

Aggregating Alternative Extensions of Abstract Argumentation Frameworks: Preservation Results for Quota Rules

Weiwei Chen ^{a,b}, Ulle Endriss ^b

^a*Institute of Logic and Cognition and Department of Philosophy
Sun Yat-sen University, China*

^b*ILLC, University of Amsterdam, The Netherlands*

Abstract. When confronted with the same abstract argumentation framework, specifying a set of arguments and an attack-relation between them, different agents may disagree on which arguments to accept, i.e., they may choose different extensions. In the context of designing systems to support collective argumentation, we may then wish to aggregate such alternative extensions into a single extension that appropriately reflects the views of the group as a whole. Focusing on a conceptually and computationally simple family of aggregation rules, the quota rules, we analyse under what circumstances relevant properties of extensions shared by all extensions reported by the individual agents will be preserved under aggregation. The properties we consider are the classical properties of argumentation semantics, such as being a conflict-free, a complete, or a preferred extension. We show that, while for some properties there are quota rules that guarantee their preservation, for the more demanding properties it is impossible to do so in general.

Keywords. Abstract argumentation, collective argumentation, social choice theory.

1. Introduction

The study of *collective argumentation* deals with questions that arise when we need to reason about a group of agents who each take an individual view on the merits of a number of arguments [5]. Relevant application scenarios include the moderation of online discussion fora and the design of support systems for collective editing efforts (e.g., for Wikipedia). Collective argumentation can (and should) be studied for any number of formalisms for modelling argumentation. Here we do so in the context of *abstract argumentation* [14], a particularly simple and mathematically elegant formalism, as this allows us to focus on fundamental principles. An abstract argumentation framework is simply a set of arguments with a binary attack-relation defined on this set. When presented with such an argumentation framework, each agent will choose her own *extension*—the subset of arguments she accepts. We can then apply an *aggregation rule*, e.g., the *majority rule* that accepts an argument if more than half of the agents do, to obtain a single collective extension that, hopefully, represents a good compromise between the individual extensions reported. The question we ask in this paper is whether certain high-level proper-

ties of extensions that all individual agents agree on will be *preserved* under aggregation. For example, if all agents report extensions that are *conflict-free*, will the collective extensions returned by the majority rule be conflict-free as well?

In studying such questions, we are taking the perspective of *social choice theory*. Specifically, we build on known results in *judgment aggregation* [19], a branch of social choice theory that deals with the aggregation of individual points of view in the presence of integrity constraints expressed in propositional logic. Regarding aggregation rules, we focus on *quota rules*, which have been studied in depth in judgment aggregation [13]. Regarding properties to be preserved, we focus on the classical properties of argumentation semantics [14], such as conflict-freeness, admissibility, or stability. To be able to analyse the behaviour of these properties in the context of judgment aggregation, we exploit known encodings of argumentation semantics in propositional logic [4].

Our results establish under what circumstances quota rules guarantee the preservation of fundamental semantic properties of extensions. This includes both *possibility* and *impossibility results*. For instance, the majority rule always preserves conflict-freeness, while no quota rule can guarantee the preservation of either admissibility or stability unless we impose demanding restrictions on the argumentation framework in question.

The paper is organised as follows. Section 2 reviews basic notions of abstract argumentation and their logical encoding. Then, Section 3 formally defines the problem of preservation we study and Section 4 presents our results. While this is the first systematic study of the preservation of semantic properties of extensions by quota rules, there is significant prior work on related questions that combines ideas from abstract argumentation and social choice theory, which we briefly review in Section 5. Section 6 concludes.

2. Abstract Argumentation

In this section, we recall some of the basic concepts of the theory of abstract argumentation first introduced by Dung [14]. We also show how to represent the constraints defining some of the classical semantics for abstract argumentation using propositional logic.

2.1. Argumentation Frameworks and Argumentation Semantics

We begin by recalling some basic terminology and notation [14]. An *argumentation framework* is a pair $AF = \langle Arg, \rightarrow \rangle$, where Arg is a finite set of arguments and \rightarrow is a binary relation on Arg . If $A \rightarrow B$ holds for two arguments $A, B \in Arg$, then we say that A *attacks* B . For $\Delta \subseteq Arg$ and $B \in Arg$, we say that Δ *attacks* B in case $A \rightarrow B$ for at least one argument $A \in \Delta$. For $\Delta \subseteq Arg$ and $C \in Arg$ we say that Δ *defends* C in case Δ attacks all arguments $B \in Arg$ with $B \rightarrow C$. We write 2^{Arg} for the powerset of Arg .

We are going to require some additional notation to describe high-level features of the topology of an argumentation framework AF . First, we write $\text{MaxAtt}(AF)$ for the maximum number of attackers on any one argument in AF (i.e., this is the maximum in-degree of the graph). Second, we write $\text{MaxDef}(AF)$ for the maximum number of attackers of an argument that itself is the source of an attack. That is, this is a measure of the number of attempts made at defending against a given attack.

$$\text{MaxAtt}(AF) = \max_{B \in Arg} |\{A \mid A \rightarrow B\}| \quad \text{MaxDef}(AF) = \max_{C \in Arg} \max_{B \rightarrow C} |\{A \mid A \rightarrow B\}|$$

Given an argumentation framework $AF = \langle Arg, \rightarrow \rangle$, the question arises which subset Δ of the set of arguments Arg one should accept. Any such set $\Delta \subseteq Arg$ is called an *extension* of AF . Different criteria have been put forward for choosing an extension.

Definition 1 (Argumentation semantics). *Let $AF = \langle Arg, \rightarrow \rangle$ be an argumentation framework and let $\Delta \subseteq Arg$ be an extension of AF . We adopt the following terminology:*

- Δ is called *conflict-free* if there are no arguments $A, B \in \Delta$ such that $A \rightarrow B$.
- Δ is called *self-defending* if $\Delta \subseteq \{C \mid \Delta \text{ defends } C\}$.
- Δ is called *reinstating* if $\{C \mid \Delta \text{ defends } C\} \subseteq \Delta$.
- Δ is called *admissible* if it is both conflict-free and self-defending.
- Δ is called *complete* if it is conflict-free, self-defending, and reinstating.
- Δ is called *preferred* if it is \subseteq -maximal amongst the admissible extensions.
- Δ is called *grounded* if it is \subseteq -minimal amongst the complete extensions.
- Δ is called *stable* if it is conflict-free and $\Delta \cup \{B \mid \Delta \text{ attacks } B\} = Arg$.

All of these alternative definitions for a suitable semantics of abstract argumentation frameworks are explained, motivated, and criticised in depth in the literature [3].

Given $AF = \langle Arg, \rightarrow \rangle$, we can think of a *property* σ of extensions, such as admissibility, as a set $\sigma \subseteq 2^{Arg}$. A property σ is called *I-maximal* (short for *inclusion-maximal*) if $\Delta_1 \subset \Delta_2$ for no two extensions $\Delta_1, \Delta_2 \in \sigma$. For instance, both the set of all preferred extensions and that of all stable extensions are I-maximal for any choice of AF [2].

2.2. Encoding Argumentation Semantics in Propositional Logic

It can be useful to be able to represent the properties of Definition 1 in a purely syntactic manner, using a logical language. So fix an argumentation framework $AF = \langle Arg, \rightarrow \rangle$, think of Arg as a set of propositional variables, and let \mathcal{L}_{AF} be the corresponding propositional language. Now extensions $\Delta \subseteq Arg$ correspond to models of formulas in \mathcal{L}_{AF} :

- $\Delta \models A$ for $A \in Arg$ if and only if $A \in \Delta$
- $\Delta \models \neg\varphi$ if and only if $\Delta \models \varphi$ is not the case
- $\Delta \models \varphi \wedge \psi$ if and only if both $\Delta \models \varphi$ and $\Delta \models \psi$

Thus, for example, $\Delta \models A \wedge \neg B$ if and only if $A \in \Delta$ and $B \notin \Delta$. The semantics of other propositional connectives (such as disjunction and implication) is defined accordingly.

Given a formula φ , we use $\text{Mod}(\varphi) = \{\Delta \subseteq Arg \mid \Delta \models \varphi\}$ to denote the set of all models of φ . Every formula φ identifies a property of extensions of AF , namely $\sigma = \text{Mod}(\varphi)$. When using a formula φ to describe such a property of extensions, we usually refer to φ as an *integrity constraint*. The following simple results, all of which are implicit in the work of Besnard and Doutre [4], characterise some of the basic properties of extensions defined earlier in terms of integrity constraints expressed in \mathcal{L}_{AF} .¹

Proposition 1 (Besnard and Doutre, 2004). *Let $AF = \langle Arg, \rightarrow \rangle$ be an argumentation framework and let $\Delta \subseteq Arg$ be an extension. Then Δ is conflict-free if and only if:*

$$\Delta \models \text{IC}_{CF}^{AF} \text{ where } \text{IC}_{CF}^{AF} = \bigwedge_{\substack{A, B \in Arg \\ A \rightarrow B}} (\neg A \vee \neg B)$$

¹While the formulas we use to encode extension properties all have a very simple syntactic structure, they are very long. This is fine for our purposes. More compact encodings, based on modal logic, exist [8, 20].

In other words, Proposition 1 states that $\text{Mod}(\text{IC}_{CF}^{AF}) = \{\Delta \subseteq \text{Arg} \mid \Delta \text{ is conflict-free}\}$.

Proposition 2 (Besnard and Doutre, 2004). *Let $AF = \langle \text{Arg}, \rightarrow \rangle$ be an argumentation framework and let $\Delta \subseteq \text{Arg}$ be an extension. Then Δ is self-defending if and only if:*

$$\Delta \models \text{IC}_{SD}^{AF} \quad \text{where} \quad \text{IC}_{SD}^{AF} = \bigwedge_{C \in \text{Arg}} [C \rightarrow \bigwedge_{\substack{B \in \text{Arg} \\ B \rightarrow C}} \bigvee_{\substack{A \in \text{Arg} \\ A \rightarrow B}} A]$$

Proposition 3 (Besnard and Doutre, 2004). *Let $AF = \langle \text{Arg}, \rightarrow \rangle$ be an argumentation framework and let $\Delta \subseteq \text{Arg}$ be an extension. Then Δ is reinstating if and only if:*

$$\Delta \models \text{IC}_{RI}^{AF} \quad \text{where} \quad \text{IC}_{RI}^{AF} = \bigwedge_{C \in \text{Arg}} [(\bigwedge_{\substack{B \in \text{Arg} \\ B \rightarrow C}} \bigvee_{\substack{A \in \text{Arg} \\ A \rightarrow B}} A) \rightarrow C]$$

We can now use the integrity constraints defined above to construct integrity constraints for the properties of being admissible and complete:²

- Δ is admissible if and only if $\Delta \models \text{IC}_{CF}^{AF} \wedge \text{IC}_{SD}^{AF}$.
- Δ is complete if and only if $\Delta \models \text{IC}_{CF}^{AF} \wedge \text{IC}_{SD}^{AF} \wedge \text{IC}_{RI}^{AF}$.

Example 1. Consider the argumentation framework $AF = \langle \{A, B\}, \{A \rightarrow B, B \rightarrow A\} \rangle$, consisting of two arguments that attack each other. Then $\text{IC}_{CF}^{AF} = (\neg A \vee \neg B) \wedge (\neg B \vee \neg A)$, which simplifies to $\neg A \vee \neg B$. The models of IC_{CF}^{AF} are \emptyset , $\{A\}$, and $\{B\}$. Indeed, these are the only conflict-free extensions of AF . Furthermore, $\text{IC}_{SD}^{AF} = (A \rightarrow A) \wedge (B \rightarrow B)$ is a tautology, so admissible and conflict-free extensions coincide for AF . \triangle

3. Aggregating Alternative Extensions

In this section, we formally introduce the scenario we study in this paper and define the central concept of the *preservation* of a property of extensions under aggregation.

Fix an argumentation framework $AF = \langle \text{Arg}, \rightarrow \rangle$. Let $N = \{1, \dots, n\}$ be a finite set of *agents*. Suppose each agent $i \in N$ supplies us with an extension $\Delta_i \subseteq \text{Arg}$, reflecting her individual views of what constitutes an acceptable set of arguments in the context of AF . Thus, we are supplied with a *profile* $\mathbf{\Delta} = (\Delta_1, \dots, \Delta_n)$, a vector of extensions, one for each agent. An *aggregation rule* is a function $F : (2^{\text{Arg}})^n \rightarrow 2^{\text{Arg}}$, mapping any given profile of extensions to a single extension. An example is the majority rule.

3.1. Quota Rules

Our focus in this paper is on a family of aggregation rules known as the *quota rules*, which generalise the idea underlying the definition of the majority rule.

²Stable extensions can be characterised in an analogous manner (but we do not require such a characterisation for our present purposes). Characterising preferred or grounded extensions is more difficult, due to the notions of maximality and minimality inherent in their definitions. But we emphasise that doing so would be possible in principle. Indeed, due to the functional completeness of the propositional calculus, for every property $\sigma \subseteq 2^{\text{Arg}}$ of extensions of AF there exists an integrity constraint IC expressed in \mathcal{L}_{AF} such that $\text{Mod}(\text{IC}) = \sigma$.

Definition 2 (Quota rules). Let $AF = \langle Arg, \rightarrow \rangle$ be an argumentation framework, let N be a set of n agents, and let $q \in \{1, \dots, n\}$. The quota rule F_q with quota q is defined as the aggregation rule mapping any given profile $\Delta = (\Delta_1, \dots, \Delta_n) \in (2^{Arg})^n$ of extensions to the extension including exactly those arguments accepted by at least q agents:


$$F_q(\Delta) = \{A \in Arg \mid \#\{i \in N \mid A \in \Delta_i\} \geq q\}$$

Such quota rules are sometimes more precisely referred to as *uniform* quota rules, to stress the fact that the acceptance of each argument is subject to *the same* quota q . Three quota rules with specific quotas will be of special interest to us:

- The (strict) *majority rule* is the quota rule F_q with quota $q = \lceil \frac{n+1}{2} \rceil$.
- The *nomination rule* is the quota rule F_q with quota $q = 1$.
- The *unanimity rule* is the quota rule F_q with quota $q = n$.

Quota rules are very natural—albeit simple—rules to consider when contemplating mechanisms to perform aggregation. They have been studied in depth in judgment aggregation [13]. One attractive feature of quota rules is their low computational complexity: computing outputs is straightforward. For comparison, more sophisticated distance-based rules (another natural class to consider) give rise to intractable optimisation problems [21, 17]. Quota rules also satisfy some appealing normative properties (known as “axioms” in social choice theory). For instance, in the context of judgment aggregation they are known to be *monotonic* (ensuring that additional support for a collectively accepted view never results in that view getting rejected) and *strategyproof* (meaning that—under certain assumptions on agents’ preferences—no agent can benefit from reporting false views).³ Having said this, there also is a significant disadvantage to using quota rules and that is the fact that they can produce seemingly paradoxical outcomes.

Example 2. Suppose three agents evaluate the following argumentation framework:



They report the extensions $\{A, C\}$, $\{A, D\}$, and $\{A, E\}$, respectively, all of which are admissible. But applying the majority rule yields $\{A\}$, which is not admissible! \triangle

3.2. Preservation of Properties of Extensions

Example 2 shows that an aggregation rule may not always *preserve* the properties shared by the extensions in the profile. Ideally, when aggregating a profile $\Delta = (\Delta_1, \dots, \Delta_n)$ in which every individual Δ_i satisfies a given property σ , we would like the output $F(\Delta)$ to satisfy σ as well. Let us now formally define this central concept of preservation.

Definition 3 (Preservation). Let $AF = \langle Arg, \rightarrow \rangle$ be an argumentation framework and let $\sigma \subseteq 2^{Arg}$ be a property of extensions of AF . Then an aggregation rule $F : (2^{Arg})^n \rightarrow 2^{Arg}$ for n agents is said to *preserve* σ if $F(\Delta) \in \sigma$ for every profile $\Delta = (\Delta_1, \dots, \Delta_n) \in \sigma^n$.

³We refer to the literature on judgment aggregation for precise statements of these results [13, 12, 16].

4. Preservation Results

In this section, we present our results on the preservation of extension properties. Our possibility results concern both simple properties and argumentation frameworks with simple topologies. Our impossibility results apply in case of more demanding scenarios.

We investigate the question of preservation for all of the properties of Definition 1, except for groundedness, which is preserved vacuously by every quota rule and for every argumentation framework. To see this, note that it is well-known that every argumentation framework has a *unique* grounded extension [14]—so if agents report grounded extensions they in fact must all report *the same* extension—and all quota rules F_q with quota $q \in \{1, \dots, n\}$ are *unanimous* in the sense that $F_q(\Delta, \dots, \Delta) = \Delta$ for every extension Δ .

4.1. Adaptation of Results from Binary Aggregation with Integrity Constraints

In our analysis, we are going to make use of results regarding a variant of judgment aggregation known as *binary aggregation with integrity constraints* [18, 19]. These results concern the conditions under which an aggregation rule will preserve (the property corresponding to) an integrity constraint. When adapted to our setting, the main result we require, due to Grandi and Endriss [19, Corollary 31], reads as follows.

Lemma 4 (Grandi and Endriss, 2013). *Let $AF = \langle Arg, \rightarrow \rangle$ be an argumentation framework and let φ be a clause in \mathcal{L}_{AF} with k_1 positive literals and k_2 negative literals. Then a quota rule F_q for n agents preserves the property $\text{Mod}(\varphi)$ if and only if:*

$$q \cdot (k_2 - k_1) > n \cdot (k_2 - 1) - k_1$$

This lemma fully characterises the conditions under which quota rules can guarantee the preservation of properties corresponding to integrity constraints that are *clauses* (i.e., disjunctions of literals). Unfortunately, none of the integrity constraints we have put forward in Section 2.2 to characterise relevant properties of extensions have this very simple structure. However, some are conjunctions of clauses and, of course, all can be translated into conjunctions of clauses. Can we still apply Lemma 4? Yes, if we are interested in the “possibility direction” of Lemma 4 (if q satisfies the inequality, then we get preservation). No, if we are interested in the “impossibility direction” (only if q satisfies the inequality can we guarantee preservation). To see this, let us first consider the possibility direction. The following result is immediate given the relevant definitions and it also is a direct corollary to a more general result by Grandi and Endriss [19, Lemma 3].

Lemma 5 (Grandi and Endriss, 2013). *Let $AF = \langle Arg, \rightarrow \rangle$ be an argumentation framework, let φ_1 and φ_2 be integrity constraints in \mathcal{L}_{AF} , and let F be an aggregation rule that preserves both $\text{Mod}(\varphi_1)$ and $\text{Mod}(\varphi_2)$. Then F also preserves $\text{Mod}(\varphi_1 \wedge \varphi_2)$.*

Now, if we combine (the possibility direction of) Lemma 4 and Lemma 5, we see that, given some clauses $\varphi_1, \dots, \varphi_\ell$, the quota rule F_q preserves $\text{Mod}(\varphi_1 \wedge \dots \wedge \varphi_\ell)$ if q satisfies the constraints specified in Lemma 4 for *all* clauses φ_i . However, the converse does not hold, i.e., we cannot use Lemma 4 to prove impossibility results for integrity constraints that are conjunctions of clauses. To see this, consider the following example.

Example 3. Consider an argumentation framework with $Arg = \{A, B\}$ and the integrity constraints $\varphi = (\neg A \vee \neg B)$ and $\psi = A$. Note that $\varphi \wedge \psi \equiv (A \wedge \neg B)$. By Lemma 4, φ is preserved by F_q only if $q \cdot (2 - 0) > n \cdot (2 - 1) - 0$, i.e., only if $q > \frac{n}{2}$. Furthermore, again by applying Lemma 4, ψ is preserved by every quota rule (because the condition $q \cdot (0 - 1) > n \cdot (0 - 1) - 1$ reduces to $n > q - 1$, which is always true). So one might be tempted to assume that the conjunction of these two integrity constraints, $\varphi \wedge \psi$, also is preserved only if $q > \frac{n}{2}$. But this clearly is false: $\varphi \wedge \psi$ is preserved by every quota rule, as under every profile in which all individual extensions satisfy $\varphi \wedge \psi \equiv (A \wedge \neg B)$ every agent in fact must be reporting the same extension $\{A\}$. \triangle

4.2. Preserving Conflict-Freeness

The following result shows that the (strict) majority rule preserves conflict-freeness, as does every quota rule with an even higher quota. On the downside, no quota rule with a quota of $q = \frac{n}{2}$ or lower will work.

Theorem 6. *Let AF be any argumentation framework with at least one attack between two arguments that do not attack themselves. Then a quota rule F_q for n agents preserves conflict-freeness for AF if and only if $q > \frac{n}{2}$.*

Proof. First, pick any quota $q > \frac{n}{2}$ and consider an arbitrary argumentation framework AF (for this direction, we do not require the restriction on AF). We need to show that F_q preserves conflict-freeness for AF . Recall from Proposition 1 that the integrity constraint IC_{CF}^{AF} characterising conflict-freeness is a conjunction of clauses of the form $\neg A \vee \neg B$. Most of these clauses are 2-clauses (with 0 positive literals and 2 negative literals each), although in case of self-attacks there may also be 1-clauses (with 1 negative literal only). In case there is not a single attack in AF , we get $IC_{CF}^{AF} = \top$. Thus, by Lemmas 4 and 5, F_q will preserve $\text{Mod}(IC_{CF}^{AF})$ provided the following three constraints are satisfied:

- 2-clauses: $q \cdot (2 - 0) > n \cdot (2 - 1) - 0$, simplifying to $q > \frac{n}{2}$
- 1-clauses: $q \cdot (1 - 0) > n \cdot (1 - 1) - 0$, simplifying to $q > 0$ (which always holds)
- tautology: $q \cdot (0 - 0) > n \cdot (0 - 1) - 0$, simplifying to $n > 0$ (which always holds)

As $q > \frac{n}{2}$ by assumption, F_q will indeed preserve $\text{Mod}(IC_{CF}^{AF})$, so we are done.

Second, pick any quota $q \leq \frac{n}{2}$. We need to find an argumentation framework $AF = \langle Arg, \rightarrow \rangle$ with arguments $A, B \in Arg$ such that $A \rightarrow B$ but neither $A \rightarrow A$ nor $B \rightarrow B$, and for which F_q does not preserve conflict-freeness. Suppose $\lceil \frac{n}{2} \rceil$ agents report the conflict-free extension $\{A\}$ and the remaining $\lfloor \frac{n}{2} \rfloor$ agents report the conflict-free extension $\{B\}$. We certainly have $\lceil \frac{n}{2} \rceil \geq q$, but as q must be an integer we also have $\lfloor \frac{n}{2} \rfloor \geq q$ (even if q is odd). Thus, F_q returns the set $\{A, B\}$, which is not conflict-free. \square

Theorem 6 has both a possibility and an impossibility component: if $q > \frac{n}{2}$, then it is possible to preserve conflict-freeness (whatever the argumentation framework), while it is impossible to find a quota rule F_q with $q \leq \frac{n}{2}$ that preserves conflict-freeness—at least when AF includes at least one attack between arguments that do not attack themselves.

4.3. Preserving Admissibility and Completeness

To analyse the preservation of admissibility, we first consider its second constituent property, namely that of being self-defending. We start with a technical lemma.

Lemma 7. *A quota rule F_q for n agents preserves the property of being self-defending for an argumentation framework AF if $q \cdot (\text{MaxDef}(AF) - 1) < \text{MaxDef}(AF)$.*

Proof. Let $AF = \langle \text{Arg}, \rightarrow \rangle$ be the argumentation framework under consideration. Recall that IC_{SD}^{AF} is a conjunction of formulas of the form $C \rightarrow \bigwedge_{\substack{B \in \text{Arg} \\ B \rightarrow C}} \bigvee_{\substack{A \in \text{Arg} \\ A \rightarrow B}} A$, which we can rewrite as $\bigwedge_{\substack{B \in \text{Arg} \\ B \rightarrow C}} (\neg C \vee \bigvee_{\substack{A \in \text{Arg} \\ A \rightarrow B}} A)$. Let us consider one such clause $\neg C \vee \bigvee_{\substack{A \in \text{Arg} \\ A \rightarrow B}} A$. Its number of negative literals is 1. Its number of positive literals depends on both C and B , so let us refer to it as $k_{C,B}$. By Lemma 4, a quota rule F_q will preserve this clause if and only if $q \cdot (1 - k_{C,B}) > n \cdot (1 - 1) - k_{C,B}$, i.e., if and only if $q \cdot (k_{C,B} - 1) < k_{C,B}$. As IC_{SD}^{AF} is equivalent to a conjunction of such clauses, by Lemma 5, we need to satisfy this inequality for all relevant $k_{C,B}$. This requirement is most demanding for large values of $k_{C,B}$. Observe that the largest value of $k_{C,B}$ is $\text{MaxDef}(AF)$, i.e., we satisfy all inequalities (and thus preserve all clauses) in case $q \cdot (\text{MaxDef}(AF) - 1) < \text{MaxDef}(AF)$. \square

Applying Lemma 7, we immediately obtain the following two possibility results (by noting that $q = 1$ for the nomination rule and by recalling Theorem 6, respectively).


Proposition 8. *The nomination rule preserves the property of being self-defending.*

Proposition 9. *Every quota rule F_q for n agents with a quota $q > \frac{n}{2}$ preserves admissibility for all argumentation frameworks AF with $\text{MaxDef}(AF) \leq 1$.*

But what if we cannot or do not want to make any *a priori* assumptions regarding the structure of the argumentation framework (such as $\text{MaxDef}(AF) \leq 1$)? Lemma 7 is a positive result, but it also hints at a problem: it shows that preservation is possible for *low* quotas, while Theorem 6 requires *high* quotas. While Lemma 7 only states a sufficient and not a necessary condition (so while there could be hope in principle), the following impossibility result shows that this apparent conflict cannot be resolved in general.

Theorem 10. *No quota rule preserves admissibility for all argumentation frameworks.*

Proof. It suffices to show that there exists a specific argumentation framework AF for which no quota rule can preserve admissibility. So let AF be defined as follows:



First, consider any quota rule F_q with a quota $q \geq 2$. Suppose one agent reports $\{A, C\}$, $q - 1$ agents report $\{A, D\}$, and the remaining agents report \emptyset . All of these extensions are admissible. Yet, the outcome returned by F_q is $\{A\}$, which is not admissible.

It remains for us to show that the claim holds also for the nomination rule, i.e., for $q = 1$. For the profile in which one agent reports $\{A, C\}$ and all others report $\{A, D\}$ (both of which are admissible), the nomination rule returns the inadmissible $\{A, C, D\}$. \square

We remark that Theorem 10 is conceptually weaker than the impossibility direction of Theorem 6. Theorem 10 only states that no quota rule will work for *all* argumentation frameworks, while Theorem 6 shows that even if we know the argumentation framework in advance we cannot design a quota rule with a low quota that will work.

Next, we provide an analogous analysis of the preservation of complete extensions. Recall that an extension is complete if it is both admissible and reinstating.

Lemma 11. *A quota rule F_q for n agents preserves the property of being reinstating for an argumentation framework AF if $q \cdot (\text{MaxAtt}(AF) - 1) > n \cdot (\text{MaxAtt}(AF) - 1) - 1$.*

Proof. Let $AF = \langle \text{Arg}, \rightarrow \rangle$ be the argumentation framework under consideration. Recall that IC_{RI}^{AF} is a conjunction of formulas of the form $(\bigwedge_{\substack{B \in \text{Arg} \\ B \rightarrow C}} \bigvee_{\substack{A \in \text{Arg} \\ A \rightarrow B}} A) \rightarrow C$. For a given C , let $\{B_1, \dots, B_m\}$ be the set of attackers of C . We can rewrite as follows:

$$\left(\bigwedge_{\substack{i=1 \\ A \in \text{Arg} \\ A \rightarrow B_i}}^m \bigvee A \right) \rightarrow C \quad \equiv \quad \left(\bigvee_{\substack{i=1 \\ A \in \text{Arg} \\ A \rightarrow B_i}}^m \bigwedge \neg A \right) \vee C \quad \equiv \quad \bigwedge_{\substack{A_1 \in \text{Arg} \\ A_1 \rightarrow B_1}} \cdots \bigwedge_{\substack{A_m \in \text{Arg} \\ A_m \rightarrow B_m}} (\neg A_1 \vee \cdots \vee \neg A_m \vee C)$$

Thus, we obtain a conjunction of $(m + 1)$ -clauses, with m negative literals and one positive literal each (recall that m is the number of attackers of C). By Lemma 4, we can preserve this part of the integrity constraint by ensuring $q \cdot (m - 1) > n \cdot (m - 1) - 1$. Doing so becomes harder as m increases. Hence, by Lemma 5, we can ensure IC_{RI}^{AF} will be preserved if this inequality holds for the maximal value of m , which is $\text{MaxAtt}(AF)$. \square

Applying Lemma 11 immediately yields the following two possibility results. To prove the first, note that $q = n$ for the unanimity rule. To prove the second, recall Proposition 9 and observe that $\text{MaxAtt}(AF) \leq 1$ implies $\text{MaxDef}(AF) \leq 1$.

Proposition 12. *The unanimity rule preserves the property of being reinstating.*

Proposition 13. *Every quota rule F_q for n agents with a quota $q > \frac{n}{2}$ preserves completeness for all argumentation frameworks AF with $\text{MaxAtt}(AF) \leq 1$.*

Finally, the proof of Theorem 10 can be adapted to obtain the following impossibility result (given that all the admissible extensions mentioned in that proof are also complete).

Theorem 14. *No quota rule preserves completeness for all argumentation frameworks.*

4.4. Preserving I-Maximal Properties of Extensions

Rather than focusing on the remaining two classical argumentation semantics, the preferred and the stable semantics, we provide an analysis that covers all extension properties that are I-maximal. We obtain a sweeping impossibility result.

Theorem 15. *Let $AF = \langle \text{Arg}, \rightarrow \rangle$ be an argumentation framework, let $\sigma \subseteq 2^{\text{Arg}}$ be an I-maximal property of extensions of AF with $|\sigma| \geq 2$, and let n be the number of agents. If n is even, then no quota rule preserves σ for AF . If n is odd, then no quota rule different from the majority rule preserves σ for AF .*

Proof. Consider two distinct extensions $\Delta_1, \Delta_2 \in \sigma$. As $\Delta_1 \neq \Delta_2$ and $\Delta_1 \not\subseteq \Delta_2$ (which follows from I-maximality), we must have $\Delta_1 \cap \Delta_2 \subset \Delta_1$ and $\Delta_2 \subset \Delta_1 \cup \Delta_2$. Thus, neither $\Delta_1 \cap \Delta_2$ nor $\Delta_1 \cup \Delta_2$ can belong to the I-maximal σ .

Now suppose n is even. If exactly half of the agents report Δ_1 and the other half report Δ_2 , then the outcome will be either $\Delta_1 \cap \Delta_2$ (for quotas $q > \frac{n}{2}$) or $\Delta_1 \cup \Delta_2$ (for $q \leq \frac{n}{2}$), i.e., σ will not be preserved under any quota rule.

Next, suppose n is odd. Consider a profile in which $\lfloor \frac{n}{2} \rfloor$ agents report Δ_1 and the remaining $\lceil \frac{n}{2} \rceil$ agents report Δ_2 . For quotas $q > \lceil \frac{n}{2} \rceil$ the outcome is $\Delta_1 \cap \Delta_2$ and for $q \leq \lfloor \frac{n}{2} \rfloor$ it is $\Delta_1 \cup \Delta_2$. Thus, for no quota rule with a quota different from $q = \lceil \frac{n}{2} \rceil = \lceil \frac{n+1}{2} \rceil$ (corresponding to the majority rule for odd n) will σ be preserved under aggregation. \square

Thus, for instance, for even n no quota rule can preserve the property of being a preferred extension for argumentation frameworks without a unique preferred extension.

Theorem 15 applies if $|\sigma| \geq 2$. For $|\sigma| = 0$ every aggregation rule vacuously preserves σ , and for $|\sigma| = 1$ every unanimous rule (and thus every quota rule) does. This leaves the case of $|\sigma| \geq 2$ for odd n and the majority rule. For the case of *exactly* two extensions we obtain a simple possibility result (not just for I-maximal properties).

Proposition 16. *Let $AF = \langle Arg, \rightarrow \rangle$ and let $\sigma \subseteq 2^{Arg}$ be a property of extensions of AF with $|\sigma| = 2$. Then the majority rule for an odd number of agents preserves σ .*

Proof. Let $\sigma = \{\Delta_1, \Delta_2\}$. In any profile satisfying σ for an odd number of agents, there will either be a strict majority for Δ_1 or a strict majority for Δ_2 , but not both. Thus, the outcome under the majority rule must be either Δ_1 or Δ_2 . \square

For $|\sigma| \geq 3$ there are examples where the majority rule (for odd n) preserves σ and others where it does not. Together with Theorem 15, the latter kind of example can be used to obtain impossibility results like Theorems 10 and 14 also for preferred and stable extensions. Finally, using similar techniques as in Section 4.3, one can identify classes of topologically simple argumentation frameworks for which the majority rule (for odd n) preserves certain I-maximal properties. For instance, we have been able to show that it preserves stability in case $\text{MaxAtt}(AF) \leq 1$. We omit all details for lack space.

5. Related Work

Our work contributes to the research agenda on collective argumentation recently surveyed by Bodanza et al. [5]. Most prior work applying the methodology of social choice theory in the context of abstract argumentation has dealt with the problem of aggregating alternative argumentation frameworks [see, e.g., 24, 15, 11, 10]. In contrast, here we are concerned with the aggregation of alternative extensions of the same argumentation framework. In this section, we briefly review recent contributions to the literature that, like ours, deal with the problem of aggregating individual views on which arguments to accept when confronted with a fixed argumentation framework. In all of them the stances of individual agents are represented using Caminada's three-valued labelling approach [7] rather than Dung's extensions [14]. Therefore, results are not directly comparable at a technical level, but there clearly are conceptual connections.

The work most closely related to ours is that of Rahwan and Tohmé [23]. These authors analyse the "argument-wise plurality rule" (which is similar in spirit to the majority rule) and show that it has a number of desirable axiomatic properties. These results are closely related to well-known results for the quota rules in judgment aggregation [13]. Rahwan and Tohmé also observe that this rule does not preserve the completeness of labellings and then generalise this observation to an impossibility result in the spirit of the List-Pettit Theorem in judgment aggregation [22]. The latter result may be considered a conceptual—albeit not technical—generalisation of our Theorem 14. This line of

work has later been continued by Awad et al. [1], who also identify special classes of argumentation frameworks that permit more favourable results.

Other prior work regarding the aggregation of alternative labellings of a given argumentation framework focuses on the design of new aggregation rules with desirable properties. For instance, the “sceptical outcome” of Caminada and Pigozzi [9] is computed by first applying a rule akin to the unanimity rule to establish which arguments all agents agree on and then “correcting” that preliminary outcome to ensure a meaningful final result. Booth et al. [6] introduce an approach for defining aggregation rules that generalises both the argument-wise plurality rule studied by Rahwan and Tohmé [23] and the rules put forward by Caminada and Pigozzi [9].

6. Conclusion

We have established a range of results on the capacity of quota rules to preserve some of the classical properties of interest to the semantics of abstract argumentation frameworks. For the most basic properties, we have obtained possibility results: conflict-freeness is preserved by the majority rule (amongst others), self-defense by the nomination rule, and reinstatement by the unanimity rule. However, most argumentation semantics involve combinations of two or more of such basic properties and we have seen that this typically gives rise to impossibility results: no quota rule can guarantee the preservation of such properties in general. While for some argumentation semantics possibility results can still be proved for specific argumentation frameworks with particularly simple topologies, for I-maximal properties even this usually is impossible.

For our possibility results, we have exploited a simple encoding of argumentation semantics in propositional logic [4] and built on prior work in judgment aggregation [19]. This systematic approach greatly simplifies the task of deriving proofs. This illustrates—once more [5]—how ideas originating in social choice theory can be fruitfully applied to the analysis of scenarios of collective argumentation.

There are at least three natural directions for future work. First, one may want to consider further properties of extensions proposed in the literature on abstract argumentation [3]. Second, one may want to consider other aggregation rules besides the quota rules. In particular, rather than analysing specific rules we may wish to consider classes of rules characterised by axiomatic properties [10, 16, 23]. Third, while abstract argumentation is attractive due to its mathematical simplicity, it also has obvious shortcomings when it comes to modelling complex real-world scenarios. Therefore, one may want to follow a similar research agenda also for other types of argumentation formalisms.

Acknowledgments. We would like to thank Umberto Grandi and the anonymous COMMA-2018 reviewers for their helpful suggestions and feedback.

References

- [1] E. Awad, R. Booth, F. Tohmé, and I. Rahwan. Judgement aggregation in multi-agent argumentation. *Journal of Logic and Computation*, 27(1):227–259, 2017.
- [2] P. Baroni and M. Giacomin. On principle-based evaluation of extension-based argumentation semantics. *Artificial Intelligence*, 171(10):675–700, 2007.

- [3] P. Baroni, M. Caminada, and M. Giacomin. An introduction to argumentation semantics. *Knowledge Engineering Review*, 26(4):365–410, 2011.
- [4] P. Besnard and S. Doutre. Checking the acceptability of a set of arguments. In *Proceedings of the 10th International Workshop on Non-Monotonic Reasoning (NMR)*, 2004.
- [5] G. A. Bodanza, F. A. Tohmé, and M. R. Auday. Collective argumentation: A survey of aggregation issues around argumentation frameworks. *Argument & Computation*, 8(1):1–34, 2017.
- [6] R. Booth, E. Awad, and I. Rahwan. Interval methods for judgment aggregation in argumentation. In *Proceedings of the 14th International Conference on Principles of Knowledge Representation and Reasoning (KR)*, 2014.
- [7] M. Caminada. On the issue of reinstatement in argumentation. In *Proceedings of the 10th European Conference on Logics in Artificial Intelligence (JELIA)*. Springer, 2006.
- [8] M. Caminada and D. M. Gabbay. A logical account of formal argumentation. *Studia Logica*, 93(2/3):109–145, 2009.
- [9] M. Caminada and G. Pigozzi. On judgment aggregation in abstract argumentation. *Journal of Autonomous Agents and Multiagent Systems*, 22(1):64–102, 2011.
- [10] W. Chen and U. Endriss. Preservation of semantic properties during the aggregation of abstract argumentation frameworks. In *Proceedings of the 16th Conference on Theoretical Aspects of Rationality and Knowledge (TARK)*, 2017.
- [11] J. Delobelle, S. Konieczny, and S. Vesic. On the aggregation of argumentation frameworks. In *Proceedings of the 24th International Joint Conference on Artificial Intelligence (IJCAI)*, 2015.
- [12] F. Dietrich and C. List. Strategy-proof judgment aggregation. *Economics & Philosophy*, 23(3):269–300, 2007.
- [13] F. Dietrich and C. List. Judgment aggregation by quota rules: Majority voting generalized. *Journal of Theoretical Politics*, 19(4):391–424, 2007.
- [14] P. M. Dung. On the acceptability of arguments and its fundamental role in nonmonotonic reasoning, logic programming and n -person games. *Artificial Intelligence*, 77(2):321–358, 1995.
- [15] P. E. Dunne, P. Marquis, and M. Wooldridge. Argument aggregation: Basic axioms and complexity results. In *Proceedings of the 4th International Conference on Computational Models of Argument (COMMA)*. IOS Press, 2012.
- [16] U. Endriss. Judgment aggregation. In F. Brandt, V. Conitzer, U. Endriss, J. Lang, and A. D. Procaccia, editors, *Handbook of Computational Social Choice*, chapter 17, pages 399–426. Cambridge University Press, 2016.
- [17] U. Endriss, U. Grandi, and D. Porello. Complexity of judgment aggregation. *Journal of Artificial Intelligence Research (JAIR)*, 45:481–514, 2012.
- [18] U. Grandi and U. Endriss. Binary aggregation with integrity constraints. In *Proceedings of the 22nd International Joint Conference on Artificial Intelligence (IJCAI)*, 2011.
- [19] U. Grandi and U. Endriss. Lifting integrity constraints in binary aggregation. *Artificial Intelligence*, 199–200:45–66, 2013.
- [20] D. Grossi. On the logic of argumentation theory. In *Proceedings of the 9th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*. IFAAMAS, 2010.
- [21] S. Konieczny, J. Lang, and P. Marquis. DA^2 merging operators. *Artificial Intelligence*, 157(1–2):49–79, 2004.
- [22] C. List and P. Pettit. Aggregating sets of judgments: An impossibility result. *Economics and Philosophy*, 18(1):89–110, 2002.
- [23] I. Rahwan and F. A. Tohmé. Collective argument evaluation as judgement aggregation. In *Proceedings of the 9th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*. IFAAMAS, 2010.
- [24] F. A. Tohmé, G. A. Bodanza, and G. R. Simari. 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, 2008.