The median rule in judgement aggregation

Klaus Nehring¹ · Marcus Pivato²

Received: 27 April 2020 / Accepted: 29 January 2021 / Published online: 19 March 2021
© The Author(s) 2021

Abstract

A *judgement aggregation rule* takes the views of a collection of voters over a set of interconnected issues and yields a logically consistent collective view. The *median rule* is a judgement aggregation rule that selects the logically consistent view which minimizes the *average distance* to the views of the voters (where the “distance” between two views is the number of issues on which they disagree). In the special case of preference aggregation, this is called the *Kemeny rule*. We show that, under appropriate regularity conditions, the median rule is the unique judgement aggregation rule which satisfies three axioms: *Ensemble Supermajority Efficiency*, *Reinforcement*, and *Continuity*. Our analysis covers aggregation problems in which the consistency restrictions on input and output judgements may differ. We also allow for issues to be weighted, and provide numerous examples in which issue weights arise naturally.

Keywords Judgement aggregation · Majoritarianism · Reinforcement · Consistency · Median

JEL Classification D71

Versions of the paper have been presented at the 2016 Meeting of Society for Social Choice and Welfare (Lund), the Workshop on Decision Making and Contest Theory (Kibbutz Ein Gedi, 2016), D-TEA (Paris 2017), and various seminars. We are grateful to the participants at these presentations for their valuable suggestions. We also thank Florian Brandl, Jerome Lang and Clemens Puppe for helpful comments. Finally, we thank two reviewers for very perceptive reports and valuable suggestions. M. Pivato: An early version of this paper was written while Pivato was at the Department of Mathematics of Trent University in Canada. This research was supported by NSERC Grant #262620-2008 and Labex MME-DII (ANR11-LBX-0023-01).

✉ Klaus Nehring
kdnehring@ucdavis.edu
Marcus Pivato
marcupivato@gmail.com

¹ Department of Economics, UC Davis, California, USA

² THEMA, CNRS UMR 8184, CY Cergy Paris Université, Cergy-Pontoise Cedex, France

1 Introduction

In judgement aggregation, a group is faced with a joint decision; frequently, the members of the group disagree about which decision the group should take and/or the grounds for the decision. Complex decisions can often be described as an interrelated set of judgements on a set of binary issues subject to some admissibility constraints. Admissibility constraints may be logical, normative or physical.

Judgement aggregation theory attempts to determine normative criteria on how to best resolve the disagreement at hand. Such “resolution of disagreement” can be understood in two ways: *compromise* or *consensus*. By “consensus”, we mean a well-supported inference from the position of a hypothetical impartial observer. By “compromise” we mean the best accommodation of each member’s views from their own perspective. In a slogan, the consensus perspective aims to maximize *plausibility*, while the compromise perspective aims to maximize *concordance*.

Judgement aggregation pertains both to groups that act as sovereigns (“electorates”), as in democratic elections and referenda, and to groups that act as delegate bodies (“committees”), such as multi-member courts, corporate boards, central banks. A key rationale for the existence of many committees is the production of decisions that are sound from an independent third-party perspective; in those cases, the judgement aggregation framework seems especially apt, arguably often more so than the more common and established preference aggregation framework.

Given a profile of views by the group members (henceforth: “voters”), which view should the group adopt? Which view enjoys the highest “support” (“plausibility” or “concordance”)? In earlier work (Nehring et al. 2014, 2016; Nehring and Pivato 2014, 2019), we have explored a “majoritarian” approach to this question. Its hallmark is to evaluate support issue by issue in terms of the sign and size of issue-wise majorities. In simple cases in which the issue-wise majorities happen to produce a jointly feasible view, on the majoritarian approach, this view enjoys the highest support, and should thus be adopted by the group. But issue-wise majorities may well not yield a consistent view. In the context of preference aggregation, this is the well-known Condorcet paradox. Analogous inconsistencies are common in judgement aggregation (Guilbaud 1952; Kornhauser and Sager 1986; List and Pettit 2002), and have stimulated an intensive investigation in economic theory, philosophy, and computer science; see e.g. List and Puppe (2009), Mongin (2012) and Lang et al. (2017) for surveys.

A satisfactory normative account thus needs to be able to deal with the Condorcet-inconsistent cases. In contrast to the focus on impossibility results in much of the literature just referenced, in this paper we make a case for a particular judgement aggregation procedure: the *median rule*. The median rule maximizes the total numerical support (number of votes) for a view, summed over all issues. Equivalently, the median rule minimizes the average distance to the views of the voters (where the “distance” between two views is measured by the number of issues on which they differ).

Our characterization of the median rule is based on three axioms. First, majoritarianism over multiple issues is encoded in an axiom called “Ensemble Supermajority Efficiency” (ESME), which itself is an extension of the normatively more basic principle of “Supermajority Efficiency” (SME). The SME principle says that if, in the

comparison of two admissible views $\mathbf{x}$ and $\mathbf{y}$, $\mathbf{x}$ agrees with the majority on more issues than $\mathbf{y}$, then $\mathbf{y}$ is inferior as a group view, and thus should not be adopted by the group. Furthermore, and more demanding, for any fixed ‘quorum’ q , we count how many supermajorities of size at least q agree with $\mathbf{x}$ or with $\mathbf{y}$. If $\mathbf{x}$ agrees with at least as many size- q supermajorities as $\mathbf{y}$ for all values of q , and $\mathbf{x}$ agrees with *more* size- q supermajorities than $\mathbf{y}$ for *some* value of q , then again $\mathbf{y}$ is inferior as a group view, and should not be adopted by the group. Thus, SME takes into account the size of the supporting supermajorities in a manner analogous to first-order stochastic dominance in decision theory. This is a simple and natural extension of Condorcet’s majoritarian approach to social choice. It acquires extra force when it is applied simultaneously to combinations of different, logically independent judgement aggregation problems. This additional power can be harnessed in a larger framework, with the help of an appropriate additional Combination axiom (Nehring and Pivato 2019); however, in this paper we stick to the standard framework and obtain these implications by leveraging the SME axiom to a more elaborate ESME axiom. The ESME axiom extends SME to “ensembles” of judgement aggregation problems consisting of multiple instances of the same type of judgement aggregation problem but with different profiles.

Second, the axiom of Reinforcement says that if two subpopulations independently choose the same view under the rule, then the combined population should also choose this view under the rule. It is a standard, highly versatile axiom originally due to Smith (1973) and Young (1974). Finally, Continuity asserts that the group judgement is robust under small perturbations of the distribution of input judgements. Our first main result, Theorem 1, characterizes the median rule as the unique judgement aggregation rule satisfying ESME, Reinforcement and Continuity. Theorem 1 is based on a weak minimal richness on the input space called “thickness”.

The main comparison result in the literature is the remarkable characterization of the median rule in the aggregation of linear orderings (“rankings”) by Young and Levenglick (1978). In this setting (mathematically equivalent to the setting of Arrovian preference aggregation), the median rule is also known as the Kemeny (1959) rule.¹

When applied to the aggregation of ordinal rankings, our axiomatic framework is broadly similar to that of Young and Levenglick (in particular, its Condorcetian foundation) but differs in the specifics, our axioms being on the whole stronger. We thus do not claim that when applied to the aggregation of rankings, our Theorem 1 improves on Young and Levenglick’s result. Our aims are simply different. While their result relies heavily on particular features of the combinatorial geometry of the space of ordinal rankings, Theorem 1 is a general-purpose result that covers a wide range of judgement aggregation problems. (See the end of Sect. 4 for a more detailed comparison.)

Theorem 1 treats all issues symmetrically. This is warranted in many standard applications whose structure is sufficiently symmetric, such as the aggregation of rankings (linear or weak orderings), of classifiers (equivalence relations) or in multi-winner choice problems. But other applications lack these symmetries; furthermore, different issues, whether or not formally symmetric, may be given different “importance”.

¹ The Kemeny rule has perhaps more applications beyond social choice than any other concept in social choice (databases, search engines, social rankings etc.) and has been studied more than any other voting rule by theoretical computer scientists; see Fischer et al. (2016).

For example, consider truth-functional aggregation, which was the focus of much of the early literature in judgment aggregation inspired by Kornhauser and Sager's (1986) "doctrinal paradox". In truth-functional aggregation, one or more "conclusion judgments" are logically ("truth-functionally") determined by a number of "premise judgments". Condorcet inconsistency takes the form of the "discursive dilemma": issuewise aggregation of majorities on the premises may well determine (by truth-functional implication) judgments on the conclusions that differ from the majority judgment on these conclusions. The discursive dilemma can be resolved via the median rule by trading off majority overrides on the premises against majority overrides on the conclusions. However, in view the structural and conceptual asymmetry between premises and conclusions, they have different standing, and it would appear quite arbitrary to give them equal weight.

In Sect. 5, we thus generalize the analysis to *weighted* judgment contexts in which different issues have different weights. The definitions of SME, ESME and the median rule generalize naturally. However, as shown by counterexample, Theorem 1 does not carry over in full generality. The characterization for weighted judgement contexts in Theorem 3 must invoke not just richness conditions on the space of admissible *input* judgments, but also restrictions on the "combinatorial geometry" of the space of admissible *output* judgments; these have no precedent in the extant literature. Most previous work on judgment aggregation assumes the input and output spaces to be the same.² But we allow them to differ. This adds useful additional generality at very modest cost in execution.

Due to the abstraction and generality of our judgment aggregation framework, it has a broad and diverse range of applications. We thus illustrate our concepts and results in a number of examples, including applications to approval voting on committees with composition constraints, assignment problems, uniform treatment of heterogeneous cases, missing information, and multiple criteria. For related contributions in various application settings, see the references in Sects. 3 and 5. See also Nehring and Puppe (2007), Nehring and Pivato (2011, 2019) or Nehring et al. (2014) for many more examples.

At a formal level, the scope of potential applications of Theorems 1 and 3 is extremely broad. But it would be rash to suppose the theorems to be normatively compelling in all cases where they are formally applicable. Instead, the content and appeal of ESME will depend on the particular setting and the way in which issues are framed, and need to be judged accordingly. In some judgement contexts, it may be possible to define a notion of the "average" opinion; for some applications, judgement aggregation rules based on such "averages" might be more normatively compelling than majoritarian rules. But even in such cases, majoritarian rules are likely to enjoy significant second-best advantages, such as superior resistance to strategic manipulation.³

This paper belongs to a larger project exploring multi-issue majoritarianism in judgement aggregation. In particular, Theorems 1 and 3 rely on the main result of Nehring and Pivato (2019), which shows that judgement aggregation rules that satisfy

² A recent exception is Endriss (2018).

³ See, e.g. Bossert and Sprumont (2014) for the Kemeny rule in the ranking problem.

ESME and Continuity are representable as “additive majority rules”. Additive majority rules can be viewed as non-linear generalizations of the median rule; they evaluate views not simply in terms of the (weighted) sum of numeric issue-wise majorities, but in terms of the (weighted) sum of issue-wise majority *gains*, which are possibly non-linear transformations of these numeric issue-wise majorities. Theorems 1 and 3 are much more focused results: augmenting the other two axioms with Reinforcement yields an additive majority rule with a *linear* gain function.

While plausible on its face, the connection between Reinforcement and linearity of the gain function is technically far from trivial. Difficulties arise in particular because the underlying representation obtained in Nehring and Pivato (2019) is hyperreal-valued rather than real-valued, and due to the need of combinatorial structural conditions in the weighted case. Moreover, conceptually, from a compromise perspective, the step from general additive majority rules to the median rule is fundamental. From this perspective, Reinforcement is quite compelling, and enables an interpretation of the optimal aggregation rule as minimizing the “aggregate burden of compromise”, as measured by the average distance between individual and group judgments. By contrast, from a consensus perspective, Reinforcement is somewhat less compelling—indeed, non-linear gain-functions offer useful degrees of freedom in distinguishing the effective strength of majority in the overall balance of argument from its numerical strength. At the same time, it is useful also from this perspective to be able to single out a particular gain-function as a default, avoiding recourse to *ad hoc* conventions. To this purpose, we introduce an axiom of Judgement Consistency, and show its functional equivalence to Reinforcement (Proposition 2). This makes Theorems 1 and 3 applicable to the consensus perspective as well.

The rest of this paper is organized as follows. Section 2 sets up the formal framework. Section 3 introduces the axioms of ESME and Continuity, and explains that additive majority rules are the only rules satisfying these axioms. Section 4 introduces the Reinforcement and Judgement Consistency axioms, and states the main result. Section 5 introduces *weighted* judgement contexts, and Sect. 6 extends our axiomatic characterization to such contexts. Appendix A discusses conditions under which the equal-weights assumption of Theorem 1 is appropriate. Appendix B contains the proofs of some auxiliary results. Appendix C contains the proofs of the main results.

2 Judgement aggregation

Let $\mathcal{K}$ be a finite set of logical propositions or issues, each of which can be either affirmed or denied. A *view* is an assignment of an assertoric (Yes-No) value to each issue, represented by an element of $\{\pm 1\}^{\mathcal{K}}$. A *judgement space* is a collection of views—that is, a subset $\mathcal{X}$ of $\{\pm 1\}^{\mathcal{K}}$ —determined by certain constraints. These constraints can arise in several ways: as a matter of logical consistency (as in truth-functional aggregation problems), as a matter of “rational coherence” (as in transitivity conditions on orderings) or as a matter of mere “feasibility” (as in multi-winner choice problems).

Example 1 (*Aggregation of rankings*) Let $\mathcal{A}$ be a finite set of alternatives. We can represent the set of all strict ordinal rankings of $\mathcal{A}$ as a judgement space $\mathcal{X}_{\mathcal{A}}^{\text{rk}} \subset \{\pm 1\}^{\mathcal{K}}$, where elements of $\mathcal{K}$ represent assertions of the form “ $a > b$ ” for some $a, b \in \mathcal{A}$, and admissibility is given by transitivity. To be precise, let $\mathcal{K} \subset \mathcal{A} \times \mathcal{A}$ be a subset such that $(a, a) \notin \mathcal{K}$ for any $a \in \mathcal{A}$, and for each distinct $a, b \in \mathcal{A}$, exactly one of the pairs (a, b) or (b, a) is in $\mathcal{K}$. Any complete, antisymmetric binary relation $>$ on $\mathcal{A}$ can be represented by a unique element $\mathbf{x}$ of $\{\pm 1\}^{\mathcal{K}}$ by setting $x_{ab} = 1$ if $a > b$, whereas $x_{ab} = -1$ if $a < b$. Let $\mathcal{X}_{\mathcal{A}}^{\text{rk}}$ be the set of all elements of $\{\pm 1\}^{\mathcal{K}}$ corresponding to ordinal rankings of $\mathcal{A}$.⁴ Judgement aggregation on $\mathcal{X}_{\mathcal{A}}^{\text{rk}}$ is formally equivalent to classical Arrovian preference aggregation.⁵ $\square$

Example 2 (*Classifier aggregation*) We can represent the set of all equivalence relations on $\mathcal{A}$ as a judgement space $\mathcal{X}_{\mathcal{A}}^{\text{eq}} \subset \{\pm 1\}^{\mathcal{K}}$, where the elements of $\mathcal{K}$ represent assertions of the form “ $a \approx b$ ”, for some $a, b \in \mathcal{A}$. To be precise, let $\mathcal{K}$ be the set of all two-element subsets of $\mathcal{A}$. Any symmetric, reflexive binary relation $\sim$ on $\mathcal{A}$ is represented by a unique element $\mathbf{x}$ of $\{\pm 1\}^{\mathcal{K}}$ by setting $x_{ab} = 1$ if and only if $a \sim b$. Let $\mathcal{X}_{\mathcal{A}}^{\text{eq}}$ be the set of all elements of $\{\pm 1\}^{\mathcal{K}}$ corresponding to equivalence relations on $\mathcal{A}$.⁶ Judgement aggregation on $\mathcal{X}_{\mathcal{A}}^{\text{eq}}$ arises when each voter has her own way of classifying the elements of $\mathcal{A}$ into equivalence classes, and the group must agree on a common classification system. $\square$

Other judgement spaces represent common collective decision problems such as resource allocation, committee selection, or taxonomic classification. One particularly well-known class of examples are the so-called *truth-functional* aggregation problems. In this case, the issues in $\mathcal{K}$ are divided into two classes: “premises” and “conclusions”, and the truth-values of the conclusions are logically entailed by the truth-values of the premises. The space $\mathcal{X}$ is then the set of all logically consistent assignments of truth values to the premises and conclusions. See Nehring and Puppe (2007), Nehring and Pivato (2011) or Nehring et al. (2014) for many more examples.

Judgement aggregation rules map profiles of views to a group view or set of views. Typically, both outputs and inputs are subject to feasibility or logical consistency constraints, which are encoded by two judgement spaces $\mathcal{X}$ and $\mathcal{Y}$, respectively. In many cases, the restrictions on inputs and outputs are the same (so that $\mathcal{X} = \mathcal{Y}$), but they need not be. For example, one might require output views to be fully rationally coherent (e.g. transitive), but allow input views that are not, for example to accommodate bounded rationality in voters. Or output views may take into account feasibility considerations, while input views do not, as in single- or multi-winner approval voting. We thus define a *judgement context* to be a triple $\mathcal{C} := (\mathcal{K}, \mathcal{X}, \mathcal{Y})$, where $\mathcal{K}$ is a (finite) set of issues and $\mathcal{X} \subseteq \mathcal{Y} \subseteq \{\pm 1\}^{\mathcal{K}}$, with $\mathcal{Y}$ being the “input space”, and $\mathcal{X}$ being the “output space”. We assume that $\mathcal{X} \subseteq \mathcal{Y}$; i.e. output views are required to

⁴ Formally: $\mathcal{X}_{\mathcal{A}}^{\text{rk}}$ is the set of all $\mathbf{x} \in \{\pm 1\}^{\mathcal{K}}$ satisfying the following transitivity constraint, for all distinct $a, b, c \in \mathcal{A}$: if $(x_{ab} = 1 \text{ or } x_{ba} = -1)$, and $(x_{bc} = 1 \text{ or } x_{cb} = -1)$, then $(x_{ac} = 1 \text{ or } x_{ca} = -1)$.

⁵ As in much of the literature on preference aggregation, we represent rankings as linear orderings; a modified specification would allow to represent weak orders as a single negatively transitive relation.

⁶ Formally: $\mathcal{X}_{\mathcal{A}}^{\text{eq}}$ is the set of all $\mathbf{x} \in \{\pm 1\}^{\mathcal{K}}$ satisfying the transitivity constraint $(x_{ab} = x_{bc} = 1) \implies (x_{ac} = 1)$, for all distinct $a, b, c \in \mathcal{A}$.

satisfy all admissibility restrictions that input views do, and maybe more.⁷ A *profile* is a function $\mu : \mathcal{Y} \rightarrow [0, 1]$ such that $\sum_{\mathbf{y} \in \mathcal{Y}} \mu(\mathbf{y}) = 1$. This represents a population of weighted voters; for each $\mathbf{y} \in \mathcal{Y}$, $\mu(\mathbf{y})$ is the total weight of the voters who hold the view $\mathbf{y}$. If all voters have the same weight, then $\mu(\mathbf{y})$ is simply the *proportion* of the electorate which holds the view $\mathbf{y}$. But we allow the possibility that different voters have different weights, e.g. because of different levels of expertise or different stakes in the outcome. By summarizing the voters' views with a function $\mu : \mathcal{Y} \rightarrow [0, 1]$, we abstract from the exact number of voters, and we render them anonymous, except for their weights: voters with the same weight are indistinguishable in our model.

If $\mathcal{Y}$ is a judgement space, we define $\Delta(\mathcal{Y})$ to be the set of all profiles on $\mathcal{Y}$. A *judgement problem* is an ordered pair $(\mathcal{C}, \mu)$, where $\mathcal{C} = (\mathcal{K}, \mathcal{X}, \mathcal{Y})$ is a judgement context, and $\mu \in \Delta(\mathcal{Y})$. *Judgement aggregation* is the process of converting such a judgement problem into a view (or set of views) in $\mathcal{X}$. A *judgement aggregation rule* on $\mathcal{C}$ is a correspondence $F : \Delta(\mathcal{Y}) \rightrightarrows \mathcal{X}$ which produces a nonempty (usually singleton) subset $F(\mu) \subseteq \mathcal{X}$ for any profile $\mu \in \Delta(\mathcal{Y})$. (Sometimes, we will write “ $F(\mathcal{X}, \mu)$ ” instead of “ $F(\mu)$ ”).

The *median rule* is a particularly attractive judgement aggregation rule. To define it, we need some notation. Recall that $\mathcal{X} \subseteq \mathcal{Y} \subseteq \{\pm 1\}^{\mathcal{K}} \subset \mathbb{R}^{\mathcal{K}}$. Thus, each view $\mathbf{y} \in \mathcal{Y}$ can be regarded as a vector in $\mathbb{R}^{\mathcal{K}}$. For any profile $\mu \in \Delta(\mathcal{Y})$, we define its *majority vector*

$$\tilde{\mu} := \sum_{\mathbf{y} \in \mathcal{Y}} \mu(\mathbf{y}) \mathbf{y} \in [-1, 1]^{\mathcal{K}}. \quad (1)$$

For all $k \in \mathcal{K}$, we have $\tilde{\mu}_k > 0$ if a (weighted) majority of voters affirm or support the issue k , whereas $\tilde{\mu}_k < 0$ if a majority deny or oppose k . The *majority ideal* is the element $\mathbf{x}^\mu \in \{\pm 1\}^{\mathcal{K}}$ defined by setting $x_k^\mu := \text{sign}(\tilde{\mu}_k)$ for all $k \in \mathcal{K}$.⁸ However, for many profiles $\mu \in \Delta(\mathcal{Y})$ it turns out that $\mathbf{x}^\mu \notin \mathcal{X}$. (This can happen even when $\mathcal{Y} = \mathcal{X}$.) In other words, it is frequently impossible to agree with the μ -majority in *every* issue in $\mathcal{K}$, while respecting the underlying logical constraints which define the judgement space $\mathcal{X}$.

Informally, the median rule maximizes the *average agreement* with μ -majorities across all the issues in $\mathcal{K}$. Formally, for all $\mu \in \Delta(\mathcal{Y})$, we define

$$\text{Median}(\mathcal{X}, \mu) := \underset{\mathbf{x} \in \mathcal{X}}{\text{argmax}} \left(\sum_{k \in \mathcal{K}} x_k \tilde{\mu}_k \right). \quad (2)$$

(This maximizer is unique for all μ in an open, dense subset of $\Delta(\mathcal{Y})$.) As we noted in the introduction, in the special case of the aggregation of rankings (i.e. when $\mathcal{X} = \mathcal{Y} = \mathcal{A}^{\text{rk}}$), the median rule is equivalent to the Kemeny rule.

⁷ There are also social decisions in which the output views are *less* constrained than input views —i.e. where $\mathcal{Y} \subset \mathcal{X}$. (One example is the aggregation of transitive individual preferences into a quasitransitive social preference.) We do not consider this case.

⁸ For simplicity, we assume in this paragraph that $\tilde{\mu}_k \neq 0$ for all $k \in \mathcal{K}$; this is not essential.

There is another way to define and motivate the median rule via a natural notion of distance due to Kemeny (1959); see Monjardet (2008) for a broad survey. For any $\mathbf{x}, \mathbf{y} \in \{\pm 1\}^{\mathcal{K}}$, we define their *Hamming distance* by $d(\mathbf{x}, \mathbf{y}) := \#\{k \in \mathcal{K}; x_k \neq y_k\}$. It is easy to see that the median rule selects the view(s) in $\mathcal{X}$ that minimize the *average Hamming distance* to the views of the voters; in the terminology of Miller and Osherson (2009) and Lang et al. (2011), it is a *distance-minimizing rule*.⁹

This metric interpretation is particularly appealing when our goal is to find an optimal compromise. The Hamming distance of the collective view from the view of any voter is a natural measure of the “burden of compromise” imposed on that voter; thus, the median view(s) are those that minimize the aggregate burden of compromise.

3 Additive majority rules

In this section and the next, we will focus on the special and simpler case in which all issues are treated on par and are assigned equal weight. Such symmetric treatment of issues is naturally motivated in applications in which the judgment context exhibits appropriate global symmetries. For example, many spaces of binary relations (linear or weak orders as in Example 1, partial orders or acyclic relations, interval orders, equivalence relations as in Example 2, graphs with certain qualitative properties) are symmetric under relabeling of alternatives/vertices.¹⁰ By consequence, only constant weight vectors are invariant under these symmetries. Global symmetries are sufficient but not necessary condition for the use of equal weights; see Appendix A for further discussion. In Sects. 5 and 6, our approach will be generalized to comparisons of sets of issues in terms of arbitrary real-valued weights.

Supermajority Efficiency. A judgement aggregation problem is *Condorcet inconsistent* if there is no admissible view that agrees with the majority view on every issue. It is then necessary to trade off majorities on some issues against majorities on others. To motivate our approach to such trade-offs, consider a very special case in which two feasible views $\mathbf{x}, \mathbf{y} \in \mathcal{X}$ are compared at some profile μ such that the majority margin on all issues distinguishing $\mathbf{x}, \mathbf{y}$ is equal in absolute amount. So, within this comparison, one must trade off one set of issues (on which $\mathbf{x}$ but not $\mathbf{y}$ is aligned with the majority) against another another set (on which $\mathbf{y}$ is aligned with the majority but not $\mathbf{x}$). Since the majority margins on these issues are equal by assumption, the choice between $\mathbf{x}$ and $\mathbf{y}$ is naturally thought of as turning on which of these sets is “larger” or “more important”. In this section, for simplicity we will focus on the case in which all issues are treated equally; sets of issues are thus compared by their cardinality, i.e. by counting. To illustrate, consider the ranking problem on four alternatives.

Example 1 continued. Let $\mathcal{A} = \{a, b, c, d\}$, let $\mathcal{K} := \{(a, b), (a, c), (a, d), (b, c), (b, d), (c, d)\}$, and consider a profile $\mu \in \Delta(\mathcal{X}_{\mathcal{A}}^{\text{rk}})$ with majority margins $\tilde{\mu}_{ab} =$

⁹ Miller and Osherson (2009) call it *Prototype*, while Lang et al. (2011) call it $R^{dH, \Sigma}$. Distance-minimizing rules are themselves a subclass of *scoring rules*, which have been studied by Dietrich (2014).

¹⁰ In fact, all of the examples in Barthélémy and Monjardet’s (1981) classical survey belong to this category.

$\tilde{\mu}_{ac} = \tilde{\mu}_{bc} = \tilde{\mu}_{bd} = \tilde{\mu}_{cd} = \epsilon > 0$, and $\tilde{\mu}_{ad} = -\epsilon$ for some $\epsilon > 0$.¹¹ Note that the linear order $abcd$ differs from the majority ideal on a single issue only, namely the issue (a, d) ; by contrast, all other orderings differ from the majority ideal by at least two issues. Hence, based on the counting criterion of size, $abcd$ beats any other linear orderings in pairwise comparison. $\square$

Thus equipped with a notion of comparative size of sets of issues, one can derive further normative implications for general profiles by invoking the basic majoritarian principle that “many heads are better than few”. Here we apply this principle not just within issues, but across issues as well. This leads to the first key normative axiom of this paper, Supermajority Efficiency. The basic idea is to compare views at various supermajority thresholds $q \in [0, 1]$. If $\mathbf{x}$ is aligned with the majority with a margin of at least q on a (weakly, and at least once strictly) larger set of issues than $\mathbf{y}$ is, then $\mathbf{x}$ supermajority dominates $\mathbf{y}$, rendering $\mathbf{y}$ inadmissible. Note that both the size of a set of issues and of the size of a majority are here employed in a strictly ordinal sense.

Formally, for any $\mu \in \Delta(\mathcal{Y})$, $\mathbf{x} \in \mathcal{X}$ and $q \in [0, 1]$, we define $\gamma_{\mathbf{x}}^{\mu}(q) := \#\{k \in \mathcal{K}; x_k \tilde{\mu}_k \geq q\}$; this yields a non-increasing function $\gamma_{\mathbf{x}}^{\mu} : [0, 1] \rightarrow \mathbb{R}$. We say $\mathbf{x} \in \mathcal{X}$ is *supermajority efficient* (SME) for the judgement problem $(\mathcal{C}, \mu)$ if there does not exist any $\mathbf{z} \in \mathcal{X}$ such that $\gamma_{\mathbf{z}}^{\mu}(q) \geq \gamma_{\mathbf{x}}^{\mu}(q)$ for all $q \in [0, 1]$, with strict inequality for some q . Let $\text{SME}(\mathcal{C}, \mu)$ be the set of such views. If $\mathbf{x} \in \text{SME}(\mathcal{C}, \mu)$, then it is impossible to change some coordinates of $\mathbf{x}$ to capture one more μ -supermajority of size q , without either losing at least one μ -supermajority of size $q' \geq q$, or losing at least two μ -supermajorities of size $q' \leq q$. In the Condorcet consistent case, i.e. if the majority ideal $\mathbf{x}^{\mu}$ is in $\mathcal{X}$, then $\text{SME}(\mathcal{C}, \mu) = \{\mathbf{x}^{\mu}\}$. We will say that a judgement aggregation rule $F : \Delta(\mathcal{Y}) \rightrightarrows \mathcal{X}$ is *supermajority efficient* (SME) if $F(\mu) \subseteq \text{SME}(\mathcal{C}, \mu)$ for any choice of $\mu \in \Delta(\mathcal{Y})$.

In simple cases, the SME criterion is decisive by itself. To illustrate, consider the ranking problem on three alternatives, with $\mathcal{A} = \{a, b, c\}$. Let μ be a profile concentrated on the three rankings abc , bca , and cab , with frequency $\alpha > \beta > \gamma$, respectively, and $\alpha < \frac{1}{2}$. A majority ranks a above b , b above c , and c above a , a classic Condorcet cycle. Among those three comparisons, the majority ranking of $c \succ a$ is supported by the smallest majority margin, namely $\beta + \gamma - \alpha$. Since the ordering abc overrides the majority on this issue only, it is the unique supermajority efficient view. Extending this argument, it is easily seen the SME criterion generically selects a unique view when ranking three alternatives.¹²

In more complex settings, generic uniqueness is easily lost. This happens, for instance, in order aggregation on at least four alternatives. Return to the profile in Example 1, and modify it by allowing the majority margin $\tilde{\mu}_{ad}$ on (a, d) to differ from $-\epsilon$.¹³ If $\tilde{\mu}_{ad} > -\epsilon$, then the ordering $abcd$ is still the unique SME view. By contrast, if $\tilde{\mu}_{ad} < -\epsilon$, other views such as $dabc$ are SME as well. A tradeoff arises:

¹¹ If ϵ is small enough, then such a μ exists. This follows from McGarvey's (1953) Theorem: for any vector $\mathbf{v}$ in $\mathbb{R}^{\mathcal{K}}$ sufficiently close to 0, there is a profile $\mu \in \Delta(\mathcal{X}_{\mathcal{A}}^{\text{rk}})$ with $\tilde{\mu} = \mathbf{v}$ (see also Nehring and Pivato, 2011). For brevity, we will not build such a profile explicitly; it is enough to know that it exists.

¹² By “generically”, we mean on an open and dense set of profiles $\mu \in \Delta(\mathcal{Y})$ in the Euclidean topology. Furthermore, any non-uniqueness is due to a supermajority “tie” in an obvious sense.

¹³ Again, this is possible by McGarvey's Theorem; see Footnote 11.

while the ordering $abcd$ overrides the issue-wise majority on only one issue (with a large majority margin), $dabc$ overrides the issue-wise majority on two issues, but with a smaller majority margin.

To further illustrate the content of SME as well as the scope of judgement aggregation in contexts, consider Voting on Committees.

Example 3 (Voting on Committees) Suppose a committee of L members must be chosen. Let $\mathcal{K}$ be the set of candidates. For any $\mathbf{x} \in \{\pm 1\}^{\mathcal{K}}$, let $\#\mathbf{x} := \#\{k \in \mathcal{K} : x_k = 1\}$. Then the set of feasible committees is $\mathcal{X}_L^{\mathcal{K}} = \{\mathbf{x} \in \{\pm 1\}^{\mathcal{K}} : \#\mathbf{x} = L\}$. This encoding breaks down the selection of a committee into $|\mathcal{K}|$ binary issues, each of the form “should candidate k be a member of the committee?”

There are two natural input spaces. On the one hand, with $\mathcal{Y} = \mathcal{X}$, each voter’s input consists of a feasible committee, naturally interpreted as her own view of the best committee. On the other hand, with $\mathcal{Y} = \{\pm 1\}^{\mathcal{K}}$, a voter’s input consists of an independent judgement on each candidate; this could be interpreted as a judgement of “competency” or “merit” of the candidate. In line with the literature, we will refer to such judgements as judgements of “approval”, and the domain $\{\pm 1\}^{\mathcal{K}}$ as the *approval domain* on $\mathcal{K}$.

Consider any profile $\mu \in \Delta(\mathcal{Y})$; to sidestep ties, assume for simplicity that all majority margins $\tilde{\mu}_k$ are different. In such cases, the SME criterion singles out a unique committee as optimal, namely the committee composed of the L candidates with the highest majority margins $\tilde{\mu}_k$ (regardless of whether or not these margins are positive). With the approval domain and $L = 1$, SME thus selects the candidate k with the highest $\tilde{\mu}_k$; this is just the approval voting rule of Brams and Fishburn (1983). With $L > 1$, SME yields “multi-winner approval voting”. Multi-winner approval voting is rationalized here via SME which evaluates views (i.e. committees) issue by issue (i.e. candidate by candidate). $\square$

Multi-winner approval voting evaluates committees by separately assessing the merit of each member. In many situations, there is also an interest in the overall composition of the committee, for example to ensure broad expertise or broad representation of perspectives or stake-holders. One can incorporate such considerations through exogenous *composition constraints* which restrict the set of admissible committees $\mathcal{X}$ to a strict subset of $\mathcal{X}_L^{\mathcal{K}}$.

Example 4 In the selection of a university-wide committee, one may want to ensure an adequate representation of types of disciplines $j \in \mathcal{J}$ (such as humanities, social sciences, natural sciences, engineering) by imposing a minimum size L_j on the representation of group j . With $\{\mathcal{K}_j\}_{j \in \mathcal{J}}$ describing the partition of candidates according to their disciplinary type, this composition constraint yields an output space $\mathcal{X} = \{\mathbf{x} \in \mathcal{X}_L^{\mathcal{K}} : \#\{k \in \mathcal{K}_j : x_k = 1\} \geq L_j \text{ for all } j \in \mathcal{J}\}$. Again, except for possible ties, the SME criterion yields unique optimal committees: first, for each group $j \in \mathcal{J}$, select the L_j candidates with the highest majority support in that group; second, fill the remaining $L - \sum_{j \in \mathcal{J}} L_j$ candidates with those among the not-yet selected candidates with the highest majority support across all disciplines.

As an empirical example, for many years the German Green Party has adopted a rule according to which two persons must be elected as a ‘leadership duo’, of which

at least one is female. (So $L = 2$; $L_f = 1$, and $L_m = 0$.) According to this rule, the composition constraint is applied to both the input and the output space, i.e. $\mathcal{X} = \mathcal{Y}$. The voting rule employed agrees with the SME criterion and is thus, in this case, equivalent to the median rule. $\square$

In more complex situations, more than one type of classification may be considered relevant, such as academic rank, gender, or ethnicity, yielding additional, overlapping composition requirements. In those cases, SME might no longer select a uniquely optimal committee; that selection might now depend on the particular rule. The median rule in particular selects those committees $\mathbf{x}$ with the highest overall approval $\sum_{k: x_k=1} \tilde{\mu}_k$; it thus yields a natural formulation of *multi-winner approval voting under composition constraints*.¹⁴

Example 5 Suppose a committee of size two must satisfy *two* composition constraints. Modifying the Green Party example a bit, suppose that one person of each gender is to be elected, and one person of each basic political outlook (center left vs. left). By SME, attention can be restricted to those candidates with the largest number of votes for each combination of characteristics. Thus, the choice reduces to two “duos”, one with a female centrist and a male leftist, and one with male centrist and a female leftist. Suppose that the candidates fc, ml, mc , and fl garner 45%, 10%, 25% and 25% of the votes, respectively.¹⁵ Then both duos are supermajority efficient. The first one is chosen by the median rule. $\square$

Additive Majority Rules. In contexts in which SME does not generically select unique views, further considerations must be brought into play. First, one must ensure the consistency of SME choices across profiles; this yields the class of *additive majority rules*, via the axiomatic characterization given in Nehring and Pivato (2019) and briefly summarized in the next two subsections. Second, one must single out the median rule among the additive majority rules, by imposing additional axioms; this is the first main result of the present paper.

Like the median rule, additive majority rules try to maximize the “total agreement with majorities”, where the “total” is taken by summing over all issues in $\mathcal{K}$, and where the “agreement with majorities” is measured by applying an increasing function to the coordinates of the majority vector $\tilde{\mu}$. Formally, let ϕ be an increasing and odd function (called the *gain function*) defined on the interval $[-1, +1]$.¹⁶ The *additive majority rule* (AMR) $F_\phi : \Delta(\mathcal{Y}) \rightrightarrows \mathcal{X}$ is defined as follows:

$$\text{For all } \mu \in \Delta(\mathcal{Y}), \quad F_\phi(\mu) := \operatorname{argmax}_{\mathbf{x} \in \mathcal{X}} \left(\sum_{k \in \mathcal{K}} x_k \phi(\tilde{\mu}_k) \right). \quad (3)$$

¹⁴ Optimal committees under composition constraints are studied for example in Bredereck et al. (2018); here we axiomatically derive a particular objective function grounded in the aggregation of individual optimality judgments.

Interpreted in the context of the allocation of scarce resources, multiple overlapping composition constraints are akin to the reserve systems proposed by Pathak et al. (2020). In this context, the median rule constitutes a simple procedure of how the actual allocation is determined reflecting a profile of judgments by a panel of ethical experts or stakeholders.

¹⁵ This is possible if $\mathcal{Y}$ is the approval space or if $\mathcal{Y} = \mathcal{X}$ and there are enough candidates.

¹⁶ A function ϕ is *odd* if $\phi(-r) = -\phi(r)$ for all $r \in [-1, 1]$.

In particular, the median rule is an additive majority rule. To see this, let $\phi(r) := r$ for all $r \in [-1, 1]$; then formula (3) reduces to the formula (2).

A non-linear gain function ϕ allows larger majorities (especially unanimous or almost unanimous majorities) to carry a disproportionately greater weight than smaller majorities. The added generality appears potentially useful especially from a consensus perspective, from which non-linearities in the gain function can naturally be interpreted as reflecting non-linearities in the plausibility of (evidential support for) judgements as a function of the balance of majorities supporting them.

To illustrate, return to Example 5. Recall that the median rule selected a two-member committee with voter support for the two members of 45% and 10% over another committee with voter support of 25% and 25%. The superiority of the first could be questioned on the grounds that the first contains a very weakly supported member (not supported or even disapproved by 90% of the voters), while the second does not (to a comparable extent). From a consensus perspective, this may lead to disproportionate doubts about the merits of this candidate compared to any of the two candidates in the second committee; the second committee may thus be seen to be the safer, more defensible choice. This line of reasoning can be captured by a non-linear gain function for which $\frac{\phi(r)}{r}$ increases with $|r|$. The axioms invoked here (Reinforcement and Judgement Consistency) preclude this possibility.

In defining AMRs, we have so far omitted an important technical detail, namely the codomain of the gain function. While prototypical gain functions are real-valued, not all are. Indeed, the representation results of Nehring and Pivato (2019) do not yield a real-valued gain function, because they allow for possibly infinite and/or infinitesimal gains as well. This is achieved by allowing ϕ to take values in the hyperreal numbers ${}^*\mathbb{R}$. This may sound exotic, but there is no reason for concern; for the purposes of this paper, AMRs mainly serve as vehicle to prove the main result, with gain functions serving as an efficient computation device. In fact, the only properties of the hyperreals used in the proof are the standard properties of addition and a compatible ordering relation.¹⁷

Ensemble Supermajority Efficiency. In applying the SME criterion to judgement aggregation rules, not just single profiles, one can get additional leverage by considering *ensembles* of judgement problems. Such ensembles consist of N instances of the same judgement context with potentially different profiles $\mu_1, \dots, \mu_N$. To picture such ensembles concretely, an academic electorate may need to simultaneously appoint committees with different tasks but the same structure (composition constraint). Or, in a sequential version of the same idea, it may need to annually elect a committee for a given task and structure, but with different candidates every year.¹⁸

To apply SME formally to ensembles of judgement contexts, one must represent these ensembles themselves as judgement contexts. To do so, we enhance the issue space to the N -fold disjoint union of (copies of) $\mathcal{K}$ written as $N \cdot \mathcal{K}$, and let the output and

¹⁷ In mathematical terminology, all that is needed is the fact that ${}^*\mathbb{R}$ is a *linearly ordered abelian group*. That is, ${}^*\mathbb{R}$ is a commutative group $(\mathcal{G}, +)$ endowed with a linear order relation $>$ compatible with the group operation “+”, i.e. such that, for all $a, b, c \in \mathcal{G}$, $a > b$ iff $a + c > b + c$.

¹⁸ From a normative standpoint, it suffices to assume that these ensembles are meaningful as hypotheticals; they do not need to be actual features of the judgement problem at hand. Their role is thus quite similar to the role of counterfactual profiles in multi-profile restrictions.

input spaces be N -fold Cartesian powers $\mathcal{X}^N := \mathcal{X} \times \dots \times \mathcal{X}$ and $\mathcal{Y}^N := \mathcal{Y} \times \dots \times \mathcal{Y}$. Thus, we obtain a new judgement context $\mathcal{C}^N := (N \cdot \mathcal{K}, \mathcal{X}^N, \mathcal{Y}^N)$. Given a profile $\mu \in \Delta(\mathcal{Y}^N)$, we define its n th marginal $\mu^n \in \Delta(\mathcal{Y})$ to be the profile such that, for any $\mathbf{x} \in \mathcal{Y}$,

$$\mu^n(\mathbf{x}) := \sum_{\substack{(\mathbf{y}_1, \dots, \mathbf{y}_N) \in \mathcal{Y}^N \\ \text{with } y_n = \mathbf{x}}} \mu(\mathbf{y}_1, \dots, \mathbf{y}_N). \quad (4)$$

An N -tuple of profiles $\{\mu_1, \dots, \mu_N\}$ can be represented by any profile $\mu \in \Delta(\mathcal{Y}^N)$ such that, for each $n \in [1 \dots N]$, $\mu^n = \mu_n$. Consider thus an “ensemble problem” $(\mathcal{C}^N, \mu)$. Denote the output for this problem by $\tilde{F}(\mathcal{C}^N, \mu)$. It seems reasonable that this output should be obtained by applying F factor by factor. In other words, for all $\mu \in \Delta(\mathcal{Y}^N)$,

$$\tilde{F}(\mathcal{C}^N, \mu) = F(\mu^1) \times \dots \times F(\mu^N). \quad (5)$$

Since $(\mathcal{C}^N, \mu)$ is itself a well-defined judgement aggregation problem, SME applies to it, just as it does to ‘ordinary’ judgement aggregation problems.

One can formalize this line of argument by defining a judgement aggregation rule $\tilde{F}$ on the family of contexts $\{\mathcal{C}^N\}_{N \in \mathbb{N}}$, requiring $\tilde{F}$ to be SME on each $\mathcal{C}^N$ and to impose the product condition (5) as an axiom. In Nehring and Pivato (2019), we refer to (5) as the Combination axiom, and argue that it is a normatively appealing requirement on judgement aggregation, quite independently of any majoritarian considerations.

Here, we have chosen to stick with the more conventional single-context framework, primarily because the median rule is well-defined and well-motivated in this framework already. We will capture the conjunction of two distinct axioms—Supermajority Efficiency and Combination—by introducing a single ‘synthetic’ axiom: Ensemble Supermajority Efficiency.

ESME. For any number of instances $N \in \mathbb{N}$ and profile $\mu \in \Delta(\mathcal{Y}^N)$, any element of $F(\mu^1) \times \dots \times F(\mu^N)$ is SME in the judgement aggregation problem $(\mathcal{C}^N, \mu)$.¹⁹

ESME strengthens SME and can be viewed as a “reduced form” axiom that melds together the two more fundamental axioms into one single-context multi-profile axiom. Conceptually, while SME requires the efficient alignment of views with majorities across issues, ESME requires the efficient alignment of views with majorities across issues *and profiles*.

ESME is substantially stronger than SME, because an ensemble view $\mathbf{x} \in \mathcal{X}^N$ that is SME instance-by-instance need not be SME overall. To illustrate, return to Example 5. Let μ be any profile in which the candidates of type *fc*, *ml*, *mc*, and *fl* garner 45%, 10%, 25% and 25% of the votes, respectively, as above. Likewise, let μ' be a different

¹⁹ μ also contains information about the joint distribution of views over the different instances of $\mathcal{Y}$. But this extra information is immaterial to ESME: it is used in neither $F^N(\mathcal{X}^N, \mu)$ nor $\text{SME}(\mathcal{X}^N, \mu)$.

profile in which the candidates of type fc, ml, mc , and fl garner 50%, 15%, 20% and 20% of the votes, respectively. Thus, the candidates of the types fc, ml receive more votes in μ' whereas the candidates of the types fl and mc receive fewer votes; μ' is understood to describe voters' views in a different (e.g. subsequent) election, in which the candidates themselves (may) differ. SME is consistent with a choice of the duo $\mathbf{x} = (fc, ml)$ at μ , and a choice of the duo $\mathbf{y} = (fl, mc)$ at μ' . However, this is ruled out by ESME since the sequential choice $[\mathbf{x}, \mathbf{y}]$ is supermajority dominated in the ensemble problem by the sequential choice $[\mathbf{y}, \mathbf{x}]$, as is checked easily. The inferiority of $[\mathbf{x}, \mathbf{y}]$ to $[\mathbf{y}, \mathbf{x}]$ is intuitive as the latter is better aligned with the support for candidates of each type than the former.

If a rule F satisfies ESME, then it is a refinement of some additive majority rule G . If F also satisfies Continuity, then in fact $F = G$.

Continuity. For every profile $\mu \in \Delta(\mathcal{Y})$, and every sequence $\{\mu_n\}_{n=1}^\infty \subset \Delta(\mathcal{Y})$ with $\lim_{n \rightarrow \infty} \mu_n = \mu$, if $\mathbf{x} \in F(\mu_n)$ for all $n \in \mathbb{N}$, then $\mathbf{x} \in F(\mu)$.

This axiom says that the correspondence F is upper hemicontinuous with respect to the usual, Euclidean topology on $\Delta(\mathcal{Y})$. This means that F is robust against small perturbations or errors in the specification of μ . It also means that, if a very large population of voters is mixed with a much smaller population, then the views of the large population essentially determine the outcome of the rule. The following result is an adaptation of Theorems 1 and 2 of Nehring and Pivato (2019).

Proposition 1 *Let F be a judgement aggregation rule on a judgement context $\mathcal{C} = (\mathcal{K}, \mathcal{X}, \mathcal{Y})$.*

- (a) *F satisfies ESME if and only if there is an additive majority rule G such that $F(\mu) \subseteq G(\mu)$ for all $\mu \in \Delta(\mathcal{Y})$.*
- (b) *If F satisfies ESME and Continuity on $\Delta(\mathcal{Y})$, then F is an additive majority rule.*

We will use of the second part of Proposition 1. Note that any AMR with a real-valued, continuous gain function satisfies Continuity. But the converse does not hold.

4 Axiomatic characterization of the median rule in the unweighted case

To characterize the median rule, we will need one more axiom in addition to the two which appeared in Proposition 1: *Reinforcement*. Let $\mu_1, \mu_2 \in \Delta(\mathcal{Y})$ be two profiles, describing two subpopulations of size S_1 and S_2 . Let $c_1 = S_1/(S_1 + S_2)$ and $c_2 = S_2/(S_1 + S_2)$. Then $\mu = c_1\mu_1 + c_2\mu_2$ is the profile of the combined population. If each subpopulation separately endorses some view $\mathbf{x} \in \mathcal{X}$, then the combined population presumably should also endorse this view. The next axiom formalizes this desideratum.

Reinforcement. For any profiles $\mu_1, \mu_2 \in \Delta(\mathcal{Y})$ with $F(\mu_1) \cap F(\mu_2) \neq \emptyset$, and any $c_1, c_2 \in (0, 1)$ with $c_1 + c_2 = 1$, if $\mu = c_1\mu_1 + c_2\mu_2$, then $F(\mu) = F(\mu_1) \cap F(\mu_2)$. In other words, for any view $\mathbf{x} \in \mathcal{X}$, we have $\mathbf{x} \in F(\mu_1) \cap F(\mu_2)$ if and only if $\mathbf{x} \in F(\mu)$.

In the present setting, Reinforcement is appealing especially from a compromise perspective. If a particular view $\mathbf{x}$ minimizes the aggregate “burden of compromise” within some subpopulation μ_1 and the same view $\mathbf{x}$ happens to minimize the aggregate burden of compromise within a disjoint subpopulation μ_2 , it stands to reason that $\mathbf{x}$ *ipso facto* minimizes the aggregate burden of compromise within the combined population μ . This argument for Reinforcement parallels a standard argument for Reinforcement as an Extended Pareto condition in preference aggregation (Dhillon and Mertens 1999).

Reinforcement seems less compelling *prima facie* from the consensus perspective which treats the input judgements as the ‘evidential basis’ for an outside observer; in particular, the consensus perspective bars an ‘extended Pareto’ argument for Reinforcement. For example, consider a situation in which there is unanimous agreement on some issue k in one subpopulation with profile μ while there is a near tie in the other subpopulation μ' . At both profiles, the same view $\mathbf{x}$ happens to be selected as ‘most plausible’ according to F . In the combined population, say $\mu'' = \frac{1}{2}\mu + \frac{1}{2}\mu'$, there is a clear majority on k , but it is far from unanimity. So the profile μ'' is materially distinct as evidence from either μ or μ' , and it may very well be sensible to select a view $\mathbf{y}$ different from $\mathbf{x}$ as ‘most plausible’ given the evidence μ'' . This may well happen, for example, at some profiles under an additive majority rule with a non-linear gain function ϕ .

While the case for a linear gain function may not be as compelling from a consensus perspective as it is from a compromise perspective (via Reinforcement), there is still a good case to be made on the basis of a “default principle” of sorts. The next axiom serves as an axiomatic expression of such a default principle. In contrast to Reinforcement, which is a variable-population axiom, the next axiom compares profiles within a fixed population of voters. In a nutshell, it considers how the rule F should respond to a change of opinion in one sub-population while the opinion of the complementary sub-population remains fixed. It says that F should always respond to a given opinion change in the same way, independent of the opinion of the complementary sub-population; it is conceptually analogous to the axiom of *Tradeoff Consistency* in decision theory.

Judgement Consistency. For any $c_1, c_2 \in (0, 1)$ with $c_1 + c_2 = 1$, and any profiles $\mu, \mu', v, v' \in \Delta(\mathcal{Y})$, and any views $\mathbf{x}, \mathbf{y} \in \mathcal{X}$, if (a) $\mathbf{y} \in F(c_1 \mu + c_2 v)$, and (b) $\mathbf{x} \in F(c_1 \mu' + c_2 v)$, and (c) $\mathbf{y} \notin F(c_1 \mu' + c_2 v)$, and (d) $\mathbf{x} \in F(c_1 \mu + c_2 v')$, then (e) $\mathbf{y} \notin F(c_1 \mu' + c_2 v')$,

The idea here is that F should only select the views in $\mathcal{X}$ which have the greatest ‘plausibility’ in light of the input judgements. In the above axiom, μ, μ' are profiles describing two possible distributions of opinions for a subpopulation S_1 making up a proportion c_1 of the total population, while v, v' are profiles describing two possible distributions of opinions for the complementary subpopulation S_2 (making up the proportion $c_2 = 1 - c_1$ of the total population). Hypotheses (a), (b), and (c) say that the shift in the distribution of opinions from μ to μ' shifts the balance of plausibility from $\mathbf{y}$ to $\mathbf{x}$, when the S_2 subpopulation has profile v . Thus, if we start with another profile $(c_1 \mu + c_2 v')$ where $\mathbf{x}$ is *already* weakly more plausible than $\mathbf{y}$ (hypothesis (d)), then the same shift of opinion from μ to μ' among subpopulation S_1 should *again* make the $\mathbf{x}$ strictly more plausible than $\mathbf{y}$ (conclusion (e)).

It is easy to verify that the median rule satisfies Judgement Consistency. The next result describes the logical relationship between the last three axioms.

Proposition 2 *If an aggregation rule satisfies Continuity and Judgement Consistency, then it satisfies Reinforcement.*²⁰

To obtain our axiomatization of the median rule, we need a weak structural condition on $\mathcal{Y}$. Recall that $\mathcal{Y} \subseteq \{\pm 1\}^{\mathcal{K}} \subset \mathbb{R}^{\mathcal{K}}$. Let $\text{conv}(\mathcal{Y})$ be the convex hull of $\mathcal{Y}$ in $\mathbb{R}^{\mathcal{K}}$. We say that $\mathcal{Y}$ is *thick* if $\text{conv}(\mathcal{Y})$ has dimension $|\mathcal{K}|$. To interpret this, note that $\text{conv}(\mathcal{Y})$ is the set of *majority vectors* associated with admissible profiles μ ; that is, $\text{conv}(\mathcal{Y}) = \{\tilde{\mu} : \mu \in \Delta(\mathcal{Y})\}$. Thus, thickness just says that the domain of profiles is sufficiently “rich”.

Most interesting judgement aggregation spaces are thick. For examples, spaces of rankings (Example 1) are thick, as are spaces of classifiers (Example 2). Evidently, the approval spaces $\{\pm 1\}^{\mathcal{K}}$ involved in committee selection (Example 3) are thick. See Nehring and Pivato (2011) for many other examples. Here is our first main result.

Theorem 1 *Let $\mathcal{C} = (\mathcal{K}, \mathcal{X}, \mathcal{Y})$ be a judgement context where $\mathcal{Y}$ is thick and $\mathcal{X} \subseteq \mathcal{Y}$. Let $F : \Delta(\mathcal{Y}) \rightrightarrows \mathcal{X}$ be a judgement aggregation rule. Then F satisfies ESME, Continuity, and Reinforcement if and only if F is the median rule.*

Note that the thickness restriction applies only to the input space $\mathcal{Y}$, while the output space $\mathcal{X}$ is left entirely unrestricted. For example the committee spaces $\mathcal{X}_L^{\mathcal{K}}$ from Example 3 are not thick; since it is defined by an affine feasibility restriction, it has dimension $|\mathcal{K}| - 1$, not $|\mathcal{K}|$. Nevertheless, Theorem 1 still applies to Multi-Winner Approval Voting under Constraints, because the input space is $\mathcal{Y} = \{\pm 1\}^{\mathcal{K}}$. In the traditional setting in which $\mathcal{X} = \mathcal{Y}$, the thickness assumption obviously applies to the output space as well. In this case, Theorem 1 simplifies as follows.

Corollary 1 *Let $\mathcal{X}$ be a thick judgement space. An aggregation rule $F : \Delta(\mathcal{X}) \rightrightarrows \mathcal{X}$ satisfies ESME, Continuity, and Reinforcement if and only if F is the median rule.*

Thickness is needed for the conclusion of these results. For instance, it fails for the context $\mathcal{C}'$ given by $\mathcal{K} = 4$ and $\mathcal{X} = \mathcal{Y} = \{(1, 0, 0, 0), (0, 1, 0, 0), (0, 0, 1, 1)\}$. And indeed, the conclusion of Theorem 1 is false for this space.²¹

By Proposition 2, the statements of Theorem 1 and Corollary 1 remain true if Reinforcement is replaced by Judgement Consistency. All three axioms are needed for the characterization:

- If $\phi : [-1, 1] \rightarrow \mathbb{R}$ is increasing and continuous, then the AMR F_ϕ satisfies ESME and Continuity. But F_ϕ does not satisfy Reinforcement unless it is the median rule.

²⁰ Proposition 2 suggests that Judgement Consistency is logically stronger than Reinforcement. But Reinforcement depends on a *variable* population of voters, whereas Judgement Consistency can still be applied when the population of voters is *fixed*. Since our framework assumes a variable population from the beginning, it somewhat obscures this distinction.

²¹ To see this, note that the third and fourth issues are “clones”: they have the same value in every feasible view. Thus, in the terminology of Sect. 5, $\mathcal{C}'$ is equivalent to a *weighted* context $\mathcal{C}''$ where these two issues are “merged” into a single issue with double weight. But $\mathcal{C}''$ belongs to a class of contexts that violate the conclusion of Theorem 1, by Proposition 5 below. (To be precise, $\mathcal{X} = \mathcal{X}_{1,1}^3$ and $\lambda = (1, 1, 2)$.)

- Scoring rules (Myerson 1995; Dietrich 2014) satisfy Reinforcement and Continuity, but typically not ESME. For example, suppose $\mathcal{X} = \mathcal{Y}$, and consider the *plurality rule*, which selects the element(s) of $\mathcal{X}$ with the most votes —i.e. $F(\mu) = \operatorname{argmax}(\mu)$. This rule satisfies Reinforcement and Continuity, but not SME.
- Let $>$ be an arbitrary strict order on $\mathcal{X}$. One can construct a single-valued refinement of the median rule which satisfies ESME and Reinforcement and breaks any tie by choosing the $>$ -maximal element. But this rule does not satisfy Continuity.

Comparison to the literature. There is considerable discussion on various versions of the median rule in the mathematically oriented literature; see, for example, Chapter 5 of the monograph of Day and McMorris (2003), Monjardet (2008) and Barthélemy and Monjardet (1981). The latter is a broad, classical survey of the early interdisciplinary literature until the 1970s.²² Axiomatizations appear largely confined to *median spaces* defined in Nehring and Puppe (2007) as an adaption of “median graphs” to judgement aggregation spaces. In particular, McMorris et al. (2000) provide a characterization of the median rule in median graphs/spaces based on a local Condorcet condition and Reinforcement.²³

In the special case where $\mathcal{X} = \mathcal{Y} = \mathcal{X}_A^{\text{rk}}$, the space of linear orderings from Example 1, Theorem 1 yields a counterpart to the seminal contribution by Young and Levenglick (1978). As we already mentioned in the introduction, Young and Levenglick characterize the median rule for such spaces by three assumptions: Condorcet Consistency, Neutrality and Reinforcement. While their Reinforcement axiom is exactly the same as ours, the other axioms are not quite comparable. Condorcet Consistency is somewhat weaker than SME (hence *a fortiori* weaker than ESME) but not entirely, since it also deals with majority ties.²⁴ Neutrality is the standard axiom of a symmetry in alternatives. Any additive majority rule is neutral on $\mathcal{X}_A^{\text{rk}}$. So in view of Proposition 1, Neutrality is implied here by ESME plus Continuity. It is not quite implied by ESME alone, since non-neutral selections from AMRs would satisfy ESME as well. Conceptually, a lot of Neutrality is built into ESME via its symmetric treatment of issues.

The three axioms in the Young-Levenglick (YL) theorem are meaningful for general judgement aggregation contexts, with Neutrality understood as invariance to any symmetries of the context (input and output spaces) under permutations of issues. It is thus an interesting research challenge to provide YL-like axiomatizations for particular contexts. By disentangling the majoritarian (Condorcet) and neutrality ingredients in SME, such axiomatizations avoid the need to introduce a counting measure on issues. However, judging from Young and Levenglick’s own work, this task seems quite

²² Barthélemy and Monjardet (1981) focus on binary relations; this can be seen as a special case of our framework where issues correspond to ordered pairs (a, b) as in our treatment of linear orderings or equivalence relations (Examples 1 and 2). Many of the concepts and issues introduced there have more general application, but axiomatics are not discussed. The paper also allows input and output spaces to differ.

²³ However, in median spaces, the median rule is characterized by Condorcet consistency alone. Indeed, as shown in Nehring and Puppe (2007), the median spaces are exactly the spaces in which this is the case.

²⁴ Say that a view $\mathbf{x}$ is *Condorcet dominant* if, for all $k \in K$, $\tilde{\mu}_k \cdot x_k \geq 0$. A rule F is *Condorcet consistent* if F is equal to the set of all Condorcet dominant views whenever that set is non-empty.

challenging, since highly specialized and intricate proofs may be required. Indeed, in many context of interests, no YL-like results may be available due to lack of structure, especially lack of symmetries.²⁵

Mathematically, the two results and their proofs are very different. Young and Levenglick's proof is a tour de force that strongly exploits the special combinatorial features of the permutation polytope $\text{conv}(\mathcal{X}_A^{\text{rk}})$. By contrast, our proof of Theorem 1 must effectively sidestep the combinatorial structure of the context. Even equipped with Proposition 1, this requires significant work because the intended generality precludes the use of arguments that exploit special properties of the combinatorial structure of a particular context.

To illustrate the leverage one can gain from information about the particular context, we provide a custom-tailored and very direct proof of our main result for the ranking problem with rational-valued profiles, just as in the original Young and Levenglick result.²⁶

Proposition 3 *A judgment aggregation rule on the domain of rational-valued profiles in $\mathcal{X}_A^{\text{rk}}$ is the median rule if and only if it satisfies ESME, Reinforcement and Continuity.*

Here is the proof sketch; the details are provided in the appendix. First, we explicitly construct pairs of profiles μ and μ' such that all AMRs yield the same output and such that $F(\mu)$ and $F(\mu')$ have non-empty intersection. Reinforcement thus prescribes the output on any mixture $c\mu + (1 - c)\mu'$. This information can be used to infer that any gain function ϕ such that $F = F_\phi$ must be linear on some range $(-R, R) \cap \mathbb{Q}$. Thus F agrees with the median rule for all profiles μ with majority margins $\tilde{\mu}_k \in (-R, R)$.

Now consider an arbitrary profile μ . Let μ^0 denote any profile with tied majority margins on each issue (i.e. $\tilde{\mu}_k^0 = 0$ for all $k \in \mathcal{K}$). By SME, $F(\mu^0) = \mathcal{X}_A^{\text{rk}}$. By Reinforcement, $F(c\mu + (1 - c)\mu^0) = F(\mu)$. If $c < R$, by the first part of the proof, F agrees with the median rule on $c\mu + (1 - c)\mu^0$, hence, as is easy to see, on μ as well. (Note that the proof completely sidesteps the question of whether ϕ is real- or hyperreal-valued).

5 Extension to weighted judgement contexts

The formulation of judgement aggregation in Sect. 2 implicitly gave the same weight to all issues in the aggregation of the voters' opinions. But sometimes such equal weight-

²⁵ Even in cases where a YL-type result is available, this need not render an (E)SME-based result redundant. First, a general result like Theorem 1 provides additional assurance because it shows that the underlying normative logic is fairly insensitive to the particular combinatorial features of the aggregation problem. Second, in some interesting cases, SME alone suffices to characterize the median rule (with continuity to deal with ties), without appeal to Reinforcement, and without need to leverage it to ESME. Examples are the ranking problem with three alternatives, and the committee selection problem with at most one composition constraint. In such cases, the normative rationale for the median rule via SME is simple and direct, being based on a fixed-profile, fixed-population argument.

²⁶ A rational-valued profile corresponds to a profile of *equally weighted* voters. The proof of Proposition 3 on this restricted domain is no more difficult than the corresponding result on the richer domain of real-valued profiles. In contrast, the proof of Theorem 1 makes essential use of the topological connectedness of the real-valued domain.

ing is not appropriate. For example, in a truth-functional aggregation problem, one may wish to give a higher weight to the voters' opinions about the premises than their opinions about the conclusions. The most extreme form of this is the "premise-based" aggregation rule, which aggregates the voters' views on each premise by majority vote, but completely ignores their opinions about the conclusions; instead, the collective opinion about each conclusion is logically derived from the majoritarian opinions on the premises (Mongin 2008; Dietrich and Mongin 2010). At the opposite extreme is the "conclusion-based" aggregation rule, which aggregates the voters' opinions on each conclusion by majority vote, and mostly ignores their opinions about the premises, except when these opinions can be aggregated in a manner which is logically consistent with the majority opinions about the conclusions (Pigozzi et al. 2009). Between these extremes, there are rules which give greater or lesser weight to the voters' views on different premises and conclusions (Dietrich 2015).

There are other judgement aggregation problems where one might want to assign different weight to the voters' opinions on different issues. Indeed, only if a problem had a sufficiently high degree of "symmetry" (e.g. aggregation of rankings) would there be a strong *a priori* reason to assign the same weight to all issues. For this reason, we now introduce a *weight vector* $\lambda = (\lambda_k)_{k \in \mathcal{K}}$, where $\lambda_k > 0$ is the "weight" which we assign to the voters' opinions on issue k . Roughly speaking, λ_k would be large if we were very unwilling to overrule the majority opinion in issue k . Conversely, λ_k would be small if we were quite ready to overrule this opinion, if this was necessary to achieve a coherent collective view. A *weighted judgement context* is a quadruple $\mathcal{C} := (\mathcal{K}, \lambda, \mathcal{X}, \mathcal{Y})$, where $\mathcal{K}$ is a (finite) set of issues, $\lambda \in \mathbb{R}_+^{\mathcal{K}}$ is a weight vector, and $\mathcal{X} \subseteq \mathcal{Y} \subseteq \{\pm 1\}^{\mathcal{K}}$, with $\mathcal{X}$ being the set of "admissible collective views", and $\mathcal{Y}$ being the set of "admissible individual views". (In particular, an *unweighted judgement context* of the kind considered in Sects. 2 to 4 can be represented by setting $\lambda = (1, 1, \dots, 1)$.) A *judgement problem* is an ordered pair $(\mathcal{C}, \mu)$, where $\mathcal{C}$ is a (weighted) judgement context, and $\mu \in \Delta(\mathcal{Y})$ is a profile. For any $\mathbf{x} \in \mathcal{X}$ and $q \in [0, 1]$, we now define

$$\gamma_{\mu, \mathbf{x}}^{\lambda}(q) := \sum \{\lambda_k ; k \in \mathcal{K} \text{ and } x_k \tilde{\mu}_k \geq q\}. \quad (6)$$

This yields a non-increasing function $\gamma_{\mu, \mathbf{x}}^{\lambda} : [0, 1] \rightarrow \mathbb{R}$. If $\mathcal{C}$ is an *unweighted judgement context* (i.e. $\lambda = (1, \dots, 1)$), then formula (6) reduces to the definition of $\gamma_{\mathbf{x}}^{\mu}$ from Sect. 3. We say $\mathbf{x} \in \mathcal{X}$ is *supermajority efficient* (SME) for the judgement problem $(\mathcal{C}, \mu)$ if there does not exist any $\mathbf{z} \in \mathcal{X}$ such that $\gamma_{\mu, \mathbf{z}}^{\lambda}(q) \geq \gamma_{\mu, \mathbf{x}}^{\lambda}(q)$ for all $q \in [0, 1]$, with strict inequality for some q . A judgement aggregation rule $F : \Delta(\mathcal{Y}) \rightrightarrows \mathcal{X}$ is *supermajority efficient* on $\mathcal{C}$ if, for any $\mu \in \Delta(\mathcal{Y})$, every element of $F(\mu)$ is supermajority efficient for $(\mathcal{C}, \mu)$.

If $\phi : [-1, 1] \rightarrow {}^*\mathbb{R}$ is a gain function, then the *additive majority rule* on $\mathcal{C}$ is the correspondence $F_{\phi} : \Delta(\mathcal{Y}) \rightrightarrows \mathcal{X}$ defined as follows:

$$\text{for all } \mu \in \Delta(\mathcal{Y}), \quad F_{\phi}(\mu) := \operatorname{argmax}_{\mathbf{x} \in \mathcal{X}} \left(\sum_{k \in \mathcal{K}} \lambda_k \phi(x_k \tilde{\mu}_k) \right). \quad (7)$$

In particular, the *median rule* on $\mathcal{C}$ is defined by

$$\text{Median}(\mathcal{C}, \mu) := \operatorname{argmax}_{\mathbf{x} \in \mathcal{X}} \left(\sum_{k \in \mathcal{K}} \lambda_k x_k \tilde{\mu}_k \right), \quad \text{for all } \mu \in \Delta(\mathcal{Y}). \quad (8)$$

If $\mathcal{C}$ is an unweighted judgement context (i.e. $\lambda = (1, \dots, 1)$), then this reduces to the “unweighted” median rule defined by formula (2). For any $\mathbf{x}, \mathbf{y} \in \{\pm 1\}^{\mathcal{K}}$, we define their λ -weighted distance by $d_\lambda(\mathbf{x}, \mathbf{y}) := \sum \{\lambda_k; k \in \mathcal{K} \text{ and } x_k \neq y_k\}$. The median rule (8) can be equivalently defined as selecting the view(s) in $\mathcal{X}$ minimizing the average λ -weighted distance to the views of the voters. In Sect. 6, we will provide an axiomatic characterization of the weighted median rule (8), broadly similar to Theorem 1, but requiring additional structural conditions on the context. First, we will motivate the theory of weighted judgement contexts in general—and the weighted median rule in particular—with two interesting applications: *assignment problems* and *uniform decisions*.²⁷ But an impatient reader can skip directly to Sect. 6 without loss of logical continuity.

5.1 Assignment problems

Consider a group of voters who need to assign different candidates to different positions, such as the positions in a cabinet. There is a set $\mathcal{A} := \{1, \dots, A\}$ of “candidates” and a set $\mathcal{B} := \{1, \dots, B\}$ of “positions”, with $A \geq B$.²⁸ (As in the matching literature, there are many different possible interpretations; for example, “positions” could be unique resources such as organ transplants, and “candidates” could be possible transplant recipients. Judgment aggregation might be required when different group members entertain different standards of fair allocation).

Assignments can be described in terms of an issue space $\mathcal{K} = \mathcal{A} \times \mathcal{B}$, with the issue (a, b) addressing the question: “should candidate a hold position b ”? Feasibility requires that any position be filled by exactly one candidate, and that any candidate can fill at most one position. Feasible assignments can thus be described as B -tuples $((a_1, \dots, a_B))$ saying that candidate a_b is assigned to position b . More explicitly in issue space, the tuple $((a_1, \dots, a_B))$ refers to the view $\mathbf{x} \in \{\pm 1\}^{\mathcal{A} \times \mathcal{B}}$ such that $x_{ab} = 1$ iff $a = a_b$, and $x_{ab} = -1$ iff $a \neq a_b$. The feasible output space is the set of all such judgments $\mathcal{X}_{\mathcal{A} \times \mathcal{B}}^{\text{asgn}}$. By contrast, we will allow input judgments to be unrestricted approval judgments for each position; thus $\mathcal{Y} = \{\pm 1\}^{\mathcal{A} \times \mathcal{B}}$.

Generally, different positions will differ in their importance, so it will be natural to assign weights $\lambda_{a,b}$ of the form $\lambda_{a,b} = \bar{\lambda}_b$, where $\bar{\lambda}_b$ reflects the importance of position b . This defines a judgement context $\mathcal{C} = (\mathcal{A} \times \mathcal{B}, \lambda, \mathcal{X}_{\mathcal{A} \times \mathcal{B}}^{\text{asgn}}, \{\pm 1\}^{\mathcal{A} \times \mathcal{B}})$. In such a

²⁷ There are many other potential applications. In the context of the allocation of a scarce resource (which generalizes the committee selection problem) weights are naturally defined by costs (Talmon and Faliszewski 2019) or, for tasks under a time constraint, by duration (Pascual et al. 2018).

²⁸ This problem has been considered in particular by Emerson (2016). The median rule can be viewed as an Approval Voting counterpart of sorts to Emerson’s “matrix vote”. See also <http://www.deborda.org/faq/voting-systems/what-is-the-matrix-vote.html>.

context, the median rule selects the assignment $\mathbf{x} = ((a_1, \dots, a_B))$ that maximizes the weighted sum

$$\sum_{b \in \mathcal{B}} \bar{\lambda}_b \tilde{\mu}_{(a_b, b)}.$$

To illustrate the role of the weights, consider, as a simple example, a profile of input views μ all of which assess candidate quality as independent of position, i.e. $\tilde{\mu}_{a,b} = \tilde{\mu}_{a,b'}$ for all $a \in \mathcal{A}$, and $b, b' \in \mathcal{B}$. In such profiles, supermajority efficient rules such as the weighted median rule (8) will assign the better candidates to the more important positions. Likewise, in arbitrary profiles, if the weight for some position b^* is much larger than that for any other position, the median rule fills that position by the candidate a with the highest majority support $\tilde{\mu}_{(a,b^*)}$ (up to, possibly, near ties).

5.2 Uniform decisions

The following “uniform decision” model is in fact a scheme of examples generating more complex judgment aggregation contexts from simpler ones. We first present the formal scheme, and then illustrate three types of applications referred to as *Heterogeneous Cases*, *Missing Information*, and *Multiple Criteria*.

Let $\mathcal{C} = (\mathcal{K}, \lambda, \mathcal{X}, \mathcal{Y})$ be a “base” judgement context. Let $\mathcal{S}$ be an abstract set of “instances”; in the three applications, the elements of $\mathcal{S}$ will be interpreted as “cases”, “states” or “criteria”. Let $\bar{\lambda} := (\bar{\lambda}_s)_{s \in \mathcal{S}}$ be a vector assigning a “weight” to each instance; in the three applications, these weights will be interpreted as the frequencies of the cases, the probabilities of the states, or the relative importance of the criteria. (We will provide more detail below.) While input judgments are made instance-wise, these form the basis for a single output judgment that governs all instances uniformly. This situation can be described formally by a “uniform” judgement context $\hat{\mathcal{C}} = (\hat{\mathcal{K}}, \hat{\lambda}, \hat{\mathcal{X}}, \hat{\mathcal{Y}})$ defined as follows.

- (i) $\hat{\mathcal{K}} := \mathcal{K} \times \mathcal{S}$;
- (ii) $\hat{\mathcal{X}} := \{(\mathbf{x}, \dots, \mathbf{x}) ; \mathbf{x} \in \mathcal{X}\}$;
- (iii) $\hat{\mathcal{Y}} := \mathcal{Y} \times \dots \times \mathcal{Y}$;
- (iv) For all $k \in \mathcal{K}$ and $s \in \mathcal{S}$, $\hat{\lambda}_{k,s} := \bar{\lambda}_s \lambda_k$.

Part (iii) allows input judgments to vary independently across instances, while (ii) says that output judgments must be constant across instances. Part (iv) says that the weight of a composite issue (k, s) rescales the base weight λ_k by the instance weight $\bar{\lambda}_s$.

5.2.1 Heterogeneous cases

To flesh out the first interpretation in terms of heterogeneous instances, consider the very simplest judgment aggregation problem given by a single yes-no issue, i.e. $|\mathcal{K}| = 1$.

For example, a group needs to decide on whether to establish a rule which permits or forbids some types of action or behaviour. Concrete examples of such rules include traffic laws, workplace codes of conduct, or safety regulations at a public swimming

pool. More contentious examples include restrictions on speech acts on grounds of libel, hate speech, obscenity, or incitement.²⁹ For this rule to be simple, unambiguous, and enforceable in a non-arbitrary and independently verifiable way, it must be based on a relatively coarse description of the action. At the same time, in each particular case, the “correct” or “just” decision may depend on some finer details of the action and the surrounding context. But even if these details are observable in principle, it is often not feasible to explicitly condition the actual decision on them, e.g. for reasons of complexity or non-verifiability.

So the best the group can do in such situations is to get it right “on balance” across comparable cases $s \in \mathcal{S}$. Voters give their judgments about the right decision in each case. The uniform group decision is then to be taken on the basis of the entire vector of majorities $(\tilde{\mu}_s)_{s \in \mathcal{S}}$. The cases could be actual or hypothetical. For actual cases, the weights $\bar{\lambda}_s$ would naturally reflect their frequency of occurrence. For hypothetical cases, the weights $\bar{\lambda}_s$ would naturally reflect their ‘relevance’ or ‘representativeness’. Additive majority rules yield a positive uniform decision just in case

$$\sum_{s \in \mathcal{S}} \bar{\lambda}_s \phi(\tilde{\mu}_s) \geq 0.$$

Hence, the median rule in particular yields a positive uniform decision iff

$$\sum_{s \in \mathcal{S}} \bar{\lambda}_s \tilde{\mu}_s \geq 0,$$

i.e. just in case the weighted average of majorities across comparable cases is non-negative.

5.2.2 Missing information

The uniform decision model can also be applied to situations in which the group is missing information relevant to the judgment task. For example, in foreign policy, macroeconomic management, and environmental regulation, the right course of action may depend on information which is unavailable at the moment the decision must be made. For example, during its 2014 independence referendum, Scotland confronted key uncertainties about the future petroleum prices and about its ability to join the E.U. as an independent state. It is quite legitimate—and may lead to better decisions—to make this missing information explicit in the collective decision procedure by representing it as an unknown “state of nature” $s \in \mathcal{S}$, so that each voter submits her judgment contingent on the state. In a binary (single-issue) decision problem, analogously to Example 5.2.1 above, the median rule would base the decision on the *probability-weighted majority margin* $\sum_{s \in \mathcal{S}} p_s \tilde{\mu}_s$.

The relevant probability weights $\bar{\lambda}_s = p_s$ could be obtained in various ways. For example, they might be obtained by some judgment aggregation rule from the voters

²⁹ See Miller (2013) for a contribution to this theme in the judgment aggregation literature.

themselves. Or, the group may delegate this judgment to an ‘outside authority’, for example to betting markets.³⁰

This can be extended beyond a single issue, for instance to the ranking of more than two possible courses of action. The base context would then be given by the space of rankings $\mathcal{X}_A^{\text{rk}}$ from Example 1. The median rule would select the ranking(s) $\succ$ with the highest support *ex ante*, as measured by the sum

$$\sum_{\substack{a, b \in A \\ a \succ b}} \hat{\mu}_{ab}, \quad (9)$$

where, for all $a, b \in A$, we defined

$$\hat{\mu}_{ab} := \sum_{s \in \mathcal{S}} p_s \tilde{\mu}_{ab,s}. \quad (10)$$

The median rule thus yields an *extension of the Kemeny rule to uncertainty*. Note that the *ex ante* comparison depends as much on the size of the majorities as on their sign.³¹

5.2.3 Multiple criteria

Applied to elections for public office, the standard ranking model can be interpreted as trying to determine the (impartially best) candidate based on an overall comparative evaluation of candidates in terms of their “suitability for office”. It might be argued that this gives too much room for subjective impressions of personal appeal of candidates, and that “suitability for office” is better construed as a combination of rather distinct, identifiable qualities such as leadership, integrity, judgment, etc. to be evaluated and aggregated separately, and then amalgamated into one judgement of overall suitability. Such a multi-criterion conception of candidate merit can be captured by the uniform decision model as follows. Let the base context again be the standard ranking context $\mathcal{X}_A^{\text{rk}}$ from Example 1, and let each s represent a different criterion. An input judgment $\mathbf{y}$ consists of an $\mathcal{S}$ -tuple of rankings $(\mathbf{y}_s)_{s \in \mathcal{S}}$, with $\mathbf{y}_s$ representing the ranking of candidates in terms of criterion s . The uniform output judgment $\mathbf{x}$ represents the overall group ranking to be determined. It is based on weight vector $\bar{\lambda} \in \mathbb{R}_+^{\mathcal{S}}$ describing the relative importance of these criteria. These can be determined in different ways: they could be determined concurrently by the group itself, by a separate committee or at an earlier ‘constitutional’ stage at which the general requirements for the office were determined.³²

³⁰ For a provocative and ambitious advocacy of using betting markets to “vote on values, but bet on beliefs”, see Hanson (2013).

³¹ For a discussion of the Kemeny rule under uncertainty from a different angle, see Procaccia and Shah (2016).

³² Switching the setting, a hiring committee or university department may be tasked to evaluate an applicant for an open faculty position in terms of research, teaching and service, and the weights of these might be predetermined by standing university policy.

In this setting, the median rule selects the ranking with highest overall majority support (9), where again the overall majority support $\widehat{\mu}_{ab}$ for ranking of a over b is given the weighted average (10). In this manner, the median rule thus yields an *extension of the Kemeny rule to multiple criteria*.

The general approach to multi-criterion evaluation just outlined here is not restricted to the particular, comparative format evaluation in terms of rankings; an alternative format of interest is non-comparative in terms of “grades”. In this vein, Balinski and Laraki (2011, ch. 21) introduce a multi-criteria majority grading rule³³ which—setting aside the treatment of tied grades—is equivalent to the weighted median rule in the uniform decision model for a base context that describes grading as follows.

Let $(G, >)$ be a finite, linearly ordered set, each element being interpreted as a grade. Grading can be defined as a judgment context $(G, \mathcal{X}_G, \mathcal{X}_G, \lambda)$, where the issues $g \in G$ are interpreted as whether or not the target of evaluation achieves at least grade g , and the set of admissible grade assignments is given by $\mathcal{X}_G = \{\mathbf{x} \in \{\pm 1\}^G; x_g \geq x_{g'} \text{ whenever } g \leq g'\}$, with typical element $(1, \dots, 1, 0, \dots, 0)$.³⁴

6 Characterization of the weighted median rule

To obtain a weighted generalization of Theorem 1, we must consider weighted combinations of weighted contexts. These are given by weighted contexts of the form $\widehat{\mathcal{C}} = (N \cdot \mathcal{K}, \widehat{\lambda}, \mathcal{X}^N, \mathcal{Y}^N)$, where $\widehat{\lambda} = (\widehat{\lambda}_1, \dots, \widehat{\lambda}_N) \in \mathbb{R}^{N \cdot \mathcal{K}}$ is *proportional* to λ ; that is, for for each $n \in [1, \dots, N]$, $\widehat{\lambda}_n = c_n \lambda$, for some $c_n > 0$. So, for a proportional weight vector $\widehat{\lambda}$, the relative weights *within* the set of basic issues $\mathcal{K}$ are the same (original ones) in each instance, while different instances may be assigned different relative weights reflected in the scaling factor c_n . The potential differences in the scaling factor may have different origins, as in the uniform decision model of Sect. 5.2. In particular, they could reflect differences in frequency, probability or relative importance. We are now in a position to state the axiom of Weighted Ensemble Supermajority Efficiency.

WESME. For any set of instances N , any vector of weights $\widehat{\lambda}$ proportional to λ , and any profile $\mu \in \Delta(\mathcal{Y}^N)$, any element of $F(\mu^1) \times \dots \times F(\mu^N)$ is SME for $(\widehat{\mathcal{C}}, \mu)$.

Here is the extension of Proposition 1 to weighted contexts, obtained from Theorem A.1 of Nehring and Pivato (2019).

Proposition 4 *Let F be a judgement aggregation rule on a weighted judgement context $\mathcal{C}$. If F satisfies WESME and Continuity, then F is an additive majority rule like (7).*

Note that the consideration of proportional weights is needed only to deal with the case of irrational-valued weight vectors λ in the base context. If all weights are rational, then it is sufficient to confine attention to weight vectors $\widehat{\lambda}$ such that $\widehat{\lambda}_n = \lambda$ for all $n \in [1 \dots N]$. (See Nehring and Pivato (2019) for more information.)

³³ Their term is “multi-criteria majority judgment”.

³⁴ The output of the median rule depends on the criteria weights $\bar{\lambda}_g$, but is easily seen to be independent of the issue weights λ_g . This follows from the fact that the underlying context $(G, \mathcal{X}_G, \mathcal{X}_G, \lambda)$, can be viewed as a median space.

In view of Proposition 4, one might expect Theorem 1 to generalize verbatim. In fact, one can mimic rational weights by cloning issues in proportion to their weight and apply equal weights to the cloned description of the problem. However, cloning destroys thickness, so Theorem 1 cannot be applied. And, indeed, the following simple counterexample shows that this generalization fails; it implies that additional, non-trivial structural conditions are necessary. To construct the counterexample, consider the following class of additive majority rules. For any $\alpha > 0$, we define $\phi^\alpha : [-1, 1] \rightarrow \mathbb{R}$ by $\phi_\alpha(r) := \text{sign}(r) \cdot |r|^\alpha$. The corresponding additive majority rule F_{ϕ^α} is called a *homogeneous* rule. Note that the median rule is just the homogeneous rule with $\alpha = 1$. Now let $\mathcal{K}$ be a finite set, and define $\mathcal{X}_{L,M}^{\mathcal{K}} := \{\mathbf{x} \in \{\pm 1\}^{\mathcal{K}}; L \leq \#\mathbf{x} \leq M\}$. If $L < M$, this space is thick.³⁵ But the intended generalization of Theorem 1 fails on this space. Its normative assumptions fail to single out the median rule uniquely, but are satisfied by any homogeneous rule.

Proposition 5 *Let $K \geq M \geq L \geq 0$, let $\lambda \in \mathbb{R}_+^{\mathcal{K}}$ be a weight vector, and let $\mathcal{C} := (\mathcal{K}, \lambda, \mathcal{X}_{L,M}^{\mathcal{K}}, \mathcal{X}_{L,M}^{\mathcal{K}})$. For any $\alpha > 0$, the rule F^α satisfies Reinforcement, Continuity, and WESME on $\mathcal{C}$. If λ is the uniform weight vector, then F^α is equivalent to the median rule on $\mathcal{C}$. But if λ not uniform and $\alpha \neq 1$, then F^α is not the median rule on $\mathcal{C}$.*

Proposition 5 opens a rather surprising gap between a weighted and an unweighted sum representation. Can it be fixed, and if so, how? To locate the source of the trouble, we first state a general result applicable to a broad range of contexts. We will require the judgement context and the aggregation rule to satisfy one of two hypotheses. We say that a judgement context $\mathcal{C} = (\mathcal{K}, \lambda, \mathcal{X}, \mathcal{Y})$ has *balanced weights* if for all $\mathbf{x}, \mathbf{y} \in \mathcal{X}$ with only two issues $i, j \in \mathcal{K}$ such that $x_i \neq y_j$ and $x_j \neq y_i$ (so that $d(\mathbf{x}, \mathbf{y}) = 2$), we have $\lambda_i = \lambda_j$. On the other hand, we say that $\mathcal{C}$ is *suitable* for a judgement aggregation rule F if there is some $\mu \in \Delta(\mathcal{Y})$ such that $F(\mathcal{C}, \mu) = \{\mathbf{x}, \mathbf{y}\}$ for some $\mathbf{x}, \mathbf{y} \in \mathcal{X}$ with $d(\mathbf{x}, \mathbf{y}) \geq 3$.

Theorem 2 *Let $\mathcal{C} = (\mathcal{K}, \lambda, \mathcal{X}, \mathcal{Y})$ be a weighted judgement context where $\mathcal{Y}$ is thick, and let $F : \Delta(\mathcal{Y}) \Rightarrow \mathcal{X}$ be an additive majority rule. Suppose that either $\mathcal{C}$ has balanced weights or $\mathcal{C}$ is suitable for F . Then F satisfies Continuity and Reinforcement if and only if it is the median rule (8).*

Note that any unweighted judgement context automatically has balanced weights. Thus, Theorem 2 immediately entails Theorem 1 above. For the case of generic (unbalanced) weights, Theorem 2 is of limited value on its own, since suitability is both hard to verify and without obvious normative significance. To obtain a satisfactory result, we must find structural conditions on the context $\mathcal{C}$ which, together with the axioms imposed on F , ensure its suitability. To these we now turn.

For any $\mu \in \Delta(\{\pm 1\}^{\mathcal{K}})$, let $\mathbf{x}^\mu$ be the *majority ideal*, as defined below formula (1) above. A judgement space $\mathcal{Y}$ is *McGarvey* if, for all $\mathbf{x} \in \{\pm 1\}^{\mathcal{K}}$, there is some $\mu \in \Delta(\mathcal{Y})$ such that $\mathbf{x}^\mu = \mathbf{x}$.³⁶ For example, the spaces $\mathcal{X}_{\mathcal{A}}^{\text{rk}}$ and $\mathcal{X}_{\mathcal{A}}^{\text{eq}}$ (Examples 1

³⁵ $\mathcal{X}_{L,M}^{\mathcal{K}}$ is often used to represent committee selection problems, as in Example 3: $\mathcal{K}$ is a set of potential “candidates”, and the committee in question must have at least L and most M members. It also arises in certain resource allocation problems. But these interpretations are not relevant for Proposition 5.

³⁶ Equivalently, $\mathcal{Y}$ is McGarvey if and only if the zero vector $\mathbf{0}$ lies in the topological interior of $\text{conv}(\mathcal{Y})$.

and 2) are McGarvey, as are many other commonly occurring judgement spaces; see Nehring and Pivato (2011) for a systematic study and many more examples. Clearly, any McGarvey space is thick, and any superset of a McGarvey space is also McGarvey.

Let $\mathbf{x}, \mathbf{y}, \mathbf{z} \in \{\pm 1\}^{\mathcal{K}}$. We say that $\mathbf{y}$ is *between* $\mathbf{x}$ and $\mathbf{z}$ if, for any $k \in \mathcal{K}$ such that $x_k = z_k$, we also have $y_k = z_k$ (and hence, $y_k = x_k$). Now let $\mathcal{X} \subseteq \{\pm 1\}^{\mathcal{K}}$ be a judgement space. Say that $\mathbf{x}, \mathbf{z} \in \{\pm 1\}^{\mathcal{K}}$ are *near* if there is no $\mathbf{y} \in \mathcal{X} \setminus \{\mathbf{x}, \mathbf{z}\}$ such that $\mathbf{y}$ is between $\mathbf{x}$ and $\mathbf{z}$. We will say that $\mathcal{X}$ is *distal* if there exist some $\mathbf{x}, \mathbf{z} \in \mathcal{X}$ such that $\mathbf{x}$ is near $\mathbf{z}$ and $d(\mathbf{x}, \mathbf{z}) \geq 3$. Heuristically, this means that the elements of $\mathcal{X}$ are not packed tightly together everywhere. For any $\mathbf{z} \in \{\pm 1\}^{\mathcal{K}} \setminus \mathcal{X}$, let $\mathcal{X}_{\mathbf{z}}$ be the set of views in $\mathcal{X}$ that are near to $\mathbf{z}$; heuristically, these are the “best admissible approximations” to $\mathbf{z}$. We say that $\mathcal{X}$ is *rugged* if there exists $\mathbf{z} \in \{\pm 1\}^{\mathcal{K}} \setminus \mathcal{X}$ and $\mathbf{x}, \mathbf{y} \in \mathcal{X}_{\mathbf{z}}$ such that $d(\mathbf{x}, \mathbf{z}) \neq d(\mathbf{y}, \mathbf{z})$.

Both distality and ruggedness are conditions on the combinatorial geometry of $\mathcal{X}$, and typically easy to verify. We illustrate them in the spaces of rankings and of classifiers.

1. If $|\mathcal{A}| \geq 4$, the space of rankings $\mathcal{X}_{\mathcal{A}}^{rk}$ is rugged. To see this, let $\mathcal{A} = \{a, b, c, d\}$ and $\mathcal{K} := \{(a, b), (a, c), (a, d), (b, c), (b, d), (c, d)\}$ as in “Example 1 continued”. Let $\mathbf{z} \in \{\pm 1\}^{\mathcal{K}} = (1, 1, -1, 1, 1, 1)$; $\mathbf{z}$ represents an asymmetric relation with a 4-cycle. Let $\mathbf{x}$ be the linear order $abcd$, and $\mathbf{y}$ the linear order $dabc$. Evidently, both $\mathbf{x}$ and $\mathbf{y}$ are near $\mathbf{z}$. This verifies ruggedness since $d(\mathbf{x}, \mathbf{z}) = 1$ while $d(\mathbf{y}, \mathbf{z}) = 2$.
2. If $|\mathcal{A}| \geq 4$, the space of classifiers $\mathcal{X}_{\mathcal{A}}^{eq}$ is rugged. To see this, let $\mathbf{z}$ represent the classifier by the three equivalences (edges) $\{ab, bc, cd\}$. Let $\mathbf{x}$ be the equivalence relation $\{ab, ac, bc\}$, and $\mathbf{y}$ the total equivalence relation $\{ab, ac, ad, bc, bd, cd\}$. Both $\mathbf{x}$ and $\mathbf{y}$ are near $\mathbf{z}$. Again, this verifies ruggedness since $d(\mathbf{x}, \mathbf{z}) = 2$ while $d(\mathbf{y}, \mathbf{z}) = 3$.
3. This also establishes the distality of $\mathcal{X}_{\mathcal{A}}^{eq}$, since $\mathbf{x}$ is near $\mathbf{y}$ but $d(\mathbf{x}, \mathbf{y}) = 3$.
4. By contrast, any two linear orders near each other differ only in one comparison, falsifying distality of $\mathcal{X}_{\mathcal{A}}^{rk}$ for any $|\mathcal{A}|$.

Ruggedness seems to be satisfied in the great majority of cases, while distality is somewhat more restrictive. Indeed, distality “almost” implies ruggedness. That is, suppose that there exist some $\mathbf{x}, \mathbf{y} \in \mathcal{X}$ such that $\mathbf{x}$ is near $\mathbf{y}$ and $d(\mathbf{x}, \mathbf{y}) \geq 3$, and assume additionally that some such $d(\mathbf{x}, \mathbf{y})$ is odd. Consider any $\mathbf{z} \in \{\pm 1\}^{\mathcal{K}} \setminus \{\mathbf{x}, \mathbf{y}\}$ such that $\mathbf{z}$ is between $\mathbf{x}$ and $\mathbf{y}$. Since $\mathbf{x}$ is near $\mathbf{y}$, it follows that $\mathbf{z} \in \{\pm 1\}^{\mathcal{K}} \setminus \mathcal{X}$. Furthermore, $\mathbf{x} \in \mathcal{X}_{\mathbf{z}}$, because any element of $\mathcal{X}$ between $\mathbf{x}$ and $\mathbf{z}$ would also be between $\mathbf{x}$ and $\mathbf{y}$, contradicting the assumed nearness $\mathbf{x}$ and $\mathbf{y}$. Likewise $\mathbf{y} \in \mathcal{X}_{\mathbf{z}}$. But $d(\mathbf{x}, \mathbf{y}) = d(\mathbf{x}, \mathbf{z}) + d(\mathbf{z}, \mathbf{y})$ (because $\mathbf{z}$ is between $\mathbf{x}$ and $\mathbf{y}$), and $d(\mathbf{x}, \mathbf{y})$ is odd, so $d(\mathbf{x}, \mathbf{z}) \neq d(\mathbf{z}, \mathbf{y})$, verifying the ruggedness of $\mathcal{X}$.

We now come to the second main result of this paper.

Theorem 3 *Let $(\mathcal{K}, \lambda, \mathcal{X}, \mathcal{Y})$ be a weighted judgement context, and suppose that either (a) $\mathcal{X}$ is rugged and $\mathcal{Y}$ is McGarvey, or (b) $\mathcal{X}$ is distal and $\mathcal{Y}$ is thick. Let $F : \Delta(\mathcal{Y}) \rightrightarrows \mathcal{X}$ be a judgement aggregation rule. Then F satisfies WESME, Continuity, and Reinforcement if and only if F is the median rule (8).*

The presuppositions of Theorem 3 are satisfied in many applications of interest, but they are more restrictive than those of Theorem 1. We illustrate their broad applicability in a variety of examples before providing an example where they do not apply.

1. As already shown, the spaces of rankings (Example 1) and of classifiers (Example 2) are rugged and McGarvey. So Theorem 3 applies.
2. In the assignment problems of Sect. 5.1, the input space is trivially McGarvey as it is an approval space. For $B \geq 3$, it is rugged. To see this, consider the assignments $\mathbf{z} := ((a_1, a_1, a_2))$, $\mathbf{x} := ((a_1, a_3, a_2))$ and $\mathbf{y} := ((a_1, a_2, a_3))$. Note that $\mathbf{x}$ and $\mathbf{y}$ are feasible, while $\mathbf{z}$ is not. Evidently, $\mathbf{x}$ and $\mathbf{y}$ are adjacent to $\mathbf{z}$; the triple of views verifies ruggedness since $d(\mathbf{z}, \mathbf{x}) = 2$ while $d(\mathbf{z}, \mathbf{y}) = 4$.
3. In the uniform decision problems of Sect. 5.2, it follows from elementary linear algebra to see that $\widehat{\mathcal{Y}}$ is thick if and only if $\mathcal{Y}$ is thick. On the other hand, for any $\mathbf{x}, \mathbf{y} \in \mathcal{X}$, $d((\mathbf{x}, \dots, \mathbf{x}), (\mathbf{y}, \dots, \mathbf{y})) = |S| d(\mathbf{x}, \mathbf{y})$. Thus $\mathcal{X}$ is distal whenever $|S| \geq 3$.
4. By contrast, the structural assumptions of Theorem 3 fail for $\mathcal{X}_{L,M}^K$. To see that $\mathcal{X} = \mathcal{X}_{0,M}^K$ is not rugged, for example, one simply notes that, for any $\mathbf{z} \notin \mathcal{X}$, $\mathbf{x}$ is near $\mathbf{z}$ iff $\#\mathbf{x} = M$ and if, for all $k \in \mathcal{K}$, $x_k = 1$ implies $z_k = 1$. Thus, for any $\mathbf{x}$ near $\mathbf{z}$, we have $d(\mathbf{x}, \mathbf{z}) = \#\mathbf{z} - M$, contradicting ruggedness. The failure of distality is verified easily as well.

Discussion

The meaning of weights

Theorems 2 and 3 do not attempt to *justify* the use of weights; they assume weights as given and settled by the aggregator. The meaning of the weight vector λ is given by their use in the WESME axiom. That meaning is extensive—i.e., λ serves to compare the “size” of sets of issues. The result endows weights with a second, “intensive”, meaning in the context of the median rule; in that second meaning, one can meaningfully say that the votes on one issue count twice as much (or 1.414 times as much, etc.) as the votes on some other issue. The path from the first to the second meaning takes two steps: the ensemble construction, which allows one to meaningfully distinguish between any real-valued weight vectors λ , and the Reinforcement axiom resulting in linearity of the gain-function, hence constant “exchange rates” between votes on different issues.

The need for SME

As exemplified by Young and Leventick’s classical characterization of the Kemeny rule, in particular settings one can replace SME by the arguably simpler conjunction of axioms of Condorcet consistency and Neutrality. As discussed in Nehring and Pivato (2019), this may work for other special contexts such as, for example, spaces of classifiers, but such a result would need to rely heavily on appropriate symmetries of the context. Dropping the Neutrality axiom, one might attempt to “derive” weights as part of a representation result but we know of no example of such a result.

On the other hand, taking equal or unequal weights as given, it stands to reason that SME is in fact the simplest way to relate weights to the output of a majoritarian aggregation rule. Note also that, with SME alone rather than (W)ESME, one cannot

hope to obtain a result analogous to Theorem 2, even for special and well-behaved context. The reason is simple: if two weight vectors λ and λ' induce the same ordinal size comparisons of subsets of issues, then they entail the same SME-axiom.

A sketch of the proofs

Theorems 1 and 3 both follow from Theorem 2; the bulk of the work is in proving this underlying result. From Propositions 1 and 4, we know that any rule satisfying Continuity and (W)ESME is an additive majority rule F_ϕ , for some gain function ϕ . Thus, the key task in the proof of Theorem 2 is to show that the identity function (or any linear function) on $[-1, 1]$ is among these gain functions.

A possible strategy would be to try to show that any representing gain function was linear. If ϕ was real-valued, it would then be enough to show that ϕ satisfies the Cauchy functional equation, and hence deduce linearity. But this straightforward-looking strategy fails for two reasons. First of all, Propositions 1 and 4 do not guarantee that ϕ is real-valued, so one cannot appeal to the Cauchy functional equation. Second, the conclusion of Theorem 2 does not require that *all* representing gain functions ϕ are linear, only that *some* of them are. Indeed, Approval Voting on size- L committees provides a counterexample, in which the assumptions of Theorem 1 are satisfied, but any AMR agrees with the median rule, whatever the gain functions.³⁷

Hence this strategy is a nonstarter. Instead, the proof of Theorem 2 shows in a couple of steps that any representing ϕ has the “linearity-like” property that $\phi(r + \epsilon) - \phi(r) = \phi(s + \epsilon) - \phi(s)$ for any r and s , and any sufficiently small ϵ ; in effect, this says that ϕ has a “constant slope” property. This is enough to show that F_ϕ is the median rule, because it implies that a gain (or loss) of ϵ in the majority support on one issue can be exactly offset by a gain (or loss) of ϵ in the support on another issue.³⁸

To show in turn that ϕ satisfies this “constant slope” property, we must study “perfect tie” profiles $\mu \in \Delta(\mathcal{J})$ such that $F(\mu) = \{\mathbf{x}, \mathbf{y}\}$ for some $\mathbf{x}, \mathbf{y} \in \mathcal{X}$. If the set of such “perfect tie” profiles is not empty, then it is a relatively open subset of an affine hyperplane (this is a consequence of Continuity and Reinforcement); we can then demonstrate the “constant slope” property by looking at how the values of $\phi(\tilde{\mu})$ change as μ moves around in this hyperplane. However, we also need $\mathbf{x}$ and $\mathbf{y}$ to differ in at least three coordinates—in effect, this is because we need the freedom to manipulate the r , s , and ϵ variables independently while preserving the perfect tie. This is the reason for the hypothesis of *suitability* in Theorem 2.

Theorem 2 just *assumes* that $\mathcal{C}$ is suitable for F . So to derive Theorems 1 and 3 from Theorem 2, we must provide conditions on $\mathcal{C}$ that ensure that it is suitable for any judgment aggregation rule F satisfying the three axioms of these theorems. This

³⁷ For a counterexample under the structural assumptions of Theorem 3, return to Example 5 with two dichotomous composition constraints. Here, whatever the profile, there are at most two supermajority efficient views (ignoring trivial ties) $\mathbf{x}, \mathbf{y}$; these must differ on exactly four issues. It follows that, for any $a > 0$, the AMR F_ϕ given by $\phi(r) = r + a \cdot \text{sign}(r)$ coincides with the median rule for these contexts. Note that these gain functions are affine, not linear, and, for most contexts, yield different outputs from the median rule.

³⁸ This assumes uniform weights. In the case of nonuniform weights, the relevant statement is that a gain (or loss) of ϵ/λ_j in the majority support on issue j can be exactly offset by a gain (or loss) of ϵ/λ_k in the majority support on issue k .

is done in Theorem C.1 in Appendix C, via a constructive condition of “frangibility” on contexts.

If $\mathcal{X}$ is rugged and $\mathcal{Y}$ is McGarvey, then frangibility of the context follows from Lemma C.4 in Appendix C. Meanwhile, if $\mathcal{X}$ is distal and $\mathcal{Y}$ is thick, then frangibility of the context follows from Lemma C.5. Finally, Theorem 3 follows by combining Theorem C.1 with Lemmas C.4 and C.5.

For *unweighted* judgement contexts, a second route is available. If $\mathcal{C}$ is suitable for F , then we can argue as above. Otherwise, that is if $d(\mathbf{x}, \mathbf{y}) \leq 2$ whenever $F(\mu) = \{\mathbf{x}, \mathbf{y}\}$, then we can use a straightforward direct argument (along with the three axioms) to show that F is the median rule. This is why Theorem 1 does not require the auxiliary structural conditions of Theorem 3.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

A On symmetry and equal weights

We have assumed that the aggregation is to be based on some ‘given’ assignment of weights to the importance of different issues in aligning the output judgment with voters input judgments on the corresponding issue. The evaluation of such alignment is the hallmark of majoritarian judgment aggregation. To be relevant, our results do not require that majoritarian judgment aggregation encoded in the (E)SME axiom is normatively compelling in all formally well-posed judgment aggregation problems. Instead, we have presented a range of applications in which the axiom is arguably well-motivated with weights arising naturally from the application. For a further discussion of some of the issues arising from the assumption of given issues and given weights, see Nehring and Pivato (2019).

Of special interest is the case of equal weights. As pointed out in Sect. 3, such symmetric treatment of issues is naturally motivated when the context is globally symmetric. But while global symmetries provide strong sufficient grounds for equal weighting, they seem quite far from necessary for this conclusion. In particular, appropriate local symmetries can be sufficient as well. To illustrate, return to the application to committee selection under a single composition constraint, as in Example 3. Asymmetries across types of candidates arise from different minimum requirements L_i and/or different cardinalities in the number of available candidates of a given type. Such asymmetries might but need not diminish the appeal of SME.

In particular, one can argue that asymmetries of the output space $\mathcal{X}$ remain immaterial as long as the input space $\mathcal{Y}$ is symmetric, for instance if $\mathcal{Y}$ is the approval space $2^{\mathcal{K}}$. Indeed, the SME requirement relies on a dominance comparison among pairs

of (feasible) views; that comparison does not depend on the structure of the entire feasible set of views $\mathcal{X}$ to which the pairs belong; hence it should not matter what $\mathcal{X}$ looks like, in particular: which symmetries $\mathcal{X}$ exhibits. In Example 3, for instance, SME compares two feasible committees in terms of the ‘estimated’ overall merit. The validity of that comparison is, arguably, not affected by which other committees are also feasible.

By contrast, suppose that the input space coincides with the output space and thus exhibits the same asymmetric composition constraints, as Example 4 above. Here, the case for SME with equal weights is less straightforward, but a good argument can still be given on the basis of local symmetries. To see how there might be an issue, consider an instance of the Green leadership duo election example in which there are 7 male and 2 female candidates. Consider a profile in which the vote on the female candidates is fairly evenly split, and some voters give both votes to these candidates, so that the two female candidates receive a vote tally of 65% and 45%. These will constitute the leadership duo if and only if the top male candidate receives less than 45% of the vote out of the remaining 90%; in other words, for a male candidate to be elected, he must receive an absolute majority within the male pool. One might question the output of this rule normatively and object that a vote for a male and a female candidate are not equally strong signals of candidate merit in this case. But note that *if* the designers of the voting rule had wished transparent comparability of votes for all candidates in terms of “merit content”, they could have chosen the approval space as the input space. Yet, in point of fact, the designers adopted an asymmetric design of the input spaces; any asymmetry in how votes effectively count is thus reasonably viewed as a feature of the judgment context, not a bug. So the putative objection is weak.

Somewhat more formally, one can justify the application of SME with equal weights as follows. For each $j \in J$, select a subset $\mathcal{K}_j$ of cardinality L_j of the candidates, and consider the subdomain $\mathcal{D}$ of profiles in which voters unanimously vote for each candidate in $\bigcup \mathcal{K}_j$. Take the election of these candidates as a given. Then, for the resulting subdomain, the problem reduces to filling the remaining $L - \sum L_j$ slots, *without any composition constraints*. This reduced problem is completely symmetric in issues (candidates), and thus warrants an application of SME with equal weights over the set $\mathcal{K} - \bigcup \mathcal{K}_j$. Since the subsets $\mathcal{K}_j$ were selected arbitrarily, if the judgment rule on the original domain $\Delta(\mathcal{V})$ is supermajority efficient with respect to some vector of weights, these weights must be equal.³⁹

B Proof of Propositions 2, 3 and 5

The proof of Proposition 2 will require some technical preliminaries. Let $(\mathcal{K}, \lambda, \mathcal{X}, \mathcal{V})$ be a judgement context, and let $F : \Delta(\mathcal{V}) \rightrightarrows \mathcal{X}$ be a weighted judgement aggregation rule. For any $\mu, \nu \in \Delta(\mathcal{V})$, define $\mu \oplus \nu := \frac{1}{2}\mu + \frac{1}{2}\nu$. Consider the following, weaker versions of Reinforcement and Judgement Consistency.

³⁹ As it stands, the argument works only if there are at least two ‘free’ slots, i.e. if $L - \sum L_j \geq 2$. With some additional twist, it can be extended to the case $L - \sum L_j = 1$; the details need not concern us here, since the main point is to illustrate the general idea of an appeal to local symmetries, rather than to nail this particular application.

Even Reinforcement: For any $\mu_1, \mu_2 \in \Delta(\mathcal{Y})$, if $F(\mu_1) \cap F(\mu_2) \neq \emptyset$, then $F(\mu_1 \oplus \mu_2) = F(\mu_1) \cap F(\mu_2)$.

Even Judgement Consistency: For any $\mu, \mu', v, v' \in \Delta(\mathcal{Y})$, and any views $\mathbf{x}, \mathbf{y} \in \mathcal{X}$, if (a) $\mathbf{y} \in F(\mu \oplus v)$, and (b) $\mathbf{x} \in F(\mu' \oplus v)$, and (c) $\mathbf{y} \notin F(\mu' \oplus v)$, and (d) $\mathbf{x} \in F(\mu \oplus v')$, then (e) $\mathbf{y} \notin F(\mu' \oplus v')$.

Lemma B.1 *Let $\mathcal{C} = (\mathcal{K}, \lambda, \mathcal{X}, \mathcal{Y})$ be any judgement context, and let $F : \Delta(\mathcal{Y}) \rightrightarrows \mathcal{X}$ be any judgement aggregation rule. If F satisfies Even Judgement Consistency, then F satisfies Even Reinforcement.*

Proof Let $\mu_1, \mu_2 \in \Delta(\mathcal{Y})$.

Claim 1: *Let $\mathbf{y} \in \mathcal{X}$. If $\mathbf{y} \in F(\mu_1)$ and $\mathbf{y} \in F(\mu_2)$, then $\mathbf{y} \in F(\mu_1 \oplus \mu_2)$.*

Proof Set $\mu = v = \mu_1$ and $\mu' = v' = \mu_2$. Then for any $\mathbf{x} \in \mathcal{X} \setminus \{\mathbf{y}\}$, Even Judgement Consistency says: if (a) $\mathbf{y} \in F(\mu_1)$, and (bd) $\mathbf{x} \in F(\mu_1 \oplus \mu_2)$, and (c) $\mathbf{y} \notin F(\mu_2 \oplus \mu_1)$, then $\mathbf{y} \notin F(\mu_2)$.

Taking the contrapositive, if $\mathbf{y} \in F(\mu_2)$, then one of the hypotheses (a), (bd), or (c) must be false. In particular, if $\mathbf{y} \in F(\mu_2)$ and $\mathbf{y} \in F(\mu_1)$, then

$$\text{either } \mathbf{x} \notin F(\mu_1 \oplus \mu_2) \text{ or } \mathbf{y} \in F(\mu_1 \oplus \mu_2). \quad (\text{B1})$$

This holds for all $\mathbf{x} \in \mathcal{X} \setminus \{\mathbf{y}\}$.

Now, by contradiction, suppose $\mathbf{y} \in F(\mu_2)$ and $\mathbf{y} \in F(\mu_1)$, but $\mathbf{y} \notin F(\mu_1 \oplus \mu_2)$. Applying (B1), we obtain $\mathbf{x} \notin F(\mu_1 \oplus \mu_2)$ for all $\mathbf{x} \in \mathcal{X} \setminus \{\mathbf{y}\}$, which means that $F(\mu_1 \oplus \mu_2) = \emptyset$. This is a contradiction. The claim follows. $\diamond$ Claim 1

Claim 2: *Let $\mathbf{y} \in \mathcal{X}$. If $F(\mu_1) \cap F(\mu_2) \neq \emptyset$, and $\mathbf{y} \notin F(\mu_2)$, then $\mathbf{y} \notin F(\mu_1 \oplus \mu_2)$.*

Proof Let $\mathbf{x} \in F(\mu_1) \cap F(\mu_2)$. If we set $\mu' = v = \mu_2$ and $\mu = v' = \mu_1$, then Even Judgement Consistency says: if (a) $\mathbf{y} \in F(\mu_1 \oplus \mu_2)$, and (b) $\mathbf{x} \in F(\mu_2)$, and (c) $\mathbf{y} \notin F(\mu_2)$, and (d) $\mathbf{x} \in F(\mu_1)$, then (e) $\mathbf{y} \notin F(\mu_2 \oplus \mu_1)$.

Thus, $\mathbf{y} \notin F(\mu_2 \oplus \mu_1)$, as claimed. $\diamond$ Claim 2

Combining Claims 1 and 2, we conclude that, if $F(\mu_1) \cap F(\mu_2) \neq \emptyset$, then $F(\mu_1 \oplus \mu_2) = F(\mu_1) \cap F(\mu_2)$. $\square$

Lemma B.2 *Let $\mathcal{C} = (\mathcal{K}, \lambda, \mathcal{X}, \mathcal{Y})$ be any judgement context, and suppose $F : \Delta(\mathcal{Y}) \rightrightarrows \mathcal{X}$ satisfies Continuity. If F satisfies Even Reinforcement, then F satisfies Reinforcement.*

Proof Let $\mu_1, \mu_2 \in \Delta(\mathcal{Y})$, and suppose $F(\mu_1) \cap F(\mu_2) \neq \emptyset$. Let $\mathbb{Q}_2$ be the set of dyadic rationals (that is, $\mathbb{Q}_2 := \{\frac{n}{2^k}; n \text{ and } k \text{ integers, } k > 0\}$).

Claim 1: *$F(q\mu_1 + (1-q)\mu_2) = F(\mu_1) \cap F(\mu_2)$ for all $q \in \mathbb{Q}_2 \cap [0, 1]$.*

Proof Even Reinforcement implies that $F(\mu_1 \oplus \mu_2) = F(\mu_1) \cap F(\mu_2)$. It follows that $F(\mu_1 \oplus \mu_2) \cap F(\mu_1) = F(\mu_1) \cap F(\mu_2) \neq \emptyset$ and $F(\mu_1 \oplus \mu_2) \cap F(\mu_2) = F(\mu_1) \cap F(\mu_2) \neq \emptyset$. Thus, applying Even Reinforcement again, we deduce that $F[\mu_1 \oplus (\mu_1 \oplus \mu_2)] = F(\mu_1) \cap F(\mu_2)$ and $F[(\mu_1 \oplus \mu_2) \oplus \mu_2] = F(\mu_1) \cap F(\mu_2)$. But $\mu_1 \oplus (\mu_1 \oplus \mu_2) = \frac{3}{4}\mu_1 + \frac{1}{4}\mu_2$, while $(\mu_1 \oplus \mu_2) \oplus \mu_2 = \frac{1}{4}\mu_1 + \frac{3}{4}\mu_2$. Iterating this argument yields the claim, by induction. $\diamond$ Claim 1

Now, $\mathbb{Q}_2$ is dense in $\mathbb{R}$. Thus, Continuity and Claim 1 imply that

$$F(r\mu_1 + (1-r)\mu_2) \supseteq F(\mu_1) \cap F(\mu_2), \quad \text{for all } r \in [0, 1]. \quad (\text{B.2})$$

Claim 2: $F(r\mu_1 + (1-r)\mu_2) \subseteq F(\mu_1) \cap F(\mu_2)$ for all $r \in [0, 1]$.

Proof Suppose $\mathbf{y} \notin F(\mu_1)$. We must show that $\mathbf{y} \notin F(r\mu_1 + (1-r)\mu_2)$ for all $r \in [0, 1]$.

By contradiction, let $\mathcal{R}_{\mathbf{y}} := \{r \in [0, 1]; \mathbf{y} \in F(r\mu_1 + (1-r)\mu_2)\}$, and suppose $\mathcal{R}_{\mathbf{y}} \neq \emptyset$. By Continuity, 1 is not a cluster point of $\mathcal{R}_{\mathbf{y}}$ (because $\mathbf{y} \notin F(\mu_1)$). Thus, if $R := \sup(\mathcal{R})$, then $R < 1$. Now let $r \in \mathcal{R}_{\mathbf{y}}$. Find some $s \in (R, 1)$ such that $r/s = q \in \mathbb{Q}_2$.

Let $v := s\mu_1 + (1-s)\mu_2$. Then $\mathbf{y} \notin F(v)$, by definition of R . But $F(v) \supseteq F(\mu_1) \cap F(\mu_2)$, by statement (B.2). Thus, $F(v) \cap F(\mu_2) \neq \emptyset$. Thus, Claim 1 implies that $F(qv + (1-q)\mu_2) = F(v) \cap F(\mu_2)$. But $qv + (1-q)\mu_2 = r\mu_1 + (1-r)\mu_2$ (because $r = qs$). Thus, we conclude that $F(r\mu_1 + (1-r)\mu_2) = F(v) \cap F(\mu_2)$. Hence $\mathbf{y} \notin F(r\mu_1 + (1-r)\mu_2)$, which contradicts the fact that $r \in \mathcal{R}_{\mathbf{y}}$.

To avert the contradiction, we must have $\mathcal{R}_{\mathbf{y}} = \emptyset$. This argument works for any $\mathbf{y} \notin F(\mu_1)$, and likewise any $\mathbf{y} \notin F(\mu_2)$. The claim follows. $\diamond_{\text{Claim 2}}$

Statement (B.2) and Claim 2 imply $F(r\mu_1 + (1-r)\mu_2) = F(\mu_1) \cap F(\mu_2)$ for all $r \in [0, 1]$. Thus, F satisfies Reinforcement. $\square$

Proof of Proposition 2 If F satisfies Judgement Consistency, then it satisfies Even Judgement Consistency. Then Lemma B.1 says F satisfies Even Reinforcement. If F also satisfies Continuity then Lemma B.2 says it satisfies Reinforcement. $\square$

We now introduce some notation that will be used in the proofs of Proposition 3 and Theorem 1. For any $\mathbf{x} \in \mathcal{X}$, let $\mathbf{x} \bullet \tilde{\mu} := \sum_{k \in \mathcal{K}} x_k \tilde{\mu}_k$. Then we can rewrite (2) more simply:

$$\text{Median}(\mathcal{X}, \mu) := \operatorname{argmax}_{\mathbf{x} \in \mathcal{X}} (\mathbf{x} \bullet \tilde{\mu}), \quad \text{for all } \mu \in \Delta(\mathcal{Y}). \quad (\text{B.3})$$

If ϕ is a gain function and $\tilde{\mu} = (\tilde{\mu}_k)_{k \in \mathcal{K}} \in \mathbb{R}^{\mathcal{K}}$, then we define $\phi(\tilde{\mu}) := (\phi(\tilde{\mu}_k))_{k \in \mathcal{K}}$. For any $\mathbf{x} = (x_k)_{k \in \mathcal{K}} \in \{\pm 1\}^{\mathcal{K}}$, we define $\mathbf{x} \bullet \phi(\tilde{\mu}) := \sum_{k \in \mathcal{K}} x_k \phi(\tilde{\mu}_k)$.

Thus, formula (3) becomes:

$$F_{\phi}(\mu) = \operatorname{argmax}_{\mathbf{x} \in \mathcal{X}} (\mathbf{x} \bullet \phi(\tilde{\mu})), \quad \text{for all } \mu \in \Delta(\mathcal{Y}). \quad (\text{B.4})$$

Proof of Proposition 3 Let $\mathcal{A}$ be a finite set. If $|\mathcal{A}| = 2$, then every AMR (including the median rule) coincides with simple majority rule. If $|\mathcal{A}| = 3$, then it is easily verified that all AMRs (including the median rule) agree.⁴⁰ So it remains to consider the case

⁴⁰ If the majority tournament is transitive, then it is selected by any AMR. If the majority tournament is intransitive—say, $a > b > c > a$, then any AMR selects the transitive ranking(s) obtained by reversing whichever of the three comparisons $a > b$, $b > c$ or $c > a$ is supported by the *smallest* majority, as explained in Sect. 3, just prior to Example 3.

$|\mathcal{A}| \geq 4$. Let F be an AMR on the domain of rational-valued profiles in $\mathcal{X}_{\mathcal{A}}^{\text{rk}}$. Let $\phi : [-1, 1] \rightarrow {}^*\mathbb{R}$ be an odd, increasing function such that $F = F_{\phi}$.

Claim 1: *There exists $Q \in (0, 1) \cap \mathbb{Q}$ such that ϕ is linear on $[-Q, Q] \cap \mathbb{Q}$.*

Proof It suffices to deal with the case $|\mathcal{A}| = 4$. So, suppose $\mathcal{A} = \{a, b, c, d\}$. Let $\mathcal{K} = \{ab, ac, ad, bc, bd, cd\}$, where ab for example represents the proposition “ $a \succ b$ ”. Thus, $\mathbf{x}_{cdab} = (1, -1, -1, -1, -1, 1)$ is the element of $\mathcal{X}_{\mathcal{A}}^{\text{rk}}$ representing the preference order “ $c \succ d \succ a \succ b$ ”, for example.

There exists $Q \in (0, 1) \cap \mathbb{Q}$ such that $[-Q, Q]^{\mathcal{K}} \subset \text{conv}(\mathcal{X}_{\mathcal{A}}^{\text{rk}})$ (Nehring and Pivato 2011, Example 3.3). This means that, for any $\mathbf{r} \in [-Q, Q]^{\mathcal{K}}$, there exists a profile $\mu \in \Delta(\mathcal{X}_{\mathcal{A}}^{\text{rk}})$ such that $\tilde{\mu} = \mathbf{r}$. In particular, for any nonzero $q \in (0, Q] \cap \mathbb{Q}$, we can construct profiles $\mu, \mu' \in \Delta(\mathcal{X}_{\mathcal{A}}^{\text{rk}})$ such that

$$\begin{aligned} \tilde{\mu} &= (q, & ab & & ac & ad & bc & bd & & cd \\ \tilde{\mu}' &= (0, & & & q, & -q, & 0, & 0, & & q), \end{aligned} \quad (\text{B.5})$$

and thus, $c\tilde{\mu} + (1-c)\tilde{\mu}' = (cq, (1-c)q, -q, 0, cq, (1-c)q)$,

for all $c \in [0, 1] \cap \mathbb{Q}$. Formula (3) implies that $F_{\phi}[\mathcal{X}_{\mathcal{A}}^{\text{rk}}, \mu]$ consists of all preference orders which agree with any two out of the three assertions “ $a \succ b$ ”, “ $b \succ d$ ”, and “ $d \succ a$ ”. Likewise, $F_{\phi}[\mathcal{X}_{\mathcal{A}}^{\text{rk}}, \mu']$ consists of all preference orders which agree with any two out of the three assertions “ $a \succ c$ ”, “ $c \succ d$ ”, and “ $d \succ a$ ”. In other words,

$$\begin{aligned} F_{\phi}[\mathcal{X}_{\mathcal{A}}^{\text{rk}}, \mu] &= \{\mathbf{x}_{cabd}, \mathbf{x}_{acbd}, \mathbf{x}_{abcd}, \mathbf{x}_{abdc}, \mathbf{x}_{cbda}, \mathbf{x}_{bcda}, \mathbf{x}_{bdca}, \mathbf{x}_{bdac}, \mathbf{x}_{cdab}, \mathbf{x}_{dcab}, \\ &\quad \mathbf{x}_{dacb}, \mathbf{x}_{dabc}\}, \\ F_{\phi}[\mathcal{X}_{\mathcal{A}}^{\text{rk}}, \mu'] &= \{\mathbf{x}_{bacd}, \mathbf{x}_{abdc}, \mathbf{x}_{acbd}, \mathbf{x}_{acdb}, \mathbf{x}_{bcda}, \mathbf{x}_{cbda}, \mathbf{x}_{cdba}, \mathbf{x}_{cdab}, \mathbf{x}_{bdac}, \mathbf{x}_{dbac}, \\ &\quad \mathbf{x}_{dabc}, \mathbf{x}_{dacb}\}. \end{aligned}$$

Let $\mu'' := c\mu + (1-c)\mu'$. Then Reinforcement yields

$$\begin{aligned} F_{\phi}(\mathcal{X}_{\mathcal{A}}^{\text{rk}}, \mu'') &= F_{\phi}[\mathcal{X}_{\mathcal{A}}^{\text{rk}}, \mu] \cap F_{\phi}[\mathcal{X}_{\mathcal{A}}^{\text{rk}}, \mu'] \\ &= \{\mathbf{x}_{acbd}, \mathbf{x}_{abcd}, \mathbf{x}_{cbda}, \mathbf{x}_{bcda}, \mathbf{x}_{bdac}, \mathbf{x}_{cdab}, \mathbf{x}_{dacb}, \mathbf{x}_{dabc}\}. \end{aligned}$$

In particular, $\{\mathbf{x}_{abcd}, \mathbf{x}_{cdab}\} \subset F_{\phi}(\mathcal{X}_{\mathcal{A}}^{\text{rk}}, \mu'')$. Thus, $\phi(\tilde{\mu}'') \bullet \mathbf{x}_{abcd} = \phi(\tilde{\mu}'') \bullet \mathbf{x}_{cdab}$. But $\mathbf{x}_{abcd} = (1, 1, 1, 1, 1, 1)$ and $\mathbf{x}_{cdab} = (1, -1, -1, -1, -1, 1)$; thus equation (B.5) yields

$$\phi(\tilde{\mu}'') \bullet \mathbf{x}_{abcd} = -\phi(q) + 2\phi(cq) + 2\phi((1-c)q) \quad \text{and} \quad \phi(\tilde{\mu}'') \bullet \mathbf{x}_{cdab} = \phi(q).$$

Thus, $\phi(q) = -\phi(q) + 2\phi(cq) + 2\phi((1-c)q)$, which means that

$$\phi(q) = \phi(cq) + \phi((1-c)q). \quad (\text{B.6})$$

Now, let $N \in \mathbb{N}$, and let $q := Q/N$. Then (B.6) yields $\phi(2q) = \phi(q) + \phi(q) = 2\phi(q)$. Then (B.6) yields $\phi(3q) = \phi(q) + \phi(2q) \stackrel{(*)}{=} \phi(q) + 2\phi(q) = 3\phi(q)$, where

(*) is by the previous sentence. By induction, we get

$$\phi(nq) = n\phi(q) \quad \text{for all } n \in [0 \dots N]. \quad (\text{B.7})$$

Thus, $\phi(Q) = N\phi(q)$ (since $Nq = Q$), so $\phi(q) = \phi(Q)/N$. Putting this into (B.7),

$$\phi\left(\frac{n}{N}Q\right) = \frac{n}{N}\phi(Q) \quad \text{for all } n \in [0 \dots N].$$

This holds for all $N \in \mathbb{N}$. Thus, $\phi(pQ) = p\phi(Q)$ for all $p \in [0, 1] \cap \mathbb{Q}$. Since ϕ is odd,

$$\phi(pQ) = p\phi(Q), \quad \text{for all } p \in [-1, 1] \cap \mathbb{Q}. \quad (\text{B.8})$$

Now let $q \in [-Q, Q] \cap \mathbb{Q}$, and let $p := q/Q$. Then $pQ = q$, so (B.8) yields $\phi(q) = q\phi(Q)/Q$. Let $r := \phi(Q)/Q$; then we get $\phi(q) = rq$ for all $q \in [-Q, Q] \cap \mathbb{Q}$.

◇ Claim 1

In view of Claim 1,

$$F_\phi(\mathcal{X}_{\mathcal{A}}^{\text{rk}}, \mu) = \text{Median}(\mathcal{X}_{\mathcal{A}}^{\text{rk}}, \mu), \quad \text{for any } \mu \in \Delta(\mathcal{X}_{\mathcal{A}}^{\text{rk}}) \text{ such that } \tilde{\mu} \in \mathbb{Q}^{\mathcal{K}} \cap [-Q, Q]^{\mathcal{K}}. \quad (\text{B.9})$$

Now, let μ_0 be a “completely tied” profile, so that $\tilde{\mu}_0$ is the zero vector. Thus, $F(\mathcal{X}_{\mathcal{A}}^{\text{rk}}, \mu_0) = \mathcal{X}_{\mathcal{A}}^{\text{rk}}$. Let $\mu \in \Delta(\mathcal{X}_{\mathcal{A}}^{\text{rk}})$, and let $\mu' := c\mu + (1-c)\mu_0$ for some $c \in (0, 1)$. Then $\tilde{\mu}' = c\tilde{\mu}$, so that

$$\text{Median}(\mathcal{X}_{\mathcal{A}}^{\text{rk}}, \mu') = \text{Median}(\mathcal{X}_{\mathcal{A}}^{\text{rk}}, \mu). \quad (\text{B.10})$$

Meanwhile, Reinforcement implies that

$$F(\mathcal{X}_{\mathcal{A}}^{\text{rk}}, \mu') = F(\mathcal{X}_{\mathcal{A}}^{\text{rk}}, \mu_0) \cap F(\mathcal{X}_{\mathcal{A}}^{\text{rk}}, \mu) = F(\mathcal{X}_{\mathcal{A}}^{\text{rk}}, \mu). \quad (\text{B.11})$$

Finally, by choosing c sufficiently small, we can ensure that $\tilde{\mu}' \in \mathbb{Q}^{\mathcal{K}} \cap [-Q, Q]^{\mathcal{K}}$; thus, equation (B.9) says that

$$F(\mathcal{X}_{\mathcal{A}}^{\text{rk}}, \mu') = \text{Median}(\mathcal{X}_{\mathcal{A}}^{\text{rk}}, \mu'). \quad (\text{B.12})$$

Combining equations (B.10) - (B.12), we obtain $F(\mathcal{X}_{\mathcal{A}}^{\text{rk}}, \mu) = \text{Median}(\mathcal{X}_{\mathcal{A}}^{\text{rk}}, \mu)$. □

Proof of Proposition 5 It is easy to verify that F^α satisfies WESME and Continuity. (Or see Lemma D.7 from Nehring and Pivato (2019).) Thus, it remains to show that F^α satisfies Reinforcement.

For any weight vector $\lambda = (\lambda_k)_{k \in \mathcal{K}}$, and any exponent $\alpha > 0$, we define the correspondence $F_\lambda^\alpha : \Delta(\mathcal{X}_{L,M}^K) \rightrightarrows \mathcal{X}_{L,M}^K$ by setting

$$F_\lambda^\alpha(\mu) := \operatorname{argmax}_{\mathbf{x} \in \mathcal{X}} \left(\sum_{k \in \mathcal{K}} \lambda_k \phi_\alpha(x_k \tilde{\mu}_k) \right), \quad \text{for all } \mu \in \Delta(\mathcal{X}_{L,M}^K). \quad (\text{B.13})$$

We also define $\lambda^\alpha := (\lambda_k^\alpha)_{k \in \mathcal{K}}$.

Claim 1: Let $\alpha, \beta > 0$, and let $\lambda, \kappa \in \mathbb{R}_+^{\mathcal{K}}$ be weight vectors. If $\lambda^\beta = \kappa^\alpha$, then $F_\lambda^\alpha(\mathcal{X}_{L,M}^K, \mu) = F_\kappa^\beta(\mathcal{X}_{L,M}^K, \mu)$ for all $\mu \in \Delta(\mathcal{X}_{L,M}^K)$.

Proof Let $\mu \in \Delta(\mathcal{X}_{L,M}^K)$, and let $\mathcal{J} := \{k \in \mathcal{K}; \tilde{\mu}_k > 0\}$. There are now three cases.

Case 1. If $L \leq |\mathcal{J}| \leq M$, then $F_\lambda^\alpha(\mathcal{X}_{L,M}^K, \mu) = F_\kappa^\beta(\mathcal{X}_{L,M}^K, \mu) = \{\mathbf{x}\}$, where $\mathbf{x} \in \mathcal{X}_{L,M}^K$ is the view with $x_j = 1$ for all $j \in \mathcal{J}$ and $x_k = -1$ for all $k \in \mathcal{K} \setminus \mathcal{J}$.

Case 2. Suppose $|\mathcal{J}| > M$. Define $\eta := \lambda^{1/\alpha}$. Then we also have $\eta = \kappa^{1/\beta}$ (because $\eta^\beta = \lambda^{\beta/\alpha} = (\lambda^\beta)^{1/\alpha} = (\kappa^\alpha)^{1/\alpha} = \kappa$, because $\lambda^\beta = \kappa^\alpha$ by hypothesis). By reordering the elements of $\mathcal{K} = [1 \dots K]$ if necessary, we can assume without loss of generality that

$$\eta_1 \tilde{\mu}_1 \geq \eta_2 \tilde{\mu}_2 \geq \dots \geq \eta_J \tilde{\mu}_J > 0 \geq \eta_{J+1} \tilde{\mu}_{J+1} \geq \dots \geq \eta_K \tilde{\mu}_K. \quad (\text{B.14})$$

(Thus, $\mathcal{J} = [1 \dots J]$.) Now, for any $k \in \mathcal{K}$, observe that $\lambda_k \phi_\alpha(x_k \tilde{\mu}_k) = \operatorname{sign}(x_k) \lambda_k |\tilde{\mu}_k|^\alpha = \operatorname{sign}(x_k) |\eta_k \tilde{\mu}_k|^\alpha$. Thus,

$$\sum_{k \in \mathcal{K}} \lambda_k \phi_\alpha(x_k \tilde{\mu}_k) = \sum_{k \in \mathcal{K}} \operatorname{sign}(x_k) |\eta_k \tilde{\mu}_k|^\alpha. \quad (\text{B.15})$$

Suppose $\eta_M \tilde{\mu}_M > \eta_{M+1} \tilde{\mu}_{M+1}$. Then the unique maximizer in $\mathcal{X}_{L,M}^K$ of the sum (B.15) is the element $\mathbf{x} \in \mathcal{X}_{L,M}^K$ such that $x_m = 1$ for all $m \in [1 \dots M]$, while $x_k = -1$ for all $k \in [M+1 \dots K]$. Thus, definition (B.13) yields $F_\lambda^\alpha(\mathcal{X}_{L,M}^K, \mu) = \{\mathbf{x}\}$. But by an identical argument, $\kappa_k \phi_\beta(x_k \tilde{\mu}_k) = \operatorname{sign}(x_k) |\eta_k \tilde{\mu}_k|^\beta$, so that

$$\sum_{k \in \mathcal{K}} \kappa_k \phi_\beta(x_k \tilde{\mu}_k) = \sum_{k \in \mathcal{K}} \operatorname{sign}(x_k) |\eta_k \tilde{\mu}_k|^\beta, \quad (\text{B.16})$$

so this sum is *also* uniquely maximized by $\mathbf{x}$, so $F_\kappa^\beta(\mathcal{X}_{L,M}^K, \mu) = \{\mathbf{x}\}$ also. Thus, $F_\lambda^\alpha(\mathcal{X}_{L,M}^K, \mu) = F_\kappa^\beta(\mathcal{X}_{L,M}^K, \mu)$, as claimed.

On the other hand, suppose $\eta_{N-1} \tilde{\mu}_{N-1} > \eta_N \tilde{\mu}_N = \eta_{N+1} \tilde{\mu}_{N+1} = \dots = \eta_P \tilde{\mu}_P > \eta_{P+1} \tilde{\mu}_{P+1}$ for some N, P with $N \leq M \leq P$. In this case, the sums (B.15) and (B.16) have more than one maximizer.⁴¹ Even in this case, however,

⁴¹ To be precise, they have $\binom{P-N+1}{M-N+1}$ maximizers.

it is easy to see that they have exactly the *same* set of maximizers, so once again $F_{\lambda}^{\alpha}(\mathcal{X}_{L,M}^K, \mu) = F_{\kappa}^{\beta}(\mathcal{X}_{L,M}^K, \mu)$.

Case 3. Suppose $|\mathcal{J}| < L$. The argument is similar to *Case 2*. Again, assume without loss of generality that (B.14) holds. If $\eta_L \tilde{\mu}_L > \eta_{L+1} \tilde{\mu}_{L+1}$, then by invoking equations (B.15) and (B.16), we see that $F_{\lambda}^{\alpha}(\mathcal{X}_{L,M}^K, \mu) = F_{\kappa}^{\beta}(\mathcal{X}_{L,M}^K, \mu) = \{\mathbf{x}\}$, where $\mathbf{x} \in \mathcal{X}_{L,M}^K$ is defined by $x_{\ell} := 1$ for all $\ell \in [1 \dots L]$, while $x_k := -1$ for all $k \in [L+1 \dots K]$. If $\eta_L \tilde{\mu}_L = \eta_{L+1} \tilde{\mu}_{L+1}$, then the sums (B.15) and (B.16) have more than one maximizer, but they have the same maximizers, so that $F_{\lambda}^{\alpha}(\mathcal{X}_{L,M}^K, \mu) = F_{\kappa}^{\beta}(\mathcal{X}_{L,M}^K, \mu)$.

◇ Claim 1

Now, fix $\lambda \in \mathbb{R}_{+}^K$, and consider the judgement context $\mathcal{C} := (\mathcal{K}, \lambda, \mathcal{X}_{L,M}^K, \mathcal{X}_{L,M}^K)$. Let $\kappa := \lambda^{1/\alpha}$. Then

$$F^{\alpha}(\mathcal{C}, \mu) \stackrel{(*)}{=} F_{\lambda}^{\alpha}(\mathcal{X}_{L,M}^K, \mu) \stackrel{(\dagger)}{=} F_{\kappa}^1(\mathcal{X}_{L,M}^K, \mu), \quad \text{for all } \mu \in \Delta(\mathcal{X}_{L,M}^K). \quad (\text{B.17})$$

Here, $(*)$ is obtained by comparing equations (7) and (B.13), while $(\dagger)$ follows from Claim 1. Now, F_{κ}^1 is just the $(\kappa$ -weighted) median rule; thus, F_{κ}^1 satisfies Reinforcement on $\mathcal{X}_{L,M}^K$. Thus, F_{λ}^{α} also satisfies Reinforcement on $\mathcal{X}_{L,M}^K$. This proves the first assertion of Proposition 5.

If $\lambda = (1, 1, \dots, 1)$, then $\kappa = \lambda$. Thus $F_{\kappa}^1 = F_{\lambda}^1$. Thus, statement (B.17) implies that $F^{\alpha}(\mathcal{C}, \mu) = F^1(\mathcal{C}, \mu)$ for all $\mu \in \Delta(\mathcal{X}_{L,M}^K)$ —in other words, F^{α} itself is the median rule on $\mathcal{C}$. This proves the second assertion of Proposition 5.

On the other hand, if $\lambda \neq (1, 1, \dots, 1)$, then $\kappa \neq \lambda$. Thus, F_{λ}^1 and F_{κ}^1 will not always agree on $\mathcal{X}_{L,M}^K$. (Discrepancies between these two rules can be constructed using reasoning similar to the proof of Claim 1.) Thus, statement (B.17) implies that F^{α} is *not* the median rule on $\mathcal{C}$. This proves the third assertion of Proposition 5. □

We end this appendix with the following result, which we will need later.

Lemma B.3 *Any additive majority rule is SME on any judgement context.*

Proof See Lemma D.7 from Nehring and Pivato (2019). □

C Proofs of the main results

Before proving Theorem 2, we introduce some notation. Let $F : \Delta(\mathcal{Y}) \rightrightarrows \mathcal{X}$ be any additive majority rule, and let $\mathcal{C} := \text{conv}(\mathcal{Y})$. For any $\mathbf{c} \in \mathcal{C}$, there exists some $\mu \in \Delta(\mathcal{Y})$ such that $\tilde{\mu} = \mathbf{c}$. We then define $F(\mathbf{c}) := F(\mu)$. By inspection of defining formulae (3) and (7), it is clear that this definition is independent of the choice of μ . Thus, we can define a correspondence $F : \mathcal{C} \rightrightarrows \mathcal{X}$. We will make use of this convention frequently in what follows. For any $\mathbf{x} \in \mathcal{X}$, define

$$\mathcal{C}_{\mathbf{x}}^F := \{\mathbf{c} \in \mathcal{C} ; \mathbf{x} \in F(\mathbf{c})\} \quad \text{and} \quad {}^*\mathcal{C}_{\mathbf{x}}^F := \{\mathbf{c} \in \mathcal{C} ; F(\mathbf{c}) = \{\mathbf{x}\}\}. \quad (\text{C1})$$

Let ${}^\circ\mathcal{C}$ be the topological interior of $\mathcal{C}$ as a subset of $\mathbb{R}^{\mathcal{K}}$. (Note that ${}^\circ\mathcal{C} \neq \emptyset$ because $\mathcal{Y}$ is thick.) For any $\mathbf{x}, \mathbf{y} \in \mathcal{X}$, we define

$$\mathcal{B}_{\mathbf{x}, \mathbf{y}}^F := \mathcal{C}_{\mathbf{x}}^F \cap \mathcal{C}_{\mathbf{y}}^F = \{\mathbf{c} \in \mathcal{C} ; \mathbf{x}, \mathbf{y} \in F(\mathbf{c})\}, \quad (\text{C2})$$

$$\text{and } {}^*\mathcal{B}_{\mathbf{x}, \mathbf{y}}^F := \{\mathbf{c} \in {}^\circ\mathcal{C} ; F(\mathbf{c}) = \{\mathbf{x}, \mathbf{y}\}\}. \quad (\text{C3})$$

For any weight vector $\lambda = (\lambda_k)_{k \in \mathcal{K}}$, view $\mathbf{x} \in \{\pm 1\}^{\mathcal{K}}$, and profile $\mu \in \Delta(\mathcal{Y})$, we define

$$\mathbf{x} \bullet_{\lambda} \tilde{\mu} := \sum_{k \in \mathcal{K}} \lambda_k x_k \tilde{\mu}_k.$$

Thus, for any judgement context $\mathcal{C} := (\mathcal{K}, \lambda, \mathcal{X}, \mathcal{Y})$ and $\mu \in \Delta(\mathcal{Y})$, formula (8) becomes:

$$\text{Median}(\mathcal{C}, \mu) = \operatorname{argmax}_{\mathbf{x} \in \mathcal{X}} (\mathbf{x} \bullet_{\lambda} \tilde{\mu}), \quad \text{for all } \mu \in \Delta(\mathcal{Y}).$$

This generalizes formula (B.3). For any gain function $\phi : [-1, 1] \rightarrow {}^*\mathbb{R}$, formula (7) becomes

$$F_{\phi}(\mu) = \operatorname{argmax}_{\mathbf{x} \in \mathcal{X}} (\mathbf{x} \bullet_{\lambda} \phi(\tilde{\mu})),$$

where $\phi(\tilde{\mu}) := (\phi(\tilde{\mu}_k))_{k \in \mathcal{K}}$. This generalizes (B.4). Let $\mathcal{K}(\mathbf{x}, \mathbf{y}) := \{k \in \mathcal{K}; x_k \neq y_k\}$. Thus,

$$\text{for any } \mu \in \Delta(\mathcal{Y}), \quad (\mathbf{x} - \mathbf{y}) \bullet_{\lambda} \phi(\tilde{\mu}) = \sum_{k \in \mathcal{K}(\mathbf{x}, \mathbf{y})} \lambda_k (x_k - y_k) \phi(\tilde{\mu}_k), \quad (\text{C4})$$

because $(x_k - y_k) = 0$ for all $k \in \mathcal{K} \setminus \mathcal{K}(\mathbf{x}, \mathbf{y})$. In particular, if $\tilde{\mu} \in \mathcal{B}_{\mathbf{x}, \mathbf{y}}^F$ —that is, if $\{\mathbf{x}, \mathbf{y}\} \subseteq F(\mathcal{X}, \mu)$ —then the sum (C4) must be zero.

Proof of Theorem 2 It is easy to verify that the median rule satisfies the Reinforcement and Continuity. It remains to verify the converse. So, let $F : \Delta(\mathcal{Y}) \rightrightarrows \mathcal{X}$ be an additive majority rule with gain function $\phi : [-1, 1] \rightarrow {}^*\mathbb{R}$.

Claim 1: Let $\mathbf{x}, \mathbf{y} \in \mathcal{X}$, and let $\mathbf{z} := \frac{1}{2}(\mathbf{x} + \mathbf{y})$. Then $\mathbf{z} \in \mathcal{B}_{\mathbf{x}, \mathbf{y}}^F$.

Proof Observe that $z_k = x_k = y_k$ for all $k \in \mathcal{K} \setminus \mathcal{K}(\mathbf{x}, \mathbf{y})$, while $z_k = 0$ for all $k \in \mathcal{K}(\mathbf{x}, \mathbf{y})$. Let $\epsilon > 0$ and let $\mathbf{z}^{\epsilon} := \epsilon \mathbf{x} + (1 - \epsilon) \mathbf{z}$. Then $\text{sign}(z_k^{\epsilon}) = x_k$ for all $k \in \mathcal{K}$; thus, $F(\mathbf{z}^{\epsilon}) = \{\mathbf{x}\}$ by supermajority efficiency (because of Lemma B.3). Likewise, if $\mathbf{z}^{-\epsilon} := \epsilon \mathbf{y} + (1 - \epsilon) \mathbf{z}$, then $\text{sign}(z_k^{-\epsilon}) = y_k$ for all $k \in \mathcal{K}$; thus, $F(\mathbf{z}^{-\epsilon}) = \{\mathbf{y}\}$ by supermajority efficiency. However, clearly $\lim_{\epsilon \rightarrow 0} \mathbf{z}^{\epsilon} = \mathbf{z} = \lim_{\epsilon \rightarrow 0} \mathbf{z}^{-\epsilon}$.

Thus, $\mathbf{z}$ is a cluster point of both $\mathcal{C}_{\mathbf{x}}^F$ and $\mathcal{C}_{\mathbf{y}}^F$, so $\mathbf{z} \in \mathcal{B}_{\mathbf{x}, \mathbf{y}}^F$ because F satisfies Continuity.

◇ Claim 1

Claim 2: Let $\mathbf{x}, \mathbf{y} \in \mathcal{X}$, and let $\mathbf{b} \in \mathcal{B}_{\mathbf{x}, \mathbf{y}}^F$. For all $s \in [0, 1)$, we have $s\mathbf{b} + (1-s)\mathbf{x} \in {}^*\mathcal{C}_{\mathbf{x}}^F$ and $s\mathbf{b} + (1-s)\mathbf{y} \in {}^*\mathcal{C}_{\mathbf{y}}^F$.

Proof Lemma B.3 says that F is supermajority efficient (SME); thus, $F(\mathbf{x}) = \{\mathbf{x}\}$. By hypothesis, $\{\mathbf{x}, \mathbf{y}\} \subseteq F(\mathbf{b})$. Thus, for all $s \in [0, 1)$ we have $F(s\mathbf{b} + (1-s)\mathbf{x}) = F(\mathbf{x}) \cap F(\mathbf{b}) = \{\mathbf{x}\}$, by Reinforcement. Thus, $s\mathbf{b} + (1-s)\mathbf{x} \in {}^*\mathcal{C}_{\mathbf{x}}^F$. By the same argument, $s\mathbf{b} + (1-s)\mathbf{y} \in {}^*\mathcal{C}_{\mathbf{y}}^F$. $\diamond$ Claim 2

Claim 3: For all $\mathbf{x}, \mathbf{y} \in \mathcal{X}$, if there exists $\mu \in \Delta(\mathcal{Y})$ with $F(\mathcal{C}, \mu) = \{\mathbf{x}, \mathbf{y}\}$, then ${}^*\mathcal{B}_{\mathbf{x}, \mathbf{y}}^F$ is nonempty, and it is a relatively open subset of some affine hyperplane in $\mathbb{R}^{\mathcal{K}}$.

Proof Let $K := |\mathcal{K}|$. For all $\mathbf{x} \in \mathcal{X}$, Reinforcement implies that $\mathcal{C}_{\mathbf{x}}^F$ is a convex subset of $\mathcal{C}$. Continuity implies that $\mathcal{C}_{\mathbf{x}}^F$ is closed. Thus, $\mathcal{B}_{\mathbf{x}, \mathbf{y}}^F$ is closed and convex, because $\mathcal{B}_{\mathbf{x}, \mathbf{y}}^F = \mathcal{C}_{\mathbf{x}}^F \cap \mathcal{C}_{\mathbf{y}}^F$. But Claim 2 implies that $\mathcal{B}_{\mathbf{x}, \mathbf{y}}^F$ has empty interior; thus, it is a convex subset of $\mathbb{R}^{\mathcal{K}}$ with dimension at most $K - 1$.

SME implies that $F(\mathbf{x}) = \{\mathbf{x}\}$. Thus, Continuity implies that $F(\mathbf{c}) = \{\mathbf{x}\}$ for all $\mathbf{c} \in \mathcal{C}$ in some open ball around $\mathbf{x}$. Thus, $\mathcal{C}_{\mathbf{x}}^F$ contains an open ball around $\mathbf{x}$. Thus, $\mathcal{C}_{\mathbf{x}}^F$ itself has nonempty interior, so it is a closed, convex set of dimension K .

Now, let $\mathbf{b} := \tilde{\mu}$. Then $F(\mathbf{b}) = \{\mathbf{x}, \mathbf{y}\}$. Thus, Continuity yields some $\epsilon > 0$ such that, if $\mathcal{D}$ is the ball of radius ϵ around $\mathbf{b}$ in $\mathcal{C}$, then $F(\mathbf{d}) \subseteq \{\mathbf{x}, \mathbf{y}\}$ for all $\mathbf{d} \in \mathcal{D}$. Define the convex sets:

$$\mathcal{D}_{\mathbf{x}} := \mathcal{D} \setminus \mathcal{C}_{\mathbf{y}}^F, \quad \mathcal{D}_{\mathbf{y}} := \mathcal{D} \setminus \mathcal{C}_{\mathbf{x}}^F, \quad \text{and} \quad \mathcal{D}_0 := \mathcal{D} \cap \mathcal{B}_{\mathbf{x}, \mathbf{y}}^F.$$

Thus, $\mathcal{D}_{\mathbf{x}} \subset {}^*\mathcal{C}_{\mathbf{x}}^F$, and $\mathcal{D}_{\mathbf{y}} \subset {}^*\mathcal{C}_{\mathbf{y}}^F$, and $\mathcal{D} = \mathcal{D}_{\mathbf{x}} \sqcup \mathcal{D}_0 \sqcup \mathcal{D}_{\mathbf{y}}$. Also, $\mathcal{D}_{\mathbf{x}}$ and $\mathcal{D}_{\mathbf{y}}$ are nonempty, because Claim 2 implies that the line segment from $\mathbf{b}$ to $\mathbf{x}$ lies in $\mathcal{D}_{\mathbf{x}}$, while the line segment from $\mathbf{b}$ to $\mathbf{y}$ lies in $\mathcal{D}_{\mathbf{y}}$.

Claim 3A: The set $\mathcal{D} \setminus \mathcal{D}_0$ is path-disconnected.

Proof Note that $\mathcal{D} \setminus \mathcal{D}_0 = \mathcal{D}_{\mathbf{x}} \sqcup \mathcal{D}_{\mathbf{y}}$. Let $\mathbf{d}_{\mathbf{x}} \in \mathcal{D}_{\mathbf{x}}$ and $\mathbf{d}_{\mathbf{y}} \in \mathcal{D}_{\mathbf{y}}$. Let $\mathcal{P} \subset \mathcal{D}$ be any path from $\mathbf{d}_{\mathbf{x}}$ to $\mathbf{d}_{\mathbf{y}}$. Then $F(\mathbf{p}) \subseteq \{\mathbf{x}, \mathbf{y}\}$ for all $\mathbf{p} \in \mathcal{P}$. Continuity implies that $F(\mathbf{p}) = \{\mathbf{x}\}$ for all points $\mathbf{p} \in \mathcal{P}$ close to $\mathbf{d}_{\mathbf{x}}$ and $F(\mathbf{p}) = \{\mathbf{y}\}$ for all points $\mathbf{p} \in \mathcal{P}$ close to $\mathbf{d}_{\mathbf{y}}$. Thus, Continuity yields some $\mathbf{p}_0 \in \mathcal{P}$ such that $F(\mathbf{p}_0) = \{\mathbf{x}, \mathbf{y}\}$; thus $\mathbf{p}_0 \in \mathcal{D}_0$. Thus, $\mathbf{d}_{\mathbf{x}}$ and $\mathbf{d}_{\mathbf{y}}$ lie in different path components of $\mathcal{D} \setminus \mathcal{D}_0$. ∇ Claim 3A

Claim 3A implies that $\mathcal{D}_0$ meets the interior of $\mathcal{D}$, which means it meets ${}^\circ\mathcal{C}$ (because $\text{int}(\mathcal{D}) \subseteq {}^\circ\mathcal{C}$). But clearly $\mathcal{D}_0 \cap {}^\circ\mathcal{C} \subseteq {}^*\mathcal{B}_{\mathbf{x}, \mathbf{y}}^F$; thus, we deduce that ${}^*\mathcal{B}_{\mathbf{x}, \mathbf{y}}^F \neq \emptyset$, as claimed.

Furthermore, $\mathcal{D}_0$ is a convex subset of $\mathcal{D}$ of dimension at most $K - 1$, which cuts the ball $\mathcal{D}$ into at least two disconnected pieces (by Claim 3A). The only way this can happen is if $\mathcal{D}_0 = \mathcal{D} \cap \mathbb{H}$ for some affine hyperplane $\mathbb{H} \subset \mathbb{R}^{\mathcal{K}}$. Now, $\mathcal{B}_{\mathbf{x}, \mathbf{y}}^F$ is a convex subset of $\mathbb{R}^{\mathcal{K}}$, and we have just found an open ball $\mathcal{D} \subset \mathbb{R}^{\mathcal{K}}$ such that $\mathcal{B}_{\mathbf{x}, \mathbf{y}}^F \cap \mathcal{D} \subset \mathbb{H}$; thus, $\mathcal{B}_{\mathbf{x}, \mathbf{y}}^F \subset \mathbb{H}$. Thus, ${}^*\mathcal{B}_{\mathbf{x}, \mathbf{y}}^F \subset \mathbb{H}$.

Finally, for any point $\mathbf{b} \in {}^*\mathcal{B}_{\mathbf{x}, \mathbf{y}}^F$, we can repeat the above construction to obtain an $\mathbb{H}$ -relatively open neighbourhood $\mathcal{D}_0$ around $\mathbf{b}$ in ${}^*\mathcal{B}_{\mathbf{x}, \mathbf{y}}^F$; thus, ${}^*\mathcal{B}_{\mathbf{x}, \mathbf{y}}^F$ is a relatively open subset of $\mathbb{H}$. $\diamond$ Claim 3

Claim 4: (a) For any $\mathbf{x} \in \mathcal{X}$, $\mathcal{C}_{\mathbf{x}}^F$ is a closed, convex polyhedron in $\mathbb{R}^{\mathcal{K}}$, and

$$(\partial \mathcal{C}_{\mathbf{x}}^F) \cap \mathcal{C} = \bigcup_{\mathbf{y} \in \mathcal{X} \setminus \{\mathbf{x}\}} \mathcal{B}_{\mathbf{x}, \mathbf{y}}^F.$$

(b) For all $\mathbf{y} \in \mathcal{X}$, if ${}^*\mathcal{B}_{\mathbf{x}, \mathbf{y}}^F \neq \emptyset$, then it is contained a codimension-1 face of $(\partial \mathcal{C}_{\mathbf{x}}^F) \cap \mathcal{C}$.

(c) Conversely, for each each codimension-1 face $\mathcal{F}$ of $(\partial \mathcal{C}_{\mathbf{x}}^F) \cap \mathcal{C}$, there is some $\mathbf{y} \in \mathcal{X}$ such that $\emptyset \neq {}^*\mathcal{B}_{\mathbf{x}, \mathbf{y}}^F \subseteq \mathcal{F}$.

Proof (a) $\mathcal{C}_{\mathbf{x}}^F$ is closed by Continuity, and convex by Reinforcement. For any $\mathbf{y} \in \mathcal{X}$, $\mathcal{B}_{\mathbf{x}, \mathbf{y}}^F \subset \partial \mathcal{C}_{\mathbf{x}}^F$ by Claim 2. This proves that $(\partial \mathcal{C}_{\mathbf{x}}^F) \cap \mathcal{C} \supseteq \bigcup_{\mathbf{y} \in \mathcal{X} \setminus \{\mathbf{x}\}} \mathcal{B}_{\mathbf{x}, \mathbf{y}}^F$. To see the opposite inclusion, let $\mathbf{b} \in (\partial \mathcal{C}_{\mathbf{x}}^F) \cap \mathcal{C}$. Then $\mathbf{b}$ is a cluster point of $\mathcal{C}_{\mathbf{x}}^F$, but $\mathbf{b}$ is also a cluster point of $\mathcal{C}_{\mathbf{y}}^F$ for some $\mathbf{y} \neq \mathbf{x}$. Thus, $\mathbf{b} \in \mathcal{C}_{\mathbf{x}}^F$ and $\mathbf{b} \in \mathcal{C}_{\mathbf{y}}^F$ by Continuity, so $\mathbf{b} \in \mathcal{B}_{\mathbf{x}, \mathbf{y}}^F$. Thus, $(\partial \mathcal{C}_{\mathbf{x}}^F) \cap \mathcal{C} \subseteq \bigcup_{\mathbf{y} \in \mathcal{X} \setminus \{\mathbf{x}\}} \mathcal{B}_{\mathbf{x}, \mathbf{y}}^F$.

(b) If ${}^*\mathcal{B}_{\mathbf{x}, \mathbf{y}}^F \neq \emptyset$, then Claim 3 says it is a nonempty open subset of some hyperplane, hence of codimension 1, hence contained in some codimension-1 face of $\partial \mathcal{C}_{\mathbf{x}}^F$.

(c) Let $\mathcal{F}$ be a codimension-1 face of $(\partial \mathcal{C}_{\mathbf{x}}^F) \cap \mathcal{C}$; we claim that ${}^*\mathcal{B}_{\mathbf{x}, \mathbf{y}}^F \subseteq \mathcal{F}$ for some $\mathbf{y} \in \mathcal{X}$. To see this, first note that $\mathcal{F} \subseteq \bigcup_{\mathbf{y} \in \mathcal{X} \setminus \{\mathbf{x}\}} \mathcal{B}_{\mathbf{x}, \mathbf{y}}^F$, by part (a). Since this is a finite collection of sets, there must be some $\mathbf{y} \in \mathcal{X}$ such that $\mathcal{B}_{\mathbf{x}, \mathbf{y}}^F \cap \mathcal{F}$ is a subset of $\mathcal{F}$ with nonempty relative interior. Call this relative interior set $\mathcal{B}_0$.

We claim that $\mathcal{B}_0 \subset {}^*\mathcal{B}_{\mathbf{x}, \mathbf{y}}^F$. To see this, consider the set $\mathcal{D}_{\mathbf{y}} := \{s\mathbf{b} + (1-s)\mathbf{y}; \mathbf{b} \in \mathcal{B}_0 \text{ and } s \in (0, 1)\}$. This is an open cone in $\mathbb{R}^{\mathcal{K}}$ with base $\mathcal{B}_0$, and Reinforcement says that $\mathcal{D}_{\mathbf{y}} \subset {}^*\mathcal{C}_{\mathbf{y}}^F$. Likewise, if we define $\mathcal{D}_{\mathbf{x}} := \{s\mathbf{b} + (1-s)\mathbf{x}; \mathbf{b} \in \mathcal{B}_0 \text{ and } s \in (0, 1)\}$, then $\mathcal{D}_{\mathbf{x}}$ is an open cone with base $\mathcal{B}_0$, and Reinforcement says that $\mathcal{D}_{\mathbf{x}} \subset {}^*\mathcal{C}_{\mathbf{x}}^F$. Note that $\mathcal{D}_{\mathbf{y}}$ and $\mathcal{D}_{\mathbf{x}}$ are disjoint. Since both have $\mathcal{B}_0$ as their base, it follows that the set $\mathcal{O} := \mathcal{D}_{\mathbf{y}} \sqcup \mathcal{B}_0 \sqcup \mathcal{D}_{\mathbf{x}}$ is an open subset of $\mathbb{R}^{\mathcal{K}}$, containing $\mathcal{B}_0$.

Now let $\mathbf{b} \in \mathcal{B}_0$, and suppose by contradiction that $\mathbf{b} \notin {}^*\mathcal{B}_{\mathbf{x}, \mathbf{y}}^F$. Then $F(\mathbf{b}) \supseteq \{\mathbf{x}, \mathbf{y}, \mathbf{z}\}$ for some distinct $\mathbf{z} \in \mathcal{X} \setminus \{\mathbf{x}, \mathbf{y}\}$. Thus, if we define $\mathcal{L} := \{s\mathbf{b} + (1-s)\mathbf{z}; s \in (0, 1]\}$, then $\mathcal{L}$ is a line-segment with one end at $\mathbf{b}$, and Reinforcement says that $\mathcal{L} \subset {}^*\mathcal{C}_{\mathbf{z}}^F$. But $\mathcal{O}$ is an open neighbourhood of $\mathbf{b}$, so $\mathcal{L}$ must pass through $\mathcal{O}$ to reach $\mathbf{b}$. Thus, ${}^*\mathcal{C}_{\mathbf{z}}^F \cap \mathcal{O}$ is nonempty. But since it is an intersection of two open sets, ${}^*\mathcal{C}_{\mathbf{z}}^F \cap \mathcal{O}$ itself is open; thus, it must contain points in either $\mathcal{D}_{\mathbf{x}}$ or $\mathcal{D}_{\mathbf{y}}$. But this is impossible, because these are subsets of ${}^*\mathcal{C}_{\mathbf{x}}^F$ and ${}^*\mathcal{C}_{\mathbf{y}}^F$, which are disjoint from ${}^*\mathcal{C}_{\mathbf{z}}^F$ by definition. To avoid contradiction, we must have $\mathbf{b} \in {}^*\mathcal{B}_{\mathbf{x}, \mathbf{y}}^F$.

From this, it follows that $\mathcal{F} \cap {}^*\mathcal{B}_{\mathbf{x}, \mathbf{y}}^F \neq \emptyset$ (because it contains $\mathcal{B}_0$). However, part (b) tells us that ${}^*\mathcal{B}_{\mathbf{x}, \mathbf{y}}^F$ is entirely contained in some codimension-1 face of $\mathcal{C}$; thus, we must have ${}^*\mathcal{B}_{\mathbf{x}, \mathbf{y}}^F \subseteq \mathcal{F}$. ◇ Claim 4

The strategy of the proof is now as follows. Claim 4(a) tells us that the sets $\{\mathcal{C}_{\mathbf{x}}^F\}_{\mathbf{x} \in \mathcal{X}}$ partition $\mathcal{C}$ into closed, convex polyhedra, which overlap only on their boundaries. By a similar argument, the median rule also partitions $\mathcal{C}$ into closed, convex polyhedra $\{\mathcal{C}_{\mathbf{x}}^{\text{med}}\}_{\mathbf{x} \in \mathcal{X}}$, which overlap only on their boundaries. We will show that these two

partitions are identical. To do this, it suffices to show that the codimension-1 faces of $\mathcal{C}_{\mathbf{x}}^F$ and $\mathcal{C}_{\mathbf{x}}^{\text{med}}$ are the same, for each $\mathbf{x} \in \mathcal{X}$. Claim 4(b,c) tells us that the codimension-1 faces of $\mathcal{C}_{\mathbf{x}}^F$ can be identified with the “boundary” sets ${}^*\mathcal{B}_{\mathbf{x},\mathbf{y}}^F$ for $\mathbf{y} \in \mathcal{X}$; thus, it suffices to show that these boundary sets coincide with those of the median rule.

The boundary face ${}^o\mathcal{B}_{\mathbf{x},\mathbf{y}}^{\text{med}}$ lies in the hyperplane $\mathbb{H}_{\mathbf{x},\mathbf{y}} = \{\mathbf{r} \in \mathbb{R}^{\mathcal{K}}; \mathbf{r} \bullet_{\lambda} (\mathbf{x} - \mathbf{y}) = 0\}$. We will show that ${}^*\mathcal{B}_{\mathbf{x},\mathbf{y}}^F$ also lies in this hyperplane (see Claim 10 below). To do this, we will show that, if we start with a point in ${}^*\mathcal{B}_{\mathbf{x},\mathbf{y}}^F$ and perturb it slightly by a vector parallel to $\mathbb{H}_{\mathbf{x},\mathbf{y}}$, then it remains in ${}^*\mathcal{B}_{\mathbf{x},\mathbf{y}}^F$. But to perform such a perturbation analysis on points in ${}^*\mathcal{B}_{\mathbf{x},\mathbf{y}}^F$, we must perform a corresponding perturbation analysis on the gain function ϕ . To do this, we need some control over “ ϕ -increments” of the form $\phi(r + \delta) - \phi(r)$, where $r \in [-1, 1]$ and δ is a “small” perturbation. To achieve this, Claims 5 to 9 establish more and more precise control over these ϕ -increments. For all $i \in \mathcal{K}$, let $\mathbf{e}_i := (0, 0, \dots, 0, 1, 0, \dots, 0)$, where the 1 appears in the i th coordinate.

Claim 5: Let $\mathbf{x}, \mathbf{y} \in \mathcal{X}$, with $d(\mathbf{x}, \mathbf{y}) \geq 2$ and let $\mathbf{b} \in {}^*\mathcal{B}_{\mathbf{x},\mathbf{y}}^F$. Let $i, j \in \mathcal{K}(\mathbf{x}, \mathbf{y})$ with $i \neq j$.

(a) For any $\delta_i, \delta_j \in \mathbb{R}$,

$$\begin{aligned} (\mathbf{x} - \mathbf{y}) \bullet_{\lambda} \phi(\mathbf{b} + \delta_i \mathbf{e}_i + \delta_j \mathbf{e}_j) &= \lambda_i (x_i - y_i) (\phi(b_i + \delta_i) - \phi(b_i)) \\ &\quad + \lambda_j (x_j - y_j) (\phi(b_j + \delta_j) - \phi(b_j)). \end{aligned}$$

(b) Let $r := b_i$ and $s = b_j$. There exists $\epsilon_{rs} > 0$ and a unique constant $c_{rs} > 0$ (which is determined by $\mathbf{x}, \mathbf{y}, i$, and j) such that, for any $\epsilon \in (-\epsilon_{rs}, \epsilon_{rs})$, we have

$$\phi(r + \epsilon) - \phi(r) = \frac{\lambda_j}{\lambda_i} (\phi(s + c_{rs} \epsilon) - \phi(s)).$$

Proof (a) $(\mathbf{x} - \mathbf{y}) \bullet_{\lambda} \phi(\mathbf{b} + \delta_i \mathbf{e}_i + \delta_j \mathbf{e}_j)$

$$\begin{aligned} &\stackrel{(\diamond)}{=} \sum_{k \in \mathcal{K}(\mathbf{x}, \mathbf{y}) \setminus \{i, j\}} \lambda_k (x_k - y_k) \phi(b_k) + \lambda_i (x_i - y_i) \phi(b_i + \delta_i) + \lambda_j (x_j - y_j) \phi(b_j + \delta_j) \\ &= \sum_{k \in \mathcal{K}(\mathbf{x}, \mathbf{y})} \lambda_k (x_k - y_k) \phi(b_k) + \lambda_i (x_i - y_i) (\phi(b_i + \delta_i) - \phi(b_i)) \\ &\quad + \lambda_j (x_j - y_j) (\phi(b_j + \delta_j) - \phi(b_j)) \\ &\stackrel{(\diamond)}{=} (\mathbf{x} - \mathbf{y}) \bullet_{\lambda} \phi(\mathbf{b}) + \lambda_i (x_i - y_i) (\phi(b_i + \delta_i) - \phi(b_i)) \\ &\quad + \lambda_j (x_j - y_j) (\phi(b_j + \delta_j) - \phi(b_j)) \\ &\stackrel{(\dagger)}{=} \lambda_i (x_i - y_i) (\phi(b_i + \delta_i) - \phi(b_i)) + \lambda_j (x_j - y_j) (\phi(b_j + \delta_j) - \phi(b_j)). \end{aligned}$$

Here, both $(\diamond)$ are because $(\mathbf{x} - \mathbf{y})_k = (x_k - y_k) = 0$ for all $k \in \mathcal{K} \setminus \mathcal{K}(\mathbf{x}, \mathbf{y})$, and $(\dagger)$ is because $\mathbf{b} \in {}^*\mathcal{B}_{\mathbf{x},\mathbf{y}}^F$, so that $(\mathbf{x} - \mathbf{y}) \bullet_{\lambda} \phi(\mathbf{b}) = 0$.

(b) By negating the i coordinate and/or j coordinate of $\mathcal{X}$ and $\mathcal{Y}$ if necessary, we can assume without loss of generality that $x_i = y_j = 1$ and $x_j = y_i = -1$. Claim 3 yields some vector $\mathbf{v} \in \mathbb{R}^{\mathcal{K}}$ and some constant $a \in \mathbb{R}$ such that ${}^*\mathcal{B}_{\mathbf{x},\mathbf{y}}^F$ is a relatively open

subset of the affine hyperplane $\mathbb{H} := \{\mathbf{r} \in \mathbb{R}^{\mathcal{K}}; \mathbf{v} \bullet \mathbf{r} = a\}$. Let $c_{rs} := -v_i/v_j$. For any $\epsilon \in \mathbb{R}$, let $\mathbf{b}_\epsilon := \mathbf{b} + \epsilon \mathbf{e}_i + c_{rs} \epsilon \mathbf{e}_j$. Then

$$\mathbf{b}_\epsilon \bullet \mathbf{v} = \mathbf{b} \bullet \mathbf{v} + \epsilon \mathbf{e}_i \bullet \mathbf{v} - \frac{v_i}{v_j} \epsilon \mathbf{e}_j \bullet \mathbf{v} = a + \epsilon v_i - \frac{v_i}{v_j} \epsilon v_j = a. \quad (\text{C5})$$

Thus, $\mathbf{b}_\epsilon \in \mathbb{H}$. Thus, if $|\epsilon|$ is small enough, then $\mathbf{b}_\epsilon \in {}^*\mathcal{B}_{\mathbf{x},\mathbf{y}}^F$. Thus,

$$0 = (\mathbf{x} - \mathbf{y}) \bullet_\lambda \phi(\mathbf{b}_\epsilon) \stackrel{(*)}{=} 2 \left[\lambda_i (\phi(r + \epsilon) - \phi(r)) - \lambda_j (\phi(s + c_{rs} \epsilon) - \phi(s)) \right],$$

where $(*)$ is by setting $\delta_i := \epsilon$ and $\delta_j := c_{rs} \epsilon$ in part (a), and noting that $(x_i - y_i) = 2$ while $(x_j - y_j) = -2$. Thus, we conclude that $\phi(r + \epsilon) - \phi(r) = \frac{\lambda_j}{\lambda_i} (\phi(s + c_{rs} \epsilon) - \phi(s))$, as desired. Finally, observe that $c_{rs} > 0$ and is unique, because ϕ is strictly increasing. $\diamond$ Claim 5

Claim 6: Let $\mathbf{x}, \mathbf{y} \in \mathcal{X}$ be such that $d(\mathbf{x}, \mathbf{y}) \geq 3$ and ${}^*\mathcal{B}_{\mathbf{x},\mathbf{y}}^F \neq \emptyset$. Let $\mathbf{b} \in {}^*\mathcal{B}_{\mathbf{x},\mathbf{y}}^F$. Let $i, j \in \mathcal{K}(\mathbf{x}, \mathbf{y})$ with $i \neq j$. Let $r := b_i$ and $s = b_j$.

- (a) There is an open interval $\mathcal{T}_r \subseteq [-1, 1]$ containing r such that, for all $t \in \mathcal{T}_r$, we have $c_{ts} = c_{rs}$, where c_{rs} and c_{ts} are as in Claim 5(b).
- (b) For all $t \in \mathcal{T}_r$, there is an open interval $\mathcal{E}_{tr}$ containing 0 such that, for all $\epsilon \in \mathcal{E}_{tr}$, we have $\phi(r + \epsilon) - \phi(r) = \phi(t + \epsilon) - \phi(t)$.

Proof (a) Since $d(\mathbf{x}, \mathbf{y}) \geq 3$, there is a third coordinate $k \in \mathcal{K}(\mathbf{x}, \mathbf{y}) \setminus \{i, j\}$.

Claim 3 yields some vector $\mathbf{v} \in \mathbb{R}^{\mathcal{K}}$ and some constant $a \in \mathbb{R}$ such that ${}^*\mathcal{B}_{\mathbf{x},\mathbf{y}}^F$ is a relatively open subset of the affine hyperplane $\mathbb{H} := \{\mathbf{r} \in \mathbb{R}^{\mathcal{K}}; \mathbf{v} \bullet \mathbf{r} = a\}$. For any $\eta \in \mathbb{R}$, let $\mathbf{b}^\eta := \mathbf{b} + \eta \mathbf{e}_i - (v_i/v_k) \eta \mathbf{e}_k$. Then $\mathbf{b}^\eta \in \mathbb{H}$, by an argument identical to equation (C5). Thus, there is some $\bar{\eta} > 0$ such that $\mathbf{b}^\eta \in {}^*\mathcal{B}_{\mathbf{x},\mathbf{y}}^F$ for all $\eta \in (-\bar{\eta}, \bar{\eta})$. Let $\mathcal{T}_r := (r - \bar{\eta}, r + \bar{\eta})$. Let $t \in \mathcal{T}_r$, and let $\eta := t - r$. Then $\mathbf{b}^\eta \in {}^*\mathcal{B}_{\mathbf{x},\mathbf{y}}^F$, and by construction we have $b_i^\eta = b_i + \eta = r + \eta = t$, while $b_j^\eta = b_j = s$. Then, repeating the construction in Claim 5(b) using $\mathbf{b}^\eta$ in place of $\mathbf{b}$, we get $\phi(t + \epsilon) - \phi(t) = \frac{\lambda_j}{\lambda_i} (\phi(s + c_{ts} \epsilon) - \phi(s))$, for all sufficiently small ϵ , where $c_{ts} = -v_i/v_j$. This works for all $t \in \mathcal{T}_r$. But $r \in \mathcal{T}_r$; thus, in particular $c_{rs} = -v_i/v_j$. Thus, $c_{ts} = c_{rs}$ for all $t \in \mathcal{T}_r$, as claimed.

(b) Let $t \in \mathcal{T}_r$, and let $\epsilon_t := \min\{\epsilon_{rs}, \epsilon_{ts}\}$, where these are defined as in Claim 5(b). Then $\epsilon_t > 0$, and for all $\epsilon \in (-\epsilon_t, \epsilon_t)$, we have

$$\phi(r + \epsilon) - \phi(r) \stackrel{(*)}{=} \frac{\lambda_j}{\lambda_i} (\phi(s + c_{r,s} \epsilon) - \phi(s)) \stackrel{(\dagger)}{=} \phi(t + \epsilon) - \phi(t),$$

where $(*)$ is by Claim 5(b), and $(\dagger)$ is by Claim 5(b) and part (a). $\diamond$ Claim 6

Let $\bar{R} := \sup\{b_i; \mathbf{b} \in {}^*\mathcal{B}_{\mathbf{x},\mathbf{y}}^F \text{ and } i \in \mathcal{K}(\mathbf{x}, \mathbf{y}) \text{ for some } \mathbf{x}, \mathbf{y} \in \mathcal{X} \text{ with } d(\mathbf{x}, \mathbf{y}) \geq 3\}$. Note that $\bar{R} > 0$ if and only if $\mathcal{C}$ is suitable for F . Thus, Claims 7, 8 and 9 (below) are non-vacuous if $\mathcal{C}$ is suitable for F .

Claim 7: (Assuming compatibility) For all $r \in (0, \bar{R})$, there is an open interval $\mathcal{T}_r$ containing r , and for all $t \in \mathcal{T}_r$, there is an open interval $\mathcal{E}_{tr}$ containing 0, such that, for any $\epsilon \in \mathcal{E}_{tr}$, we have $\phi(t + \epsilon) - \phi(t) = \phi(r + \epsilon) - \phi(r)$.

Proof Since $r \leq \bar{R}$, there exist $\mathbf{x}, \mathbf{y} \in \mathcal{X}$ with $d(\mathbf{x}, \mathbf{y}) \geq 3$ and some $\mathbf{b} \in {}^*\mathcal{B}_{\mathbf{x}, \mathbf{y}}^F$ and $i \in \mathcal{K}(\mathbf{x}, \mathbf{y})$ such that $r \leq b_i \leq \bar{R}$. Let $s = r/b_i$; thus, $s \in [0, 1]$ and $r = s b_i$. Let $\mathbf{z} := \frac{1}{2}(\mathbf{x} + \mathbf{y})$, and then let $\mathbf{b}^s := s \mathbf{b} + (1-s)\mathbf{z}$. Then $F(\mathbf{b}^s) = F(\mathbf{b}) \cap F(\mathbf{z}) = \{\mathbf{x}, \mathbf{y}\}$, because F satisfies Reinforcement and $F(\mathbf{z}) \supseteq \{\mathbf{x}, \mathbf{y}\}$ by Claim 1. Thus, $\mathbf{b}^s \in {}^*\mathcal{B}_{\mathbf{x}, \mathbf{y}}^F$. Note that $b_i^s = s b_i = r$. Thus, Claim 6 yields some neighbourhood $\mathcal{T}_r$ around r , and for all $t \in \mathcal{T}_r$, an open interval $\mathcal{E}_{tr}$ containing 0, such that, for all $t \in \mathcal{T}_r$ and $\epsilon \in \mathcal{E}_{tr}$ we have $\phi(t + \epsilon) - \phi(t) = \phi(r + \epsilon) - \phi(r)$. $\diamond$ Claim 7

Claim 8: (Assuming compatibility) For all $r, s \in (0, \bar{R})$, there exists $\bar{\epsilon} = \bar{\epsilon}(r, s) > 0$ containing 0, such that, for any $\epsilon \in (-\bar{\epsilon}, \bar{\epsilon})$, we have $\phi(s + \epsilon) - \phi(s) = \phi(r + \epsilon) - \phi(r)$.

Proof Without loss of generality, suppose $r < s$. For all $q \in [r, s]$, let $\mathcal{T}_q$ be as in Claim 7. The family $\{\mathcal{T}_q\}_{q \in [r, s]}$ is an open cover of the compact set $[r, s]$, so it has a finite subcover, say $\{\mathcal{T}_{q_0}, \dots, \mathcal{T}_{q_N}\}$, where $r = q_0 \leq q_1 < q_2 < \dots < q_N \leq s = q_N$. By dropping to a subsequence of $\{q_0, q_1, \dots, q_N\}$ if necessary, we can ensure that $q_n \in \mathcal{T}_{q_{n-1}}$ for all $n \in [1 \dots N]$ (because these intervals cover $[r, s]$). For all $n \in [1 \dots N]$, let $\mathcal{E}_{q_n, q_{n-1}}$ be the open interval around 0 defined in Claim 7. Let $\mathcal{E} := \mathcal{E}_{q_1, q_0} \cap \mathcal{E}_{q_2, q_1} \cap \dots \cap \mathcal{E}_{q_N, q_{N-1}}$; then $\mathcal{E}$ is an open interval containing 0, so there is some $\bar{\epsilon} > 0$ such that $(-\bar{\epsilon}, \bar{\epsilon}) \subseteq \mathcal{E}$. For all $\epsilon \in (-\bar{\epsilon}, \bar{\epsilon})$, we have

$$\begin{aligned} \phi(q_N + \epsilon) - \phi(q_N) &= \phi(q_{N-1} + \epsilon) - \phi(q_{N-1}) = \dots = \phi(q_1 + \epsilon) - \phi(q_1) \\ &= \phi(q_0 + \epsilon) - \phi(q_0), \end{aligned}$$

where each equality is an invocation of Claim 7. In other words, $\phi(s + \epsilon) - \phi(s) = \phi(r + \epsilon) - \phi(r)$, as claimed. $\diamond$ Claim 8

Claim 9: (Assuming compatibility) For all $r \in (0, \bar{R})$, there is some $\delta = \delta(r) > 0$ such that for any $q \in \mathbb{Q} \cap [-1, 1]$, we have $\phi(r + q\delta) - \phi(r) = q \cdot [\phi(r + \delta) - \phi(r)]$.

Proof Find $\delta > 0$ such that $0 < r - \delta < r + \delta < \bar{R}$. Thus, if we define $\mathcal{S} := [r - \delta, r + \delta]$, then $\mathcal{S} \subset (0, \bar{R})$. Thus, for all $s \in \mathcal{S}$, Claim 8 yields some $\epsilon_s := \bar{\epsilon}(r, s) > 0$ such that $\phi(s + \epsilon) - \phi(s) = \phi(r + \epsilon) - \phi(r)$ for all $\epsilon \in (-\epsilon_s, \epsilon_s)$. It is easily verified that the function $\mathcal{S} \ni s \mapsto \epsilon_s \in \mathbb{R}_+$ is continuous. The interval $\mathcal{S}$ is compact. Thus, there exists some $\epsilon' > 0$ such that $\epsilon_s \geq \epsilon'$ for all $s \in \mathcal{S}$.

Now, let $M_0 \in \mathbb{N}$ be large enough that $\delta/M_0 < \epsilon'$. For any rational number $q \in [-1, 1]$, we can write $q = N/M$ for some $N \in [-M \dots M]$ and some $M \in \mathbb{N}$ with $M \geq M_0$. Thus, if we define $\epsilon := \delta/M$, then $\epsilon < \epsilon'$, and $q\delta = N\epsilon$, and we have

$$\begin{aligned} \phi(r + q\delta) - \phi(r) &= \phi(r + N\epsilon) - \phi(r) = \sum_{n=1}^N (\phi(r + n\epsilon) - \phi(r + (n-1)\epsilon)) \\ &\stackrel{(*)}{=} \sum_{n=1}^N (\phi(r + \epsilon) - \phi(r)) = N \cdot (\phi(r + \epsilon) - \phi(r)) \end{aligned}$$

$$= N \cdot \left(\phi \left(r + \frac{1}{M} \delta \right) - \phi(r) \right), \quad (\text{C6})$$

where $(*)$ is by N applications of Claim 8 (which applies because for all $n \in [1 \dots N]$, we have $s := r + (n - 1) \epsilon \in \mathcal{S}$ and $\epsilon < \epsilon' \leq \epsilon_s$). In particular, if $q = 1$ (so that $N = M$), then (C6) yields

$$\phi(r + \delta) - \phi(r) = M \cdot \left(\phi \left(r + \frac{\delta}{M} \right) - \phi(r) \right),$$

which means that

$$\phi \left(r + \frac{1}{M} \delta \right) - \phi(r) = \frac{1}{M} (\phi(r + \delta) - \phi(r)). \quad (\text{C7})$$

For any $q = N/M \in [-1, 1]$, if we substitute (C7) into (C6), we get

$$\phi(r + q \delta) - \phi(r) = \frac{N}{M} \cdot (\phi(r + \delta) - \phi(r)) = q \cdot (\phi(r + \delta) - \phi(r)),$$

as desired. $\diamond$ Claim 9

Claim 10: Let $\mathbf{x}, \mathbf{y} \in \mathcal{X}$ be distinct, and let $\mathbb{H}_{\mathbf{x}, \mathbf{y}} := \{\mathbf{r} \in \mathbb{R}^{\mathcal{K}}; (\mathbf{x} - \mathbf{y}) \bullet_{\lambda} \mathbf{r} = 0\}$. If ${}^*\mathcal{B}_{\mathbf{x}, \mathbf{y}}^F \neq \emptyset$, then ${}^*\mathcal{B}_{\mathbf{x}, \mathbf{y}}^F \subset \mathbb{H}_{\mathbf{x}, \mathbf{y}}$.

Proof First suppose $d(\mathbf{x}, \mathbf{y}) = 1$, and let $\mathcal{K}(\mathbf{x}, \mathbf{y}) = \{i\}$. Then $\mathbb{H}_{\mathbf{x}, \mathbf{y}} = \{\mathbf{r} \in \mathbb{R}^{\mathcal{K}}; r_i = 0\}$. By negating the i th coordinate of $\mathcal{X}$ and $\mathcal{Y}$ if necessary, we can assume without loss of generality that $x_i = 1$ and $y_i = -1$. Then for any $\mu \in \Delta(\mathcal{Y})$, if $\tilde{\mu}_i > 0$ then $\mathbf{y} \notin \text{SME}(\mathcal{C}, \mu)$ (because $\gamma_{\mathbf{x}, \mu}(q) \geq \gamma_{\mathbf{y}, \mu}(q)$ for all $q \in [0, 1]$, and $\gamma_{\mathbf{x}, \mu}(q) \geq \gamma_{\mathbf{y}, \mu}(q) + \lambda_i$ for all $q \in [0, \tilde{\mu}_i]$.) Likewise, if $\tilde{\mu}_i < 0$ then $\mathbf{x} \notin \text{SME}(\mathcal{C}, \mu)$. Thus, if $\{\mathbf{x}, \mathbf{y}\} \subseteq \text{SME}(\mathcal{C}, \mu)$, then we must have $\tilde{\mu}_i = 0$. Since F is SME, it follows that $\mathcal{B}_{\mathbf{x}, \mathbf{y}}^F \subset \mathbb{H}_{\mathbf{x}, \mathbf{y}}$.

Now suppose $d(\mathbf{x}, \mathbf{y}) \geq 2$. There are now two cases: Either $\mathcal{C}$ is suitable for F , or it is *not* suitable for F , but has balanced weights.

Case A. Suppose $\mathcal{C}$ is not suitable for F , but has balanced weights. If $\mathbf{x}, \mathbf{y} \in \mathcal{X}$, and ${}^*\mathcal{B}_{\mathbf{x}, \mathbf{y}}^F \neq \emptyset$, then we must have $d(\mathbf{x}, \mathbf{y}) \leq 2$ (because $\mathcal{C}$ is not suitable for F). We have already dealt with the case $d(\mathbf{x}, \mathbf{y}) = 1$, so suppose $d(\mathbf{x}, \mathbf{y}) = 2$. Let $\mathcal{K}(\mathbf{x}, \mathbf{y}) = \{i, j\}$. By balanced weights, we must have $\lambda_i = \lambda_j$. For simplicity, suppose $\lambda_i = \lambda_j = 1$. By negating the i th and j th coordinate of $\mathcal{X}$ and $\mathcal{Y}$ if necessary, we can assume without loss of generality that $x_i = y_j = 1$ and $y_i = x_j = -1$. Thus, $\mathbb{H}_{\mathbf{x}, \mathbf{y}} = \{\mathbf{r} \in \mathbb{R}^{\mathcal{K}}; r_i = r_j\}$. Now, for any $\mu \in \Delta(\mathcal{Y})$, if $\tilde{\mu}_i > \tilde{\mu}_j$ then $\mathbf{y} \notin \text{SME}(\mathcal{C}, \mu)$ (because $\gamma_{\mathbf{x}, \mu}(q) \geq \gamma_{\mathbf{y}, \mu}(q)$ for all $q \in [0, 1]$, and $\gamma_{\mathbf{x}, \mu}(q) \geq \gamma_{\mathbf{y}, \mu}(q) + 1$ for all $q \in (\tilde{\mu}_j, \tilde{\mu}_i]$.) Likewise, if $\tilde{\mu}_i < \tilde{\mu}_j$ then $\mathbf{x} \notin \text{SME}(\mathcal{C}, \mu)$. Thus, if $\{\mathbf{x}, \mathbf{y}\} \subseteq \text{SME}(\mathcal{C}, \mu)$, then we must have $\tilde{\mu}_i = \tilde{\mu}_j$. Since F is SME, it follows that $\mathcal{B}_{\mathbf{x}, \mathbf{y}}^F \subset \mathbb{H}_{\mathbf{x}, \mathbf{y}}$.

Case B. Suppose $\mathcal{C}$ is suitable for F . Then $\bar{R} > 0$. Let $\mathbf{x}, \mathbf{y} \in \mathcal{X}$, and suppose $d(\mathbf{x}, \mathbf{y}) \geq 2$. Let $\mathcal{J} := \mathcal{K}(\mathbf{x}, \mathbf{y})$ and let $\mathcal{L} := \mathcal{K} \setminus \mathcal{J}$. By negating the i and/or j

coordinates of $\mathcal{X}$ and $\mathcal{Y}$ if necessary, we can assume without loss of generality that $x_j = 1$ and $y_j = -1$ for all $j \in \mathcal{J}$. Let $\mathbf{z} := \frac{1}{2}(\mathbf{x} + \mathbf{y})$. Let $\mathbf{b} \in {}^*\mathcal{B}_{\mathbf{x},\mathbf{y}}^F$. For all $u \in [0, 1]$, let $\mathbf{b}^u := u\mathbf{b} + (1-u)\mathbf{z}$. Then $\mathbf{b}^u \in {}^*\mathcal{B}_{\mathbf{x},\mathbf{y}}^F$ by Claim 1 and Reinforcement, as explained above. If u is small enough, then we have $b_i^u \in (0, \bar{R})$ for all $i \in \mathcal{K}(\mathbf{x}, \mathbf{y})$.

Claim 10A: For any distinct $i, j \in \mathcal{K}(\mathbf{x}, \mathbf{y})$, there exists $\epsilon > 0$ such that

$$\mathbf{b}^u + \epsilon(\lambda_j \mathbf{e}_i - \lambda_i \mathbf{e}_j) \in {}^*\mathcal{B}_{\mathbf{x},\mathbf{y}}^F.$$

Proof Without loss of generality, suppose $\lambda_i \geq \lambda_j$. For any $\delta > 0$ and $\theta \in [0, 1]$, let $\mathbf{b}_\delta^\theta := \mathbf{b}^u + \theta\delta \mathbf{e}_i - \delta \mathbf{e}_j$. Since $F(\mathbf{b}^u) = \{\mathbf{x}, \mathbf{y}\}$ and F satisfies Continuity, there exists some $\bar{\epsilon} > 0$ such that, for all $\epsilon_1, \epsilon_2 \in (-\bar{\epsilon}, \bar{\epsilon})$, we have $F(\mathbf{b}^u + \epsilon_1 \mathbf{e}_i + \epsilon_2 \mathbf{e}_j) \subseteq \{\mathbf{x}, \mathbf{y}\}$. In particular, for any $\delta \in (0, \bar{\epsilon})$, and $\theta \in [0, 1]$, we have $F(\mathbf{b}_\delta^\theta) \subseteq \{\mathbf{x}, \mathbf{y}\}$. Without loss of generality, suppose $\bar{\epsilon} \leq \bar{\epsilon}(b_i^u, b_j^u)$, where $\bar{\epsilon}(b_i^u, b_j^u)$ is as in Claim 8. Let $\delta(b_i^u)$ be defined as in Claim 9, and find some small enough $q \in \mathbb{Q} \cap [0, 1]$ such that, if $\delta := q\delta(b_i^u)$, then $\delta \in (0, \bar{\epsilon})$.

Let $\theta := \frac{\lambda_j}{\lambda_i}$. Then $\theta \in [0, 1]$. We will show that $\mathbf{b}_\delta^\theta \in {}^*\mathcal{B}_{\mathbf{x},\mathbf{y}}^F$. If we knew that ϕ was linear, then we could deduce that $\phi(b_i^u + \theta\delta) - \phi(b_i^u) = \theta \cdot [\phi(b_i^u + \delta) - \phi(b_i^u)]$, and from here, use Claims 8 and 5(a) to obtain $(\mathbf{x} - \mathbf{y}) \bullet_\lambda \phi(\mathbf{b}_\delta^\theta) = 0$ and hence $F(\mathbf{b}_\delta^\theta) = \{\mathbf{x}, \mathbf{y}\}$. But we *don't* know that ϕ is linear; instead, Claim 9 only tells us that ϕ is “locally $\mathbb{Q}$ -linear”. Thus, we must approximate θ with rational numbers.

Let $\{q_n\}_{n=1}^\infty$ be a decreasing sequence in $\mathbb{Q} \cap [0, 1]$ with $\lim_{n \rightarrow \infty} q_n = \theta$. For all $n \in \mathbb{N}$, we have

$$\begin{aligned} \phi(b_i^u + q_n\delta) - \phi(b_i^u) &\stackrel{(*)}{=} q_n [\phi(b_i^u + \delta) - \phi(b_i^u)] \stackrel{(\dagger)}{=} q_n [\phi(b_j^u) - \phi(b_j^u - \delta)] \\ &\stackrel{(\diamond)}{\geq} \frac{\lambda_j}{\lambda_i} [\phi(b_j^u) - \phi(b_j^u - \delta)], \quad \text{and thus,} \\ \lambda_i [\phi(b_i^u + q_n\delta) - \phi(b_i^u)] &> \lambda_j [\phi(b_j^u) - \phi(b_j^u - \delta)], \end{aligned} \quad (\text{C8})$$

where $(*)$ is by setting $r := b_i^u$ in Claim 9, $(\dagger)$ is by setting $r := b_i^u, s := b_j^u - \delta$, and $\epsilon := \delta$ in Claim 8, and $(\diamond)$ is because $q_n > \theta = \frac{\lambda_j}{\lambda_i}$.

Since $(x_i - y_i) = (x_j - y_j) = 2$, the inequality (C8) yields

$$\lambda_i (x_i - y_i) [\phi(b_i^u + q_n\delta) - \phi(b_i^u)] + \lambda_j (x_j - y_j) [\phi(b_j^u - \delta) - \phi(b_j^u)] > 0. \quad (\text{C9})$$

Thus, Claim 5(a) yields $(\mathbf{x} - \mathbf{y}) \bullet_\lambda \phi(\mathbf{b}_\delta^{q_n}) > 0$. Thus, $\mathbf{y} \notin F(\mathbf{b}_\delta^{q_n})$. But we have already noted that $\emptyset \neq F(\mathbf{b}_\delta^{q_n}) \subseteq \{\mathbf{x}, \mathbf{y}\}$. Thus, we must have $F(\mathbf{b}_\delta^{q_n}) = \{\mathbf{x}\}$. But $\lim_{n \rightarrow \infty} q_n = \theta$, so $\lim_{n \rightarrow \infty} \mathbf{b}_\delta^{q_n} = \mathbf{b}_\delta^\theta$. Thus, Continuity implies that $\mathbf{x} \in F(\mathbf{b}_\delta^\theta)$ also.

Now let $\{q_n\}_{n=1}^\infty$ be an increasing sequence in $\mathbb{Q} \cap [0, 1]$ with $\lim_{n \rightarrow \infty} q_n = \theta$. For all $n \in \mathbb{N}$, we obtain $\lambda_i (x_i - y_i) [\phi(b_i^u + q_n \delta) - \phi(b_i^u)] + \lambda_j (x_j - y_j) [\phi(b_j^u - \delta) - \phi(b_j^u)] < 0$, by an argument similar to inequality (C9). Thus, Claim 5(a) yields $(\mathbf{x} - \mathbf{y}) \bullet_\lambda \phi(\mathbf{b}_\delta^{q_n}) < 0$. Thus, by an argument similar to the previous paragraph, we get $F(\mathbf{b}_\delta^{q_n}) = \{\mathbf{y}\}$, for all $n \in \mathbb{N}$. But $\lim_{n \rightarrow \infty} q_n = \theta$, so $\lim_{n \rightarrow \infty} \mathbf{b}_\delta^{q_n} = \mathbf{b}_\delta^\theta$. Thus, Continuity implies that $\mathbf{y} \in F(\mathbf{b}_\delta^\theta)$ also.

Combining these observations, we deduce that $\{\mathbf{x}, \mathbf{y}\} \subseteq F(\mathbf{b}_\delta^\theta)$. But we have already noted that $F(\mathbf{b}_\delta^\theta) \subseteq \{\mathbf{x}, \mathbf{y}\}$. Thus, $F(\mathbf{b}_\delta^\theta) = \{\mathbf{x}, \mathbf{y}\}$ —that is, $\mathbf{b}_\delta^\theta \in {}^*\mathcal{B}_{\mathbf{x}, \mathbf{y}}^F$. Now define $\epsilon := \delta/\lambda_i$; then $\mathbf{b}_\delta^\theta = \mathbf{b}^u + \epsilon (\lambda_j \mathbf{e}_i - \lambda_i \mathbf{e}_j)$, which proves the claim. ∇ Claim 10A

Claim 10B: Let $\ell \in \mathcal{L}$. If $\epsilon > 0$ is small enough, then $\mathbf{b}^u + \epsilon \mathbf{e}_\ell \in {}^*\mathcal{B}_{\mathbf{x}, \mathbf{y}}^F$.

Proof Let $\epsilon > 0$. If ϵ is small enough, then Continuity implies that $\emptyset \neq F(\mathbf{b}^u + \epsilon \mathbf{e}_\ell) \subseteq \{\mathbf{x}, \mathbf{y}\}$ (because $F(\mathbf{b}^u) = \{\mathbf{x}, \mathbf{y}\}$). But $(\mathbf{x} - \mathbf{y}) \bullet_\lambda \phi(\mathbf{b}^u + \epsilon \mathbf{e}_\ell) = (\mathbf{x} - \mathbf{y}) \bullet_\lambda \phi(\mathbf{b}^u) = 0$ (because $\phi(\mathbf{b}^u + \epsilon \mathbf{e}_\ell)_j = \phi(b_j)$ for all $j \in \mathcal{J}$, while $(\mathbf{x} - \mathbf{y})_k = 0$ for all $k \in \mathcal{L}$). Thus, we must have $F(\mathbf{b}^u + \epsilon \mathbf{e}_\ell) = \{\mathbf{x}, \mathbf{y}\}$; hence $\mathbf{b}^u + \epsilon \mathbf{e}_\ell \in {}^*\mathcal{B}_{\mathbf{x}, \mathbf{y}}^F$. ∇ Claim 10B

Claim 3 says that ${}^*\mathcal{B}_{\mathbf{x}, \mathbf{y}}^F$ is a relatively open subset of some affine hyperplane $\mathbb{H}$. Claims 10A and 10B imply that $\mathbb{H}$ is parallel to all vectors in the set $\{(\lambda_j \mathbf{e}_i - \lambda_i \mathbf{e}_j); i, j \in \mathcal{J}\} \cup \{\mathbf{e}_\ell; \ell \in \mathcal{L}\}$. But the hyperplane $\mathbb{H}_{\mathbf{x}, \mathbf{y}}$ is spanned by this set. Thus, $\mathbb{H}$ is parallel to $\mathbb{H}_{\mathbf{x}, \mathbf{y}}$. Let $\mathbf{z} = (\mathbf{x} + \mathbf{y})/2$. Then $\mathbf{z} \in \mathbb{H}_{\mathbf{x}, \mathbf{y}}$. But Claim 1 and Reinforcement imply that $\mathbf{z}$ is a cluster point of ${}^*\mathcal{B}_{\mathbf{x}, \mathbf{y}}^F$. Thus, $\mathbb{H} = \mathbb{H}_{\mathbf{x}, \mathbf{y}}$. $\diamond$ Claim 10

For any $\mathbf{x} \in \mathcal{X}$, let $\mathcal{C}_{\mathbf{x}}^{\text{med}} := \{\mathbf{c} \in \mathcal{C}; \mathbf{x} \in \text{Median}(\mathcal{X}, \mathbf{c})\}$. Then $\mathcal{C}_{\mathbf{x}}^{\text{med}}$ is a convex polyhedron whose supporting hyperplanes are the sets $\mathbb{H}_{\mathbf{x}, \mathbf{y}}$ (from Claim 10) for all $\mathbf{y} \in \mathcal{X} \setminus \{\mathbf{x}\}$, along with the supporting hyperplanes of $\mathcal{C}$ itself. Claims 4 and 10 show that every one of these supporting hyperplanes is also a supporting hyperplane of the convex polyhedron $\mathcal{C}_{\mathbf{x}}^F$. Thus, $\mathcal{C}_{\mathbf{x}}^F \subseteq \mathcal{C}_{\mathbf{x}}^{\text{med}}$ for all $\mathbf{x} \in \mathcal{X}$. However, the systems $\{\mathcal{C}_{\mathbf{x}}^F\}_{\mathbf{x} \in \mathcal{X}}$ and $\{\mathcal{C}_{\mathbf{x}}^{\text{med}}\}_{\mathbf{x} \in \mathcal{X}}$ are each partitions of $\mathcal{C}$ into closed convex polyhedra which meet only along their boundaries. Thus, they must be identical. Thus, F is the median rule. $\square$

Proof of Theorem 1 It is easy to verify that the median rule (2) satisfies ESME, Continuity, and Reinforcement; we must verify the converse. So, suppose F is a rule satisfying these axioms. Proposition 1 says that F is an additive majority rule, because it satisfies ESME and Continuity. Any unweighted judgement context obviously has balanced weights. Thus, if F also satisfies Reinforcement, then Theorem 2 says it is the median rule. $\square$

Theorem 3 is a consequence of a more general result, involving a more complicated structural condition. Let F be a judgement aggregation rule on the context $\mathcal{C} = (\mathcal{K}, \lambda, \mathcal{X}, \mathcal{Y})$, and let $\Theta \subset \Delta(\mathcal{Y})$ be a collection of profiles. Let $\mathcal{V}_\Theta^F := \{\mathbf{x} \in \mathcal{X}; \exists \mu \in \Theta \text{ such that } F(\mu) = \{\mathbf{x}\}\}$. For any $\mathbf{v}, \mathbf{w} \in \mathcal{V}_\Theta^F$, we write $\mathbf{v} \stackrel{F}{\sim} \mathbf{w}$ if $d(\mathbf{v}, \mathbf{w}) \leq 2$ and there is some $\mu \in \Theta$ with $F(\mu) = \{\mathbf{v}, \mathbf{w}\}$. Thus, $(\mathcal{V}_\Theta^F, \stackrel{F}{\sim})$ is a graph. This graph is *path-connected* if, for any $\mathbf{u}, \mathbf{w} \in \mathcal{V}_\Theta^F$, there is some path $\mathbf{u} \stackrel{F}{\sim} \mathbf{v}_1 \stackrel{F}{\sim} \mathbf{v}_2 \stackrel{F}{\sim} \dots \stackrel{F}{\sim} \mathbf{v}_N \stackrel{F}{\sim} \mathbf{w}$ connecting them in $\mathcal{V}_\Theta^F$. We will say that the judgement context $\mathcal{C}$ is *frangible* if for any additive majority rule F satisfying Continuity,

there exists an open, connected subset $\Theta \subset \Delta(\mathcal{Y})$ such that the graph $(\mathcal{V}_\Theta^F, \overset{F}{\sim})$ is not path-connected. Theorem 3 is a corollary of the next result.

Theorem C.1 *Let $\mathcal{C} = (\mathcal{K}, \lambda, \mathcal{X}, \mathcal{Y})$ be a frangible weighted judgement context such that $\mathcal{Y}$ is thick, and let $F : \Delta(\mathcal{Y}) \rightrightarrows \mathcal{X}$ be a judgement aggregation rule. Then F satisfies WESME, Continuity, and Reinforcement if and only if F is the median rule (8).*

Proof Suppose $\mathcal{C}$ is frangible and $\mathcal{Y}$ is thick, and F is an additive majority rule satisfying Continuity and Reinforcement. It suffices to show that $\mathcal{C}$ is suitable for F .

Since $\mathcal{C}$ is frangible, there is some open, connected $\Theta \subseteq \Delta(\mathcal{Y})$ and some $\mathbf{v}, \mathbf{w} \in \mathcal{V}_\Theta^F$ in different $\overset{F}{\sim}$ -connected components. Let $\tilde{\Theta} := \{\tilde{\mu}; \mu \in \Theta\}$; this is an open, connected subset of $\mathbb{R}^{\mathcal{K}}$ (because the function $\Delta(\mathcal{Y}) \ni \mu \mapsto \tilde{\mu} \in \mathbb{R}^{\mathcal{K}}$ is open and continuous). Let $\mathcal{C} = \text{conv}(\mathcal{Y})$; then $\tilde{\Theta} \subseteq \mathcal{C}$. As explained after Claim 4 in the proof of Theorem 2, the sets $\{\mathcal{C}_x^F\}_{x \in \mathcal{X}}$ partition $\mathcal{C}$ into closed, convex polyhedra, which overlap only on their boundaries. Now, $\tilde{\Theta}$ intersects $^*\mathcal{C}_v^F$ and $^*\mathcal{C}_w^F$ (because $\mathbf{v}, \mathbf{w} \in \mathcal{V}_\Theta^F$). Thus, it is possible to construct a continuous path $\alpha : [0, 1] \rightarrow \tilde{\Theta}$ with $\alpha(0) \in ^*\mathcal{C}_v^F$ and $\alpha(1) \in ^*\mathcal{C}_w^F$, such that for all $\mathbf{x}, \mathbf{y} \in \mathcal{X}$, if α crosses from $\mathcal{C}_x^F$ to $\mathcal{C}_y^F$, then it does so by passing through the codimension-1 face between $\mathcal{C}_x^F$ and $\mathcal{C}_y^F$ —call this face $\mathcal{F}_{\mathbf{x}, \mathbf{y}}$.

In fact, we will now show that we can assume without loss of generality that for all $\mathbf{x}, \mathbf{y} \in \mathcal{X}$, if α crosses from $\mathcal{C}_x^F$ to $\mathcal{C}_y^F$, then it does so by passing through the set $^*\mathcal{B}_{\mathbf{x}, \mathbf{y}}^F$. To see this, let $t \in [0, 1]$, and suppose $\alpha(t) \in \mathcal{F}_{\mathbf{x}, \mathbf{y}}$ for some $\mathbf{x}, \mathbf{y} \in \mathcal{X}$. Claim 4(b,c) in the proof of Theorem 2 says that $^*\mathcal{B}_{\mathbf{x}, \mathbf{y}}^F$ is nonempty, and is a subset of $\mathcal{F}_{\mathbf{x}, \mathbf{y}}$. So, let $\mathbf{b} \in ^*\mathcal{B}_{\mathbf{x}, \mathbf{y}}^F$, let $\epsilon > 0$, and replace α with the path α' defined:

$$\alpha'(s) := \begin{cases} \alpha(s) & \text{if } s \notin (t - \epsilon, t + \epsilon); \\ (\epsilon - |s - t|)\mathbf{b} + (1 - \epsilon + |s - t|)\alpha(s) & \text{if } s \in (t - \epsilon, t + \epsilon). \end{cases}$$

Recall that $\tilde{\Theta}$ is open; thus, if ϵ is small enough, then $\alpha'(s) \in \tilde{\Theta}$ for all $s \in [0, 1]$; furthermore, α' passes through exactly the same polyhedral cells as α , and its passage through all other faces is unchanged. However, $\alpha'(t) = \epsilon\mathbf{b} + (1 - \epsilon)\alpha(t)$. Thus, Reinforcement yields $F[\alpha'(t)] = F(\mathbf{b}) \cap F[\alpha(t)] = \{\mathbf{x}, \mathbf{y}\}$, so that $\alpha'(t) \in ^*\mathcal{B}_{\mathbf{x}, \mathbf{y}}^F$, as desired.

Thus, we can construct α such that for all $t \in [0, 1]$, either $\alpha(t) \in ^*\mathcal{C}_x^F$ for some $\mathbf{x} \in \mathcal{X}$ (so that $F[\alpha(t)] = \{\mathbf{x}\}$) or $\alpha(t) \in ^*\mathcal{B}_{\mathbf{x}, \mathbf{y}}^F$ for some $\mathbf{x}, \mathbf{y} \in \mathcal{X}$ (so that $F[\alpha(t)] = \{\mathbf{x}, \mathbf{y}\}$).

Now, suppose that $\mathcal{C}$ is not suitable for F . Then whenever $\alpha(t) \in ^*\mathcal{B}_{\mathbf{x}, \mathbf{y}}^F$ for some $\mathbf{x}, \mathbf{y} \in \mathcal{X}$, we must have $d(\mathbf{x}, \mathbf{y}) = 2$, and hence, $\mathbf{x} \overset{F}{\sim} \mathbf{y}$. Thus, the function $F \circ \alpha$ defines a path from $\mathbf{v}$ to $\mathbf{w}$ in the graph $(\mathcal{V}_\Theta^F, \overset{F}{\sim})$. But $\mathbf{v}$ and $\mathbf{w}$ are in different connected components of $(\mathcal{V}_\Theta^F, \overset{F}{\sim})$. Contradiction.

To avoid this contradiction, $\mathcal{C}$ must be suitable for F . Then Theorem 2 says that F is the median rule (8). $\square$

Case (a) of Theorem 3 follows from Theorem C.1 and the next result.

Lemma C.4 Let $\mathcal{C} = (\mathcal{K}, \lambda, \mathcal{X}, \mathcal{Y})$. If $\mathcal{X}$ is rugged and $\mathcal{Y}$ is McGarvey, then $\mathcal{C}$ is frangible.

Proof Let $\mathbf{z} \in \{\pm 1\}^{\mathcal{K} \setminus \mathcal{X}}$ and $\mathbf{x}, \mathbf{y} \in \mathcal{X}_{\mathbf{z}}$ be such that $d(\mathbf{x}, \mathbf{z}) \neq d(\mathbf{y}, \mathbf{z})$. (By ruggedness, such a $\mathbf{z}$ exists). Let $\Theta := \{\mu \in \Delta(\mathcal{Y}); \mathbf{x}^\mu = \mathbf{z}\} = \{\mu \in \Delta(\mathcal{Y}); \text{sign}(\tilde{\mu}_k) = z_k, \text{ for all } k \in \mathcal{K}\}$.

Claim 1: Θ is nonempty, open and convex (hence, connected).

Proof Θ is nonempty because $\mathcal{Y}$ is McGarvey. It is defined by a finite system of strict linear inequalities, so it is open and convex. $\diamond$ Claim 1

Now let F be any additive majority rule satisfying Continuity. We will show that $(\mathcal{V}_\Theta^F, \overset{F}{\sim})$ is disconnected. First, we need some terminology. A view $\mathbf{x} \in \mathcal{X}$ is *Condorcet admissible* for μ if there does not exist any other $\mathbf{y} \in \mathcal{X}$ such that $y_k \tilde{\mu}_k \geq x_k \tilde{\mu}_k$ for all $k \in \mathcal{K}$ —in other words, there is no view $\mathbf{y} \in \mathcal{X}$ which agrees with the majority in a strictly larger set of issues than those where $\mathbf{x}$ agrees with the majority. Let $\text{Cond}(\mathcal{X}, \mu) \subseteq \mathcal{X}$ be the set of all views that are Condorcet admissible for μ .⁴² It is easily verified that supermajority efficiency (for any λ) implies Condorcet admissibility. Thus, $\text{SME}(\mathcal{X}, \lambda, \mu) \subseteq \text{Cond}(\mathcal{X}, \mu)$.

Claim 2: $\mathcal{V}_\Theta^F = \mathcal{X}_{\mathbf{z}}$.

Proof “ $\subseteq$ ” For any $\mu \in \Theta$, we must have $F(\mu) \subseteq \text{SME}(\mathcal{X}, \mu) \subseteq \text{Cond}(\mathcal{X}, \mu)$. But $\text{Cond}(\mathcal{X}, \mu) \subseteq \mathcal{X}_{\mathbf{z}}$ by Lemma 1.5 of Nehring et al. (2014).

“ $\supseteq$ ” Let $\mathbf{x} \in \mathcal{X}_{\mathbf{z}}$. Let $\mathcal{J} := \{j \in \mathcal{K}; x_j = z_j\}$. By negating certain coordinates of $\mathcal{X}$ and $\mathcal{Y}$ if necessary, we can assume without loss of generality that $z_j = x_j = 1$ for all $j \in \mathcal{J}$. Since $\mathcal{Y}$ is McGarvey, there exists some $\mu \in \Delta(\mathcal{Y})$ such that $\tilde{\mu}_j > 0$ for all $j \in \mathcal{J}$ and $\tilde{\mu}_k = 0$ for all $k \in \mathcal{K} \setminus \mathcal{J}$ (by Footnote 36). Since $\mathbf{z}$ is near to $\mathbf{x}$, there is no other $\mathbf{y} \in \mathcal{X} \setminus \{\mathbf{x}\}$ such that $y_j = 1$ for all $j \in \mathcal{J}$. Thus, $\text{SME}(\mathcal{X}, \mu) = \{\mathbf{x}\}$. Thus, $F(\mu) = \{\mathbf{x}\}$, because F is supermajority efficient by Lemma B.3.

Let $\theta \in \Theta$ be arbitrary. For all $s \in [0, 1]$, define $\mu^s := s\theta + (1-s)\mu$. Thus, $\lim_{s \rightarrow 0} \mu^s = \mu$. Thus, Continuity implies that $F(\mu^s) = \{\mathbf{x}\}$ for all s sufficiently close to zero. But for all $s > 0$, we have $\mu^s \in \Theta$ (because for all $k \in \mathcal{K}$ we have $\text{sign}(\tilde{\mu}_k^s) = \text{sign}(s\tilde{\theta}_k + (1-s)\tilde{\mu}_k) = \text{sign}(\tilde{\theta}_k) = z_k$). Thus, $\mathbf{x} \in \mathcal{V}_\Theta^F$. This argument works for all $\mathbf{x} \in \mathcal{X}_{\mathbf{z}}$, so $\mathcal{V}_\Theta^F \supseteq \mathcal{X}_{\mathbf{z}}$, as claimed. $\diamond$ Claim 2

Claim 3: For any $\mathbf{v}, \mathbf{w} \in \mathcal{V}_\Theta^F$, if $\mathbf{v}$ and $\mathbf{w}$ are in the same $\overset{F}{\sim}$ -connected component, then $d(\mathbf{v}, \mathbf{z}) = d(\mathbf{w}, \mathbf{z})$.

Proof It suffices to prove this in the case when $\mathbf{v} \overset{F}{\sim} \mathbf{w}$; the general case follows by induction on path length. Now, if $\mathbf{v} \overset{F}{\sim} \mathbf{w}$, then $d(\mathbf{v}, \mathbf{w}) = 2$. Thus, $\mathcal{K}(\mathbf{v}, \mathbf{w}) = \{i, j\}$ for some $i, j \in \mathcal{K}$.

Claim 3A: Either $i \in \mathcal{K}(\mathbf{v}, \mathbf{x})$ or $i \in \mathcal{K}(\mathbf{w}, \mathbf{z})$, but not both. Likewise, either $j \in \mathcal{K}(\mathbf{v}, \mathbf{z})$ or $j \in \mathcal{K}(\mathbf{w}, \mathbf{x})$, but not both.

⁴² See Nehring et al. (2014, 2016) and Nehring and Pivato (2014) for an analysis of Condorcet admissibility in judgement aggregation.

Proof Since $v_i = -w_i$, we either have $v_i = -z_i$ or $w_i = -z_i$. But if $v_i = -z_i$, then evidently $w_i = z_i$. This proves the first claim. The second is similar. ∇ Claim 3A

Claim 3B: Exactly one of i or j is in $\mathcal{K}(\mathbf{v}, \mathbf{z})$.

Proof (by contradiction) If $\{i, j\} \subseteq \mathcal{K}(\mathbf{v}, \mathbf{z})$, then Claim 3A implies $\mathcal{K}(\mathbf{w}, \mathbf{z}) = \mathcal{K}(\mathbf{v}, \mathbf{z}) \setminus \{i, j\}$, which means $\mathcal{K}(\mathbf{w}, \mathbf{z}) \subsetneq \mathcal{K}(\mathbf{v}, \mathbf{z})$, which contradicts the fact that $\mathbf{v}$ is near to $\mathbf{z}$. On the other hand, if $\{i, j\}$ is disjoint from $\mathcal{K}(\mathbf{v}, \mathbf{z})$, then Claim 3A implies $\mathcal{K}(\mathbf{w}, \mathbf{z}) = \mathcal{K}(\mathbf{v}, \mathbf{z}) \sqcup \{i, j\}$, which means $\mathcal{K}(\mathbf{v}, \mathbf{z}) \subsetneq \mathcal{K}(\mathbf{w}, \mathbf{z})$, which contradicts the fact that $\mathbf{w}$ is near to $\mathbf{z}$. ∇ Claim 3B

Without loss of generality, suppose $i \in \mathcal{K}(\mathbf{v}, \mathbf{z})$ and $j \notin \mathcal{K}(\mathbf{v}, \mathbf{z})$. Thus, Claim 3A says $i \notin \mathcal{K}(\mathbf{w}, \mathbf{z})$. But by an argument similar to Claim 3B, exactly one of i or j is in $\mathcal{K}(\mathbf{w}, \mathbf{z})$. Thus, we must have $j \in \mathcal{K}(\mathbf{w}, \mathbf{z})$. At this point, we deduce that $\mathcal{K}(\mathbf{w}, \mathbf{z}) = \{j\} \sqcup \mathcal{K}(\mathbf{v}, \mathbf{z}) \setminus \{i\}$. Thus, $|\mathcal{K}(\mathbf{w}, \mathbf{z})| = |\mathcal{K}(\mathbf{v}, \mathbf{z})|$. In other words, $d(\mathbf{v}, \mathbf{z}) = d(\mathbf{w}, \mathbf{z})$, as claimed. $\diamond$ Claim 3

By the definition of $\mathbf{z}$, there exist $\mathbf{x}, \mathbf{y} \in \mathcal{X}_{\mathbf{z}}$ such that $d(\mathbf{x}, \mathbf{z}) \neq d(\mathbf{y}, \mathbf{z})$. Claim 2 says that $\mathbf{x}, \mathbf{y} \in \mathcal{V}_{\Theta}^F$, and Claim 3 implies that they must be in different $\sim^F$ -connected components. Thus, $(\mathcal{V}_{\Theta}^F, \sim^F)$ is disconnected, as desired.

This argument works for any additive majority rule (actually, any supermajority efficient rule) satisfying Continuity. Thus, $\mathcal{C}$ is frangible. $\square$

Case (b) of Theorem 3 follows from Theorem C.1 and the next result.

Lemma C.5 Let $\mathcal{C} = (\mathcal{K}, \lambda, \mathcal{X}, \mathcal{Y})$. If $\mathcal{X} \subseteq \{\pm 1\}^{\mathcal{K}}$ is distal, then $\mathcal{C}$ is frangible.

Proof By hypothesis, there exist some $\mathbf{x}, \mathbf{z} \in \mathcal{X}$ with $\mathbf{x}$ near to $\mathbf{z}$ and $d(\mathbf{x}, \mathbf{z}) \geq 3$. Let $\delta_{\mathbf{x}}$ and $\delta_{\mathbf{z}}$ be the point masses at $\mathbf{x}$ and $\mathbf{z}$, respectively, and let $\mu := \frac{1}{2}(\delta_{\mathbf{x}} + \delta_{\mathbf{z}})$. Let Θ be an open ball of small radius around μ ; Θ is obviously open and connected.

Now let F be any additive majority rule satisfying Continuity. We claim that $\mathcal{V}_{\Theta}^F \subseteq \{\mathbf{x}, \mathbf{z}\}$. To see this, let $\mathcal{J} := \mathcal{K}(\mathbf{x}, \mathbf{y})$ and let $\mathcal{L} := \mathcal{K} \setminus \mathcal{J}$. Without loss of generality, suppose $x_{\ell} = z_{\ell} = 1$ for all $\ell \in \mathcal{L}$. Then $\tilde{\mu}_{\ell} = 1$ for all $\ell \in \mathcal{L}$, while $\tilde{\mu}_j = 0$ for all $j \in \mathcal{J}$. Fix a weight vector $\lambda \in \mathbb{R}_{+}^{\mathcal{K}}$, and let $\Lambda = \sum_{\ell \in \mathcal{L}} \lambda_{\ell}$. Then for any $\mathbf{y} \in \mathcal{X}$ and any $q \in [0, 1]$, we have $\gamma_{\mu, \mathbf{x}}^{\lambda}(q) = \gamma_{\mu, \mathbf{z}}^{\lambda}(q) = \Lambda \geq \gamma_{\mu, \mathbf{y}}^{\lambda}(q)$, with equality if and only if $\mathbf{y}$ is between $\mathbf{x}$ and $\mathbf{z}$. But $\mathbf{x}$ is near to $\mathbf{z}$, so there is no $\mathbf{y} \in \mathcal{X} \setminus \{\mathbf{x}, \mathbf{y}\}$ which is between $\mathbf{x}$ and $\mathbf{z}$. Thus, $\text{SME}(\mathcal{X}, \mu) \subseteq \{\mathbf{x}, \mathbf{z}\}$. Thus, Lemma B.3 implies that $F(\mu) \subseteq \{\mathbf{x}, \mathbf{z}\}$. If we make the ball Θ small enough, then Continuity implies that $F(\theta) \subseteq \{\mathbf{x}, \mathbf{z}\}$ for all $\theta \in \Theta$; thus, $\mathcal{V}_{\Theta}^F \subseteq \{\mathbf{x}, \mathbf{z}\}$.

To see that $\mathcal{V}_{\Theta}^F = \{\mathbf{x}, \mathbf{z}\}$, define $\mu_n := \frac{n+1}{2n}\delta_{\mathbf{x}} + \frac{n-1}{2n}\delta_{\mathbf{z}}$ for all $n \in \mathbb{N}$. Then by an argument similar to the previous paragraph, supermajority efficiency implies that $F_n(\mu_n) = \{\mathbf{x}\}$ for all $n \in \mathbb{N}$. If n is large enough, then $\mu_n \in \Theta$; thus, $\mathbf{x} \in \mathcal{V}_{\Theta}^F$. By an identical argument (defining $\mu_n := \frac{n-1}{2n}\delta_{\mathbf{x}} + \frac{n+1}{2n}\delta_{\mathbf{z}}$ for all $n \in \mathbb{N}$), we obtain $\mathbf{z} \in \mathcal{V}_{\Theta}^F$. Thus, $\mathcal{V}_{\Theta}^F = \{\mathbf{x}, \mathbf{z}\}$, as claimed. But $d(\mathbf{x}, \mathbf{z}) \geq 3$; thus, $(\mathcal{V}_{\Theta}^F, \sim)$ is not path-connected, and thus, $\mathcal{C}$ is frangible. $\square$

Proof of Theorem 3 It is easy to verify that the weighted median rule (8) satisfies WESME, Continuity, and Reinforcement; we must verify the converse. If F is a rule

satisfying WESME and Continuity, then Proposition 4 says that F is an additive majority rule. Suppose F also satisfies Reinforcement. If $\mathcal{X}$ is rugged and $\mathcal{Y}$ is McGarvey, then Lemma C.4 says $\mathcal{C}$ is frangible. On the other hand, if $\mathcal{X}$ is distal and $\mathcal{Y}$ is thick, then Lemma C.5 says $\mathcal{C}$ is frangible. Either way, Theorem C.1 implies that F is the median rule. $\square$

References

- Balinski, M., Laraki, R.: Majority Judgment: Measuring, Ranking, and Electing. MIT Press, Cambridge (2011)
- Barthélemy, J.-P., Monjardet, B.: The median procedure in cluster analysis and social choice theory. *Math. Soc. Sci.* **1**(3), 235–267 (1981)
- Bossert, W., Sprumont, Y.: Strategy-proof preference aggregation: possibilities and characterizations. *Games Econ. Behav.* **85**, 109–126 (2014)
- Brams, S.J., Fishburn, P.C.: Approval Voting. Birkhäuser, Boston, MA (1983)
- Bredereck, R., Faliszewski, P., Igarashi, A., Lackner, M., Skowron, P.: Multiwinner elections with diversity constraints. *Proc. AAAI Conf. Artif. Intell.* **32**, 933–940 (2018)
- Day, W.H., McMorris, F.R.: Axiomatic Consensus Theory in Group Choice and Biomathematics. SIAM, Philadelphia (2003)
- Dhillon, A., Mertens, J.-F.: Relative utilitarianism. *Econometrica* **67**, 471–498 (1999)
- Dietrich, F.: Scoring rules for judgment aggregation. *Soc. Choice Welf.* **42**(4), 873–911 (2014)
- Dietrich, F.: Aggregation theory and the relevance of some issues to others. *J. Econ. Theory* **160**, 463–493 (2015)
- Dietrich, F., Mongin, P.: The premiss-based approach to judgment aggregation. *J. Econ. Theory* **145**(2), 562–582 (2010)
- Emerson, P.: From Majority Rule to Inclusive Politics—Electing a Power-Sharing Coalition. Springer, Berlin (2016)
- Endriss, U.: Judgment aggregation with rationality and feasibility constraints. In: Proceedings of the 17th International Conference on Autonomous Agents and MultiAgent Systems. International Foundation for Autonomous Agents and Multiagent Systems, pp. 946–954 (2018)
- Fischer, F., Hudry, O., Niedermeier, R.: Weighted tournament solutions. In: Brandt, F., Conitzer, V., Endriss, U., Lang, J., Procaccia, A.D. (eds.) *Handbook of Computational Social Choice*, pp. 85–102. Cambridge University Press, New York, NY (2016)
- Guilbaud, G.-T.: Octobre-Décembre: Les théories de l'intérêt général et le problème logique de l'aggrégation. *Economie Appliquée* **V 4**, 501–551 (1952)
- Hanson, R.: Shall we vote on values, but bet on beliefs? *J. Polit. Philos.* **21**, 151–178 (2013)
- Kemeny, J.G.: Fall: math without numbers. *Daedalus* **88**, 571–591 (1959)
- Kornhauser, L., Sager, L.: Unpacking the court. *Yale Law J.* **96**, 82–117 (1986)
- Lang, J., Pigozzi, G., Slavkovik, M., van der Torre, L., Vesic, S.: A partial taxonomy of judgment aggregation rules and their properties. *Soc. Choice Welf.* **48**(2), 327–356 (2017)
- Lang, S., Pigozzi, G., Slavkovik, M., van der Torre, L.: Judgement aggregation rules based on minimization. In: Proceedings of the 13th Conference on Theoretical Aspects of Rationality and Knowledge. ACM, New York, NY, pp. 238–246 (2011)
- List, C., Pettit, P.: Aggregating sets of judgements: an impossibility result. *Econ. Philos.* **18**, 89–110 (2002)
- List, C., Puppe, C.: Judgement aggregation: a survey. In: *Oxford Handbook of Rational and Social Choice*, pp. 457–482. Oxford University Press, Oxford (2009)
- McGarvey, D.C.: A theorem on the construction of voting paradoxes. *Econometrica* **21**, 608–610 (1953)
- McMorris, F.R., Mulder, H.M., Powers, R.C.: The median function on median graphs and semilattices. *Discrete Appl. Math.* **101**(1–3), 221–230 (2000)
- Miller, A.: A model of community standards. *J. Econ. Theory* **148**(6), 2696–2705 (2013)
- Miller, M.K., Osherson, D.: Methods for distance-based judgment aggregation. *Soc. Choice Welf.* **32**(4), 575–601 (2009)
- Mongin, P.: Factoring out the impossibility of logical aggregation. *J. Econ. Theory* **141**(1), 100–113 (2008)

- Mongin, P.: The doctrinal paradox, the discursive dilemma, and logical aggregation theory. *Theory Decis.* **73**(3), 315–355 (2012)
- Monjardet, B.: “Mathématique Sociale” and Mathematics. A case study: Condorcet’s effect and medians. *Electron. J. Hist. Probab. Stat.* **4**, 1–26 (2008)
- Myerson, R.B.: Axiomatic derivation of scoring rules without the ordering assumption. *Soc. Choice Welf.* **12**(1), 59–74 (1995)
- Nehring, K., Pivato, M.: Incoherent majorities: the McGarvey problem in judgement aggregation. *Discrete Appl. Math.* **159**, 1488–1507 (2011)
- Nehring, K., Pivato, M.: How indeterminate is sequential majority voting? A judgement aggregation perspective. In: *The Mathematics of Decisions, Elections, and Games*. Vol. 624 of *Contemp. Math.* Amer. Math. Soc., Providence, RI, pp. 55–88 (2014)
- Nehring, K., Pivato, M.: Majority rule in the absence of a majority. *J. Econ. Theory* **183**, 213–257 (2019)
- Nehring, K., Pivato, M., Puppe, C.: The Condorcet set: majority voting over interconnected propositions. *J. Econ. Theory* **151**, 268–303 (2014)
- Nehring, K., Pivato, M., Puppe, C.: Unanimity overruled: majority voting and the burden of history. *J. Theor. Polit.* **28**(4), 552–597 (2016)
- Nehring, K., Puppe, C.: The structure of strategy-proof social choice I: general characterization and possibility results on median spaces. *J. Econ. Theory* **135**, 269–305 (2007)
- Pascual, F., Rządca, K., Skowron, P.: Collective schedules: scheduling meets computational social choice. In: *Seventeenth International Conference on Autonomous Agents and Multiagent Systems (AAMAS’18)*, pp. 667–675 (2018)
- Pathak, P.A., Sönmez, T., Ünver, M.U., Yenmez, M.B.: Fair allocation of vaccines, ventilators and antiviral treatments: leaving no ethical value behind in health care rationing (2020). *arXiv preprint [arXiv:2008.00374](https://arxiv.org/abs/2008.00374)*
- Pigozzi, G., Slavkovik, M., van der Torre, L.: A complete conclusion-based procedure for judgment aggregation. In: *International Conference on Algorithmic Decision Theory*. Springer, pp. 1–13 (2009)
- Procaccia, A., Shah, N.: Optimal aggregation of uncertain preferences. *Proc. AAAI Conf. Artif. Intell.* **30**, 608–614 (2016)
- Smith, J.H.: Aggregation of preferences with variable electorate. *Econometrica* **41**, 1027–1041 (1973)
- Talmon, N., Faliszewski, P.: A framework for approval-based budgeting methods. *Proc. AAAI Conf. Artif. Intell.* **33**, 2181–2188 (2019)
- Young, H.P.: A note on preference aggregation. *Econometrica* **42**, 1129–1131 (1974)
- Young, H.P., Levenglick, A.: A consistent extension of Condorcet’s election principle. *SIAM J. Appl. Math.* **35**(2), 285–300 (1978)

Publisher’s Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.