Social Choice and Welfare
https://doi.org/10.1007/s00355-019-01230-5

ORIGINAL PAPER

Neutrality and relative acceptability in judgment aggregation

Zoi Terzopoulou[1] · **Ulle Endriss**[1]

Received: 20 March 2019 / Accepted: 28 November 2019

Abstract
One of the fundamental normative principles in social choice theory is that of neutrality. In the context of judgment aggregation, neutrality is encoded in the form of an axiom expressing that, when two possible judgments enjoy the same support amongst the individuals, then either both or neither of them should be accepted. This is a reasonable requirement in many scenarios. However, we argue that for scenarios in which individuals are asked to pass judgment on very diverse kinds of propositions, a notion of relative acceptability is better suited. We capture this notion by a new axiom that hinges on a binary "acceptability" relation A between propositions: if a given coalition accepting a proposition p entails the collective acceptance of p, then the same should be true for every other proposition q related to p via A. Intuitively, pAq means that p is at least as acceptable as q. Classical neutrality is then a special case where all propositions are equally acceptable. We show that our new axiom allows us to circumvent a classical impossibility theorem in judgment aggregation for certain scenarios of practical interest. Also, we offer a precise characterisation of all scenarios that are *safe*, in the sense that any aggregation rule respecting the relative acceptability between propositions will always return logically consistent outcomes.

1 Introduction

How can we ensure that a group of individuals, despite having diverse views, will be able to make collective decisions in a manner that is fair and that guarantees that those decisions are internally coherent? Judgment aggregation is a formal framework

The authors are grateful to the associate editor and the two reviewers for their comments and suggestions. We especially thank Franz Dietrich for first directing our attention towards the concept of relative acceptability during a presentation of this work in Paris.

✉ Zoi Terzopoulou
z.terzopoulou@uva.nl

[1] Institute for Logic, Language and Computation (ILLC), University of Amsterdam, Postbus 94242, 1090 GE Amsterdam, The Netherlands

Published online: 04 December 2019

for group decision making in logically rich domains, modelling situations where the judgments of several individuals need to be aggregated into a single collective judgment. The development of judgment aggregation was originally inspired by problems in legal theory (Kornhauser and Sager 1993). Since then, judgment aggregation has attracted the interest of scholars in a variety of disciplines at the interface of economics and philosophy (see, e.g., List and Pettit 2002; Dietrich 2006; Pauly and van Hees 2006; Nehring and Puppe 2008; List and Puppe 2009a; Dokow and Holzman 2010; List 2012), and has also been studied within the areas of computer science and artificial intelligence (see, e.g., Grossi and Pigozzi 2014; Baumeister et al. 2016; Endriss 2016a). Uniformly across all research fields where judgment aggregation can be applied, important considerations concern the properties—also known as *axioms* (Thomson 2001)—of the specific aggregation rules that are employed (List and Pettit 2002; Dietrich and List 2010). Axioms in judgment aggregation play a *normative* role, meaning that they formally define requirements that we would like an aggregation rule to satisfy, motivated by philosophical, political, or ethical concerns.

In this paper we delve deeply, both conceptually and technically, into one of the axioms that are prominent in the literature on judgment aggregation, namely *neutrality* (Dietrich and List 2007a, 2008; List 2012; Endriss 2016a).[1] Neutrality demands that all the propositions judged by the individuals should be treated symmetrically during aggregation—without any kind of differentiation. This clearly is a natural requirement for some applications, as it encodes a basic notion of fairness. But, as we are going to argue, in most domains in which judgment aggregation can be applied the propositions to be judged will be very diverse in nature and, therefore, the classical neutrality axiom is too strong a requirement. In particular, different propositions are often associated with different standards of acceptance—respecting their relative acceptability does not necessarily mean treating them symmetrically. To prepare the ground for this argument, let us consider an example.

Example 1 Suppose the governor of Arizona has to make two decisions concerning the state:

(i) Should a wall be constructed separating Arizona and Mexico?
(ii) Should the tuition fees for the state colleges be increased?

It is publicly known that in order for the wall to be built a very high budget is required, so for that budget to be covered money has to be collected (also) through more tuition. Since there is a lot at stake for the governor, he puts together a committee of counselors to help him settle the issues, and the committee happens to be quite diverse, consisting of four pure democrats and six extreme conservatives. Suppose the counselors—not deviating from their party lines—express the judgments depicted in the table that follows.

Thus, the governor observes that 60% of the committee members agree with building the wall, and that exactly the same counselors who want the wall support the increase on college tuition as well. As expected, everyone also acknowledges that they cannot accept the wall construction without a tuition raise (as can be seen from the table

[1] In fact, in the early literature on judgment aggregation, neutrality is often combined with an *independence* axiom and appears only in disguise, as part of an axiom called *systematicity*.

	Wall?	Tuition?
6 Counselors	Yes	Yes
4 Counselors	No	No

above, nobody says *yes* on wall and *no* on tuition). But given these judgments of the committee, how will the governor make a final decision? First, regarding the tuition, it makes sense that the governor is determined to simply follow the opinion of the majority of the counselors; this is an issue that will only have some local effects in the state of Arizona, so a simple majority by the committee members suffices for the governor to make a decision in favour of an increase. However, constructing a wall will have much more serious consequences to Arizona's finances and international relations with Mexico, so the governor only feels confident to decide in favour of it if he has at least a $2/3$ majority of the committee backing him on this (which he does not have in this case).[2] Overall, the governor would now like to increase tuition but not build a wall, even though the committee provided the same support on both issues.

Apparently, in situations like the one described above, propositions may vary in *meta-value:* Judgments on certain propositions may have much more severe and large-scale effects than judgments on others. Then, we arguably should treat propositions of the same kind symmetrically, while we may still want to apply very different standards across the different kinds.

Our conceptual contribution in this paper is to propose a new family of axioms that can account for such diversity between the propositions under consideration, providing a refinement of the standard notion of neutrality in judgment aggregation.[3] Our axioms capture a plethora of scenarios based on an acceptability relation, which is a binary relation expressing the relative acceptability of the different propositions in a decision-making domain. Given such an acceptability relation A, an *A-respecting* aggregation rule imposes equal treatment *only* for those propositions that are linked to each other by A. However, if one proposition φ is related to another proposition ψ via A, but the converse is not true, then an A-respecting rule merely has to ensure that the standards for accepting φ are at least as demanding as those for accepting ψ. We provide multiple application-driven examples for concrete definitions of relations.[4]

Our technical contribution is twofold. First, known *impossibility theorems* state that no aggregation rule satisfies certain combinations of desirable axioms (List and Pettit 2002, 2004; Pauly and van Hees 2006; Dietrich 2007; van Hees 2007). We investigate to what extent this pessimistic finding can be avoided by relaxing neutrality as far as the original impossibility theorem in the field, due to List and Pettit (2002),

[2] It is actually the case in many legislative systems, like the German one, that simple laws can be changed by getting the approval of a simple majority in the parliament, while constitutional laws require, for instance, a $2/3$ majority.

[3] We have previously outlined this idea in a preliminary version of this paper (Terzopoulou and Endriss 2019).

[4] In the judgment aggregation literature, binary relations between propositions have also been employed to capture priority (List 2004; Dietrich and List 2007b) or relevance (Dietrich 2015). Although our axioms formally rely on a binary relation as well, our intended interpretation of it—hinging on the propositions' standards of acceptance—is different.

is concerned. We focus on a family of simple yet powerful scenarios (so-called *conjunctive* and *disjunctive agendas*), which are the most extensively studied scenarios in the literature (Chapman 1998; Pettit 2001; Dietrich and List 2007c; Chapman 2002; Hartmann and Sprenger 2012; Terzopoulou and Endriss 2018), and we find that the classical impossibility can be circumvented for several cases of practical interest. Second, we provide a precise *characterisation* of scenarios for which it is *safe* to require individuals to use an aggregation rule that is A-respecting (given an acceptability relation A between the propositions), in the sense of never resulting in a logically inconsistent outcome. This subsumes known results regarding the special case of the standard neutrality axiom (Endriss et al. 2012).

Although we are the first to formalise the concept of weakening neutrality in judgment aggregation, the idea was already briefly mentioned more than fifteen years ago in Footnote 4 of the the seminal paper by List and Pettit (2002). Having said this, List and Pettit only discuss the option of limiting neutrality to propositions that belong to the same equivalence class, while our definition accounts for a much broader spectrum of relations that may occur, motivated by the immense diversity of possible aggregation domains. More recently, Slavkovik (2014) also noted that judgment aggregation often deals with propositions of diverse types and, to account for this, argued for considering non-neutral aggregation rules. But her focus is on proposing new rules, rather than on studying the neutrality axiom itself. Others have observed that certain intuitively reasonable aggregation rules in fact violate the standard neutrality axiom. For example, contrary to what one might expect, even though the well-known *Kemeny rule* is neutral in the context of preference aggregation, this rule—which is also known as the *median rule*—is not neutral in judgment aggregation (Costantini et al. 2016). Moreover, in many classical examples of judgment aggregation (like some versions of the doctrinal paradox due to Kornhauser and Sager 1993) propositions are distinguished, being either *premises* or *conclusions*—as Mongin (2012) observes in his Footnote 13, standard neutrality (the part of the so-called systematicity axiom discussed by Mongin that is relevant here) is not sensitive to this distinction. In conclusion, while neutrality clearly is a debatable property and while that point has been made repeatedly in the literature, relaxing this overly demanding axiom has not yet received the attention it deserves. In this paper we close this gap.

The remainder of the paper is organised as follows. In Sect. 2 we recall the basic model of judgment aggregation and review the standard neutrality axiom. Section 3 then introduces our new family of axioms, motivated by a number of typical examples demonstrating that in settings of judgment aggregation propositions will often be associated with different standards of acceptance, which we would like to respect. The next two sections contain our main technical results: we study the impact that relaxing neutrality has on the impossibility of List and Pettit (2002) in Sect. 4, and we determine under which conditions rules satisfying our new axioms are safe in Sect. 5. Finally, Sect. 6 concludes. In the Appendix we include a comparison of the definition of classical neutrality with two alternative definitions, inspired by the subtly different way in which neutrality usually is defined in preference aggregation.

2 The model

In this section we present the basic framework of judgment aggregation (List and Pettit 2002; Dietrich 2007; List and Puppe 2009a; List 2012; Grossi and Pigozzi 2014; Endriss 2016a) and introduce the relevant notation and terminology.

2.1 Preliminaries

Consider a finite *group of individuals* $\mathcal{N} = \{1, 2, \ldots, n\}$, with $n \geq 2$, who make judgments on a finite number of propositions. Each proposition is defined as a formula φ in propositional logic. An *agenda* $\Phi = \Phi^+ \cup \Phi^-$ is a collection of all such formulas, where Φ^+ contains only non-negated ones and $\Phi^- = \{\neg\varphi \mid \varphi \in \Phi^+\}$. We define $\sim\varphi$, the *complement* of φ, as follows: if $\varphi \in \Phi^+$, then $\sim\varphi = \neg\varphi$, and if $\varphi = \neg\psi$ for some $\psi \in \Phi^+$ (so $\varphi \in \Phi^-$), then $\sim\varphi = \psi$.

An individual i's *judgment set* (or simply *judgment*) $J_i \subseteq \Phi$ is the set of formulas in Φ with which she agrees. We assume that all individual judgment sets are *consistent*, i.e., consistent sets of formulas in the standard sense of logic, and *complete*, i.e., they contain one member of each formula-complement pair in Φ. A *profile* $\boldsymbol{J} = (J_1, \ldots, J_n)$ is a vector of judgment sets, one for every individual in the group. We use $N_\varphi^{\boldsymbol{J}} = \{i \in \mathcal{N} \mid \varphi \in J_i\}$ to denote the set of individuals who accept φ in $\boldsymbol{J}$. Since individual judgments are complete and consistent, it holds that $N_{\sim\varphi}^{\boldsymbol{J}} = \mathcal{N} \setminus N_\varphi^{\boldsymbol{J}}$. Finally, an *aggregation rule* F is a function that maps every profile $\boldsymbol{J}$ of individual judgments to a set of collective judgments (not necessarily complete or consistent). If $F(\boldsymbol{J})$ is a singleton independently of the input profile, then F is called *resolute*. All our formal results in this paper concern resolute rules; thus, slightly abusing notation, we will usually treat $F(\boldsymbol{J})$ as a collective judgment itself (rather than as a set containing a single collective judgment).

Example 2 Recall the scenario of Example 1. Formally, the ten counselors constitute a group $\mathcal{N} = \{1, \ldots, 10\}$ and are asked to express opinions on propositions $w, \neg w$ and $t, \neg t$ (capturing the questions about the wall and the tuition). To be precise, the counselors also judge the pair of propositions $w \to t$ and $\neg(w \to t)$ (capturing the constraint that for the wall to be built an increase on tuition is necessary), but for simplicity we omit the relevant judgments, since all individuals agree with proposition $w \to t$ in this case and it does not affect the outcome. Then, given the profile $\boldsymbol{J}$ depicted in the table of Example 1, a conservative i's and a democrat j's judgments will be $J_i = \{w, t\}$ and $J_j = \{\neg w, \neg t\}$, respectively, while the sets $N_w^{\boldsymbol{J}} = N_t^{\boldsymbol{J}}$ will contain all conservatives and the sets $N_{\neg w}^{\boldsymbol{J}} = N_{\neg t}^{\boldsymbol{J}}$ will contain all democrats.

2.2 Specific agendas and aggregation rules

Different agendas may be of relevance across applications, relating to diverse domains of decision making. For instance, a wide range of interesting scenarios are captured by *conjunctive* agendas, built from a set of *premises* and a single *conclusion* (Pettit 2001; Dietrich and List 2007c; Hartmann and Sprenger 2012). The premises are understood

to denote the *reasons* behind a final decision that has to be made, and the conclusion is true if and only if all the premises are. Formally, the premises in a conjunctive agenda are taken to be literals $p_1, \neg p_1, \ldots, p_k, \neg p_k$, with $k \geq 2$, and the conclusion is of the form $c = (\ell_1 \wedge \cdots \wedge \ell_k)$, where $\ell_j \in \{p_j, \neg p_j\}$ for every $j \in \{1, \ldots, k\}$. Analogously, the conclusion of a *disjunctive* agenda is a disjunction of premises.

It is also well-known that judgment aggregation can describe *preference aggregation* problems, including those of voting (List and Pettit 2004; Dietrich and List 2007a). It does so via *preference agendas* Φ that contain a formula for each (ranked) pair of alternatives in a given set $\mathcal{A}$. That is, $\Phi = \{p_{a \succ b} \mid a, b \in \mathcal{A}, a \neq b\}$. Omitting technical details, we can assume that the complete and consistent subsets of a preference agenda are exactly those corresponding to the linear orders over the alternatives (Dietrich and List 2007a).

Directing our attention back to aggregation rules, let us define the resolute *premise-based rule* F^{pr}, which functions on conjunctive and disjunctive agendas (Pettit 2001; Chapman 2002; Dietrich and List 2007c; Hartmann and Sprenger 2012). Given a profile $\boldsymbol{J}$ (for an odd number of individuals), F^{pr} accepts a premise p if and only if a majority of individuals do. Then for the conclusion, F^{pr} accepts either c or $\neg c$, depending on what is implied by the accepted premises. The premise-based rule is very popular not only due to its intuitive appeal, but also because it provides a natural way to resolve the *discursive dilemma* (Pettit 2001), i.e., cases where simply following the majority both on the premises and on the conclusion leads to inconsistent outcomes.

Next, an *oligarchy* $F_{\mathcal{N}'}$, associated with a non-empty subset of the individuals $\mathcal{N}' \subseteq \mathcal{N}$, is a resolute aggregation rule that does not distinguish between premises and conclusions, while it still guarantees consistent outcomes for conjunctive and disjunctive agendas (Dokow and Holzman 2009).[5] Take Φ to be a conjunctive agenda. For a profile $\boldsymbol{J}$ and for every formula $\varphi \in \Phi$, it holds that:

$$\varphi \in F_{\mathcal{N}'}(\boldsymbol{J}) \text{ if and only if } \begin{cases} N_{\varphi}^{\boldsymbol{J}} \supseteq \mathcal{N}' & \text{for } \varphi \in \Phi^{+} \\ N_{\sim\varphi}^{\boldsymbol{J}} \not\supseteq \mathcal{N}' & \text{for } \varphi \in \Phi^{-} \end{cases}$$

From the definition above we see that an oligarchy is *biased* towards negated formulas $\neg\varphi \in \Phi^{-}$. Speaking informally, the default is to accept $\neg\varphi$, and φ is accepted only if all oligarchs in $\mathcal{N}'$ accept it.[6] When $\mathcal{N}' = \mathcal{N}$, the corresponding oligarchic rule is the *unanimity rule* $F_{\mathcal{N}}$: for every pair of formulas, it accepts the positive formula if and only if all individuals do. For disjunctive agendas instead, Dokow and Holzman (2009) define an oligarchy to be biased towards non-negated formulas.

Finally, we abstract away from special agendas and define the non-resolute *Kemeny rule* F^{K} (Miller and Osherson 2009). For a profile $\boldsymbol{J} = (J_1, \ldots, J_n)$:

$$F^{K}(\boldsymbol{J}) = \underset{\substack{J \subseteq \Phi \\ J \text{ complete and consistent}}}{\operatorname{argmax}} \sum_{i \in \mathcal{N}} |J_i \cap J|$$

[5] Following Dokow and Holzman (2009), we present a particular definition of an oligarchy that treats positive and negative formulas differently.

[6] For a study of (un)biased judgment aggregation, we refer to Dietrich and List (2010).

In words, the Kemeny rule tries to maximise agreement with all the individual judgment sets, with the constraint of producing only complete and consistent outcomes.

2.3 Axioms

A resolute aggregation rule F may or may not satisfy a number of normatively desirable properties, known as axioms. We say that F is:

- *Independent* if and only if for all profiles $\boldsymbol{J}, \boldsymbol{J'}$ and formulas $\varphi \in \Phi$, if $N_\varphi^{\boldsymbol{J}} = N_\varphi^{\boldsymbol{J'}}$, then $\varphi \in F(\boldsymbol{J}) \Leftrightarrow \varphi \in F(\boldsymbol{J'})$. This means that F treats each proposition separately.
- *Anonymous* if and only if for all permutations $\pi : \mathcal{N} \to \mathcal{N}$ and profiles $\boldsymbol{J} = (J_1, \ldots, J_n)$, $F(J_1, \ldots, J_n) = F(J_{\pi(1)}, \ldots, J_{\pi(n)})$. That is, the names of the individuals do not play a role for F.
- *Unanimous* if and only if for all profiles $\boldsymbol{J} = (J_1, \ldots, J_n)$ and formulas $\varphi \in \Phi$, it is the case that $\varphi \in F(\boldsymbol{J})$ whenever $\varphi \in J_i$ for all $i \in \mathcal{N}$. Thus, F respects all unanimous judgments of the group.
- *Complement-free* if and only if for all profiles $\boldsymbol{J}$, there is no formula $\varphi \in \Phi$ such that both $\varphi \in F(\boldsymbol{J})$ and $\sim\varphi \in F(\boldsymbol{J})$.
- *Collectively rational* if and only if for all profiles $\boldsymbol{J}$, the collective judgment $F(\boldsymbol{J})$ is complete and consistent.

Next, the idea that a *neutral* aggregation rule advocates is the following: what should matter for the collective outcome is not a proposition itself (for example its semantic value), but only who agrees with it. We provide the dominant definition of this property in the judgment aggregation literature to date (Grandi and Endriss 2013; Grossi and Pigozzi 2014; Endriss 2016a). Note also that, in the literature, neutrality can often be found not as an axiom itself, but as a property that functions in addition to independence, defining the new axiom of *systematicity* (Dietrich and List 2007a; List and Polak 2010).

Definition 1 The aggregation rule F is **neutral** if and only if for all profiles $\boldsymbol{J}$ and all formulas $\varphi, \psi \in \Phi$, it is the case that:

$$N_\varphi^{\boldsymbol{J}} = N_\psi^{\boldsymbol{J}} \text{ implies that } \varphi \in F(\boldsymbol{J}) \Rightarrow \psi \in F(\boldsymbol{J})$$

By symmetry, Definition 1 entails the more common formulation with $\varphi \in F(\boldsymbol{J}) \Leftrightarrow \psi \in F(\boldsymbol{J})$ on the right-hand side of the formal condition.

We include a discussion about a few other reasonable approaches with respect to the definition of the *neutrality* property and we compare them in logical terms in the Appendix.

3 Relaxing neutrality by respecting relative acceptability

In this section we explore a promising direction towards the relaxation of neutrality in judgment aggregation that incorporates the notion of *relative acceptability* of different

propositions, captured by an acceptability relation A (that is, a binary relation on the agenda). We first motivate and then formalise the new family of resulting axioms.

3.1 Motivation

In the introduction we already discussed that various aggregation situations may involve decisions on propositions that are different in nature. We now shed light on several more such cases.

To start off, the propositions that the individuals are asked to judge in a given scenario may be either *subjective* or *objective* in character, meaning that they are or are not associated with a ground truth (Slavkovik 2014). For instance, propositions may express *preferences* rather than *opinions*, as in a case where friends decide together what kind of food to have for dinner. Preferences are subjective. On the other hand, when a committee of doctors need to evaluate whether a patient has cancer or not, their collective judgment will obviously be either right or wrong, depending on reality. It is reasonably fair to ask that when two propositions of the same kind receive the same support from a group, they should be treated equally by an aggregation rule. However, this is not so straightforward when their kinds differ.

Besides what the propositions in a judgment aggregation scenario stand for conceptually, within the formal model they are associated with specific logical formulas. We can thus look at their structural similarities and differences. For instance, propositions of the same kind may be considered those that are logically equivalent or those that make use of the same syntactical connectives (e.g., formulas that are conjunctions of literals, or formulas that do not involve negation). An example of this kind can be found in the work of Porello and Endriss (2014), who discuss the case of judgment aggregation for description logics, where the agenda is naturally separated into two kinds of formulas (referred to as *Tbox* and *Abox*).[7] Moreover, in a conjunctive or disjunctive agenda, the premises also form a natural collection of propositions whose role during aggregation is much more similar to each other than to the conclusion.

Finally, distinctions may be made with respect to the *amount of evidence* different propositions require for their acceptance. For instance, a proposition can be very easy to agree with because the logical formula representing the proposition is true in a large number of models, while another proposition may be easy to reject, because it is true in a smaller set of models. In many cases it makes sense that easily acceptable propositions require at least as much support as easily rejectable ones to be collectively accepted; but the other way around can be reasonable too.

Example 3 Think of a group of movie experts judging the quality of films based on three criteria. Proposition (p_j) "film j passes at least one criterion" asks for less evidence to be satisfied than proposition (q_j) "film j passes all criteria". Suppose now that a subset S of the experts agree with propositions p_a and q_b for two films a and b, and based on the support received by S, the whole group announces that

[7] More generally, *agenda separability* is discussed by Lang et al. (2016). In their framework, it could be sensible to consider as formulas of the same kind those that belong to the same subagenda in an agenda decomposition. Note though that agenda separability relates to the axiom of independence, which is orthogonal to that of neutrality.

film a passes at least one criterion (p_a). But q_b gathered exactly the same support, and, most essentially, the evidence needed for that support was harder to obtain. So, one could argue, the group should also announce that film b passes all criteria, as a matter of fairness for the two films. On the other hand, it could also be argued that for the group to make a logically stronger announcement like q_b, it should be harder (formally requiring more support) than a weaker announcement like p_a. Our new axioms will be able to capture both points of view.

3.2 Formalisation

By looking at different acceptability relations $A \subseteq \Phi \times \Phi$, we will be able to formally describe the various scenarios that triggered our interest, as explained at the beginning of this section. The main definition of this paper follows.

Definition 2 The aggregation rule F is **A-respecting** for a given relation $A \subseteq \Phi \times \Phi$ if and only if for all profiles $\boldsymbol{J}$ and all formulas $\varphi, \psi \in \Phi$ with $\varphi A \psi$, it is the case that:

$$N_{\varphi}^{\boldsymbol{J}} = N_{\psi}^{\boldsymbol{J}} \text{ implies that } \varphi \in F(\boldsymbol{J}) \Rightarrow \psi \in F(\boldsymbol{J})$$

In words, for two propositions φ and ψ, $\varphi A \psi$ should be read as "φ is associated with higher (or equal) standards of acceptance than ψ". Then, an A-respecting aggregation rule F complies with the relative acceptability of different propositions as follows: provided that φ and ψ receive the same support in a given profile, F includes φ in the outcome only if ψ is included as well. Definition 2 directly modifies the standard neutrality axiom (Definition 1) by limiting the scope of its main quantifier. More precisely, instead of talking about *all* formulas $\varphi, \psi \in \Phi$, we now only refer to formulas with $\varphi A \psi$.

Fact 1 *The following are true:*

- *for any relation $A \subseteq A_{id}$ (where $A_{id} = \{(\varphi, \varphi) \mid \varphi \in \Phi\}$ is the identity relation), an A-respecting aggregation rule does not impose any restriction on the treatment of different propositions.*
- *for the total relation $A = \Phi \times \Phi$, an aggregation rule is A-respecting if and only if it is classically neutral.*

The second statement of Fact 1 sheds light on a way of thinking about classical neutrality that is new to the literature: neutrality can now be seen as presupposing the existence of one single relation between all formulas in an agenda (which is not always a well-founded assumption).

Some further fundamental properties of A-respecting aggregation rules are derived from Definition 2 (note that the last one implies the first two):

Fact 2 *For all relations $A_1, A_2 \subseteq \Phi \times \Phi$,*

- *if $A_1 \supseteq A_2$, then every A_1-respecting rule is A_2-respecting.*
- *if a rule is A_1-respecting, then it is also $A_1^{\circlearrowleft}$-respecting, where $A_1^{\circlearrowleft}$ is the reflexive closure of A_1.*

– *if a rule is both A_1-respecting and A_2-respecting, then it is also $(A_1 \cup A_2)$-respecting.*

The second statement (in combination with the first one) of Fact 2 allows us to restrict attention to the reflexive closure of given relations, whenever needed, since this can be done without any impact on the induced axiom. It is worth stressing here that the transitive closure of a relation A does not necessarily induce the same axiom. However, almost all relations corresponding to interesting applications will already be transitive themselves. Note, furthermore, that when the relation A is symmetric and transitive, the corresponding axiom operates on a partition of the agenda and states that formulas in the same equivalence class should be treated equally.

Next, Lemmas 1 and 2 show that the definition of an A-respecting aggregation rule is wide enough to even include the axiom of unanimity, which is simulated for a specific relation A and for any agenda that contains some formula that is a tautology (we denote an arbitrary such formula by $\top$).

Lemma 1 *Consider an agenda Φ with $\top \in \Phi$ and the relation $A = \{(\top, \varphi) \mid \varphi \in \Phi\}$. Then, every collectively rational aggregation rule is unanimous if and only if it is A-respecting.*

Proof Let $A = \{(\top, \varphi) \mid \varphi \in \Phi\}$. A collectively rational aggregation rule F is A-respecting if and only if for every profile $\boldsymbol{J}$ and formula $\varphi \in \Phi$, $\top \in F(\boldsymbol{J})$ implies $\varphi \in F(\boldsymbol{J})$, whenever $N_{\top}^{\boldsymbol{J}} = N_{\varphi}^{\boldsymbol{J}}$. Equivalently (since both the individual and the collective judgments are complete and consistent), F is A-respecting if and only if $\mathcal{N} = N_{\varphi}^{\boldsymbol{J}}$ implies that $\varphi \in F(\boldsymbol{J})$, which is the case if and only if F is unanimous. □

Lemma 2 *Consider an agenda Φ with $\top \in \Phi$ and a relation A with $A \supseteq \{(\top, \varphi) \mid \varphi \in \Phi\}$. Then, every collectively rational A-respecting aggregation rule is unanimous.*

Proof Follows from the first statement of Fact 2 and Lemma 1. □

So, for any agenda that contains a tautology, Lemma 2 tells us that standard neutrality implies unanimity for every collectively rational rule.[8]

Let us now define a *type* $T \subseteq \Phi$ as a subset of the agenda Φ. Then, the (symmetric) relation $A_T = \{(\varphi, \psi) \mid \varphi, \psi \in T\}$ relates all formulas that are of type T and an aggregation rule being A_T-respecting simply means that all formulas of type T will be treated equally. The formulas in an agenda may belong (simultaneously) to different types. Defining suitable types T and making use of A_T-respecting rules, we can model all scenarios described in Sect. 3.1, except for the last one, concerning semantically weaker and stronger formulas. For instance, for all formulas $\varphi_1, \ldots, \varphi_m$ in the agenda Φ, we can take the equivalence relations $A_{\leftrightarrow}^{\varphi_j}$, with $1 \leq j \leq m$, induced by the type $T_{\leftrightarrow}^{\varphi_j} = \{\psi \in \Phi \mid \vDash \varphi_j \leftrightarrow \psi\}$.[9] Then, taking the union of these relations, we can obtain the bigger relation $A_{\leftrightarrow} = A_{\leftrightarrow}^{\varphi_1} \cup \cdots \cup A_{\leftrightarrow}^{\varphi_m}$, according to which any

[8] Endriss et al. (2010) observe a similar fact, namely that in such agendas neutrality together with independence imply unanimity for every complete and complement-free rule (for which at least one collective outcome contains the tautology).

[9] We write $\vDash \varphi \leftrightarrow \psi$ when φ is logically equivalent to ψ.

two logically equivalent formulas in Φ are related.[10] Analogously, we can define A_{pr}, which relates all premises in an agenda to each other using the type $T_{pr} = \{\varphi \in \Phi \mid \varphi \text{ is a premise}\}$.

Finally, by employing a non-symmetric relation, we can also capture the last scenario discussed in Sect. 3.1, where formulas are related via their varying semantical strength. The *semantics-based* relations hinge on comparisons between the (number of) models of the formulas involved. We consider eight such relations. Given an agenda Φ, we denote with $\mathcal{M}_\varphi^\Phi$ the set of models of φ, taking into account all valuations on propositional symbols that appear in formulas of Φ. Whenever the relevant agenda is clear from the context, we will omit the superscript and simply write $\mathcal{M}_\varphi$. For $\star \in \{\geq, >, \leq, <\}$ and $\bullet \in \{\supseteq, \supset, \subseteq, \subset\}$, we define:

$$A_\star = \{(\varphi, \psi) \in \Phi \times \Phi \mid |\mathcal{M}_\varphi| \star |\mathcal{M}_\psi|\}$$
$$A_\bullet = \{(\varphi, \psi) \in \Phi \times \Phi \mid \mathcal{M}_\varphi \bullet \mathcal{M}_\psi\}$$

For example, we have that $A_\geq = \{(\varphi, \psi) \in \Phi \times \Phi \mid |\mathcal{M}_\varphi| \geq |\mathcal{M}_\psi|\}$.

3.3 Examples

To conclude this section, we demonstrate how A-respecting rules—more than neutral rules—are relevant in specific scenarios of practical interest.

Example 4 Suppose that a governmental agenda Φ contains both formulas with important potential consequences (starting a war) and ones with less significance (what the president will have for dinner). Using equivalence relations A we can partition the agenda into classes of formulas of the same type. Then, an aggregation rule that accepts a formula in a given class if its support exceeds a fixed (class-specific) quota is A-respecting but not neutral.

Example 5 Consider a conjunctive agenda Φ with a set of premises $\Phi_{pr} = \{p, q, r, z, \neg p, \neg q, \neg r, \neg z\}$, and the relation $A_{pr} = \{(\varphi, \psi) \mid \varphi, \psi \in \Phi_{pr}\}$. The premise-based rule is A_{pr}-respecting, but it is not neutral in general. Here is a scenario violating neutrality:

	p	q	r	z		$c = p \wedge q \wedge r \wedge z$
Agent 1	Yes	No	Yes	Yes	$\Rightarrow$	No
Agent 2	Yes	Yes	No	Yes	$\Rightarrow$	No
Agent 3	Yes	Yes	Yes	No	$\Rightarrow$	No

In the above profile premise p receives the same support as $\neg c$. However, the premise-based rule will accept p and all other premises, so it will also accept the conclusion c instead of $\neg c$.

[10] The relation $A_\leftrightarrow$ can also be defined more straightforwardly, as the set $\{(\varphi, \psi) \in \Phi \times \Phi \mid \vDash \varphi \leftrightarrow \psi\}$. All collectively rational aggregation rules are in fact $A_\leftrightarrow$-respecting.

Example 6 This example is about the Kemeny rule, which is not resolute. For the purposes of illustration, we consider the following (strong) neutrality axiom (Slavkovik 2014):[11] A non-resolute rule F is neutral if for all profiles $\boldsymbol{J}$ and formulas $\varphi, \psi \in \Phi$, $N_{\varphi}^{\boldsymbol{J}} = N_{\psi}^{\boldsymbol{J}}$ implies that $\varphi \in J \Rightarrow \psi \in J$ for all $J \in F(\boldsymbol{J})$. Remarkably, although its voting counterpart is trivially neutral (Zwicker 2016), in judgment aggregation the Kemeny rule is not neutral even for preference agendas. Consider the following scenario:

$$\begin{aligned}
&\text{Agent 1 } a_1 \succ a_2 \succ a_3 \succ a_4 \succ a_5\\
&\text{Agent 2 } a_1 \succ a_2 \succ a_3 \succ a_4 \succ a_5\\
&\text{Agent 3 } a_2 \succ a_1 \succ a_4 \succ a_3 \succ a_5\\
&\text{Agent 4 } a_2 \succ a_1 \succ a_5 \succ a_4 \succ a_3
\end{aligned}$$

The Kemeny outcomes are $a_1 \succ a_2 \succ a_3 \succ a_4 \succ a_5$, $a_1 \succ a_2 \succ a_4 \succ a_3 \succ a_5$, $a_2 \succ a_1 \succ a_3 \succ a_4 \succ a_5$, $a_2 \succ a_1 \succ a_4 \succ a_3 \succ a_5$. Note that in the profile $\boldsymbol{J}$ above, we have that $N^{\boldsymbol{J}}_{p_{a_1 \succ a_2}} = N^{\boldsymbol{J}}_{p_{a_3 \succ a_4}}$, but there is an outcome J with $p_{a_1 \succ a_2} \in J$ and $p_{a_3 \succ a_4} \notin J$.

Now, consider a formula φ in an agenda Φ such that neither φ nor $\sim\varphi$ are logically implied by any consistent subset of the agenda that does not contain them. Define the type T_{ind} containing all such formulas and the relation A_{ind} associated with it. Formally, $T_{ind} = \{\varphi \in \Phi \mid \varphi \text{ and } \sim\varphi \text{ are logically independent of } S \backslash \{\varphi, \sim\varphi\}, \text{ for all consistent } S \subseteq \Phi\}$. Recall that a formula φ is logically independent of a set of formulas S if both $S \cup \{\varphi\}$ and $S \cup \{\sim\varphi\}$ are logically consistent. The Kemeny rule just follows the majority's opinion on propositions in T_{ind}, trying to maximise agreement with them, and hence it must be A_{ind}-respecting.

4 (Im)possibilities

Recall Example 1 in the introduction, where the governor of Arizona had to disregard neutrality if he wanted to make a decision in accordance with his majority standards concerning building a wall and increasing college tuition. In fact, we can find numerous such problematic scenarios. To make matters worse, it is known that there exists no collectively rational aggregation rule that is simultaneously unanimous, anonymous, neutral, and independent (List and Pettit 2002).[12] So, we have to compromise on some of these properties. But collective rationality simply requires a consistent collective judgment about all propositions, and unanimity is a highly appealing axiom. Anonymity is a basic requirement we cannot compromise on in many applications either, be it out of fundamental fairness considerations or because we have no means for ascribing different levels of expertise to different individuals. Weakening the more debatable property of independence has already been examined in prior work (Mongin

[11] The axiom of neutrality for non-resolute rules certainly deserves further investigation, but this is beyond the scope of this paper.

[12] To be precise, List and Pettit (2002) proved a slightly stronger result, without the unanimity axiom.

2008; Grossi and Pigozzi 2014). In this section we put neutrality into the spotlight, and ask:

> *For a given agenda Φ, what relations A guarantee the existence of a unanimous, anonymous, A-respecting, independent, and collectively rational aggregation rule?*

We answer this question for the family of conjunctive agendas.[13]

4.1 General characterisation

We show that the relations associated with impossibility results are exactly those that have a non-empty intersection with a specific relation A^*. For a given agenda Φ, we define:

$$A^* = \{(\varphi, \psi) \in \Phi^- \times \Phi^+ \mid \varphi \wedge \psi \text{ and } \sim\varphi \wedge \sim\psi \text{ are consistent}\}$$

The relation A^* may be thought of as "mixing" between positive and negative formulas, which can cause trouble if in addition those formulas are mutually consistent (i.e., can be satisfied together). To illustrate, consider the following example starring a conjunctive agenda, inspired by the famous *doctrinal paradox* of Kornhauser and Sager (1993).

Example 7 A committee of three reviewers, Alice, Bob, and Carol, are judging the work of a colleague in terms of excellence. The work can be rendered excellent (e) if and only if it is both novel (φ) and sound (ψ). That is, the propositions φ and ψ are the premises towards the conclusion $e = (\varphi \wedge \psi)$. In this setting, we have that:

$$A^* = \{(\neg\varphi, \psi), (\neg\psi, \varphi)\}$$

Note that no other pair of formulas can be in A^*. For instance, $(\neg(\varphi \wedge \psi), \varphi) \notin A^*$ because $\varphi \wedge \psi$ is not consistent with $\neg\varphi$.

Figure 1 instantiates our findings. Before we are able to prove our main theorem, we will need a lemma characterising all situations where the unanimity rule is A-respecting:

Lemma 3 *For a conjunctive agenda Φ, the unanimity rule $F_{\mathcal{N}}$ is A-respecting if and only if $A \cap A^* = \varnothing$.*

Proof ($\Rightarrow$) We show the contrapositive. Suppose that $A \cap A^* \neq \varnothing$. Then, there are two mutually consistent formulas $\varphi \in \Phi^-, \psi \in \Phi^+$ with $\sim\varphi, \sim\psi$ also mutually consistent, such that $\varphi A \psi$. We can then construct a profile $\boldsymbol{J}$ where one individual accepts both φ and ψ and everyone else accepts both $\sim\varphi$ and $\sim\psi$. In such a case, the unanimity rule will include φ in the outcome (because $\varphi \in \Phi^-$) but not ψ. But this violates neutrality since $N_\varphi^{\boldsymbol{J}} = N_\psi^{\boldsymbol{J}}$ and $\varphi A \psi$.

[13] Our results immediately extend to disjunctive agendas as well: simply flip Φ^+ and Φ^- in the relevant statements.

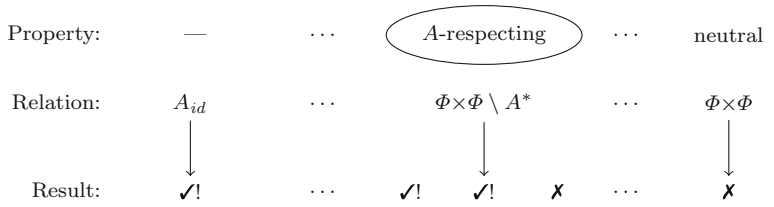

Fig. 1 Every column describes an (im)possibility result, regarding the existence of an aggregation rule that has the relevant property of the first row and is also unanimous, anonymous, independent, and collectively rational. The corresponding relations in the middle row ordered from smaller to larger according to set inclusion. The non-existence of a desirable rule is indicated by ✗, while ✓! indicates a unique existence

($\Leftarrow$) Observe that the *only* way for the unanimity rule to not be A-respecting is by having a negative and a positive formula $\varphi \in \Phi^-$ and $\psi \in \Phi^+$ with $\varphi A \psi$, that receive exactly the same support in a profile $\boldsymbol{J}$, but this support does not come from all the individuals. Then, as unanimity rule is defined, φ will belong to the collective outcome but ψ will not. Furthermore, since individual judgments are complete and consistent and $n \geq 2$, in order to be able to construct such a profile $\boldsymbol{J}$ with $N_\varphi^{\boldsymbol{J}} = N_\psi^{\boldsymbol{J}} \neq \mathcal{N}$, we must have that both pairs of formulas φ, ψ and $\sim\varphi$, $\sim\psi$ are mutually consistent. All of these conditions can be satisfied when $A \cap A^* \neq \varnothing$. □

Theorem 1 *For a conjunctive agenda, there exists a unanimous, anonymous, A-respecting, independent, and collectively rational aggregation rule if and only if $A \cap A^* = \varnothing$. Moreover, that rule is unique.*

Proof By Theorem 5.1 due to Dokow and Holzman (2009), we know that the only independent, unanimous, and collectively rational aggregation rules for a conjunctive agenda are the oligarchies. Adding anonymity, we deduce that the only rule satisfying the above properties is the unanimity rule. Now, from Lemma 3 we also have that the unanimity rule is A-respecting if and only if $A \cap A^* = \varnothing$, which completes the proof. □

Theorem 1 can be compared to the impossibility result of List and Pettit (2004). Specifically, we obtain the known impossibility (for the somewhat restricted case of conjunctive agendas) by simply replacing A with the total relation. Moreover, Theorem 1 equips us with a powerful tool we can utilise in different scenarios. For instance, suppose that a group of individuals has to make a decision on a number of propositions. They want to employ an aggregation rule that is unanimous, anonymous, independent, and of course collectively rational. But such a rule can never be neutral, even though fair treatment of the decisions to be made is important for the group as well. A solution rests in the way the propositions may be related to each other with regard to their standards of acceptance—a suitable A-respecting rule may exist, which is all the individuals need.

In the remainder of this section we will see how we can use Theorem 1 to obtain both impossibility and possibility results for specific relations of interest, namely those motivated in Sect. 3. For a given relation A, we ask whether there exists an aggregation rule that is unanimous, anonymous, A-respecting, independent, and collectively

Table 1 (Im)possibility results for aggregation rules that respect relations of practical interest (and that are also unanimous, anonymous, independent, and collectively rational)

Relation	A_{pr}	$A_{\geq}$	$A_{\leq}$	$A_{>}$	$A_{<}$	$A_{\supset}$	$A_{\supseteq}$	$A_{\subset}$	$A_{\subseteq}$
Result	✗	✗	✗	✓!	✓!	✓!	✓!	✓!	✓!

rational. Of course, even when we provide a positive answer, we have to keep in mind (from the proof of Theorem 1) that the rule satisfying all our axioms will be the—arguably disputable—unanimity rule. Note, though, that this makes our findings even more interesting: that this specific rule is the only one circumventing the impossibility is worth showing. Our results are summarised in Table 1.

4.2 Specific impossibilities

We now enlarge the corpus of impossibility theorems in the judgment aggregation literature (List and Pettit 2004; Pauly and van Hees 2006; Dietrich 2007; van Hees 2007) simply by instantiating Theorem 1.

Proposition 1 *For a conjunctive agenda, there exists no unanimous, anonymous, $A_{\geq}$(or $A_{\leq}$)-respecting, independent, and collectively rational aggregation rule.*

Proof Consider a conjunctive agenda Φ with $\neg p, q$ two of its premises. Then, take $\varphi = \neg p \in \Phi^-$ and $\psi = q \in \Phi^+$ and note that it holds that $\varphi A^* \psi$, since both $\neg p \wedge q$ and $p \wedge \neg q$ are consistent. But we also have that $\varphi A_{\geq} \psi$ (and $\varphi A_{\leq} \psi$), because $|\mathcal{M}_{\neg p}| = |\mathcal{M}_q|$. Thus, $(\neg p, q) \in A \cap A^* \neq \varnothing$, and the statement follows by Theorem 1. □

Proposition 1 makes evident that relaxing neutrality in a manner that still asks for propositions with the same number of models to be treated equally cannot circumvent the impossibility of List and Pettit (2004) associated with the standard neutrality axiom. Moreover, for any agenda Φ that contains a tautology, we obviously have that $(\top, \varphi) \in A_{\geq}$ for all formulas $\varphi \in \Phi$. Thus, by Lemma 2, we know that—for agendas with tautologies—the impossibility of Proposition 1 holds even if we drop the axiom of unanimity from our requirements. Nonetheless, in Sect. 4.3 we will see that in order to achieve a possibility result, we only need to weaken our demands one step further.

Next, Proposition 2 shows that, unfortunately, requesting a neutral aggregation rule only with respect to the premises in a conjunctive agenda—although being a very natural requirement—also leads to an impossibility.

Proposition 2 *For a conjunctive agenda, there exists no unanimous, anonymous, A_{pr}-respecting, independent, and collectively rational aggregation rule.*

Proof Analogous to the proof of Proposition 1. □

4.3 Specific possibilities

On the positive side, using Theorem 1 we can also derive several possibility results. Specifically, we will prove in Proposition 3 that we can have an A-respecting neutral

aggregation rule for all our semantics-based axioms—except for those that relate propositions with the same number of models (they are covered by Proposition 1).

Proposition 3 *For a conjunctive agenda, there exists a (unique) unanimous, anonymous, A-respecting, independent, and collectively rational aggregation rule for any relation* $A \in \{A_>, A_<, A_\geq, A_\subseteq, A_\supset, A_\subset\}$.

Proof For all relations A in the statement, the proof is similar. Consider a conjunctive agenda Φ. By Theorem 1, we need to check that $A \cap A^* = \varnothing$, or equivalently that there are no formulas $\varphi \in \Phi^-$, $\psi \in \Phi^+$ such that $\varphi A \psi$, φ, ψ are mutually consistent, and $\sim\varphi$, $\sim\psi$ also are mutually consistent. Let us sketch the proof for $A = A_>$. It suffices to note that we can only have $\varphi A_> \psi$ for formulas $\varphi \in \Phi^-$ and $\psi \in \Phi^+$ in three cases: (i): $\varphi = \neg c$ and $\psi = c$ (so $|\mathcal{M}_{\neg c}| > |\mathcal{M}_c|$); (ii) $\varphi = \neg c$ and ψ is a premise p (so $|\mathcal{M}_{\neg c}| > |\mathcal{M}_p|$); (iii): $\varphi = \neg p$ for some premise p and $\psi = c$ (so $|\mathcal{M}_{\neg p}| > |\mathcal{M}_c|$). But in the first case it is obviously not true that φ and ψ are mutually consistent. In the second case, we have that $\sim\varphi = c$ and $\sim\psi = \neg p$ are not mutually consistent, and similarly in the last case, $\varphi = \neg p$ and $\psi = c$ are not mutually consistent either (by definition of a conjunctive agenda). This means that $A_> \cap A^* = \varnothing$. □

Recall once more that the unique rule the existence of which is established by Proposition 3 is the unanimity rule.

5 Safety of the agenda

In this section we address a different problem and are interested in knowing for which exact combinations of decision-making domains and A-respecting rules paradoxical scenarios involving logical inconsistencies can *never* occur.

Consider a group of individuals that engage in collective decision making, like the members of a company's board during one of their meetings. We may be uncertain about the specific aggregation rule that the individuals will use to reach a collective judgment, and we may also be ignorant about their precise judgments in a given scenario. But what if we know that, valuing highly the fair treatment of propositions of the same type, the board will always employ some A-respecting rule? Can we then guarantee that—no matter what the individuals' individual judgments actually are—the collective outcome will always be consistent? More broadly, we ask the following question:

> *Given an aggregation domain (captured by an agenda* Φ *with an associated relation* $A \subseteq \Phi \times \Phi$*), which A-respecting rules provide a safety net against inconsistent outcomes?*

The above question corresponds to the *safety of the agenda* problem in judgment aggregation (Endriss et al. 2012).

Definition 3 An agenda Φ is called **safe** for an aggregation rule F if, for all profiles $\boldsymbol{J}$, the outcome $F(\boldsymbol{J})$ is consistent.

For the remainder, we assume that the agenda Φ does not contain any tautologies.[14] We also stress that in this section we examine all possible agendas (not only conjunctive or disjunctive ones). Let us now define some relevant notions. We say that a subset $S \subseteq \Phi$ of the agenda is *minimally inconsistent* (mi) if (i) it is inconsistent and (ii) all its strict subsets $S' \subset S$ are consistent. For example, the set $\{\varphi, \psi, \neg(\varphi \wedge \psi)\}$ is minimally inconsistent, since by removing any single formula from it, it becomes consistent.

An agenda Φ has the *Median Property* (MP) if all mi subsets of Φ are of size two (Nehring and Puppe 2007; List and Puppe 2009a). For instance, the agenda Φ of Example 7 with $\Phi^+ = \{\varphi, \psi, \varphi \wedge \psi\}$ does not have the MP, since as we saw above, it contains a mi subset of size three. Furthermore, Φ has the *Simplified Median Property* (SMP) if, for all mi subsets $\{\varphi, \psi\}$ it holds that $\vDash \varphi \leftrightarrow \neg\psi$; and Φ has the *Syntactic Simplified Median Property* (SSMP) if all its mi subsets $\{\varphi, \psi\}$ are of the form $\{\varphi, \sim\varphi\}$ for some $\varphi \in \Phi$ (Endriss et al. 2012). An example of an agenda Φ that has the MP but not the SMP is the one with $\Phi^+ = \{\varphi, \varphi \wedge \psi\}$: note that $\{\neg\varphi, \varphi \wedge \psi\}$ is a mi subset of size two but it does not hold that $\vDash \varphi \leftrightarrow (\varphi \wedge \psi)$. On the other hand, for the agenda with $\Phi^+ = \{\varphi, \varphi \wedge (\psi \vee \neg\psi)\}$, the SMP holds but the SSMP does not.

To begin with, note that if a rule is not complement-free, then it is also not consistent, for any agenda. So, it is sensible to restrict attention to complement-free rules. From Endriss et al. (2012) we learn that an agenda Φ is safe for all neutral and complement-free aggregation rules if and only if Φ has the SMP.[15] In the sequel, we will in addition provide a uniform characterisation of the agendas that are safe for all complement-free and A-respecting rules, based on a notion that we call the *A-Median Property* (A-MP).

5.1 The *A*-median property

Definition 4 The agenda Φ has the **A-Median Property (A-MP)** for some relation $A \subseteq \Phi \times \Phi$ if all its mi subsets are of the form $\{\varphi, \psi\}$, where for some $\chi \in \Phi$:

$$\vDash \varphi \leftrightarrow \chi \quad \text{and} \quad \varphi A \chi \quad \text{as well as} \quad \vDash \psi \leftrightarrow \neg\chi \quad \text{and} \quad \psi A \sim\chi.$$

In words, the A-MP says all the agenda's mi subsets have to contain exactly two formulas that are related via A to two complementary formulas.

Example 8 Consider the agenda $\Phi = \{p, p \wedge p, \neg p, \neg(p \wedge p)\}$ and denote by $c(\varphi)$ the number of connectives that appear in the formula $\varphi \in \Phi$. Then, consider the relations $A = \{(\varphi, \psi) \in \Phi \times \Phi \mid |c(\varphi) - c(\psi)| \leq 1\}$ and $A' = \{(\varphi, \psi) \in \Phi \times \Phi \mid c(\varphi) = c(\psi) = 1\}$. That is, A relates formulas that are close regarding their number of connectives, while A' simply relates formulas with exactly one connective. In practical applications these relations could capture interesting distinctions between formulas, like in the case, for example, where we would want to talk about how "easy" a formula is for a human to grasp, thereby assuming that the number of connectives in a formula

[14] This assumption is crucial for our formal results—without it, the only-if direction of our characterisation theorem would not hold.

[15] Specifically, Theorem 17 of Endriss et al. (2012) concerns all complete, anonymous, unanimous, complement-free, and neutral rules; however, their proof implies also the result we are interested in here, regarding only complement-free and neutral rules.

tells us something about how complicated the formula is (which, of course, is not always realistic). Note that A relates all semantically equivalent formulas in Φ, and hence Φ has the A-MP if and only if it has the SMP. Moreover, it is trivial to check that Φ satisfies the SMP. However, we have that the formulas p and $p \wedge p$ are semantically equivalent, but it does not hold that $pA'(p \wedge p)$ (and neither that $\neg pA'\neg(p \wedge p)$). By examining the mi set $\{p, \neg(p \wedge p)\}$, we can then conclude that Φ does not have the A'-MP.

Definition 4 implies Fact 3.

Fact 3 *For all agendas Φ and all relations $A_1, A_2 \subseteq \Phi \times \Phi$, such that $A_1 \subseteq A_2$, the A_1-MP implies the A_2-MP.*

In order to compare the A-MP to the properties of the SMP and of the SSMP, we note that for all reflexive relations A, the SSMP implies the A-MP, which implies the SMP (Fact 4): When an agenda has the SSMP, each of its mi subsets is of the form $\{\varphi, \sim\varphi\}$. So, we can take $\chi = \varphi$ and see that, for a reflexive relation A, Definition 4 holds. Even more straightforwardly, if an agenda has the A-MP, Definition 4 implies that all its mi subsets are of the form $\{\varphi, \psi\}$ with $\vDash \varphi \leftrightarrow \chi$ and $\vDash \chi \leftrightarrow \neg\psi$ for some $\chi \in \Phi$. Thus, we have that $\vDash \varphi \leftrightarrow \neg\psi$, and the SMP is satisfied.

Fact 4 *For all agendas Φ and all reflexive relations $A \subseteq \Phi \times \Phi$,*

- *the SSMP implies the A-MP.*
- *the A-MP implies the SMP.*

Moreover, by looking carefully at Definition 4, the following is immediate:

Fact 5 *For all agendas Φ, the SMP is equivalent to the $A_{\leftrightarrow}$-MP.*

Using Facts 3, 4, and 5, Corollary 1 clarifies that the A-MP differs from the SMP only for relations A that differentiate between (at least some) semantically equivalent formulas in the agenda.

Corollary 1 *For all agendas Φ and all relations $A \subseteq \Phi \times \Phi$ such that $A_{\leftrightarrow} \subseteq A$, the SMP is equivalent to the A-MP.*

Note that in general, the weaker the axioms we require for a rule F to satisfy, the more restricting the properties of the agendas that are safe for F.

5.2 The characterisation result

Having familiarised ourselves with the A-MP, we are ready to provide the main result of this section. Theorem 2 generalises existing results regarding the safety of the agenda, which only examine the extreme cases of standard neutrality and of no neutrality at all, corresponding to the agenda properties SMP and SSMP, respectively (Endriss et al. 2012).

Note that the safety of the agenda problem that we address in this paper can also be seen as a "universal agenda characterisation" problem (Endriss 2016b). That is, we ask what are the conditions on the agenda for which *all* aggregation rules satisfying

certain axioms will be consistent. A slightly different question corresponds to what is sometimes called an "existential agenda characterisation" problem. In that setting, we are interested in finding what are the conditions on the agenda for which *there exists* some aggregation rule satisfying certain axioms that is also consistent. Work on the latter problem has been done, e.g., by Dokow and Holzman (2010) and several results if this kind have been reviewed by List and Puppe (2009b).

Recall that n is the number of individuals.[16]

Theorem 2 *Consider a reflexive and transitive relation $A \subseteq \Phi \times \Phi$. An agenda Φ with $|\Phi| \leq n$ is safe for all A-respecting and complement-free rules if and only if Φ has the A-MP.*

Proof ($\Rightarrow$) We show the contrapositive. First, note that if the agenda Φ does not have the SMP, then we know that there is some neutral rule (hence also A-respecting) for which Φ is not safe (Endriss et al. 2012). Thus, let us assume that Φ has the SMP, but does not have the A-MP. This means that Definition 4 does not hold. So, there exists a mi subset $\{\varphi, \psi\} \subseteq \Phi$ for which there is no $\chi \in \Phi$ with $\vDash \varphi \leftrightarrow \chi$ and $\varphi A \chi$ and $\vDash \psi \leftrightarrow \neg\chi$ and $\psi A {\sim}\chi$.

Let us define the (possibly empty) set $X = \{\chi \in \Phi \mid \varphi A \chi \text{ and } \psi A {\sim}\chi\}$.[17] For every $\chi \in X$ it is the case that $\nvDash \varphi \leftrightarrow \chi$ or $\nvDash \psi \leftrightarrow \neg\chi$. In other words, for every $\chi \in X$ we can define the nonempty set of models $\mathcal{M}(\chi) = \{M \mid M \nvDash \varphi \leftrightarrow \chi \text{ or } M \nvDash \psi \leftrightarrow \neg\chi\}$.

Now, according to our assumption, we have a large enough group of individuals $\mathcal{N}$, with $|\mathcal{N}| \geq |X|$. We can then construct a profile $\boldsymbol{J} = (J_1, \ldots, J_i, \ldots, J_n)$ as follows. First, we consider an enumeration of $X = \{\chi_1, \chi_2, \ldots, \chi_k\}$. Then, for each individual $i \leq |X|$, we choose arbitrarily one of the models $M_i \in \mathcal{M}(\chi_i)$ and we take $J_i = \{\varphi \in \Phi \mid M_i \vDash \varphi\}$. In this way, we have established that for every formula $\chi_i \in X$ there is some individual i who cannot accept simultaneously the formulas $\varphi, \chi, {\sim}\psi$, neither the formulas ${\sim}\varphi, {\sim}\chi, \psi$. For the individuals i with $i > |X|$, we can assume they submit any judgment set—this assumption does not affect the proof.

Next, consider the aggregation rule F that maps the profile $\boldsymbol{J}$ to the set $\{\varphi, \psi\} \cup \{\chi \in \Phi \mid \varphi A \chi \text{ and } N_\varphi^{\boldsymbol{J}} = N_\chi^{\boldsymbol{J}}\} \cup \{\chi \in \Phi \mid \psi A \chi \text{ and } N_\psi^{\boldsymbol{J}} = N_\chi^{\boldsymbol{J}}\}$, and all the other profiles to the empty set. By the definition of F, and recalling that the relation A is transitive, it is straightforward to see that F is A-respecting. In addition, we can show that F is complement-free. First, note that the set $\{\varphi, \psi\}$ is already complement-free. Indeed, if (without loss of generality) we had that $\psi = {\sim}\varphi$, then we would also have that $\varphi A \varphi$, and that $\psi A {\sim}\varphi$ (because A is reflexive), $\vDash \varphi \leftrightarrow \varphi$ and $\vDash \psi \leftrightarrow \neg\varphi$. But this would imply that Φ also has the A-MP, which contradicts our assumption. Second, we can show that it is not possible to have $\{\chi, {\sim}\chi\} \subseteq F(\boldsymbol{J})$ for any $\chi \in \Phi$, since this would require that $N_\varphi^{\boldsymbol{J}} = N_\chi^{\boldsymbol{J}}$ and $N_\psi^{\boldsymbol{J}} = N_{\sim\chi}^{\boldsymbol{J}}$ for some $\chi \in X$. But by the construction of the profile $\boldsymbol{J}$, we cannot have that $N_\varphi^{\boldsymbol{J}} = N_\chi^{\boldsymbol{J}} = N_{\sim\psi}^{\boldsymbol{J}}$ (because there is some individual i with $i \leq |X|$ who cannot accept simultaneously the formulas $\varphi, \chi, {\sim}\psi$, neither the formulas ${\sim}\varphi, {\sim}\chi, \psi$). However, $F(\boldsymbol{J})$ will obviously be inconsistent because $\{\varphi, \psi\}$

[16] Theorem 2 holds under two mild assumptions that very naturally apply in a number of realistic scenarios: that the size n of the group is sufficiently large, and that the propositions are related in a transitive manner.

[17] The proof holds both for the case where $X = \varnothing$ and for the case where $X \neq \varnothing$.

is inconsistent, which means that the agenda Φ is not safe for all A-respecting and complement-free rules.

($\Leftarrow$) Assume that Φ has the A-MP and (aiming for a contradiction) that there exist an A-respecting and complement-free rule, as well as a profile $\boldsymbol{J}$ such that $F(\boldsymbol{J})$ is inconsistent. That is, $F(\boldsymbol{J}) \supseteq \{\varphi, \psi\}$, for some mi set $\{\varphi, \psi\}$. As individual judgment sets are complete and consistent, using the formula χ that appears in the definition of the A-MP, it must hold that $N_{\varphi}^{\boldsymbol{J}} = N_{\chi}^{\boldsymbol{J}}$ and $N_{\psi}^{\boldsymbol{J}} = N_{\sim\chi}^{\boldsymbol{J}}$. So F being A-respecting implies that $\chi, \sim\chi \in F(\boldsymbol{J})$, contradicting complement-freeness. □

To emphasise the significance of Theorem 2, we apply it to obtain new safety results. Interestingly, we find that relaxing neutrality to any non-strict semantics-based neutrality axiom has no effect on the conditions that make an agenda safe, and the same holds when moving from the complete absence of neutrality to any strict semantics-based axiom (Propositions 4 and 5).

Proposition 4 *Let $A \in \{A_{\geq}, A_{\leq}, A_{\supseteq}, A_{\subseteq})\}$. An agenda Φ with $|\Phi| \leq n$ is safe for all A-respecting and complement-free aggregation rules if and only if Φ has the SMP.*

Proof To be able to apply Theorem 2, we need to show that the $A_{\geq}$-MP (and accordingly for the $A_{\leq}$-MP, the $A_{\supseteq}$-MP, and the $A_{\subseteq}$-MP) is equivalent to the SMP. Since we have that $A_{\leftrightarrow} \subseteq A_{\geq}$, this follows from Corollary 1, but we will also provide a direct proof here. We know that the $A_{\geq}$-MP implies the SMP, so we only need to prove that it is also implied by it. Consider an agenda Φ with the SMP: all its mi subsets are of the form $\{\varphi, \psi\}$, where $\vDash \varphi \leftrightarrow \neg\psi$. So, we have that $\mathcal{M}_{\psi} = \mathcal{M}_{\neg\varphi}$, and hence $|\mathcal{M}_{\psi}| = |\mathcal{M}_{\sim\varphi}|$, which implies that $\psi A_{\geq} \sim\varphi$. Then, taking $\chi = \varphi$ in Definition 4, we see that the $A_{\geq}$-MP is satisfied. □

Proposition 5 *Let $A \in \{A_{>}, A_{<}, A_{\supset}, A_{\subset})\}$. An agenda Φ with $|\Phi| \leq n$ is safe for all A-respecting and complement-free aggregation rules if and only if Φ has the SSMP.*

Proof Again, in order to apply Theorem 2, we need to show that the $A_{>}$-MP (and accordingly for the $A_{<}$-MP, the $A_{\supset}$-MP, and the $A_{\subset}$-MP) implies the SSMP (since we already know that it is also implied by it). Take an agenda Φ with the $A_{>}$-MP: all its mi subsets are of the form $\{\varphi, \psi\}$, where $\vDash \varphi \leftrightarrow \chi$, $\vDash \psi \leftrightarrow \neg\chi$, $\varphi A_{>}^{\circlearrowleft} \chi$, and $\psi A_{>}^{\circlearrowleft} \sim\chi$ for some $\chi \in \Phi$. Now, $\varphi A_{>}^{\circlearrowleft} \chi$ means that either $\varphi = \chi$ or $|\mathcal{M}_{\varphi}| > |\mathcal{M}_{\chi}|$. But $\vDash \varphi \leftrightarrow \chi$ implies that $\mathcal{M}_{\varphi} = \mathcal{M}_{\chi}$, so it must hold that $\varphi = \chi$. Analogously, from $\psi A_{>}^{\circlearrowleft} \sim\chi$ and $\vDash \psi \leftrightarrow \neg\chi$ we conclude that $\psi = \sim\chi$. Hence, $\varphi = \sim\psi$ and the SSMP holds. □

Figure 2 summarises the results of this section.

6 Conclusion

We have introduced a new family of axioms in judgment aggregation that hinge on the various acceptability relations that may exist between propositions. These axioms—relaxing the classical neutrality property—are broadly applicable and provide relevant insights into a number of scenarios revolving around collective decision

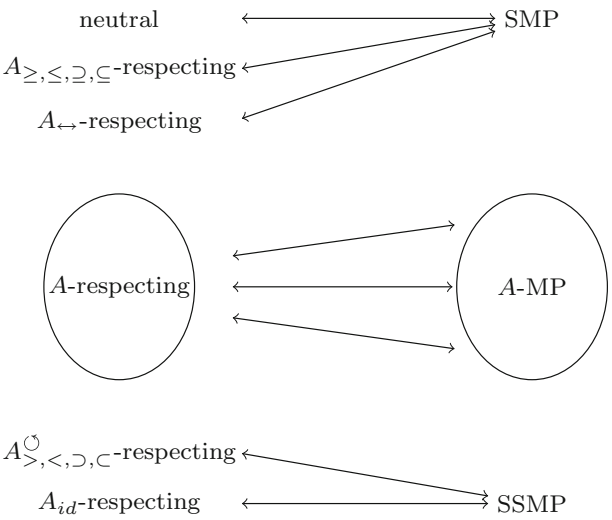

Fig. 2 On the right side we see the requirements that are imposed on the agenda if we want to guarantee its safety (i.e., consistent outcomes) for all complement-free aggregation rules that satisfy the properties appearing on the left side

making. Besides the conceptual analysis of the neutrality axiom we have conducted, we have also investigated two classes of important technical problems, providing uniform answers that generalise existing results in judgment aggregation. First, we have shown exactly how weakening neutrality affects known impossibility results, and second, we have characterised all aggregation domains for which the new axioms do not cause any inconsistencies.

Our work opens up many directions for future research. Does weakening neutrality affect other technical results as well? For instance, instead of asking what domains are safe for *all* A-respecting neutral rules, we can ask which are safe for *at least one* such rule. Are there judgment aggregation rules, beyond the ones presented in this paper, that are not neutral, but are A-respecting with regard to some interesting relation A? And more generally, what other relations can we design, and what application domains can be modelled by them? All these are intriguing questions one can explore further.

Appendix: Variants of neutrality in judgment aggregation

In this Appendix we collect and present all the alternative definitions that may be used to capture the notion of neutrality in the framework of judgment aggregation, and we discuss the logical relations between them. We start with repeating the standard definition which we have also used in this paper.

Definition 5 The aggregation rule F satisfies **neutrality (N)** if for all profiles $\boldsymbol{J}$ and all formulas $\varphi, \psi \in \Phi$, it is the case that:

$$N_{\varphi}^{\boldsymbol{J}} = N_{\psi}^{\boldsymbol{J}} \text{ implies that } \varphi \in F(\boldsymbol{J}) \Rightarrow \psi \in F(\boldsymbol{J})$$

Another appealing property is inspired by the way neutrality is usually defined in the context of voting (Brams and Fishburn 2002; Zwicker 2016), requiring invariable outcomes under permutations. Any permutation $\pi : \Phi \to \Phi$ of the agenda induces a mapping on judgment sets as usual, i.e., by letting $\pi(J_i) = \{\pi(\varphi) \mid \varphi \in J_i\}$ for all judgment sets J_i. Note that even if J_i is consistent, $\pi(J_i)$ could be inconsistent. For example, consider the judgment set $J_i = \{\varphi, \varphi \wedge \psi\}$ and the permutation π with $\pi(\varphi) = \neg\varphi$ and $\pi(\varphi \wedge \psi) = \varphi \wedge \psi$. Then, $\pi(J_i) = \{\neg\varphi, \varphi \wedge \psi\}$. For any profile $\boldsymbol{J} = (J_1, \ldots, J_n)$, let $\pi(\boldsymbol{J}) = (\pi(J_1), \ldots, \pi(J_n))$.

Definition 6 The aggregation rule F satisfies **permutation neutrality (PN)** if, for all profiles $\boldsymbol{J} = (J_1, \ldots, J_n)$ and permutations $\pi : \Phi \to \Phi$ such that $\pi(J_i)$ is consistent for every J_i, with $1 \leq i \leq n$, it holds that:

$$F(\pi(\boldsymbol{J})) = \pi(F(\boldsymbol{J}))$$

The third neutrality axiom only applies when judgment aggregation simulates preference aggregation. Consider a preference agenda Φ based on a set of alternatives $\mathcal{A}$. We say that the permutation $\pi' : \mathcal{A} \to \mathcal{A}$ *naturally induces* the permutation $\pi : \Phi \to \Phi$ if $\pi(p_{a \succ b}) = p_{\pi(a) \succ \pi(b)}$ for every $a, b \in \mathcal{A}$ and $p_{a \succ b} \in \Phi$. For example, for three alternatives a, b, c where $\pi'(a) = b$ and $\pi'(b) = c$, it must be the case that $\pi(p_{a \succ b}) = p_{b \succ c}$.

Definition 7 Consider a preference agenda Φ based on alternatives in $\mathcal{A}$. The aggregation rule F satisfies **alternative-permutation neutrality (APN)** if, for all profiles $\boldsymbol{J}$ and permutations $\pi' : \mathcal{A} \to \mathcal{A}$ that naturally induce permutations $\pi : \Phi \to \Phi$, it holds that:

$$F(\pi(\boldsymbol{J})) = \pi(F(\boldsymbol{J}))$$

How do the different neutrality axioms logically relate to each other? Figure 3 illustrates graphically our findings, and we proceed with the proofs.

Proposition 6 *PN implies N.*

Proof Consider an arbitrary rule F that satisfies PN, a profile $\boldsymbol{J}$, and two formulas $\varphi, \psi \in \Phi$ with $N_{\varphi}^{\boldsymbol{J}} = N_{\psi}^{\boldsymbol{J}}$. Suppose that $\varphi \in F(\boldsymbol{J})$. To show that N holds, we will

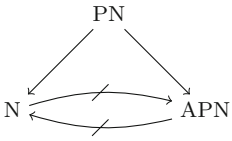

Fig. 3 Logical relations between the different variants of neutrality in judgment aggregation. An arrow points from the logically stronger axiom to the logically weaker one

deduce that $\psi \in F(\boldsymbol{J})$. Take the permutation $\pi : \Phi \to \Phi$ with $\pi(\varphi) = \psi, \pi(\psi) = \varphi$, and $\pi(\alpha) = \alpha$ for all $\alpha \neq \varphi, \psi$. Since $N^{\boldsymbol{J}}_{\varphi} = N^{\boldsymbol{J}}_{\psi}$, for every judgment set J_i in the profile $\boldsymbol{J}$ we have that $\pi(J_i) = J_i$, so $\pi(J_i)$ is consistent. Hence, F satisfying PN implies that $F(\pi(\boldsymbol{J})) = F(\boldsymbol{J})$. But $\psi = \pi(\varphi) \in \pi(F(\boldsymbol{J}))$ by definition. We conclude that $\psi \in F(\boldsymbol{J})$. □

Proposition 7 *PN implies APN (for any preference agenda Φ).*

Proof It is immediate that PN implies APN, as every permutation $\pi' : \mathcal{A} \to \mathcal{A}$ naturally induces a permutation $\pi : \Phi \to \Phi$ that satisfies the condition required by PN (intuitively, for individual preferences being linear orders, any permutation over the alternatives will retain the linear orders and consistency is preserved). □

So, permutation neutrality is logically stronger than both standard neutrality and (for preference agendas) alternative-permutation neutrality. Moreover, we show that there is no other logical relation between these axioms by constructing relevant counterexamples.[18]

Example 9 *N does not imply APN* (so *N does not imply PN*): Consider the preference agenda $\Phi = \{p_{a\succ b}, p_{b\succ a}\}$ for a set of alternatives $\mathcal{A} = \{a, b\}$ and the rule F such that for every profile $\boldsymbol{J}$:

$$F(\boldsymbol{J}) = \begin{cases} \Phi & \text{if } J_1 = \{p_{a\succ b}\} \\ \varnothing & \text{if } J_1 \neq \{p_{a\succ b}\} \end{cases}$$

F satisfies N, since for any two formulas $\varphi, \psi \in \Phi$ it is the case that $\varphi \in F(\boldsymbol{J}) \Leftrightarrow \psi \in F(\boldsymbol{J})$. However, we may consider a permutation $\pi' : \mathcal{A} \to \mathcal{A}$ with $\pi(a) = b$ and $\pi(b) = a$, and a profile $\boldsymbol{J}$ with $J_1 = \{p_{a\succ b}\}$. Then, taking $\pi : \Phi \to \Phi$ to be naturally induced by π', we will have that $\pi(J_1) = \{p_{b\succ a}\}$, so $\pi(F(\boldsymbol{J})) = \pi(\Phi) = \Phi \neq \varnothing = F(\pi(\boldsymbol{J}))$. Thus, F does not satisfy APN.

Example 10 *APN does not imply N* (so *APN does not imply PN*): Consider the preference agenda Φ based on the set of alternatives $\mathcal{A} = \{a, b, c\}$. Then, construct the rule F such that $p_{a_1\succ a_2} \in F(\boldsymbol{J})$ if $p_{a_1\succ a_2} \in J_i$ for every J_i in $\boldsymbol{J}$, but in case a linear order $a_1 \succ a_2 \succ a_3$ appears in the outcome, the pair $p_{a_1\succ a_3}$ is deleted. F satisfies APN, since renaming the alternatives would derive the same renaming in the outcome. However, taking the unanimous profile $\boldsymbol{J}$ with $J_i = (a \succ b \succ c)$ for all i, we have that $F(\boldsymbol{J}) = \{p_{a\succ b}, p_{b\succ c}\}$ (because the formula $p_{a\succ c}$ will be deleted). Thus, even though it is the case that $N^{\boldsymbol{J}}_{p_{a\succ b}} = N^{\boldsymbol{J}}_{p_{a\succ c}}$, it will hold that $p_{a\succ b} \in F(\boldsymbol{J})$ and $p_{a\succ c} \notin F(\boldsymbol{J})$, implying that F does not satisfy N.

[18] Analogous observations have been made by Endriss and Grandi (2017) in the framework of graph aggregation.

References

Baumeister D, Erdélyi G, Rothe J (2016) Judgment aggregation. In: Rothe J (ed) Economics and computation. Springer, New York

Brams SJ, Fishburn PC (2002) Voting procedures. In: Arrow KJ, Sen A, Suzumura K (eds) Handbook of social choice and welfare. Elsevier, Amsterdam

Chapman B (1998) More easily done than said: Rules, reasons and rational social choice. Oxf. J. Legal Stud. 18(2):293–329

Chapman B (2002) Rational aggregation. Polit. Philos. Econ. 1(3):337–354

Costantini M, Groenland C, Endriss U (2016) Judgment aggregation under issue dependencies. In: Proceedings of the 30th AAAI Conference on Artificial Intelligence (AAAI), pp 468–474

Dietrich F (2006) Judgment aggregation: (Im)possibility theorems. J Econ Theory 126(1):286–298

Dietrich F (2007) A generalised model of judgment aggregation. Soc Choice Welf 28(4):529–565

Dietrich F (2015) Aggregation theory and the relevance of some issues to others. J Econ Theory 160:463–493

Dietrich F, List C (2007a) Arrow's theorem in judgment aggregation. Soc Choice Welf 29(1):19–33

Dietrich F, List C (2007b) Judgment aggregation by quota rules: majority voting generalized. J Theor Polit 19(4):391–424

Dietrich F, List C (2007c) Strategy-proof judgment aggregation. Econ Philos 23(3):269–300

Dietrich F, List C (2008) Judgment aggregation without full rationality. Soc Choice Welf 31(1):15–39

Dietrich F, List C (2010) The impossibility of unbiased judgment aggregation. Theor Decis 68(3):281–299

Dokow E, Holzman R (2009) Aggregation of binary evaluations for truth-functional agendas. Soc Choice Welf 32(2):221–241

Dokow E, Holzman R (2010) Aggregation of binary evaluations. J Econ Theory 145(2):495–511

Endriss U (2016a) Judgment aggregation. In: Brandt F, Conitzer V, Endriss U, Lang J, Procaccia AD (eds) Handbook of computational social choice. Cambridge University Press, Cambridge

Endriss U (2016b) Judgment aggregation. In: Brandt F, Conitzer V, Endriss U, Lang J, Procaccia AD (eds) Handbook of computational social choice. Cambridge University Press, Cambridge chapter 17

Endriss U, Grandi U (2017) Graph aggregation. Artif Intell 245:86–114

Endriss U, Grandi U, Porello D (2010) Complexity of judgment aggregation: Safety of the agenda. In: Proceedings of the 9th International Conference on Autonomous Agents and Multiagent Systems (AAMAS), pp 359–366

Endriss U, Grandi U, Porello D (2012) Complexity of judgment aggregation. J Artif Intell Res 45:481–514

Grandi U, Endriss U (2013) Lifting integrity constraints in binary aggregation. Artif Intell 199:45–66

Grossi D, Pigozzi G (2014) Judgment aggregation: a primer, volume 8 of synthesis lectures on artificial intelligence and machine learning. Morgan & Claypool Publishers, San Rafael

Hartmann S, Sprenger J (2012) Judgment aggregation and the problem of tracking the truth. Synthese 187(1):209–221

van Hees M (2007) The limits of epistemic democracy. Soc Choice Welf 28(4):649–666

Kornhauser LA, Sager LG (1993) The one and the many: Adjudication in collegial courts. California Law Review, pp 1–59

Lang J, Slavkovik M, Vesic S (2016) Agenda separability in judgment aggregation. In: Proceedings of the 30th AAAI Conference on Artificial Intelligence (AAAI)

List C (2004) A model of path-dependence in decisions over multiple propositions. Am Polit Sci Rev 98(3):495–513

List C (2012) The theory of judgment aggregation: an introductory review. Synthese 187(1):179–207

List C, Pettit P (2002) Aggregating sets of judgments: an impossibility result. Econ Philos 18(1):89–110

List C, Pettit P (2004) Aggregating sets of judgments: two impossibility results compared. Synthese 140(1–2):207–235

List C, Polak B (2010) Introduction to judgment aggregation. J Econ Theory 145(2):441–466

List C, Puppe C (2009a) Judgment aggregation: a survey. In: Anand P, Pattanaik P, Puppe C (eds) Handbook of rational and social choice. Oxford University Press, Oxford

List C, Puppe C (2009b) Judgment aggregation: a survey. In: Brandt F, Conitzer V, Endriss U, Lang J, Procaccia AD (eds) Handbook of Rational and Social Choice. Oxford University Press, Oxford

Miller MK, Osherson D (2009) Methods for distance-based judgment aggregation. Soc Choice Welf 32(4):575–601

Mongin P (2008) Factoring out the impossibility of logical aggregation. J Econ Theory 141(1):100–113

Mongin P (2012) The doctrinal paradox, the discursive dilemma, and logical aggregation theory. Theor Decis 73(3):315–355
Nehring K, Puppe C (2007) The structure of strategy-proof social choice. Part I: general characterization and possibility results on median spaces. J Econ Theory 135(1):269–305
Nehring K, Puppe C (2008) Consistent judgement aggregation: the truth-functional case. Soc Choice Welf 31(1):41–57
Pauly M, van Hees M (2006) Logical constraints on judgement aggregation. J Philos Logic 35(6):569–585
Pettit P (2001) Deliberative democracy and the discursive dilemma. Noûs 35:268–299
Porello D, Endriss U (2014) Ontology merging as social choice: judgment aggregation under the open world assumption. J Logic Comput 24(6):1229–1249
Slavkovik M (2014) Not all judgment aggregation should be neutral. In: Proceedings of the 2nd European Conference on Social Intelligence (ECSI), pp 198–211
Terzopoulou Z, Endriss U (2018) Modelling iterative judgment aggregation. In: Proceedings of the 32nd AAAI Conference on Artificial Intelligence (AAAI)
Terzopoulou Z, Endriss U (2019) Rethinking the neutrality axiom in judgment aggregation (extended abstract). In: Proceedings of the 18th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)
Thomson W (2001) On the axiomatic method and its recent applications to game theory and resource allocation. Soc Choice Welf 18(2):327–386
Zwicker WS (2016) Introduction to the theory of voting. In: Brandt F, Conitzer V, Endriss U, Lang J, Procaccia AD (eds) Handbook of computational social choice. Cambridge University Press, Cambridge