd D$. This scenario is shown on the righthand side of Figure 5. Consider argumentation frameworks $\langle Arg, Att \cup S \rangle$ with $S \subseteq \{att_1, att_2\}$. If $S = \{att_1\}$ or $S = \{att_1, att_2\}$, the only preferred extension is $\{A, C\}$, which is also stable. If $S = \{att_2\}$, the only preferred extension is $\{A, B\}$, which again is also stable. Thus, in all three cases every preferred extension is stable, i.e., AF is coherent. On the other hand, if $S = \emptyset$, then the only preferred extension is $\{A\}$, which is not stable, as B , C , and D not attacked by A . To summarise, $P(Att \cup S)$ if and only if $S \neq \emptyset$. Hence, P is disjunctive. \square