

Aggregating Incomplete Judgments: Axiomatisations for Scoring Rules

Zoi Terzopoulou, Ulle Endriss and Ronald de Haan

Institute for Logic, Language and Computation
University of Amsterdam

Abstract

We propose a model of judgment aggregation in which agents can choose to provide a judgment for only some of the issues at stake. A natural class of aggregation rules in this context are scoring rules for which an agent's weight in the decision making process is determined by the number of issues she provides a judgment for. We formulate several appealing axioms for aggregating incomplete judgments and show how they characterise specific rules within this class of scoring rules.

1 Introduction

Judgment aggregation, going back to the seminal work of List and Pettit (2002), is a powerful framework for modelling collective decision making scenarios that involve several logically related issues. This framework has been studied in depth by scholars in Economics, Legal Theory, Philosophy, and Mathematics (see, e.g., List and Puppe, 2009), and more recently it has also attracted the interest of computer scientists (see, e.g., Grossi and Pigozzi, 2014; Endriss, 2016; Baumeister et al., 2017). Almost all prior work on judgment aggregation makes the assumption that all individual judgments are *complete*, i.e., that every agent provides a judgment on every issue at stake. In this paper, we argue that this assumption is not always justified and instead consider the problem of *aggregating incomplete judgments*. Specifically, we consider a natural class of scoring rules for this problem and characterise the rules obtained by imposing a number of normatively appealing axioms.

In judgment aggregation we are presented with a set of *issues*, which together form the *agenda* of decision making. In the model typically investigated in the literature, for every such formula φ , every agent has to either accept φ or explicitly indicate rejection of φ by accepting its complement. This is a useful model for a wide range of applications. But there are at least three types of scenarios in which it is natural to assume that agents only provide such judgments for a subset of the issues in the agenda:

- *Agents only care about a subset of the issues.* When a group decides on a large number of issues, it is not uncommon that some members of the group are personally only interested in some of the issues and thus may want to provide judgments for only those issues. A similar idea has been discussed in work on strategic considerations in judgment aggregation, where it is natural to assume that the preferences of agents depend on the outcome regarding specific issues only (Dietrich and List, 2007b).
- *Agents can only reason about a subset of the issues.* That is, agents may lack the information or the reasoning capabilities required for providing a judgment on certain issues. To model this kind of *bounded rationality*, we require a framework that allows for individual judgments that are incomplete. For example, using the well-known embedding of preference aggregation into judgment aggregation (Dietrich and List, 2007a), incomplete judgments would permit modelling partial preference orders, which are central to many decision making scenarios studied in AI (Rossi et al., 2011).

- *Agents are only asked about a subset of the issues.* For example, judgment aggregation has been suggested as a paradigm for modelling the aggregation of the information collected in a crowdsourcing exercise (Endriss and Fernández, 2013). In the context of crowdsourcing we typically have a very large number of questions we require an answer for, but each individual participant only answers a small subset of them.

In as far as prior work has considered the case of incomplete judgments, the focus has been on allowing the *output* generated by a judgment aggregation rule to be incomplete (rather than the individual judgments provided as input). Specifically, Gärdenfors (2006) has shown that dropping the requirement for an aggregation rule to decide on all issues (i.e., to return complete judgments) does not offer a convincing way out of the impossibility results one typically runs into when imposing particularly demanding axioms in judgment aggregation. His results have later been refined by both Dietrich and List (2008) and Dokow and Holzman (2010). While all of these authors also briefly discuss the idea of dropping the completeness requirement for the individual judgments, they only do so in the context of investigating how such design choices affect impossibility results. Along similar lines, van Hees (2007) has developed a model of judgment aggregation in which agents can accept a given formula to varying degrees, again concentrating on impossibility results. Somewhat closer to our approach are Slavkovik and Jamroga (2011), who investigate a family of distance-based aggregation rules for incomplete individual judgments. However, they treat abstentions in the same way as the binary (yes/no) opinions of the agents (with the underlying framework being a three-valued logic), while we interpret an abstention as the *lack of any opinion* rather than an opinion itself. To summarise, a systematic discussion of how to design and choose useful judgment aggregation rules that operate on inputs that may be incomplete has been missing from the literature to date.

In this paper we concentrate on the following simple and natural class of aggregation rules. Whenever an agent i accepts a formula φ , she implicitly assigns a score to φ . For instance, the scoring function determining this score may be constructed in such a way that an agent can choose to either have a high impact on a small number of issues or a low impact on a large number of issues.¹ The aggregation rule then selects a collective judgment—from the range of all possible judgments that are logically consistent—that maximises the sum of the scores of the formulas accepted by all the agents. These rules are closely related to the *scoring rules* introduced by Dietrich (2014), except that we permit incomplete individual judgments and we require the scoring functions to be *neutral*, in the sense of being blind to the issues themselves and only considering the number of formulas a given agent accepts.

We present three kinds of axioms that encode intuitively attractive normative requirements for aggregation rules that operate on possibly incomplete judgment sets:

- *Majoritarianism.* When deciding between accepting φ and its complement $\sim\varphi$, it is appealing to follow the choice of the majority of agents who have provided a judgment on these two issues. But, as is well known, doing so can lead to logically inconsistent outcomes—this is the infamous *doctrinal paradox* (Kornhauser and Sager, 1993). We therefore propose three weakened forms of majoritarianism that avoid this conflict between logical consistency and responsiveness to the will of the majority.
- *Splitting.* Suppose that a subgroup of the group of all agents observe that simply accepting the union of all their (disjoint) individual judgments would be logically consistent. They then might consider all submitting that union of judgments rather than their own individual judgments. It seems desirable to use an aggregation rule for which the outcome never changes when a subgroup chooses to make such a move, so we consider three axioms of varying strengths based on this fundamental idea.

¹A similar concept has been explored in the context of studying variants of the system of approval voting (Alcalde-Unzu and Vorsatz, 2009).

- *Quality over quantity.* We may wish to give more weight to agents who choose to only provide a small number of judgments. For instance, we may think that they have reflected on those judgments more carefully than another agent who provided judgments on a large number of issues. We may also consider it a matter of fairness to either give someone much influence over a few issues or little influence over many, but not much influence over many issues. We propose an axiom that takes an extreme position on this matter and requires that a single “quality” agent who provided a small number of judgments is always more powerful than any number of agents who all provided a common strictly larger set of judgments.

Our main technical results are characterisations of scoring rules (with neutral scoring functions) that satisfy the axioms sketched above. Under appropriate conditions, imposing a majoritarian axiom demands constant scores. Similarly, imposing a (weak form of the) splitting axiom causes scores to be inversely proportional to the number of issues an agent provides a judgment for. Finally, imposing the quality-over-quantity axiom forces the aggregation rule to make decisions in a lexicographic manner, with agents being ordered according to the number of issues they express an opinion upon.

The remainder of this paper is structured as follows. In Section 2 we present our model of judgment aggregation with possibly incomplete individual judgments and define the class of scoring rules with neutral scoring functions for this model. We then motivate and define our axioms for aggregating incomplete judgments in Section 3 and show how these axioms characterise some specific scoring rules in Section 4. Section 5 concludes with a brief discussion of future work. All proofs have been relegated to the appendix.

2 The Model

In this section we introduce the basic model of judgment aggregation that allows for incomplete individual judgments. Our framework is based on classical propositional logic, building on the existing literature regarding the aggregation of complete judgments (List and Pettit, 2002; List, 2012; Grossi and Pigozzi, 2014; Endriss, 2016).

2.1 The Superagenda

A *superagenda* \mathcal{A} is a finite or countably infinite set of formulas in propositional logic (that is, $|\mathcal{A}| \in \mathbb{N}$ or $|\mathcal{A}| = \aleph_0$ respectively). Intuitively, it is the set of all possible issues that a group of agents may have to decide about; for instance, it can simply be the set of all formulas of propositional logic.² Any specific aggregation scenario concerns a subset of \mathcal{A} . We define $\sim\varphi$, the *complement* of formula φ , as follows: if φ is of the form $\varphi = \neg\psi$ for some formula ψ , then $\sim\varphi = \psi$, otherwise, $\sim\varphi = \neg\varphi$. For the purposes of this paper, we will assume that the superagenda \mathcal{A} is *closed under complementation*, i.e., $\sim\varphi \in \mathcal{A}$ whenever $\varphi \in \mathcal{A}$, and that it does not contain doubly-negated formulas.

Hereafter, for any cardinal numbers $a, b \in \mathbb{N} \cup \{\aleph_0\}$, we write $a \leq b$ whenever: (i) $a, b \in \mathbb{N}$ and $a \leq b$ in the familiar way of comparing natural numbers, (ii) $a \in \mathbb{N}$ and $b = \aleph_0$, or (iii) $a = b = \aleph_0$. Moreover, we adopt the following convention: $\aleph_0/2 = \aleph_0 + 1 = \aleph_0$.

For $\lambda \in \mathbb{N}$, we say that a superagenda \mathcal{A} is λ -*constrained* if for every consistent subset $J \subseteq \mathcal{A}$ with $|J| = \lambda$, it holds that for all $\varphi \in J$:

$$J \setminus \{\varphi\} \models \varphi$$

²The notion of a superagenda captures what numerous scholars already implicitly assume in their models of judgment aggregation: that the domain of decision making may vary in different scenarios, but it is still desirable to have a uniform framework that integrates all possible cases.

A superagenda \mathcal{A} is *unconstrained* if there is no $\lambda \leq \frac{|\mathcal{A}|}{2}$ for which \mathcal{A} is λ -constrained. Unconstrainedness is a very weak property. In particular, every superagenda becomes unconstrained if we add to it a new propositional variable and its complement.³

An *agenda* A is a finite subset of the superagenda that is closed under complementation. Agendas represent the domain of decision making in specific collective choice problems.

2.2 Aggregating Incomplete Judgments

We consider a finite or countably infinite superpopulation \mathcal{N} with $|\mathcal{N}| \geq 2$. This is the set of all *potential agents* that may participate in an aggregation scenario (hence, $|\mathcal{N}| \in \mathbb{N}$ or $|\mathcal{N}| = \aleph_0$). Then, we fix a finite *group of agents* $N = \{1, \dots, n\} \subseteq \mathcal{N}$ that have to make a collective decision regarding the issues in an agenda A . During the aggregation process, every agent $i \in N$ submits a logically consistent *judgment set* (or simply *judgment*) $J_i \subseteq A$. Agent i may abstain from expressing a yes/no opinion on some issue in the agenda, that is, it may be the case that $\varphi \notin J_i$ and $\sim\varphi \notin J_i$, for some formula $\varphi \in A$.⁴ We denote by $\mathcal{J}(A)$ the set of all logically consistent subsets of the agenda A , and by $\mathcal{J}(A)^\bullet$ those that are also complete ($J \in \mathcal{J}(A)^\bullet$ if $J \in \mathcal{J}(A)$ and $\varphi \in J$ or $\sim\varphi \in J$, for all $\varphi \in A$). A *profile* $\mathbf{J} = (J_1, \dots, J_n) \in \mathcal{J}(A)^n$ captures the individual judgments of the agents in group N . We denote by $N_\varphi^{\mathbf{J}}$ and $N_{\sim\varphi}^{\mathbf{J}}$ the sets of agents that report φ and $\sim\varphi$ in profile \mathbf{J} respectively. Furthermore, we say that a formula φ is *logically independent* of a profile \mathbf{J} whenever φ is logically independent of each consistent subset of $\{\psi, \sim\psi \mid \psi \in \bigcup_{i \in N} J_i\} \setminus \{\varphi, \sim\varphi\}$. Recall that a formula φ is logically independent of a set of formulas Y if both $Y \cup \{\varphi\}$ and $Y \cup \{\sim\varphi\}$ are logically consistent.

Then, an *aggregation rule* F is a function that maps any profile of judgments $\mathbf{J} \in \mathcal{J}(A)^n$ for any agenda A and any group N to a nonempty set of collective judgment sets for the same agenda, i.e., to a nonempty subset of 2^A . Thus, there may be a tie between several “best” judgment sets and these judgments need not be consistent or complete in general.

2.3 Scoring Rules

We now define the family of *scoring rules* \mathcal{S} , following Dietrich (2014). First, a *scoring function* $s : \mathcal{A} \times 2^A \rightarrow \mathbb{R}$ assigns to each formula φ and judgment set J_i (possibly inconsistent and incomplete for the sake of this definition) a number $s_{J_i}(\varphi)$, called the *score of φ given J_i* and measuring how φ performs from the perspective of holding judgment set J_i . We will make the reasonable assumption that, if a formula does not belong to a judgment set J_i , then its relevant score is zero:

$$s_{J_i}(\varphi) = 0, \text{ if } \varphi \notin J_i \tag{1}$$

For example, for any given $c \in \mathbb{R}$, we call s^c , defined as follows, the *constant scoring function*.

$$s_{J_i}^c(\varphi) = \begin{cases} c, & \text{if } \varphi \in J_i \\ 0, & \text{if } \varphi \notin J_i \end{cases}$$

For our technical purposes, negative scores are also allowed; the requirement of positivity will be later entailed by an axiom, namely *unanimity*. Of course, one could also force the

³It is also worth noting that most superagendas used in common judgment aggregation applications, like those containing a conjunctive or disjunctive agenda with at least two premises (Dietrich and List, 2007b), a preference agenda (Dietrich and List, 2007a) or a budget agenda (Dietrich and List, 2010), are unconstrained.

⁴Gärdenfors (2006) and the authors of follow-up papers require deductive closure of (incomplete) judgment sets. However, we do not wish to make this restrictive assumption, since (according to each of the three motivations we have given in the introduction) an agent may often have no reason to report an opinion on an extra issue, even if that issue is a logical consequence of her judgment set. Our model is designed to only take into consideration the agents’ explicitly declared opinions.

scores to be positive already in the definition. The score of a set of formulas can be thought of as the sum of the scores of its members. So, we get an extended scoring function which (given agent i 's judgment set J_i) assigns a score to each set $J \subseteq \mathcal{A}$:

$$s_{J_i}(J) = \sum_{\varphi \in J} s_{J_i}(\varphi)$$

Having a profile $\mathbf{J} = (J_1, \dots, J_n)$ and a scoring function s , we say that $s_{\mathbf{J}}(J) = \sum_{i \in N} s_{J_i}(J)$ is the *score* of the judgment set J on the profile \mathbf{J} . Then, given a superagenda \mathcal{A} and a superpopulation \mathcal{N} , a scoring rule F_s determines the collective judgments for any agenda $A \subseteq \mathcal{A}$ and group of agents $N = \{1, \dots, n\} \subseteq \mathcal{N}$ by selecting the complete and consistent subsets of the agenda with the highest total score across all agents:⁵

$$F_s(\mathbf{J}) = \operatorname{argmax}_{J \in \mathcal{J}(A)^\bullet} \sum_{\varphi \in J, i \in N} s_{J_i}(\varphi) = \operatorname{argmax}_{J \in \mathcal{J}(A)^\bullet} \sum_{i \in N} s_{J_i}(J) = \operatorname{argmax}_{J \in \mathcal{J}(A)^\bullet} s_{\mathbf{J}}(J)$$

In this paper we focus on a subfamily of all scoring rules, namely the family of all rules with associated *neutral* scoring functions.⁶ Informally speaking, a neutral scoring function treats all formulas symmetrically. Concretely, the scoring function s is neutral if for every permutation $\pi : \mathcal{A} \rightarrow \mathcal{A}$ and all formulas $\varphi_1, \dots, \varphi_k, \varphi \in \mathcal{A}$ it holds that:

$$s_{\{\varphi_1, \dots, \varphi_k\}}(\varphi) = s_{\{\pi(\varphi_1), \dots, \pi(\varphi_k)\}}(\pi(\varphi))$$

Neutral scoring functions induce the same score for all formulas in a judgment set, and moreover that score depends only on the judgment set's size.⁷ Admittedly, the condition of neutrality of a scoring function is normatively questionable and certainly we do not claim that it should hold in all decision-making contexts. What our work aims to do, however, is to focus on and analyse those situations where neutrality is a reasonable and desirable assumption to make.

Additionally, one may find it useful to impose scores that are non-increasing in the judgment set's size (otherwise the agents would always prefer to submit complete opinions). This requirement can be achieved by an axiom analogous to *contraction* in the framework of size-approval voting (Alcalde-Unzu and Vorsatz, 2009). However, since all the interesting axioms we consider in Section 3 already entail contraction, we are not going to further discuss this axiom in its own right.

We call \mathcal{S}_N the class of all neutral-scoring rules, i.e., the aggregation rules with associated neutral scoring functions.⁸ In the remainder we will write s_λ as a shorthand for $s_J(\varphi)$ with $\varphi \in J$, where $0 < \lambda = |J|$. For instance, s_1 will designate the score corresponding to all singletons. Hence $s_{|J|}$ is the score of judgment set J . Restricting attention to \mathcal{S}_N , a scoring rule F_s can also be defined as:

$$F_s(\mathbf{J}) = \operatorname{argmax}_{J \in \mathcal{J}(A)^\bullet} \sum_{i \in N} s_{|J_i|} \cdot |J \cap J_i|$$

⁵Note that if an agent i does not express a yes/no judgment on any issue, that is, if $J_i = \emptyset$, then by assumption (1) this agent is completely disregarded as far as the collective outcome is concerned. In alternative frameworks, empty judgment sets could influence the collective decision (see, e.g., Slavkovik and Jamroga, 2011).

⁶The scoring rules are anonymous by definition. Moreover, we have defined scoring rules to satisfy *collective rationality*, which requires that all collective judgment sets must be complete and consistent. One could easily relax the completeness assumption, with no important technical implications for our results.

⁷The reader who finds our definition of neutrality for a scoring function too stringent can also think of an alternative—possibly more plausible—definition restricted to all *consistency-preserving* permutations, that is, permutations π such that $\{\pi(\varphi) \mid \varphi \in J\}$ is consistent whenever J is consistent. Since all judgment sets we consider in this paper are consistent, our results are invariant under the definition of neutrality we adopt.

⁸It is crucial to stress that what we call a *neutral-scoring* rule should not be confused with a *neutral* rule as studied so far in judgment aggregation (Endriss, 2016). These two notions are logically independent; for instance, the generalised Kemeny rule is a neutral-scoring rule but it is not neutral (Costantini et al., 2016).

Before continuing to the main sections of this paper, it is important to underline a detail regarding *rules* on the one hand and *scoring vectors* defining rules on the other. Given a superagenda \mathcal{A} , for a neutral-scoring rule F_s there exists a (possibly infinite) scoring vector $\mathbf{s} = (s_\lambda)_{\lambda \in \mathbb{N}, \lambda \leq |\mathcal{A}|/2}$, with $s_\lambda \in \mathbb{R}$ for every λ , such that for every agenda $A \subseteq \mathcal{A}$ only the prefix of \mathbf{s} of cardinality $|A|/2$ is relevant for F_s . Furthermore, for every neutral-scoring rule F_s there are infinitely many scoring vectors that induce F_s . For example, if we multiply all scores with some positive constant c , then the rule being induced does not change.

3 Axioms

In this section we discuss a number of reasonable features of aggregation rules for incomplete individual judgments and we formalise them in terms of axioms. For the sake of our definitions, let us consider a fixed agenda $A \subseteq \mathcal{A}$, a group of agents $N \subseteq \mathcal{N}$ of size n , and an aggregation rule F defined on profiles in $\mathcal{J}(A)^n$.

To start off, a fundamental property in the economics literature demands that whenever all agents have the same opinion, then the collective outcome should respect that opinion:

Unanimity. For every profile of the form $\mathbf{J} = (J', \dots, J') \in \mathcal{J}(A)^n$, it holds that $J' \subseteq J$ for all $J \in F(\mathbf{J})$.

Making a collective decision based on the opinion of the majority on each issue of the agenda separately is commonly considered as a desirable attribute of an aggregation rule. However, we know that majorities easily lead to inconsistencies when logical interconnections between the formulas are involved (we refer, for instance, to the *doctrinal paradox* by Kornhauser and Sager, 1993). On the other hand, there clearly are situations where the role of logic is not critical for the decision-making process, while the judgment of the majority is still considered a reliable indication concerning the judgment of the whole group. Suppose, for instance, that the jury of a court has to deliver a verdict on two distinct cases, deciding whether the defendant of the first case is guilty (proposition φ_1), and similarly for the second case (proposition φ_2). In this scenario, no matter what the verdict on φ_1 is, the jury can independently decide in favour of φ_2 or of $\sim\varphi_2$. Formally, for the agenda $A = \{\varphi_1, \varphi_2, \sim\varphi_1, \sim\varphi_2\}$ and for a jury of size n , it holds that φ_2 is logically independent of \mathbf{J} for every profile $\mathbf{J} \in \mathcal{J}(A)^n$. Now, imagine that the jury consists of five members, three of which find that the defendant of the second case is guilty. This means that a strict majority of the jury members accept proposition φ_2 . The axiom of *forward majoritarianism* says exactly that this majority ought to be respected by an aggregation rule; regardless of what the jury decides about the first case, any collective judgment of it should contain proposition φ_2 . Then, suppose that the judge of the court does not have access to the individual judgments of the jury members, but she can observe the outcome of an aggregation rule, where every collective judgment contains φ_2 . Hence, the judge cannot avoid but announce the defendant of the second case guilty, without knowing how many out of the five jury members agree with this verdict. The axiom of *backward majoritarianism* provides the judge with the guarantee that a strict majority of the jury members have deemed the defendant guilty.

For a profile $\mathbf{J} \in \mathcal{J}(A)^n$, let us define the *simple-majority set* $m(\mathbf{J}) = \{\varphi \in A \mid |N_\varphi^{\mathbf{J}}| > |N_{\sim\varphi}^{\mathbf{J}}|\}$, that is, the set of all formulas that have more advocates than adversaries in \mathbf{J} .⁹

⁹One could also want to work with the *absolute-majority set* $m(\mathbf{J}) = \{\varphi \in A \mid |N_\varphi^{\mathbf{J}}| > \frac{n}{2}\}$ (see, e.g., Lang et al., 2017). By adopting this definition for incomplete individual judgments, our axioms of forward majoritarianism and of general majoritarianism would become weaker, while backward majoritarianism would be logically stronger.

Forward majoritarianism. For every profile $\mathbf{J} \in \mathcal{J}(A)^n$ and formula $\varphi \in A$ that is logically independent of \mathbf{J} , it holds that if $\varphi \in m(\mathbf{J})$, then $\varphi \in J$ for all $J \in F(\mathbf{J})$.

Backward majoritarianism. For every profile $\mathbf{J} \in \mathcal{J}(A)^n$ and formula $\varphi \in A$ that is logically independent of \mathbf{J} , it holds that if $\varphi \in J$ for all $J \in F(\mathbf{J})$, then $\varphi \in m(\mathbf{J})$.

Regarding the general matter of majority judgments conflicting with logical constraints, a notion of minimal divergence from the majority outcome based on set-inclusion can be defined: For a set of formulas $Y \subseteq A$, $Y' \subseteq Y$ is a maximal consistent subset of Y if and only if Y' is consistent and there exists no other consistent set Y'' such that $Y' \subset Y'' \subseteq Y$. The set of maximal consistent subsets of Y is denoted by $\max(Y, \subseteq)$. Following Nehring et al. (2014), we call $\text{Con}(\mathbf{J}) = \{Y' \mid Y' \text{ complete and consistent, and } Y' \supseteq Y, \text{ for some } Y \in \max(m(\mathbf{J}), \subseteq)\}$ the *Condorcet set*. The property of *general majoritarianism* ensures that the aggregation rule induces only judgment sets that extend (possibly in an inconsistent way) judgment sets in the Condorcet set.¹⁰

General majoritarianism. For every profile $\mathbf{J} \in \mathcal{J}(A)^n$ and every judgment set $J \in F(\mathbf{J})$, there exists a judgment set $J' \in \text{Con}(\mathbf{J})$ such that $J' \subseteq J$.

By inspecting the definitions above we can see that, as far as all aggregation rules are concerned, general majoritarianism logically implies forward majoritarianism and both these axioms are independent of backward majoritarianism. However, we will show in Section 4 that the logical relations differ when we focus on all neutral-scoring rules.

Next, we introduce three novel axioms. To illustrate the idea behind them, consider the following scenario. A group of doctors has to make a decision concerning the treatment of a patient. For example, they may have to decide if treatment 1 is suitable (proposition φ_1), similarly for treatment 2 (proposition φ_2), but also if treatment 1 and 2 can be combined (proposition $\varphi_1 \wedge \varphi_2$), etc. In this context, it is reasonable to think that the doctors may have different specialisations, or that they may choose to spend their limited time investigating different issues regarding the decision they have to make. This would result in a meeting where these doctors submit judgments on different propositions. Suppose now that a subset of the doctors have the chance to eat lunch together before the meeting, and of course they discuss about their patient. Hence, they are able to share the information they collected individually and the judgments they arrived at. Then, this subset of the doctors may realise that all their judgments together form a consistent set of propositions $\{\varphi_1, \dots, \varphi_\lambda\}$. In such a case it is safe to assume that, trusting their colleagues, the doctors in this subset will report the whole set of propositions in the meeting, instead of the more restricted judgments they held before the discussion. We would like the aggregation rule employed by the group to reach the same outcome independently of whether some of the doctors choose to combine their individual judgments prior to aggregation. We present three axioms that enforce exactly this property in various degrees, ordered from stronger to weaker:

Arbitrary-splitting. For every profile $\mathbf{J} \in \mathcal{J}(A)^n$ and subgroup $\emptyset \neq N' \subseteq N$ of agents with pairwise disjoint and mutually consistent judgment sets, we have $F(\mathbf{J}) = F(\mathbf{J}')$, where \mathbf{J}' arises from \mathbf{J} by replacing the judgment set of each member of N' by the union $\bigcup_{i \in N'} J_i$.

Equal-splitting. For every profile $\mathbf{J} \in \mathcal{J}(A)^n$ and subgroup $\emptyset \neq N' \subseteq N$ of agents whose judgment sets are pairwise disjoint, mutually consistent, and of equal size, we have

¹⁰If the aggregation rule satisfies collective rationality, inducing only consistent collective judgments, general majoritarianism means that $F(\mathbf{J}) \subseteq \text{Con}(\mathbf{J})$ for every profile $\mathbf{J} \in \mathcal{J}(A)^n$.

$F(\mathbf{J}) = F(\mathbf{J}')$, where \mathbf{J}' arises from \mathbf{J} by replacing the judgment set of each member of N' by the union $\bigcup_{i \in N'} J_i$.

Single-splitting. For every profile $\mathbf{J} \in \mathcal{J}(A)^n$ and subgroup $\emptyset \neq N' \subseteq N$ of agents whose judgment sets are pairwise disjoint, mutually consistent, and singletons, we have $F(\mathbf{J}) = F(\mathbf{J}')$, where \mathbf{J}' arises from \mathbf{J} by replacing the judgment set of each member of N' by the union $\bigcup_{i \in N'} J_i$.

The following property is particularly desirable when smaller judgment sets are significantly more important to the collective decision than larger ones. For instance, in case the agents have a fixed and limited amount of energy/time/cognitive effort at their disposal, then one could expect that small judgment sets are more well-thought-out than large ones, and hence of greater value to the group. We restrict attention to profiles where only two judgments—which differ in size—are reported. The axiom of *quality-over-quantity* states that the collective judgment ought to always agree with the smaller individual judgment set, no matter how many agents adopt each opinion:¹¹

Quality-over-quantity. Consider two arbitrary judgment sets $\emptyset \neq J', J'' \in \mathcal{J}(A)$ with $|J'| < |J''|$. For every profile $\mathbf{J} \in \mathcal{J}(A)^n$ where a subgroup $\emptyset \neq N' \subset N$ of agents declare judgment J' , another subgroup $\emptyset \neq N'' \subset N$ of agents declare judgment J'' , and the rest of the agents submit an empty set, it holds that $J' \subseteq J$ for all $J \in F(\mathbf{J})$.

For the impatient reader we shall already mention that within the context of scoring rules, the axiom of quality-over-quantity does not simply require smaller judgments sets to weight more than larger ones—as we have indicated in Section 2, this would be imposed by a contraction-like property. The specific scores that the quality-over-quantity axiom generates will be fully analysed in Section 4, along with the scores related to the rest of our axioms.

4 Characterisation Results

The goal of this section is to provide axiomatic characterisations for aggregation rules within the family \mathcal{S}_N . Recall that, for a given superagenda \mathcal{A} and a given superpopulation \mathcal{N} , a scoring rule F_s induces an aggregation rule for every specific agenda $A \subseteq \mathcal{A}$ and specific group $N \subseteq \mathcal{N}$. We say that F_s satisfies a certain axiom if and only if all aggregation rules induced by F_s in this manner satisfy the axiom.

Our first lemma is straightforward and states that unanimity requires positive scores:

Lemma 1. *For any superagenda \mathcal{A} with at least one logically contingent formula, superpopulation \mathcal{N} , and neutral scoring function s , the scoring rule $F_s \in \mathcal{S}_N$ satisfies unanimity if and only if $s_\lambda > 0$ for all $\lambda \in \mathbb{N}$ with $\lambda \leq \frac{|\mathcal{A}|}{2}$.*

We begin by investigating the three majoritarian properties. For a finite superpopulation, we show that each one of the axioms of general majoritarianism and of forward majoritarianism characterises a family of rules with scores that are “close enough”, while if the superpopulation is infinite we obtain an axiomatisation of the constant-scoring rule:¹²

¹¹The property of quality-over-quantity may be judged as weak or strong, depending on different considerations. On the one hand, the axiom’s scope (i.e., the profiles that its antecedent refers to) may categorise it as weak, since it only mentions cases where just two distinct judgment sets are reported. On the other hand, the axiom may also be considered strong, because it clearly prioritises small judgment sets over large ones, not allowing for any compromise.

¹²Note that every constant scoring function s^c with $c > 0$ induces the same scoring rule.

Definition 1. Take any constant scoring function s^c with $c > 0$. We call the scoring rule F^c that is induced by s^c the **constant-scoring rule**.

In frameworks with complete individual judgments, the constant-scoring rule is known as the *distance-based* rule (Endriss et al., 2012), the *simple* scoring rule (Dietrich, 2014), the *median* rule (Nehring et al., 2014), the *Prototype* rule (Miller and Osherson, 2009), and the *generalised Kemeny* rule (Endriss, 2016).

Proposition 2. For any unconstrained superagenda \mathcal{A} , finite superpopulation \mathcal{N} , and neutral scoring function s , the scoring rule $F_s \in \mathcal{S}_{\mathcal{N}}$ satisfies general majoritarianism if and only if it satisfies forward majoritarianism if and only if $s_{\lambda} > 0$ and $\frac{s_{\lambda}}{s_{\lambda'}} < \frac{k}{k-1}$ for all $\lambda, \lambda' \in \mathbb{N}$ with $\lambda, \lambda' \leq \frac{|\mathcal{A}|}{2}$, where $k = \lceil \frac{|\mathcal{N}|}{2} \rceil$.

Theorem 3. For any unconstrained superagenda \mathcal{A} and infinite superpopulation \mathcal{N} , the only neutral-scoring rule satisfying general majoritarianism is the constant-scoring rule F^c .

Theorem 4. For any unconstrained superagenda \mathcal{A} and infinite superpopulation \mathcal{N} , the only neutral-scoring rule satisfying forward majoritarianism is the constant-scoring rule F^c .

In case, however, we wish to make no assumptions regarding the superpopulation, the constant-scoring rule F^c is directly characterised by backward majoritarianism:

Theorem 5. For any unconstrained superagenda \mathcal{A} and superpopulation \mathcal{N} , the only neutral-scoring rule satisfying backward majoritarianism is the constant-scoring rule F^c .

The above results highlight an interesting fact about the logical relations between the three majoritarian properties, when restricting attention within the family $\mathcal{S}_{\mathcal{N}}$. The axioms of forward and of general majoritarianism are now logically equivalent (even though in general the latter is stronger than the former), and logically weaker than backward majoritarianism (which in general is logically independent of both).

We proceed with exploring how the splitting axioms function within the family $\mathcal{S}_{\mathcal{N}}$. Interestingly, for a sufficiently large superpopulation we can show that the weakest axiom of the three, namely the single-splitting axiom, together with unanimity characterises the aggregation rule $F_s \in \mathcal{S}_{\mathcal{N}}$ induced by scores s_{λ} that are inversely proportional to λ (Theorem 6), and by Theorem 7 the same holds for the equal-splitting axiom. Moreover, the arbitrary-splitting axiom is proven to be too strong: it is not satisfied by any unanimous neutral-scoring rule (Proposition 8).¹³

Definition 2. Take any neutral scoring function s with $s_{\lambda} = \frac{s_1}{\lambda} > 0$ for all $\lambda \in \mathbb{N}$, $\lambda \leq \frac{|\mathcal{A}|}{2}$. We call the scoring rule F^{ee} that is induced by s the **equal-and-even-scoring rule**.

The equal-and-even-scoring rule borrows its name from its counterpart in voting theory: the equal and even cumulative voting procedure, which can be considered a special version of cumulative voting (Glasser, 1959; Alcalde-Unzu and Vorsatz, 2009).

Theorem 6. For any unconstrained superagenda \mathcal{A} and superpopulation \mathcal{N} with $|\mathcal{N}| \geq \frac{|\mathcal{A}|}{2} + 1$, the only neutral-scoring rule satisfying the single-splitting axiom simultaneously with unanimity is the equal-and-even-scoring rule F^{ee} .

Theorem 7. For any unconstrained superagenda \mathcal{A} and any superpopulation \mathcal{N} with $|\mathcal{N}| \geq \frac{|\mathcal{A}|}{2} + 1$, the only neutral-scoring rule satisfying the equal-splitting axiom simultaneously with unanimity is the equal-and-even-scoring rule F^{ee} .

¹³The results about the splitting axioms would fail without the property of unanimity in the relevant statements and the hypothesis of a sufficiently large superpopulation; dropping any one of these two assumptions would lead to a characterisation of some *family of rules* instead of a *single rule*. Similarly, assuming that $s_1 > 0$ in Definition 2 is necessary in order to guarantee the induction of a single scoring rule.

Recalling the formulations of the splitting axioms in Section 3, the reader may be curious to know how essential for our characterisation results the assumption of disjointness of the individual judgments is. The answer is “very”—without it, the splitting properties would fail for all unanimous neutral-scoring rules.

Finally, the axioms of arbitrary-splitting and of unanimity are mutually exclusive:

Proposition 8. *For any unconstrained superagenda \mathcal{A} and superpopulation \mathcal{N} with $|\mathcal{N}| \geq \frac{|\mathcal{A}|}{2} + 1 \geq 4$, no unanimous neutral-scoring rule satisfies the arbitrary-splitting axiom.*

In the sequel we introduce an original aggregation rule, which makes sense in settings where giving strict priority to agents who report smaller judgment sets is recommendable. In words, the *upward-lexicographic* rule first tries to agree as much as possible with the agents holding the smallest judgment sets; in case this process induces more than one collective judgment, the rule aims at maximising agreement with the agents with the second smallest judgment sets, and so on, until no ties can be broken anymore.

In order to proceed to a formal definition of the aggregation rule, we first need an additional technical notion, namely the one of the lexmax function. Consider a set $V = \{\mathbf{v}^1, \mathbf{v}^2, \dots, \mathbf{v}^m\}$ of equal-sized vectors of real numbers, and denote by $v_i^k \in \mathbb{R}$ the i^{th} element of vector \mathbf{v}^k . The function lexmax picks up the vector from the set V the elements of which are maximal in a lexicographic manner. More precisely, $\text{lexmax} : V \mapsto \mathbf{v}^k$, where $\mathbf{v}^k \in V$ is such that if $v_i^\ell > v_i^k$ for some $v_i^\ell \in V$, then there exists some $j < i$ for which $v_j^\ell < v_j^k$.

Let us now fix an agenda A , a group of agents N , and a profile $\mathbf{J} \in \mathcal{J}(A)^n$, and define $k_\lambda^{\mathbf{J}}(J) = \sum_{i \mid \lambda=|J_i|} |J \cap J_i|$, the “raw points” the agents with judgment sets of size λ assign to choosing J , and $K_\lambda^{\mathbf{J}}(J) = \sum_{k=1}^K k_\lambda^{\mathbf{J}}(J)$, the “raw points” the agents with judgment sets of size *up to* λ assign to choosing J . The upward-lexicographic rule selects those complete and consistent judgment sets that result in a lexicographically maximal vector (the argument function arg is defined in the standard way):

$$F^{\text{ulex}}(\mathbf{J}) = \underset{J \in \mathcal{J}(A)^\bullet}{\text{arglexmax}} \left(K_1^{\mathbf{J}}(J), \dots, K_{|\mathcal{A}|/2}^{\mathbf{J}}(J) \right)$$

We have defined the upward-lexicographic rule in terms of a procedure to be followed to compute the outcome for any given profile. We are now going to show that we can define the same rule in terms of a scoring vector. We observe that (for \mathcal{N} being finite) if a scoring function s is such that $\frac{s_{\lambda-1}}{s_\lambda} = \lambda \cdot |\mathcal{N}|$, this means that each formula of a small judgment set that appears in a profile of at most $|\mathcal{N}|$ agents has greater value for the aggregation rule F_s than all formulas in larger judgment sets; and this precisely captures the intuition behind the upward-lexicographic rule:

Definition 3. *Given a superagenda \mathcal{A} and a finite superpopulation \mathcal{N} , take the scoring function s such that $s_\lambda = \prod_{k=1}^\lambda \frac{1}{k \cdot |\mathcal{N}|}$ for all $\lambda \in \mathbb{N}$, $\lambda \leq \frac{|\mathcal{A}|}{2}$. We call the scoring rule F^{ulex} that is induced by s the **upward-lexicographic rule**.*

Recall that every scoring rule is represented by many different scoring vectors. The following lemma shows that the upward-lexicographic rule is special in the sense that, if we change the scoring vector in such a way that some or all of the ratios $\frac{s_{\lambda-1}}{s_\lambda}$ of consecutive scores increase, then we do not change the corresponding aggregation rule:

Lemma 9. *For any superagenda \mathcal{A} and finite superpopulation \mathcal{N} , let s be any neutral scoring function with $s_\lambda > 0$ and $\frac{s_{\lambda-1}}{s_\lambda} > \lambda \cdot (|\mathcal{N}| - 1)$ for all $\lambda \in \mathbb{N}$, $\lambda \leq \frac{|\mathcal{A}|}{2}$. Then F_s is the upward-lexicographic rule F^{ulex} .*

As one may have already predicted, the upward-lexicographic rule is characterised by the property of quality-over-quantity:¹⁴

Theorem 10. *For any superagenda \mathcal{A} that is closed under conjunction of literals and contains all propositional variables and for any finite superpopulation \mathcal{N} , the only neutral-scoring rule that satisfies quality-over-quantity is the upward-lexicographic rule F^{ulex} .*

After having analysed the upward-lexicographic rule, it is natural to define the *downward-lexicographic* rule, which selects those complete and consistent subsets of the agenda that result in a lexicographically maximal vector prioritising the points collected by all judgment sets in a profile. This rule breaks ties by excluding the points collected by individual judgment sets of a fixed size in rounds, moving from the largest to the smallest ones.¹⁵

$$F^{dlex}(\mathbf{J}) = \operatorname{arglexmax}_{J \in \mathcal{J}(\mathcal{A})^\bullet} \left(K_{|A|/2}^J(J), \dots, K_1^J(J) \right)$$

Even though at first sight the upward-lexicographic rule and the downward-lexicographic rule seem to satisfy some notion of duality, we will show that this is not true in general. In particular, for any infinite and unconstrained superagenda (such as a superagenda that contains all propositional variables that was used in the proof of Theorem 10), the downward-lexicographic rule is not a neutral-scoring rule:

Proposition 11. *For any infinite and unconstrained superagenda \mathcal{A} and for any superpopulation \mathcal{N} , the downward-lexicographic rule F^{dlex} does not belong to $\mathcal{S}_\mathcal{N}$.*

5 Conclusion

We have developed a general model of judgment aggregation that makes room for agents who may express an opinion on a limited number of the issues in question. An individual judgment is then assigned a weight depending on how many issues it covers, and a scoring rule tries to maximise the relevant score gained by the collective judgment. We have defined various appealing axioms in this context, concerning first the problem of dealing with clashing majority opinions, second situations of alliance-formation, and third settings where judgments referring to less issues may be deemed much more precious than those involving more issues. We have showed that each type of axiom uniquely characterises a natural scoring rule under very light conditions on the domain of decision making (i.e., the agenda) and the number of all potential agents participating in an aggregation problem.

The results presented in this paper can be seen as the beginning towards a more structured study of the possibilities that aggregating incomplete individual judgments offers, both regarding theory and applications. Our work brings into light several concrete open questions. Technically, it would be interesting to examine whether the downward-lexicographic rule we introduced corresponds to a neutral-scoring rule under the assumption of a finite superagenda, and if so, what kind of axiom(s) could characterise it. Of course, further formal properties related to incomplete individual judgments could also be considered and analysed within the family of neutral-scoring rules. In addition, it would be intriguing to

¹⁴The axiomatisation of Theorem 10 holds for a finite superagenda \mathcal{A} too, under the condition that for every $\lambda \leq \frac{|A|}{2}$ there exists a set of propositional variables $\{p_1, \dots, p_\lambda\} \subseteq \mathcal{A}$ such that $\neg p_1 \wedge \dots \wedge \neg p_\lambda \in \mathcal{A}$. However, if we take an infinite superpopulation \mathcal{N} instead of a finite one, then the characterisation result does not hold. Specifically, the upward-lexicographic rule cannot be defined as a neutral-scoring rule anymore, and the quality-over-quantity axiom is not satisfied by any rule in $\mathcal{S}_\mathcal{N}$.

¹⁵Observe that the definition of the upward-lexicographic rule is not affected if we substitute all K_λ^J by k_λ^J , while for the downward-lexicographic rule it is (and the new rule we obtain is not intuitively attractive).

provide complete axiomatisations for the aggregation rules we discussed in this paper, without restricting ourselves to the class of neutral-scoring rules—inspiration could be drawn from the recent work of Nehring and Pivato (2018). Finally, in more philosophical terms, it would be worth clarifying which exact characteristics of the agents’ reasoning and motivations are meant to be captured by an aggregation framework of incomplete judgments and spell out how these considerations play a role when designing useful aggregation rules.

Acknowledgments. We thank the COMSOC-2018 reviewers for the extensive and helpful feedback received.

References

- Alcalde-Unzu, J. and Vorsatz, M. (2009). Size approval voting. *Journal of Economic Theory*, 144(3):1187–1210.
- Baumeister, D., Rothe, J., and Selker, A.-K. (2017). Strategic behavior in judgment aggregation. In Endriss, U., editor, *Trends in Computational Social Choice*, pages 145–168. AI Access.
- Costantini, M., Groenland, C., and Endriss, U. (2016). Judgment aggregation under issue dependencies. In *Proceedings of the 30th AAAI Conference on Artificial Intelligence (AAAI)*, pages 468–474.
- Dietrich, F. (2014). Scoring rules for judgment aggregation. *Social Choice and Welfare*, 42(4):873–911.
- Dietrich, F. and List, C. (2007a). Arrow’s theorem in judgment aggregation. *Social Choice and Welfare*, 29(1):19–33.
- Dietrich, F. and List, C. (2007b). Strategy-proof judgment aggregation. *Economics and Philosophy*, 23(03):269–300.
- Dietrich, F. and List, C. (2008). Judgment aggregation without full rationality. *Social Choice and Welfare*, 31(1):15–39.
- Dietrich, F. and List, C. (2010). The problem of constrained judgment aggregation. In Stadler, F., editor, *The Present Situation in the Philosophy of Science*, pages 125–139. Springer.
- Dokow, E. and Holzman, R. (2010). Aggregation of binary evaluations with abstentions. *Journal of Economic Theory*, 145(2):544–561.
- Endriss, U. (2016). Judgment aggregation. In Brandt, F., Conitzer, V., Endriss, U., Lang, J., and Procaccia, A. D., editors, *Handbook of Computational Social Choice*, pages 399–426. Cambridge University Press.
- Endriss, U. and Fernández, R. (2013). Collective annotation of linguistic resources: Basic principles and a formal model. In *Proceedings of the 51st Annual Meeting of the Association for Computational Linguistics (ACL)*, pages 539–549.
- Endriss, U., Grandi, U., and Porello, D. (2012). Complexity of judgment aggregation. *Journal of Artificial Intelligence Research*, 45(1):481–514.
- Gärdenfors, P. (2006). A representation theorem for voting with logical consequences. *Economics and Philosophy*, 22(2):181–190.

- Glasser, G. (1959). Game theory and cumulative voting for corporate directors. *Management Science*, 5(2):151–156.
- Grossi, D. and Pigozzi, G. (2014). *Judgment Aggregation: A Primer*. Synthesis Lectures on Artificial Intelligence and Machine Learning. Morgan and Claypool Publishers.
- van Hees, M. (2007). The limits of epistemic democracy. *Social Choice and Welfare*, 28(4):649–666.
- Kornhauser, L. A. and Sager, L. G. (1993). The one and the many: Adjudication in collegial courts. *California Law Review*, 81(1):1–59.
- Lang, J., Pigozzi, G., Slavkovik, M., van der Torre, L., and Vesic, S. (2017). A partial taxonomy of judgment aggregation rules and their properties. *Social Choice and Welfare*, 48(2):327–356.
- List, C. (2012). The theory of judgment aggregation: An introductory review. *Synthese*, 187(1):179–207.
- List, C. and Pettit, P. (2002). Aggregating sets of judgments: An impossibility result. *Economics and Philosophy*, 18(1):89–110.
- List, C. and Puppe, C. (2009). Judgment aggregation: A survey. In Anand, P., Pattanaik, P., and Puppe, C., editors, *Handbook of Rational and Social Choice*. Oxford University Press.
- Miller, M. K. and Osherson, D. (2009). Methods for distance-based judgment aggregation. *Social Choice and Welfare*, 32(4):575–601.
- Nehring, K. and Pivato, M. (2018). The median rule in judgment aggregation. Technical report, Social Science Research Network (SSRN).
- Nehring, K., Pivato, M., and Puppe, C. (2014). The condorcet set: Majority voting over interconnected propositions. *Journal of Economic Theory*, 151:268–303.
- Rossi, F., Venable, K. B., and Walsh, T. (2011). *A Short Introduction to Preferences: Between Artificial Intelligence and Social Choice*. Synthesis Lectures on Artificial Intelligence and Machine Learning. Morgan & Claypool Publishers.
- Schechter, E. (1996). *Handbook of Analysis and its Foundations*. Academic Press.
- Slavkovik, M. and Jamroga, W. (2011). Distance-based judgment aggregation of three-valued judgments with weights. In *Proceedings of the IJCAI Workshop on Social Choice and Artificial Intelligence*, pages 81–87.

Zoi Terzopoulou
 ILLC, University of Amsterdam
 Amsterdam, The Netherlands
 Email: z.terzopoulou@uva.nl

Ulle Endriss
 ILLC, University of Amsterdam
 Amsterdam, The Netherlands
 Email: ulle.endriss@uva.nl

Ronald de Haan
 ILLC, University of Amsterdam
 Amsterdam, The Netherlands
 Email: r.dehaan@uva.nl

Appendix: Proofs

Proof of Proposition 2. We will demonstrate the characterisation of the scores only for the property of general majoritarianism, since the proof for forward majoritarianism is similar. Given an unconstrained superagenda \mathcal{A} and a finite superpopulation \mathcal{N} , we address the case for an even number $|\mathcal{N}| = 2k$ (when $|\mathcal{N}|$ is odd the proof is totally analogous). Consider an arbitrary neutral scoring function s and its induced scoring rule F_s .

Suppose that F_s satisfies general majoritarianism and consider two arbitrary $\lambda, \lambda' \in \mathbb{N}$. We first show that $s_\lambda > 0$: Since \mathcal{A} is unconstrained, there exists a consistent set $J = \{\varphi_1, \varphi_2, \dots, \varphi_\lambda\} \subseteq \mathcal{A}$ such that $J' = \{\sim\varphi_1, \varphi_2, \dots, \varphi_\lambda\}$ is also consistent. We take the agenda $A = \{\varphi_1, \dots, \varphi_\lambda, \sim\varphi_1, \dots, \sim\varphi_\lambda\} \subseteq \mathcal{A}$ and some number $n \leq |\mathcal{N}|$ of agents and we construct the unanimous profile $\mathbf{J} = (J, \dots, J)$, where all agents report the judgment set J . Obviously it holds that $\text{Con}(\mathbf{J}) = \{J\}$, and by general majoritarianism of F_s we must have that $F_s(\mathbf{J}) = \{J\}$. However, if it was the case that $s_{|J|} = s_\lambda \leq 0$, we could deduce that $F_s(\mathbf{J}) \neq \{J\}$, because then it would be $s_{\mathbf{J}}(J) = n \cdot \lambda \cdot s_\lambda \leq n \cdot (\lambda - 1) \cdot s_\lambda = s_{\mathbf{J}}(J')$ and $J' \in \mathcal{J}(A)^\bullet$. Hence, it must be true that $s_\lambda > 0$.

Then, if $s_\lambda = s_{\lambda'}$, we have that $\frac{s_\lambda}{s_{\lambda'}} < \frac{k}{k-1}$. So, assume that $s_\lambda > s_{\lambda'}$ and $\lambda > \lambda'$ (if $\lambda < \lambda'$ the proof is analogous). Consider J, J' as above and take J'' to be a consistent subset of J' with $\sim\varphi_1 \in J''$ and $|J''| = \lambda'$. Consider the following profile, corresponding to a group of $|\mathcal{N}|$ agents:

$$\mathbf{J} = (\emptyset, \overbrace{J, \dots, J}^{k-1}, \overbrace{J'', \dots, J''}^k)$$

Since $|N_{\sim\varphi_1}^{\mathbf{J}}| > |N_{\varphi_1}^{\mathbf{J}}|$ and $J'' \subset J' = (J \setminus \{\varphi_1\}) \cup \{\sim\varphi_1\}$, it is true that $\text{Con}(\mathbf{J}) = \{J'\}$. Then, F_s satisfying general majoritarianism implies that $F_s(\mathbf{J}) = \{J'\}$, so φ_1 must score strictly less than $\sim\varphi_1$ in \mathbf{J} . That is, $(k-1) \cdot s_\lambda < k \cdot s_{\lambda'}$ or $\frac{s_\lambda}{s_{\lambda'}} < \frac{k}{k-1}$.

For the other direction we will show the contrapositive. Concretely, we will prove that if F_s does not satisfy general majoritarianism, then $s_\lambda < 0$ for some λ or $\frac{s_\lambda}{s_{\lambda'}} \geq \frac{k}{k-1}$ for some λ, λ' . We start with assuming that F_s does not satisfy general majoritarianism. Since F_s —being a scoring rule—produces only complete and consistent collective judgments, failing general majoritarianism means that $F_s(\mathbf{J}) \not\subseteq \text{Con}(\mathbf{J})$ for some agenda $A \subseteq \mathcal{A}$ and some profile $\mathbf{J} \in \mathcal{J}(A)^n$ with $n \leq |\mathcal{N}|$. This means that there exists a judgment set $J \in F_s(\mathbf{J})$ and $J \notin \text{Con}(\mathbf{J})$. The fact that $J \notin \text{Con}(\mathbf{J})$ implies that there are some formulas $\varphi_1, \dots, \varphi_\ell \in J$ such that $|N_{\sim\varphi_j}^{\mathbf{J}}| > |N_{\varphi_j}^{\mathbf{J}}|$ for all $j \in \{1, \dots, \ell\}$, and $J' = (J \setminus \{\varphi_1, \dots, \varphi_\ell\}) \cup \{\sim\varphi_1, \dots, \sim\varphi_\ell\}$ is consistent. Since $J \in F_s(\mathbf{J})$, it must hold that $\sum_{i \in N} s_{|J_i|} \cdot |J \cap J_i| \geq \sum_{i \in N} s_{|J_i|} \cdot |J' \cap J_i|$. This means that $\sum_{j \in \{1, \dots, \ell\}} \sum_{i \in N_{\varphi_j}^{\mathbf{J}}} s_{|J_i|} \geq \sum_{j \in \{1, \dots, \ell\}} \sum_{i \in N_{\sim\varphi_j}^{\mathbf{J}}} s_{|J_i|}$. Consequently, $(|N_{\varphi_1}^{\mathbf{J}}| + \dots + |N_{\varphi_\ell}^{\mathbf{J}}|) \cdot \max_{|J_i|} s_{|J_i|} \geq (|N_{\sim\varphi_1}^{\mathbf{J}}| + \dots + |N_{\sim\varphi_\ell}^{\mathbf{J}}|) \cdot \min_{|J_i|} s_{|J_i|}$. Now, if $s_\lambda < 0$ for some λ , we are done. Otherwise, it must hold that $\min_{|J_i|} s_{|J_i|} > 0$. Then, we have that $|N_{\varphi_1}^{\mathbf{J}}| + \dots + |N_{\varphi_\ell}^{\mathbf{J}}| > 0$. So $\frac{|N_{\sim\varphi_1}^{\mathbf{J}}| + \dots + |N_{\sim\varphi_\ell}^{\mathbf{J}}|}{|N_{\varphi_1}^{\mathbf{J}}| + \dots + |N_{\varphi_\ell}^{\mathbf{J}}|} \geq \frac{|N_{\varphi_1}^{\mathbf{J}}| + \dots + |N_{\varphi_\ell}^{\mathbf{J}}| + \ell}{|N_{\varphi_1}^{\mathbf{J}}| + \dots + |N_{\varphi_\ell}^{\mathbf{J}}|} = 1 + \frac{\ell}{|N_{\varphi_1}^{\mathbf{J}}| + \dots + |N_{\varphi_\ell}^{\mathbf{J}}|} \geq 1 + \frac{\ell}{\ell(k-1)} = \frac{k}{k-1}$. We conclude that $(k-1) \cdot \max_{|J_i|} s_{|J_i|} \geq k \cdot \min_{|J_i|} s_{|J_i|}$, which means that $\frac{s_\lambda}{s_{\lambda'}} \geq \frac{k}{k-1}$ for some λ, λ' . \square

Proof of Theorem 3. We give a sketch of the proof. If general majoritarianism holds for an infinite superpopulation \mathcal{N} , then it has to hold for every finite group $\{1, \dots, n\} \subseteq \mathcal{N}$ of agents. So, following the idea of Proposition 2 we can show that (i) $s_\lambda > 0$ and (ii) $\frac{s_\lambda}{s_{\lambda'}} < \frac{k}{k-1}$ for all $k \in \mathbb{N}$, implying that $s_\lambda = s_{\lambda'} > 0$ for all λ, λ' . \square

Proof of Theorem 4. Analogous to the proof of Theorem 3. \square

Proof of Theorem 5. Consider an unconstrained superagenda \mathcal{A} and an infinite superpopulation \mathcal{N} . For the first direction we take an arbitrary profile $\mathbf{J} = (J_1, \dots, J_n) \in \mathcal{J}(A)^n$ for some $A \subseteq \mathcal{A}$ and some $n \leq |\mathcal{N}|$ and a formula $\varphi \in A$ logically independent of \mathbf{J} . The independence assumption for φ implies that the constant-scoring rule F^c with $c > 0$ will simply count how many agents report the formulas φ and $\sim\varphi$ in the profile \mathbf{J} and proceed as follows: if $N_\varphi^{\mathbf{J}} = N_{\sim\varphi}^{\mathbf{J}}$, then both φ and $\sim\varphi$ will belong to some collective judgment set in $F^c(\mathbf{J})$; if $N_\varphi^{\mathbf{J}} > N_{\sim\varphi}^{\mathbf{J}}$, then only φ will belong to some collective judgment set in $F^c(\mathbf{J})$; and if $N_\varphi^{\mathbf{J}} < N_{\sim\varphi}^{\mathbf{J}}$, then only $\sim\varphi$ will belong to some collective judgment set in $F^c(\mathbf{J})$. This means exactly that if $\varphi \in J$ for all $J \in F^c(\mathbf{J})$, then $N_\varphi^{\mathbf{J}} > N_{\sim\varphi}^{\mathbf{J}}$, and thus $\varphi \in m(\mathbf{J})$. Hence backward majoritarianism holds.

For the other direction we have to consider two cases. First, the neutral-scoring rules associated with zero or negative constant scores obviously violate backward majoritarianism. Hence, it suffices to show that if backward majoritarianism is satisfied by some scoring rule $F_s \in \mathcal{S}_N$, induced by a neutral scoring function s , then it must hold that $s_\lambda = s_{\lambda'}$ for all λ, λ' . We proceed with proving the contrapositive. Take a neutral scoring function s such that $s_\lambda \neq s_{\lambda'}$ for some $\lambda \neq \lambda'$. Suppose that $s_\lambda > s_{\lambda'}$ and $\lambda > \lambda'$ (if $\lambda < \lambda'$ the proof is analogous). Then, since \mathcal{A} is unconstrained, there is some consistent judgment set $J \subseteq \mathcal{A}$ of size λ such that $\varphi \in J$ and $J' = (J \setminus \{\varphi\}) \cup \{\sim\varphi\}$ is also consistent. Consider the agenda $A \subseteq \mathcal{A}$ that contains all formulas in J and their complements, and take $J'' \subset J'$ such that $\sim\varphi \in J''$ and $|J''| = \lambda'$. Then, we construct the profile $\mathbf{J} = (\emptyset, \dots, \emptyset, J, J'')$, where by definition, φ is logically independent of \mathbf{J} . Since $s_{|J''|} = s_{\lambda'} < s_\lambda = s_{|J|}$, we will have that $\varphi \in J$ for all $J \in F_s(\mathbf{J})$. However, $|N_\varphi^{\mathbf{J}}| = |N_{\sim\varphi}^{\mathbf{J}}| = 1$, thus $\varphi \notin m(\mathbf{J})$. We conclude that backward majoritarianism fails. \square

Proof of Theorem 6. Consider given an unconstrained superagenda \mathcal{A} and a superpopulation \mathcal{N} with $|\mathcal{N}| \geq \frac{|\mathcal{A}|}{2} + 1$. For the first direction we need to show that the scoring rule F^{ee} satisfies the single-splitting axiom (since F^{ee} is induced by positive scores, we know it satisfies unanimity by Lemma 1). So we consider an arbitrary instance of the single-splitting axiom: an agenda $A \subseteq \mathcal{A}$, a group N of $n \leq |\mathcal{N}|$ agents, a non-empty subgroup $N' \subseteq N$ of agents whose judgment sets are pairwise disjoint, mutually consistent, and singleton, and two profiles $\mathbf{J} = (J_1, \dots, J_n)$ and \mathbf{J}' , where \mathbf{J}' arises from \mathbf{J} by replacing the judgment set of each member of N' by the union $\bigcup_{i \in N'} J_i$. By the definition of the scores for the rule F^{ee} we know that every formula in $\bigcup_{i \in N'} J_i$ scores exactly the same in the profiles \mathbf{J} and \mathbf{J}' , because $|J_i| = 1$ for every $i \in N'$, and $s_\lambda = \frac{s_1}{\lambda}$ for every λ . The same holds for all formulas that are not in $\bigcup_{i \in N'} J_i$ too, since they trivially appear in exactly the same judgment sets in both profiles \mathbf{J} and \mathbf{J}' . This means that $F^{ee}(\mathbf{J}) = F^{ee}(\mathbf{J}')$.

For the other direction consider an arbitrary $\lambda \leq |\mathcal{N}| - 1$ (note that we do not need to consider larger values for λ because $|\mathcal{N}| \geq \frac{|\mathcal{A}|}{2} + 1$). Since \mathcal{A} is unconstrained, we take a consistent subset $J = \{\varphi_1, \varphi_2, \dots, \varphi_\lambda\}$ such that $\{\sim\varphi_1, \varphi_2, \dots, \varphi_\lambda\}$ is also consistent. Consider the agenda A that contains all formulas in J and their complements and take two profiles $\mathbf{J} = (\{\sim\varphi_1\}, \{\varphi_1\}, \{\varphi_2\}, \dots, \{\varphi_\lambda\})$ and $\mathbf{J}' = (\{\sim\varphi_1\}, \{\varphi_1, \varphi_2, \dots, \varphi_\lambda\}, \{\varphi_1, \varphi_2, \dots, \varphi_\lambda\}, \dots, \{\varphi_1, \varphi_2, \dots, \varphi_\lambda\})$ (this can be done because $\lambda \leq |\mathcal{N}| - 1$). Then, if a scoring rule F_s induced by a neutral scoring function s satisfies the single-splitting axiom, it must be the case that $F_s(\mathbf{J}) = F_s(\mathbf{J}')$. So the scores of φ_1 and $\sim\varphi_1$ should be the same in profile \mathbf{J}' , that is, $s_1 = \lambda \cdot s_\lambda$. If moreover F_s satisfies unanimity, then by Lemma 1 it holds that $s_\lambda > 0$. Hence, by Definition 2, F_s is the equal-and-even-scoring rule. \square

Proof of Theorem 7. Trivially, if a neutral-scoring rule F_s satisfies the equal-splitting axiom (together with unanimity), then it satisfies the single-splitting axiom (together with

unanimity) too, so Theorem 6 implies that F_s must be the equal-and-even-scoring rule. For the other direction it is easy to see that for all scores of the form $s_\lambda = \frac{s_1}{\lambda} > 0$ (which induce the equal-and-even-scoring rule by Definition 2), the equal-splitting axiom is satisfied (note, however, that the arbitrary-splitting axiom is not). \square

Proof of Proposition 8. For the sake of contradiction, assume that there exists some neutral scoring function s for which the induced scoring rule F_s satisfies both the arbitrary-splitting axiom and unanimity. Then F_s should clearly also satisfy the single-splitting axiom. Hence, by Theorem 6, we should have that $s_\lambda = \frac{s_1}{\lambda} > 0$ for every λ . However, we will now show that if this is the case, then F_s can in fact not satisfy the arbitrary-splitting axiom, which stands in direct contradiction to our assumption. Consider three formulas $\varphi_1, \varphi_2, \varphi_3 \in \mathcal{A}$ such that both $\{\varphi_1, \varphi_2, \varphi_3\}$ and $\{\sim\varphi_1, \varphi_2, \varphi_3\}$ are consistent (which is possible due to unconstrainedness), and take the agenda $A = \{\varphi_1, \varphi_2, \varphi_3, \sim\varphi_1, \sim\varphi_2, \sim\varphi_3\}$. Then, consider $\mathbf{J} = (\{\sim\varphi_1\}, \{\varphi_1\}, \{\varphi_2, \varphi_3\})$ and $\mathbf{J}' = (\{\sim\varphi_1\}, \{\varphi_1, \varphi_2, \varphi_3\}, \{\varphi_1, \varphi_2, \varphi_3\})$. Since both φ_1 and $\sim\varphi_1$ score s_1 in profile \mathbf{J} , it holds that both φ_1 and $\sim\varphi_1$ belong to some judgment set in $F_s(\mathbf{J})$. Now, note that $s_{\{\varphi_1, \varphi_2, \varphi_3\}} = \frac{s_1}{3}$ (this is true given that $|\mathcal{N}| \geq 4$ and thus Theorem 6 holds). So, φ_1 scores $\frac{2s_1}{3} < s_1$ in \mathbf{J}' . This means that for every judgment set in $F_s(\mathbf{J}')$, only $\sim\varphi_1$ but not φ_1 will belong to that judgment set. To conclude, $F_s(\mathbf{J}) \neq F_s(\mathbf{J}')$, which violates the arbitrary-splitting axiom. \square

Proof of Lemma 9. By definition, the upward-lexicographic rule is induced by a neutral scoring function s with $\frac{s_{\lambda-1}}{s_\lambda} > \lambda \cdot (|\mathcal{N}| - 1)$ for all $\lambda > 1$ and is such that $s_\lambda > 0$ for all λ . It remains to be shown that all scoring functions of this kind in fact define the same aggregation rule. So consider any two neutral scoring functions s and s' with $s_\lambda > 0$ and $s'_\lambda > 0$ for all λ , as well as with $\frac{s_{\lambda-1}}{s_\lambda} > \lambda \cdot (|\mathcal{N}| - 1)$ and $\frac{s'_{\lambda-1}}{s'_\lambda} > \lambda \cdot (|\mathcal{N}| - 1)$ for all $\lambda > 1$. We need to show that $F_s = F_{s'}$. Consider an arbitrary agenda $A \subseteq \mathcal{A}$, $n \leq |\mathcal{N}|$, and profile $\mathbf{J} \in \mathcal{J}(A)^n$, as well as two arbitrary judgment sets $J, J' \in \mathcal{J}(A)^\bullet$. We are done if we can show that $s_{\mathbf{J}}(J) \geq s_{\mathbf{J}}(J')$ holds if and only if $s'_{\mathbf{J}}(J) \geq s'_{\mathbf{J}}(J')$ does. We can rewrite $s_{\mathbf{J}}(J)$ as follows:

$$s_{\mathbf{J}}(J) = \sum_{i \in \mathcal{N}} \sum_{\varphi \in J} s_{J_i}(\varphi) = \sum_{i \in \mathcal{N}} s_{|J_i|} \cdot |J \cap J_i| = \sum_{\lambda \geq 1} s_\lambda \cdot \sum_{i \mid \lambda = |J_i|} |J \cap J_i|$$

In case profile \mathbf{J} is such that all agents report judgment sets of the same size λ , our claim clearly holds, as then $s_{\mathbf{J}}(J) = s_\lambda \cdot \sum_{i \in \mathcal{N}} |J \cap J_i|$, meaning that $s_{\mathbf{J}}(J) \geq s_{\mathbf{J}}(J')$ reduces to $\sum_{i \in \mathcal{N}} |J \cap J_i| \geq \sum_{i \in \mathcal{N}} |J' \cap J_i|$ (because we know that $s_\lambda > 0$). So, without loss of generality, we may assume that at most $|\mathcal{N}| - 1$ agents report a judgment set of the same size and thus $\sum_{i \mid \lambda = |J_i|} |J \cap J_i| \leq (|\mathcal{N}| - 1) \cdot \lambda$ for all λ (and accordingly for J'). Together with our assumptions regarding the lower bounds for $\frac{s_{\lambda-1}}{s_\lambda}$ and $\frac{s'_{\lambda-1}}{s'_\lambda}$, this implies that $s_{\mathbf{J}}(J) \geq s_{\mathbf{J}}(J')$ holds if and only if the $\frac{|A|}{2}$ -vector $(\sum_{i \mid \lambda = |J_i|} |J \cap J_i|)_{\lambda \geq 1}$ lexicographically precedes $(\sum_{i \mid \lambda = |J_i|} |J' \cap J_i|)_{\lambda \geq 1}$. In other words, the relative ranking of J and J' does not depend on the exact choice of scores, provided the condition on ratios is respected, i.e., we have shown that $s_{\mathbf{J}}(J) \geq s_{\mathbf{J}}(J')$ if and only if $s'_{\mathbf{J}}(J) \geq s'_{\mathbf{J}}(J')$. \square

Proof of Theorem 10. First it is easy to verify that the upward-lexicographic rule satisfies the quality-over-quantity axiom. For the other direction consider any finite superpopulation \mathcal{N} and any superagenda \mathcal{A} of the required kind. Let F_s be a scoring rule for \mathcal{A} and \mathcal{N} that is induced by some neutral scoring function s . Since \mathcal{A} contains the set of all propositional variables $\{p_1, p_2, \dots\}$, if F_s satisfies the quality-over-quantity axiom it is easy to see that $s_\lambda > 0$ for every λ : We can construct an agenda A with $p_1, \dots, p_\lambda, p_{\lambda+1} \in A$ and the

judgment sets $J = \{p_1, \dots, p_\lambda\} \in \mathcal{J}(A)$ and $J' = \{p_1, \dots, p_\lambda, p_{\lambda+1}\} \in \mathcal{J}(A)$. Afterwards, we may consider the profile $\mathbf{J} = (\emptyset, \dots, \emptyset, J, J')$. The quality-over-quantity property implies that $J \subseteq J''$ for all $J'' \in F_s(\mathbf{J})$, which is only true if $s_{|J|} = s_\lambda > 0$.

Now consider any $\lambda > 1$. We are going to establish a lower bound on the ratio $\frac{s_{\lambda-1}}{s_\lambda}$ based on quality-over-quantity. As \mathcal{A} is closed under conjunction of literals and contains all propositional variables, we can construct an agenda A with $p_1, \dots, p_{2\lambda-2} \in A$ and $\neg p_1 \wedge \dots \wedge \neg p_\lambda \in A$. We then use the judgment sets $J = \{p_1, \dots, p_\lambda\} \in \mathcal{J}(A)$ and $J' = \{\neg p_1 \wedge \dots \wedge \neg p_\lambda, p_{\lambda+1}, \dots, p_{2\lambda-2}\} \in \mathcal{J}(A)$ to construct the profile $\mathbf{J} = (J, \dots, J, J')$ for $n = |\mathcal{N}|$ agents in which $|\mathcal{N}| - 1$ agents submit J and a single agent submits J' . As $|J| = \lambda$ and $|J'| = \lambda - 1$, the quality-over-quantity axiom requires that $J' \subseteq J''$ for all $J'' \in F_s(\mathbf{J})$. Given that the propositions $p_{\lambda+1}, \dots, p_{2\lambda-2}$ are logically independent of the profile \mathbf{J} , having that $J' \subseteq J''$ for all $J'' \in F_s(\mathbf{J})$ means that accepting $\neg p_1 \wedge \dots \wedge \neg p_\lambda$ must have yielded a higher total score than accepting all of p_1, \dots, p_λ , i.e., we must have $s_{\mathbf{J}}(\{\neg p_1 \wedge \dots \wedge \neg p_\lambda\}) > s_{\mathbf{J}}(J)$. But this is equivalent to $1 \cdot 1 \cdot s_{\lambda-1} > (|\mathcal{N}| - 1) \cdot \lambda \cdot s_\lambda$, which in turn is equivalent to $\frac{s_{\lambda-1}}{s_\lambda} > \lambda \cdot (|\mathcal{N}| - 1)$. The claim now follows from Lemma 9. \square

Proof of Proposition 11. For an infinite and unconstrained superagenda \mathcal{A} and a superpopulation \mathcal{N} , we will assume, aiming for a contradiction, that the downward-lexicographic rule F^{dlex} belongs to \mathcal{S}_N . In other words, we assume that F^{dlex} is induced by a neutral scoring function s , or equivalently, with a scoring vector $\mathbf{s} = (s_1, s_2, \dots)$. We will first prove that if this is the case, then it must hold that (i) $s_\lambda > 0$ and (ii) $s_\lambda > s_{\lambda+1}$ for all $\lambda > 0$. This means that \mathbf{s} corresponds to a decreasing and bounded sequence of real numbers, which by the monotone convergence theorem (see, e.g., Schechter, 1996) has to converge to a non-negative real number. Then, we will show that (iii) $s_\lambda - s_{\lambda+1} > s_{\lambda-1} - s_\lambda$ for all $\lambda > 1$, which in simple words means that consecutive members of the sequence \mathbf{s} must be further and further from each other as λ grows large, and hence the sequence cannot converge. To prove the statements (i) – (iii), consider an arbitrary $\lambda > 0$ and a consistent judgment set J of size λ such that $J' = (J \setminus \{\varphi\}) \cup \{\sim\varphi\}$ is also consistent, where $\varphi \in J$ (such a J exists because the superagenda is unconstrained). Take the agenda A that contains all formulas in J and their complements. For (i), we construct the unanimous profile $\mathbf{J} = (J, \dots, J)$. According to the definition of the downward-lexicographic rule, it must be the case that $F^{dlex}(\mathbf{J}) = \{J\}$, which implies that $s_\lambda > 0$ (otherwise it would hold that $J' \in F^{dlex}(\mathbf{J})$). For (ii), take J and J' as before, but now with $|J| = |J'| = \lambda + 1$. Then consider a judgment set $J'' = J' \setminus \{\psi\}$, where $\psi \neq \sim\varphi$ (that is, $|J''| = \lambda$). Using the profile $\mathbf{J}' = (\emptyset, \dots, \emptyset, J, J'')$, the definition of the downward-lexicographic rule prescribes that $F^{dlex}(\mathbf{J}') = \{J'\}$, which implies that $s_\lambda > s_{\lambda+1}$. Finally, for (iii) we assume that $\lambda > 1$ and we additionally consider the judgment set $J''' = J \setminus \{\psi', \psi''\}$, where $\psi', \psi'' \neq \varphi$, which is of size $\lambda - 1$. We construct the following profile: $\mathbf{J}''' = (\emptyset, \dots, \emptyset, J''', J'', J, J''')$. Then, $K_{\lambda+1}^{\mathbf{J}'''}(J) = 2 \cdot (\lambda - 1) + \lambda + 1 + \lambda - 1 = 2 \cdot \lambda + \lambda + \lambda - 2 = K_{\lambda+1}^{\mathbf{J}'''}(J')$ and $K_\lambda^{\mathbf{J}'''}(J) = 2 \cdot (\lambda - 1) + \lambda - 1 < 2 \cdot \lambda + \lambda - 2 = K_\lambda^{\mathbf{J}'''}(J')$, which means that $F^{dlex}(\mathbf{J}''') = \{J'\}$. Thus it must be the case that $2 \cdot s_\lambda > s_{\lambda-1} + s_{\lambda+1}$ or equivalently $s_\lambda - s_{\lambda+1} > s_{\lambda-1} - s_\lambda$.

Now, by (i) and (ii) we have that the sequence of scores in \mathbf{s} converges. Take $0 < \alpha < s_{\lambda-1} - s_\lambda$. By the definition of convergence we know that there must exist some $k' \in \mathbb{N}$ such that $s_k - s_{k+1} < \alpha$ for all $k \geq k'$. But by (iii) we see that this is impossible. We reached a contradiction and we can conclude that the downward-lexicographic rule F^{dlex} does not belong to \mathcal{S}_N . \square