

Scoring rules for judgment aggregation

Franz Dietrich¹

December 2011 (small revisions later)

Abstract

This paper introduces a new class of judgment aggregation rules, to be called ‘scoring rules’ after their famous counterparts in preference aggregation theory. A scoring rule generates the collective judgment set which reaches the highest total ‘score’ across the individuals, subject to the judgment set having to be rational. Depending on how we define ‘scores’, we obtain several (old and new) solutions to the judgment aggregation problem, such as distance-based aggregation, premise- and conclusion-based aggregation, truth-tracking rules, and a generalization of the Borda rule to judgment aggregation theory. Scoring rules are shown to generalize the classical scoring rules of preference aggregation theory.

JEL Classification: D70, D71

Keywords: judgment aggregation, social choice, scoring rules, Kemeny rule, Borda rule, distance-based aggregation

1 Introduction

The judgment aggregation problem consists in merging many individuals’ yes/no judgments on some interconnected propositions into collective yes/no judgments on these propositions. The classical example, born in legal theory, is that three jurors in a court trial disagree on which of the following three propositions are true: the defendant has broken the contract (p); the contract is legally valid (q); the defendant is liable (r). According to a universally accepted legal doctrine, r (the ‘conclusion’) is true if and only if p and q (the two ‘premises’) are both true. So, r is logically equivalent to $p \wedge q$. The simplest rule to aggregate the jurors’ judgments – namely propositionwise majority voting – may generate logically inconsistent collective judgments, as Table 1 illustrates. There are of course nu-

	premise p	premise q	conclusion r ($\Leftrightarrow p \wedge q$)
Individual 1	Yes	Yes	Yes
Individual 2	Yes	No	No
Individual 3	No	Yes	No
Majority	Yes	Yes	No

Table 1: The classical example of logically inconsistent majority judgments

merous other possible ‘agendas’, i.e., kinds of interconnected propositions a group might face. Preference aggregation is a special case with propositions of the form ‘ x is better than

¹I am grateful for very useful suggestions received from two anonymous referees and the editor. Affiliation: CNRS, Paris, France & UEA, U.K. Mail: fd@franzdietrich.net. Web: www.franzdietrich.net.

y ' (for many alternatives x and y), where these propositions are interconnected through standard conditions such as transitivity. In this context, Condorcet's classical *voting paradox* about cyclical majority preferences is nothing but another example of inconsistent majority judgments. Starting with List and Pettit's (2002) seminal paper, a whole series of contributions have explored which judgment aggregation rules can be used, depending on, firstly, the agenda in question, and, secondly, the requirements placed on aggregation, such as anonymity, and of course the consistency of collective judgments. Some theorems generalize Arrow's Theorem from preference to judgment aggregation (Dietrich and List 2007a, Dokow and Holzman 2010; both build on Nehring and Puppe 2010a and strengthen Wilson 1975). Other theorems have no immediate counterparts in classical social choice theory (e.g., List 2004, Dietrich 2006a, 2010, Nehring and Puppe 2010b, Dietrich and Mongin 2010).

It is fair to say that judgment aggregation theory has until recently been dominated by 'impossibility' findings, as is evident from the *Symposium on Judgment Aggregation* in Journal of Economic Theory (C. List and B. Polak eds., 2010, vol. 145(2)). The recent conference 'Judgment aggregation and voting' in Freudenstadt in 2011 however marks a visible shift of attention towards constructing concrete aggregation rules and finding 'second best' solutions in the face of impossibility results. The new proposals range from a first Borda-type aggregation rule (Zwicker 2011) to, among others, new distance-based rules (Duddy and Piggins 2012) and rules which approximate the majority judgments when these are inconsistent (Nehring, Pivato and Puppe 2011). The more traditional proposals include premise- and conclusion-based rules (e.g., Kornhauser and Sager 1986, Pettit 2001, List and Pettit 2002, Dietrich 2006a, Dietrich and Mongin 2010), sequential rules (e.g., List 2004, Dietrich and List 2007b), distance-based rules (e.g., Konieczni and Pino-Perez 2002, Pigozzi 2006, Miller and Osherson 2008, Eckert and Klamler 2009, Hartmann *et al.* 2010, Lang *et al.* 2011), and quota rules with well-calibrated acceptance thresholds and various degrees of collective rationality (e.g., Dietrich and List 2007b; see also Nehring and Puppe 2010a).

The present paper contributes to the theory's current 'constructive' effort by investigating a class of aggregation rules to be called *scoring rules*. The inspiration comes from classical scoring rules in preference aggregation theory. These rules generate collective preferences which rank each alternative according to the sum-total 'score' it receives from the group members, where the 'score' could be defined in different ways, leading to different classical scoring rules such as the Borda rule (see Smith 1973, Young 1975, Zwicker 1991, and for abstract generalizations Myerson 1995, Zwicker 2008, Conitzer *et al.* 2009 and Pivato 2011b). In a general judgment aggregation framework, however, there are no 'alternatives'; so our scoring rules are based on assigning scores to *propositions*, not alternatives. Nonetheless, our scoring rules are related to classical scoring rules, and generalize them, as will be shown.

The paradigm underlying our scoring rules – i.e., the maximization of total score of collective judgments – differs from standard paradigms in judgment aggregation, such as the premise-, conclusion- or distance-based paradigms. Nonetheless, it will turn out that several existing rules can be re-modelled as scoring rules, and can thus be 'rationalized' in terms of the maximization of total scores. Of course, the way scores are being assigned to propositions – the '*scoring*' – differs strongly across rules; for instance, the Kemeny rule and the premise-based rule can each be viewed as a scoring rule, but with respect to two very different scorings. This paper explores various plausible scorings. It uncovers the scorings

which implicitly underlie several well-known aggregation rules, and suggests other scorings which generate novel aggregation rules. For instance, a particularly natural scoring, to be called *reversal scoring*, will lead to a new generalization of the Borda rule from preference aggregation to judgment aggregation. The problem of how to generalize the Borda rule has been a long-lasting open question in judgment aggregation theory. Recently, an interesting (so far incomplete) proposal was made by William Zwicker (2011). Surprisingly, his and the present Borda generalizations are somewhat different, as detailed below.²

Though large, the class of scoring rules is far from universal: some important aggregation rules fall outside this class (notably the mentioned rule approximating the majority judgments, proposed by Nehring, Pivato and Puppe 2011). I will also investigate a natural generalization of scoring rules, to be called *set scoring rules*, which are based on assigning scores to entire judgment sets rather than single propositions (judgments). Such rules are examples of Zwicker’s (2008) ‘generalized scoring rules’ for an abstract aggregation problem whose inputs and outputs are arbitrary objects rather than judgment sets (see also Myerson 1995 and Pivato 2011a).³

After this introduction, Section 2 defines the general framework, Section 3 analyses various scoring rules, Section 4 goes on to analyse several set scoring rules, and Section 5 draws some conclusions about where we stand in terms of concrete aggregation procedures.

2 The framework, examples and interpretations

I now introduce the framework, following List and Pettit (2002) and Dietrich (2007).⁴ We consider a set of n (≥ 2) individuals, denoted $N = \{1, \dots, n\}$. They need to decide which of certain interconnected propositions to ‘believe’ or ‘accept’.

The agenda. The set X of propositions under consideration is called the *agenda*. It is subdivided into *issues*, i.e., pairs of a proposition and its negation, such as ‘it will rain’ and ‘it won’t rain’. Rationally, an agent accepts a proposition from each issue (‘completeness’), while respecting any logical interconnections between propositions (‘consistency’). We write ‘ $\neg p$ ’ for the negation of a proposition ‘ p ’, so that the agenda takes the form $X = \{p, \neg p, q, \neg q, \dots\}$, with issues $\{p, \neg p\}$, $\{q, \neg q\}$, etc. It is worth defining the present notion of an agenda formally:

Definition 1 *An agenda is a set X (of ‘propositions’) endowed with:*

- (a) *binary ‘issues’ $\{p, p'\}$ which partition X (where the members p and p' of an issue are the ‘negations’ of each other, written $p = \neg p'$ and $p' = \neg p$),*
- (b) *interconnections, i.e., a notion of which subsets of X are consistent, or formally, a system $\mathcal{C} \equiv \mathcal{C}_X$ of (‘consistent’) subsets.*⁵

²Conal Duddy and Ashley Piggins also have independent work in progress on ‘generalizing Borda’, and I learned from Klaus Nehring that he had ideas similar to those in the present paper.

³If the inputs and outputs of that aggregation problem are taken to be judgment sets, then Zwicker’s ‘generalized scoring rules’ become our set scoring rules (suitably translated into Zwicker’s setting of an anonymous variable population).

⁴To be precise, I use a slimmer variant of their models, since the logic in which propositions are formed is not explicitly part of the model.

⁵Algebraically, the agenda is thus the structure $(X, \mathcal{I}, \mathcal{C})$ (where \mathcal{I} is the partition into issues), or alternatively the structure (X, \neg, \mathcal{C}) . To see why the issues and the negation operator are interdefinable, note that we could alternatively start with an operator \neg on X satisfying $\neg \neg p = p$, and then define the issues as the pairs $\{p, \neg p\}$.

A simple example is the agenda given by

$$X = \{p, \neg p, q, \neg q, p \wedge q, \neg(p \wedge q)\}, \quad (1)$$

where p and q are two atomic sentences, for instance ‘it rains’ and ‘it is cold’, and $p \wedge q$ is their conjunction. The structure of this agenda – i.e., the issues and the interconnections – is obvious, as it is directly inherited from logic. For instance, the set $\{p, p \wedge q\}$ is consistent, while the set $\{\neg p, p \wedge q\}$ is not.

Given an agenda X , an individual’s *judgment set* is the set $J \subseteq X$ of propositions he accepts. It is *complete* if it contains a member of each issue $\{p, \neg p\}$, and (*fully*) *rational* if it is complete and consistent. The set of all rational judgment sets is denote by $\mathcal{J} \equiv \mathcal{J}_X$.

Under the above definition of an agenda, the set \mathcal{C} of consistent sets is a primitive, while the set \mathcal{J} is derived from it. One could proceed the over way round, by using a slightly modified definition of an agenda in which clause (b) is replaced by the following clause:

- (b) *interconnections, i.e., a notion of which judgment sets are rational, or formally, a system $\mathcal{J} \equiv \mathcal{J}_X \neq \emptyset$ of (‘rational’) judgment sets $J \subseteq X$, each containing exactly one member from any issue.*⁶

Here, \mathcal{J} is the primitive, and the system of consistent sets is defined as $\mathcal{C} = \mathcal{C}_X = \{C \subseteq X : C \subseteq J \text{ for some } J \in \mathcal{J}\}$. The main difference between the two ways to define an agenda is that under the second way the agenda – more precisely, its consistency notion – is automatically *well-behaved*, that is, satisfies the following regularity conditions (see Dietrich 2007):

- C1:** no set $\{p, \neg p\}$ is consistent (self-entailment);
- C2:** subsets of consistent sets are consistent (monotonicity);
- C3:** \emptyset is consistent and each consistent set can be extended to a complete and consistent set (completability).

Well-behavedness can be expressed equivalently by a single condition:

- $\mathcal{C} = \{C \subseteq J : J \in \mathcal{J}\} \neq \emptyset$, i.e., the consistent sets are the subsets of the fully rational sets.

An advantage of the second way to define an agenda is that well-behavedness – a common assumption of the theory up to date – comes for free and thus need not be added explicitly. An advantage of the first way is that it has no hidden assumptions, and lends itself to future relaxations of well-behavedness (for instance when studying judgment aggregation in non-monotonic logics).

Henceforth, let X be a given finite agenda with well-behaved interconnections (defined in any of the two ways). Notationally, a judgment set $J \subseteq X$ is often abbreviated by concatenating its members in any order (so, $p\neg q\neg r$ is short for $\{p, \neg q, \neg r\}$); and the negation-closure of a set $Y \subseteq X$ is denoted

$$Y^\pm \equiv \{p, \neg p : p \in Y\}.$$

We now introduce the two lead examples of this paper, the first one being isomorphic to the previous example (1).

⁶So, algebraically speaking, the agenda is the structure $(X, \mathcal{I}, \mathcal{J})$ instead of $(X, \mathcal{I}, \mathcal{C})$ (where \mathcal{I} is the partition of X into issues, replaceable by the negation operator \neg on X).

Example 1: the ‘doctrinal paradox agenda’. This agenda is

$$X = \{p, q, r\}^\pm,$$

where p , q and r are atomic sentences and where the logical interconnections are defined by classical logic relative to the external constraint $r \leftrightarrow (p \wedge q)$. So, there are four rational judgment sets:

$$\mathcal{J} = \{pqr, p\neg q\neg r, \neg pq\neg r, \neg p\neg q\neg r\}.$$

Example 2: the preference agenda. For an arbitrary, finite set of alternatives A , the *preference agenda* is defined as

$$X = X_A = \{xPy : x, y \in A, x \neq y\},$$

where the negation of a proposition xPy is of course $\neg xPy = yPx$, and where logical interconnections are defined by the usual conditions of transitivity, asymmetry and connectedness, which define a *strict linear order*. Formally, to each binary relation \succ over A uniquely corresponds a judgment set, denoted $J_\succ = \{xPy \in X : x \succ y\}$, and the set of all rational judgment sets is

$$\mathcal{J} = \{J_\succ : \succ \text{ is a strict linear order over } A\}.$$

Aggregation rules. A (*multi-valued*) *aggregation rule* is a correspondence F which to every profile of ‘individual’ judgment sets (J_1, \dots, J_n) (from some domain, usually \mathcal{J}^n) assigns a set $F(J_1, \dots, J_n)$ of ‘collective’ judgment sets. Typically, the output $F(J_1, \dots, J_n)$ is a singleton set $\{C\}$, in which case we identify this set with C and write $F(J_1, \dots, J_n) = C$. If $F(J_1, \dots, J_n)$ contains more than one judgment set, there is a ‘tie’ between these judgment sets. An aggregation rule is called *single-valued* or *tie-free* if it always generates a single judgment set. A standard (single-valued) aggregation rule is *majority rule*; it is given by

$$F(J_1, \dots, J_n) = \{p \in X : p \in J_i \text{ for more than half of the individuals } i\}$$

and generates inconsistent collective judgment sets for many agendas and profiles. If both individual and collective judgment sets are rational (i.e., in \mathcal{J}), the aggregation rule defines a correspondence $\mathcal{J}^n \rightrightarrows \mathcal{J}$, and in the case of single-valuedness a function $\mathcal{J}^n \rightarrow \mathcal{J}$.⁷

Some thoughts about the model. As an excursion for interested readers, let me add some considerations about the present model and its flexibility. Firstly, I mention three salient ways of specifying an agenda in practice. All three approaches could qualify broadly as ‘logical’:

- Under the *syntactic* approach, the propositions are sentences of formal logic. The structure of the agenda – i.e., the issues $\{p, \neg p\}$ and the interconnections – need not be specified explicitly, as it is directly inherited from the logic: it is given by the logic’s negation symbol and consistency notion.⁸ The logic needs to have the right expressive

⁷More generally, dropping the requirement of collective rationality, we have a correspondence $\mathcal{J}^n \rightrightarrows 2^X$, where 2^X is the set of *all* judgment sets, rational or not. As usual, I write ‘ \rightrightarrows ’ instead of ‘ \rightarrow ’ to indicate a *multi*-function.

⁸Formally, the agenda X is any set of logical sentences which is partitionable into pairs containing a sentence and its negation. These pairs define the issues. (We have presupposed that the logic includes the negation symbol, as all interesting logics do.)

power to adequately render the group’s decision problem: one might use standard propositional logic, or standard predicate logic, or various modal or conditional logics (see Dietrich 2007). Many real-life agendas draw on non-standard logics by involving for instance modal operators or non-material conditionals. Fortunately, most relevant logics are well-behaved, i.e., satisfy the conditions C1-C3 (now read as conditions on sentences of the logic), so that the agenda is automatically well-behaved.

- Under the *semantic* approach, the propositions are subsets of some fixed underlying set Ω of possibilities or worlds. The structure of the agenda (the issues $\{p, \neg p\}$ and the interconnections) need again not be specified explicitly, as it is inherited from set theory: the negation of a proposition p is the complement $\neg p := \Omega \setminus p$, and a set of propositions is consistent just in case its intersection is non-empty.⁹ Such a semantic agenda is automatically well-behaved.¹⁰
- Under an *algebraic* (or *abstract semantic*) approach, the agenda is a subset of an underlying Boolean algebra.¹¹ The structure of the agenda (the issues $\{p, \neg p\}$ and the interconnections) is once again automatically inherited, this time from the Boolean algebra.¹² The agenda is automatically well-behaved.

Secondly, I mention an interpretational point (orthogonal to the question of whether one works with a syntactic, semantic, algebraic or other agenda). Under the standard interpretation, we are aggregating ‘judgments’, i.e., belief-type attitudes towards propositions. But one may re-interpret the nature of the attitude, so that judgment sets become *desire* sets, or *hope* sets, or *normative approval* sets, or *intention* sets, and so on – which leads to desire aggregation, or hope aggregation, and so on. In this case we still aggregate propositional attitudes, albeit not judgments. In a more radical departure, we may consider the aggregation of arbitrary individual properties (characteristics). Here the agenda contains not propositions which someone may or may not believe (or desire, or hope etc.), but *properties* which someone may or may not *have* or *satisfy*. For instance, the agenda might contain the properties of liking peace, being successful, and so on. Each individual i has some set of properties J_i , and the goal is to derive a collective property set. This is the general *property aggregation* problem, distinct from the problem of aggregating propositional attitudes such as judgments.¹³

⁹Formally, the agenda X is any set of subsets of Ω such that $p \in X \Rightarrow \Omega \setminus p \in X$.

¹⁰Nehring and Puppe’s (2010a) *property spaces* are essentially semantically defined agendas.

¹¹A Boolean algebra is a lattice \mathcal{L} such that (i) its operations of join and meet satisfy distributivity, (ii) there exists a top element \top (tautology) and a bottom element \perp (contradiction), and (iii) every element p has a negation/complement (i.e., an element whose join with p is \top and whose meet with p is \perp). An important example is a *concrete* Boolean algebra $\mathcal{L} \subseteq 2^\Omega$ (for an underlying set of worlds Ω), in which the join is given by the union, the meet by the intersection, the top by Ω , the bottom by \emptyset , and the negation by the set-theoretic complement. In this case, the algebraic approach reduces to the standard semantic approach. Another example is the Boolean algebra generated from a logic, i.e., the set of sentences modulo logical equivalence (where the logic includes classical negation and conjunction, which induce the algebra’s join, meet and complement operations).

¹²Formally, the agenda X is any subset of the Boolean algebra which is closed under Boolean-algebraic negation/complementation (see footnote 11).

¹³Property aggregation raises the question of what it means for the collective to ‘have’ a property. Presumably, collective properties are something quite different from individual properties, just as collective judgments (or desires, hopes, ...) do not have the same status as individual ones. More generally, one may of course doubt the meaningfulness and relevance of group attributes. From a behaviourist perspective, the two ‘Humean’ collective attitudes – group beliefs (judgments) and group desires (preferences) – can be given meaning based on group behaviour, at least in principle. From my own (non-behaviourist) perspective, all sorts of non-standard group attributes can be meaningful and relevant, despite possibly being underdetermined by group action. But even if such group attributes are deemed meaningless or irrelevant, one may re-interpret them as being no more than summaries of the group members’ attributes

The mathematics built on the model does not ‘know’ which kind of interpretation or application one has in mind. Nonetheless, interpretation and context matter, since the question of how to best aggregate and which axioms to require depends on it.

3 Scoring rules

Scoring rules are particular judgment aggregation rules, defined on the basis of a so-called scoring function. A *scoring function* – or simply a *scoring* – is a function $s : X \times \mathcal{J} \rightarrow \mathbb{R}$ which to each proposition p and rational judgment set J assigns a number $s_J(p)$, called the *score* of p given J and measuring how p performs (‘scores’) from the perspective of holding judgment set J . As an elementary example, so-called *simple scoring* is given by:

$$s_J(p) = \begin{cases} 1 & \text{if } p \in J \\ 0 & \text{if } p \notin J, \end{cases} \quad (2)$$

so that all accepted propositions score 1, whereas all rejected propositions score 0. This and many other scorings will be analysed. Let us think of the score of a *set* of propositions as the sum of the scores of its members. So, the scoring s is extended to a function which (given the agent’s judgment set $J \in \mathcal{J}$) assigns to each set $C \subseteq X$ the score

$$s_J(C) \equiv \sum_{p \in C} s_J(p).$$

(In Section 4, we move to a generalized approach in which the score of a set C is a primitive and need not be additively derivable from scores of single propositions.)

A scoring s gives rise to an aggregation rule, called the *scoring rule w.r.t. s* and denoted F_s . Given a profile $(J_1, \dots, J_n) \in \mathcal{J}^n$, this rule determines the collective judgments by selecting the rational judgment set(s) with the highest sum-total score across all judgments and all individuals:

$$\begin{aligned} F_s(J_1, \dots, J_n) &= \text{judgment set(s) in } \mathcal{J} \text{ with highest total score} \\ &= \operatorname{argmax}_{C \in \mathcal{J}} \sum_{p \in C, i \in N} s_{J_i}(p) = \operatorname{argmax}_{C \in \mathcal{J}} \sum_{i \in N} s_{J_i}(C). \end{aligned}$$

By a *scoring rule simpliciter* we of course mean an aggregation rule which is a scoring rule w.r.t. some scoring. Different scorings s and s' can generate the same scoring rule $F_s = F_{s'}$, in which case they are called *equivalent*. For instance, s is equivalent to $s' = 2s$.¹⁴

3.1 Simple scoring and the Kemeny rule

We first consider the most elementary definition of scoring, namely *simple scoring* (2). Table 2 illustrates the corresponding scoring rule F_s for the case of the agenda and profile of our doctrinal paradox example. The entries in Table 2 are derived as follows. First, enter the score of each proposition ($p, \neg p, q, \dots$) from each individual (1, 2 and 3). Second, enter each individual’s score of each judgment set by taking the row-wise sum. For instance,

(rather than attributes of any separate group agent).

¹⁴More generally, certain increasing transformations have no effect. As one may show, scorings s and s' are equivalent (i.e., $F_s = F_{s'}$) whenever there are coefficients $a > 0$ and $b_p \in \mathbb{R}$ ($p \in X$) with $b_p = b_{\neg p}$ for all $p \in X$ such that s' is given by $s'_J(p) = a s_J(p) + b_p$.

							Score of...			
	p	$\neg p$	q	$\neg q$	r	$\neg r$	pqr	$p\neg q\neg r$	$\neg pq\neg r$	$\neg p\neg q\neg r$
Indiv. 1 (pqr)	1	0	1	0	1	0	3	1	1	0
Indiv. 2 ($p\neg q\neg r$)	1	0	0	1	0	1	1	3	1	2
Indiv. 3 ($\neg pq\neg r$)	0	1	1	0	0	1	1	1	3	2
Group	2	1	2	1	1	2	5*	5*	5*	4

Table 2: Simple scoring (2) for the doctrinal paradox agenda and profile

individual 1's score of pqr is $1 + 1 + 1 = 3$, and his score of $p\neg q\neg r$ is $1 + 0 + 0 = 1$. Third, enter the group's score of each proposition by taking the column-wise sum. For instance, the group's score of p is $1 + 1 + 0 = 2$. Finally, enter the group's score of each judgment set, by taking either a vertical or a horizontal sum (the two give the same result), and add a star '*' in the field(s) with maximal score to indicate the winning judgment set(s). For instance, the group's score of pqr using a vertical sum is $3 + 1 + 1 = 5$, and using a horizontal sum it is $2 + 2 + 1 = 5$. Since the judgment sets pqr , $p\neg q\neg r$ and $\neg pq\neg r$ all have maximal group score, the scoring rule delivers a tie:

$$F(J_1, J_2, J_3) = \{pqr, p\neg q\neg r, \neg pq\neg r\}.$$

This is a tie between the premise-based outcome pqr and the conclusion-based outcomes $p\neg q\neg r$ and $\neg pq\neg r$. Were we to add more individuals, the tie would presumably be broken in one way or the other. In large groups, ties are a rare coincidence.

To link simple scoring to distance-based aggregation, suppose we measure the distance between two rational judgment sets by using some *distance function* ('metric') d over \mathcal{J} .¹⁵ The most common example is the *Kemeny distance* $d = d_{\text{Kemeny}}$, which can also be attributed to the statistician Kendall. It is defined as follows (where by a 'judgment reversal' I mean the replacement of an accepted proposition by its negation):

$$\begin{aligned} d_{\text{Kemeny}}(J, K) &= \text{number of judgment reversals needed to transform } J \text{ into } K \\ &= |J \setminus K| = |K \setminus J| = \frac{1}{2} |J \Delta K|. \end{aligned}$$

For instance, the Kemeny-distance between pqr and $p\neg q\neg r$ (for our doctrinal paradox agenda) is 2.

Now the *distance-based rule* w.r.t. distance d is the aggregation rule F_d which for any profile $(J_1, \dots, J_n) \in \mathcal{J}^n$ determines the collective judgment set(s) by minimizing the sum-total distance to the individual judgment sets:

$$\begin{aligned} F_d(J_1, \dots, J_n) &= \text{judgment set(s) in } \mathcal{J} \text{ with minimal sum-distance to the profile} \\ &= \operatorname{argmin}_{C \in \mathcal{J}} \sum_{i \in N} d(C, J_i). \end{aligned}$$

The most popular example, the *Kemeny rule* $F_{d_{\text{Kemeny}}}$ (also referred to as the *Hamming rule* or the *median rule*), can be characterized as a scoring rule:

Proposition 1 *The simple scoring rule is the Kemeny rule.*

¹⁵A *distance function* or *metric* over \mathcal{J} is a function $d : \mathcal{J} \times \mathcal{J} \rightarrow [0, \infty)$ satisfying three conditions: for all $J, K, L \in \mathcal{J}$, (i) $d(J, K) = 0 \Leftrightarrow J = K$, (ii) $d(J, K) = d(K, J)$ ('symmetry'), and (iii) $d(J, L) \leq d(J, K) + d(K, L)$ ('triangle inequality').

3.2 Classical scoring rules for preference aggregation

I now show that our scoring rules generalize the classical scoring rules of preference aggregation theory. Consider the preference agenda X for a given set of alternatives A of finite size k . Classical scoring rules (such as the Borda rule) are defined by assigning scores to alternatives in A , not to propositions xPy in X . Given a strict linear order \succ over A , each alternative $x \in A$ is assigned a score $SCO_{\succ}(x) \in \mathbb{R}$. The most popular example is of course *Borda scoring*, for which the highest ranked alternative in A scores k , the second-highest $k - 1$, the third-highest $k - 2$, ..., and the lowest 1. Given a profile $(\succ_1, \dots, \succ_n)$ of individual preferences (strict linear orders), the collective ranks the alternatives $x \in X$ according to their sum-total score $\sum_{i \in N} SCO_{\succ_i}(x)$. To translate this into the judgment aggregation formalism, recall that each strict linear order \succ over A uniquely corresponds to a rational judgment set $J \in \mathcal{J}$ (given by $xPy \in J \Leftrightarrow x \succ y$); we may therefore write $SCO_J(x)$ instead of $SCO_{\succ}(x)$, and view the classical scoring SCO as a function of (x, J) in $A \times \mathcal{J}$. Formally, I define a *classical scoring* as an arbitrary function $SCO : A \times \mathcal{J} \rightarrow \mathbb{R}$, and the *classical scoring rule* w.r.t. it as the judgment aggregation rule $F \equiv F_{SCO}$ for the preference agenda which for every profile $(J_1, \dots, J_n) \in \mathcal{J}^n$ returns the rational judgment set(s) that rank an alternative x over another y whenever x has a higher sum-total score than y :¹⁶

$$F(J_1, \dots, J_n) = \{C \in \mathcal{J} : C \text{ contains all } xPy \in X \text{ s.t. } \sum_{i \in N} SCO_{J_i}(x) > \sum_{i \in N} SCO_{J_i}(y)\}.$$

Now, any given classical scoring SCO induces a scoring s in our (proposition-based) sense. In fact, there are two canonical (and, as we will see, equivalent) ways to define s : one might define s either by

$$s_J(xPy) = SCO_J(x) - SCO_J(y), \quad (3)$$

or, if one would like the lowest achievable score to be zero, by

$$s_J(xPy) = \max\{SCO_J(x) - SCO_J(y), 0\} = \begin{cases} SCO_J(x) - SCO_J(y) & \text{if } xPy \in J \\ 0 & \text{if } xPy \notin J \end{cases} \quad (4)$$

(where the last equality assumes that $SCO_J(x) > SCO_J(y) \Leftrightarrow xPy \in J$ for all x, y and J , a property that is so natural that we might have built it into the definition of a ‘classical scoring’ SCO). This allows us to characterize classical scoring rules in terms of proposition-based rather than alternative-based scoring:

Proposition 2 *In the case of the preference agenda (for any finite set of alternatives), every classical scoring rule is a scoring rule, namely one with respect to a scoring s derived from the classical scoring SCO via (3) or via (4).*

¹⁶A technical difference between the standard notion of a scoring rule in preference aggregation theory and our judgment-theoretic rendition of it arises when there happen to exist distinct alternatives with identical sum-total score. In such cases, the standard scoring rule returns collective *indifferences*, whereas our F_{SCO} returns a *tie* between *strict* preferences. From a formal perspective, however, the two definitions are equivalent, since to any weak order corresponds the set (tie) of all strict linear orders which linearize the weak order by breaking its indifferences (in any cycle-free way). The structural asymmetry between input and output preferences of scoring rules as defined standardly (i.e., the possibility of indifferences at the collective level) may have been one of the obstacles – albeit only a small, mainly psychological one – for importing scoring rules and Borda aggregation into judgment aggregation theory.

3.3 Reversal scoring and a Borda rule for judgment aggregation

Given the agent's judgment set J , let us think of the score of a proposition $p \in X$ as a measure of how 'distant' the negation $\neg p$ is from J ; so, p scores high if $\neg p$ is far from J , and low if $\neg p$ is contained in J . More precisely, let the score of a proposition p given $J \in \mathcal{J}$ be the number of judgment reversals needed to reject p , i.e., the number of propositions in J that must (minimally) be negated in order to obtain a consistent judgment set containing $\neg p$. So, denoting the judgment set arising from J by negating the propositions in a subset $R \subseteq J$ by $J_{-R} = (J \setminus R) \cup \{\neg r : r \in R\}$, so-called *reversal scoring* is defined by

$$\begin{aligned} s_J(p) &= \text{number of judgment reversals needed to reject } p \\ &= \min_{R \subseteq J: J_{-R} \in \mathcal{J} \& p \notin J_{-R}} |R| = \min_{J' \in \mathcal{J}: p \notin J'} |J \setminus J'| = \min_{J' \in \mathcal{J}: p \notin J'} d_{\text{Kemeny}}(J, J'). \end{aligned} \quad (5)$$

For instance, a rejected proposition $p \notin J$ scores zero, since J itself contains $\neg p$ so that it suffices to negate *zero* propositions ($R = \emptyset$). An accepted proposition $p \in J$ scores 1 if J remains consistent by negating p ($R = \{p\}$), and scores more than 1 otherwise ($R \supsetneq \{p\}$). Table 3 illustrates reversal scoring for our doctrinal paradox example. For instance, individual 1's judgment set pqr leads to a score of 2 for proposition p , since in order for him to reject p he needs to negate not just p (as $\neg pqr$ is inconsistent), but also r (where $\neg pq\neg r$ is consistent). The scoring rule delivers a tie between the judgment sets

	Score of...									
	p	$\neg p$	q	$\neg q$	r	$\neg r$	pqr	$p\neg q\neg r$	$\neg pq\neg r$	$\neg p\neg q\neg r$
Indiv. 1 (pqr)	2	0	2	0	2	0	6	2	2	0
Indiv. 2 ($p\neg q\neg r$)	1	0	0	2	0	2	1	5	2	4
Indiv. 3 ($\neg pq\neg r$)	0	2	1	0	0	2	1	2	5	4
Group	3	2	3	2	2	4	8	9*	9*	8

Table 3: Reversal scoring (5) for the doctrinal paradox agenda and profile

$p\neg q\neg r$ and $\neg pq\neg r$. This is a tie between two conclusion-based outcomes; the premise-based outcome pqr is rejected (unlike for simple scoring in Section 3.1).

The remarkable feature of reversal scoring rule is that it generalizes the Borda rule from preference to judgment aggregation. The Borda rule is initially only defined for the preference agenda X (for a given finite set of alternatives), namely as the classical scoring rule w.r.t. Borda scoring; see the last subsection. The key observation is that reversal scoring is intimately linked to Borda scoring:

Remark 1 *In the case of the preference agenda (for any finite set of alternatives), reversal scoring s is given by (4) with SCO defined as classical Borda scoring.*

Let me sketch the simple argument – it should sound familiar to social choice theorists. Let s be reversal scoring, X the preference agenda for a set of alternatives A of size $k < \infty$, and SCO classical Borda scoring. Consider any $xPy \in X$ and $J \in \mathcal{J}$. If $xPy \in X \setminus J$, then $\neg xPy = yPx \in J$, which implies $s_J(xPy) = 0$, as required by (4). Now suppose $xPy \in J$. Clearly, $SCO_J(x) > SCO_J(y)$. Consider the alternatives in the order \succ established by J :

$$x_k \succ x_{k-1} \succ \cdots \succ x \succ \cdots \succ y \succ \cdots \succ x_1,$$

where x_j is the alternative with $SCO_J(x_j) = j$. Step by step, we now move y up in the ranking, where each step consists in raising the position (score) of y by one. Each step

corresponds to negating one proposition in J , namely the proposition zPy where z is the alternative that is currently being ‘overtaken’ by y . After exactly $SCO_J(x) - SCO_J(y)$ steps, y has ‘overtaken’ x , i.e., xPy has been negated. So, $s_J(xPy)$ is *at most* $SCO_J(x) - SCO_J(y)$. It is *exactly* $SCO_J(x) - SCO_J(y)$, since, as the reader may check, no smaller number of judgment reversals allows y to ‘overtake’ x in the ranking.

Remark 1 and Proposition 2 imply that reversal scoring allows us to extend the Borda rule to arbitrary judgment aggregation problems:

Proposition 3 *The reversal scoring rule generalizes the Borda rule, i.e., matches it in the case of the preference agenda (for any finite set of alternatives).*

I note that one could use a perfectly equivalent variant of reversal scoring s which, in the case of the preference agenda, is related to classical Borda scoring SCO via (3) instead of (4):

Remark 2 *Reversal scoring s is equivalent (in terms of the resulting scoring rule) to the scoring s' given by*

$$s'_J(p) = s_J(p) - s_J(\neg p) = \begin{cases} s_J(p) & \text{if } p \in J \\ -s_J(\neg p) & \text{if } p \notin J, \end{cases}$$

and in the case of the preference agenda (for any finite set of alternatives) this scoring is given by

$$s'_J(xPy) = SCO_J(x) - SCO_J(y)$$

with SCO defined as classical Borda scoring.

For comparison, I now sketch Zwicker’s (2011) interesting approach to extending the Borda rule to judgment aggregation – let me call such an extension a ‘Borda-Zwicker’ rule. The motivation derives from a geometric characterization of Borda preference aggregation obtained by Zwicker (1991). Let me write the agenda as $X = \{p_1, \neg p_1, p_2, \neg p_2, \dots, p_m, \neg p_m\}$, where m is the number of ‘issues’. Each profile gives rise to a vector $\mathbf{v} \equiv (v_1, \dots, v_m)$ in \mathbb{R}^m whose j^{th} entry v_j is the *net support for p_j* , i.e., the number of individuals accepting p_j minus the same number for $\neg p_j$. Now if X is the preference agenda for any finite set of alternatives A , then each p_j takes the form xPy for certain alternatives $x, y \in A$. Each preference cycle can be mapped to a vector in \mathbb{R}^m ; for instance, if $p_1 = xPy$, $p_2 = yPz$ and $p_3 = xPz$, then the cycle $x \succ y \succ z \succ x$ becomes the vector $(1, 1, -1, 0, \dots, 0) \in \mathbb{R}^m$. The linear span of all vectors corresponding to preference cycles defines the so-called ‘cycle space’ $\mathbf{V}_{\text{cycle}} \subseteq \mathbb{R}^m$, and its orthogonal complement defines the ‘cocycle space’ $\mathbf{V}_{\text{cocycle}} \subseteq \mathbb{R}^m$. Let $\mathbf{v}_{\text{cocycle}}$ be the orthogonal projection of \mathbf{v} on the cocycle space $\mathbf{V}_{\text{cocycle}}$. Intuitively, $\mathbf{v}_{\text{cocycle}}$ contains the ‘consistent’ or ‘acyclic’ part of \mathbf{v} . The upshot is that the Borda outcome can be read off from $\mathbf{v}_{\text{cocycle}}$: for each $p_j = xPy$, the Borda group preference ranks x above (below) y if the j^{th} entry of $\mathbf{v}_{\text{cocycle}}$ is positive (negative). Zwicker’s strategy for extending the Borda rule to judgment aggregation is to define a subspace $\mathbf{V}_{\text{cycle}}$ analogously for agendas other than the preference agenda; one can then again project \mathbf{v} on the orthogonal complement of $\mathbf{V}_{\text{cycle}}$ and determine collective ‘Borda’ judgments according to the signs of the entries of this projection. This approach has proved successful for simple agendas, in which there is a natural way to define $\mathbf{V}_{\text{cycle}}$. Whether the approach is viable

for general agendas (i.e., whether \mathbf{V}_{cycle} has a useful general definition) seems to be open so far.¹⁷

A Borda-Zwicker rule is not just constructed differently from a scoring rule in our sense, but, as I conjecture, it also cannot generally be remodelled as a scoring rule, since most interesting scoring rules use information that goes beyond the information contained in the profile’s ‘net support vector’ $\mathbf{v} \in \mathbb{R}^m$. (Even more does the required information go beyond the projection of \mathbf{v} on the orthogonal complement of \mathbf{V}_{cycle} .)

In summary, there seem to exist two quite different approaches to generalizing Borda aggregation. One approach, taken by Zwicker, seeks to filter out the profile’s ‘inconsistent component’ along the lines of the just-described geometric technique. The other approach, taken here, seeks to retain the principle of score-maximization inherent in Borda aggregation (with scoring now defined at the level of propositions, not alternatives, as these do not exist outside the world of preferences). One might view the normative core of the scoring approach as being the use of the *strength* of someone’s judgments, i.e., the strength of acceptance of propositions, as measured by the score – just as Borda preference aggregation is sometimes interpreted as a way to incorporate the *strength* of someone’s preference between alternatives x and y , as measured by the score of xPy , i.e., the difference between x ’s and y ’s score.¹⁸

3.4 A generalization of reversal scoring

Recall that the reversal score of a proposition p can be characterized as the distance by which one must deviate from the current judgment set in order to reject p – where ‘distance’ is understood as Kemeny-distance. It is natural to also consider other kinds of a distance. Relative to any given distance function d over \mathcal{J} , one may define a corresponding scoring by

$$\begin{aligned} s_J(p) &= \text{distance by which one must depart from } J \text{ to reject } p & (6) \\ &= \min_{J' \in \mathcal{J}: p \notin J'} d(J, J'). \end{aligned}$$

This provides us with a whole class of scoring rules, all of which are variants of our judgment-theoretic Borda rule. In the special case of the preference agenda, we thus obtain new variants of the classical Borda rule.

For instance, we could adopt Duddy and Piggins’ (2012) distance function, i.e., we could define $d(J, J')$ as the number of *minimal consistent modifications* needed to transform J into J' .¹⁹ So, while Duddy and Piggins had introduced their distance in the context of distance-based aggregation to develop an alternative to the Kemeny rule, their distance can

¹⁷One might at first be tempted to generally define \mathbf{V}_{cycle} as the linear span of those vectors which correspond to the agenda’s *minimal inconsistent subsets*. Unfortunately, this span is often the entire space \mathbb{R}^m , an example for this being our doctrinal paradox agenda.

¹⁸By contrast, Condorcet’s rule of pairwise majority voting only cares about *whether or not* someone prefers an alternative over another, without attempting to extract strength-of-preference information from the person’s full preference relation. It is of course notoriously controversial whether strength or intensity of preference is a permissible concept – and even if it is, whether the score of xPy is an adequate proxy of the strength to which x is preferred to y . The ordinalist approach takes a sceptical stance here. Note however that Borda aggregation can be defended for reasons unrelated to strength of preference. For instance, its axiomatic characterizations offer possible normative defences.

¹⁹Judgment sets $J, J' \in \mathcal{J}$ are *minimal consistent modifications* of each other if the set $S = J \setminus J'$ of propositions in J which need to be negated to transform J into J' is non-empty and minimal (i.e., J couldn’t have been transformed into a consistent set by negating only a strict non-empty subset of S). For

also be used in our context of (reversal) scoring rules to develop an alternative to Borda aggregation.

3.5 Scoring based on logical entrenchment

We now consider scoring rules which explicitly exploit the logical structure of the agenda. Let us think of the score of a proposition p ($\in X$) given the judgment set J ($\in \mathcal{J}$) as the degree to which p is logically entrenched in the belief system J , i.e., as the ‘strength’ with which J entails p . We measure this strength by the number of ways in which p is entailed by J , where each ‘way’ is given by a particular judgment subset $S \subseteq J$ which entails p , i.e., for which $S \cup \{\neg p\}$ is inconsistent. If J does not contain p , then no judgment subset – not even the full set J – can entail p ; so the strength of entailment (score) of p is zero. If J contains p , then p is entailed by the judgment subset $\{p\}$, and perhaps also by very different judgment subsets; so the strength of entailment (score) of p is positive and more or less high.

There are different ways to formalise this idea, depending on precisely which of the judgment subsets that entail p are deemed relevant. I now propose four formalizations. Two of them will once again allow us to generalize the Borda rule from preference to judgment aggregation. These generalizations differ from that based on reversal scoring in Section 3.3.

Our first, naive approach is to count *each* judgment subset which entails p as a separate, full-fledged ‘way’ in which p is entailed. This leads to so-called *entailment scoring*, defined by:

$$\begin{aligned} s_J(p) &= \text{number of judgment subsets which entail } p & (7) \\ &= |\{S \subseteq J : S \text{ entails } p\}|. \end{aligned}$$

If $p \notin J$ then $s_J(p) = 0$, while if $p \in J$ then $s_J(p) \geq 2^{|X|/2-1}$ since p is entailed by at least all sets $S \subseteq J$ which contain p , i.e., by at least $2^{|J|-1} = 2^{|X|/2-1}$ sets. One might object that this definition of scoring involves redundancies, i.e., ‘multiple counting’. Suppose for instance p belongs to J and is logically independent of all other propositions in J . Then p is entailed by several subsets S of J – all $S \subseteq J$ which contain p – and yet these entailments are essentially identical since all premises in S other than p are irrelevant.

I now present three refinements of scoring (7), each of which responds differently to the mentioned redundancy objection. In the first refinement, we count two entailments of p as different only if they have no premise in common. This leads to what I call *disjoint-entailment scoring*, formally defined by:

$$\begin{aligned} s_J(p) &= \text{number of } \textit{pairwise disjoint} \text{ judgment subsets entailing } p & (8) \\ &= \max\{m : J \text{ has } m \text{ pairwise disjoint subsets each entailing } p\}. \end{aligned}$$

In the mentioned case where p ($\in J$) is logically independent of all other propositions in J , we now avoid ‘multiple counting’: $s_J(p)$ is only 1, as one cannot find different pairwise disjoint judgment subsets entailing p . For our doctrinal paradox agenda and profile, the scoring rule delivers a tie between the two conclusion-based outcomes $p\neg q\neg r$ and $\neg pq\neg r$, as illustrated in Table 4. For instance, individual 2 has judgment set $p\neg q\neg r$, so that p

our doctrinal paradox agenda, the judgment sets pqr and $p\neg q\neg r$ are minimal consistent modifications of each other, and hence have Duddy-Piggins-distance of 1.

							Score of...			
	p	$\neg p$	q	$\neg q$	r	$\neg r$	pqr	$p\neg q\neg r$	$\neg pq\neg r$	$\neg p\neg q\neg r$
Indiv. 1 (pqr)	2	0	2	0	2	0	6	2	2	0
Indiv. 2 ($p\neg q\neg r$)	1	0	0	2	0	2	1	5	2	4
Indiv. 3 ($\neg pq\neg r$)	0	2	1	0	0	2	1	2	5	4
Group	3	2	3	2	2	4	8	9*	9*	8

Table 4: Disjoint-entailment scoring (8) for the doctrinal paradox agenda and profile

scores 1 (it is entailed by $\{p\}$ but by no other disjoint judgment subset), $\neg q$ scores 2 (it is disjointly entailed by $\{\neg q\}$ and $\{p, \neg r\}$), $\neg r$ scores 2 (it is disjointly entailed by $\{\neg r\}$ and $\{\neg q\}$), and all rejected propositions score zero (they are not entailed by any judgment subsets).

Disjoint-entailment scoring turns out to match reversal scoring for our doctrinal paradox agenda (check that Tables 3 and 4 coincide), as well as for the preference agenda (as shown later). Is this pure coincidence? The general relationship is that the disjoint-entailment score of a proposition p is always *at most* the reversal score, as one may show.²⁰

While this refinement of naive entailment scoring (7) avoids ‘multiple counting’ by only counting entailments with pairwise disjoint sets of premises, the next two refinements use a different strategy to avoid ‘multiple counting’. The new strategy is to count only those entailments whose sets of premises are *minimal* – with minimality understood either in the sense that no premises can be removed, or in the sense that no premises can be logically weakened. To begin with the first sense of minimality, I say that a set *minimally entails* p ($\in X$) if it entails p but no strict subset of it entails p , and I define *minimal-entailment scoring* by

$$\begin{aligned}
 s_J(p) &= \text{number of judgment subsets which } \textit{minimally} \textit{ entail } p & (9) \\
 &= |\{S \subseteq J : S \textit{ minimally entails } p\}|.
 \end{aligned}$$

If for instance p is contained in J , then $\{p\}$ minimally entails p ,²¹ but strict supersets of $\{p\}$ do not and are therefore not counted. For our doctrinal paradox agenda, this scoring happens to coincide with reversal scoring and disjoint-entailment scoring. Indeed, Table 3 resp. 4 still applies; e.g., for individual 2 with judgment set $p\neg q\neg r$, p still scores 1 (it is minimally entailed only by $\{p\}$), $\neg q$ still scores 2 (it is minimally entailed by $\{\neg q\}$ and by $\{p, \neg r\}$), $\neg r$ still scores 2 (it is minimally entailed by $\{\neg r\}$ and by $\{\neg q\}$), and all rejected propositions still score zero (they are not minimally entailed by any judgment subsets).

Scoring (9) is certainly appealing. Nonetheless, one might complain that it still allows for certain redundancies, albeit of a different kind. Consider the preference agenda with set of alternatives $A = \{x, y, z, w\}$, and the judgment set $J = \{xPy, yPz, zPw, xPz, yPw, xPw\}$ ($\in \mathcal{J}$). The proposition xPw is minimally entailed by the subset $S = \{xPy, yPz, zPw\}$. While this entailment is minimal in the (set-theoretic) sense that we cannot *remove* premises, it is non-minimal in the (logical) sense that we can *weaken* some of its premises: if we replace xPy and yPz in S by their logical implication xPz , then we obtain a weaker set of premises $S' = \{xPz, zPw\}$ which still entails xPw . We

²⁰The reason is that, given m mutually disjoint judgment subsets which each entail p , the reversal score of p is at least m since one must negate at least one proposition from each of these m sets in order to consistently reject p .

²¹Assuming that p is not a tautology, i.e., that $\{\neg p\}$ is consistent. (Otherwise, \emptyset minimally entails p .)

shall say that S fails to ‘irreducibly’ entail xPw , in spite of minimally entailing it. In general, a set of propositions is called *weaker* than another one (which is called *stronger*) if the second set entails each member of the first set, but not vice versa. A set $S (\subseteq X)$ is defined to *irreducibly* (or *logically minimally*) entail p if S entails p and no subset of premises $Y \subsetneq S$ can be weakened in this entailment (i.e., for no subset $Y \subsetneq S$ there exists a weaker set $Y' \subseteq X$ such that $(S \setminus Y) \cup Y'$ still entails p). Each irreducible entailment is a minimal entailment, as is seen by taking $Y' = \emptyset$.²² In the previous example, the set $\{xPy, yPz, zPw\}$ minimally, but not irreducibly entails xPw , and the set $\{xPz, zPw\}$ irreducibly entails xPw . *Irreducible-entailment scoring* is naturally defined by

$$\begin{aligned} s_J(p) &= \text{number of judgment subsets which irreducibly entail } p & (10) \\ &= |\{S \subseteq J : S \text{ irreducibly entails } p\}|. \end{aligned}$$

This scoring matches reversal scoring and both previous scorings in the case of our doctrinal paradox example: Table 3 resp. 4 still applies. But for many other agendas these scorings all deviate from one another, resulting in different collective judgments. As for the preference agenda, we have already announced the following result:

Proposition 4 *Disjoint-entailment scoring (8) and irreducible-entailment scoring (10) match reversal scoring (5) in the case of the preference agenda (for any finite set of alternatives).*

Propositions 3 and 4 jointly have an immediate corollary.

Corollary 1 *The scoring rules w.r.t. scorings (8) and (10) both generalize the Borda rule, i.e., match it in the case of the preference agenda (for any finite set of alternatives).*

3.6 Propositionwise scoring and a way to repair quota rules with non-rational outputs

We now consider a special class of scorings: *propositionwise* scorings. This will allow us to relate scoring rules to the well-known judgment aggregation rules called *quota rules* (Dietrich and List 2007b) – in fact, to ‘repair’ these rules by rendering their outcomes rational across all profiles.

I call scoring s *propositionwise* if the score of a proposition $p \in X$ only depends on whether p is accepted, i.e., if $s_J(p) = s_K(p)$ whenever J and K (in \mathcal{J}) both contain p or both do not contain p . Equivalently, scoring is propositionwise just in case for each $p \in X$ there is a pair of real numbers $s_+(p), s_-(p)$ such that

$$s_J(p) = \begin{cases} s_+(p) & \text{for all } J \in \mathcal{J} \text{ containing } p \\ s_-(p) & \text{for all } J \in \mathcal{J} \text{ not containing } p. \end{cases} \quad (11)$$

Intuitively, $s_+(p)$ is the score of an accepted proposition p , and $s_-(p)$ is the score of a rejected proposition p . Typically, of course, $s_+(p) > s_-(p)$. An example is simple scoring: there, $s_+(p) = 1$ and $s_-(p) = 0$.

How do propositionwise scoring rules behave? They derive a proposition p ’s sum-total score ‘locally’, i.e., based only on people’s judgments about p . This property stands in

²²Assuming X contains no tautology, i.e., no p such that $\{\neg p\}$ is inconsistent.

obvious analogy to a well-studied axiom on aggregation rules, namely the axiom of *propositionwise* or *independent* aggregation, which prescribes that the collective judgment about any given proposition p is derived ‘locally’, i.e., again based only on people’s judgments about p . Can we therefore relate propositionwise scoring to independent aggregation? The paradigmatic independent aggregation rules are the *quota rules*.²³ A quota rule is a (single-valued) aggregation rule which is given by an acceptance threshold $m_p \in \{1, \dots, n\}$ for each proposition $p \in X$. The quota rule corresponding to the so-called *threshold family* $(m_p)_{p \in X}$ is denoted $F_{(m_p)_{p \in X}}$ and accepts those propositions p which are supported by at least m_p individuals: for each profile $(J_1, \dots, J_n) \in \mathcal{J}^n$,

$$F_{(m_p)_{p \in X}}(J_1, \dots, J_n) = \{p \in X : |\{i : p \in J_i\}| \geq m_p\}.$$

Special cases are unanimity rule (given by $m_p = n$ for all p), majority rule (given by the majority threshold $m_p = \lceil (n+1)/2 \rceil$ for all p), and more generally, *uniform* quota rules (given by a uniform threshold $m_p \equiv m$ for all p). A uniform quota rule is also referred to as a supermajority rule if m exceeds the majority threshold, and a submajority rule if m is below the majority threshold. Note that supermajority rules may generate incomplete collective judgment sets, while submajority rule may accept both members of a pair $p, \neg p \in X$, a drastic form of inconsistency. If one wishes that exactly one member of each pair $p, \neg p \in X$ is accepted, the thresholds of p and $\neg p$ should be ‘complements’ of each other: $m_p = n + 1 - m_{\neg p}$.

A non-trivial question is how the acceptance thresholds would have to be set to ensure that the collective judgment set satisfies some given degree of rationality, such as to be (i) consistent, or (ii) deductively closed, or (iii) consistent and deductively closed, or even (iv) fully rational, i.e., in \mathcal{J} . These questions have been settled (see Nehring and Puppe 2010a for (iv), and, subsequently, Dietrich and List 2007b for (i)-(iv)). Unfortunately, for many agendas the thresholds would have to be set at ‘extreme’ and normatively unattractive levels. Worse, often *no* thresholds achieve (iv) (see Nehring and Puppe 2010a). For our doctrinal paradox agenda $X = \{p, q, r\}^\pm$ only the extreme thresholds $m_p = m_q = m_r = n$ and $m_{\neg p} = m_{\neg q} = m_{\neg r} = 1$ achieve (iv), and for the preference agenda (with more than two alternatives) *no* thresholds achieve (iv).

Given that quota rules with ‘reasonable’ thresholds typically violate many of the conditions (i)-(iv), one may want to depart from ordinary quota rules by modifying (‘repairing’) them so that they always generate rational outputs. This can be done by using propositionwise scoring rules. Given an arbitrary quota rule with threshold family $(m_p)_{p \in X}$, one can specify a propositionwise scoring such that the scoring rule replicates the quota rule whenever the quota rule generates a rational output, while ‘repairing’ the output otherwise. How must we calibrate $s_+(p)$ and $s_-(p)$ in order to achieve this? The idea is that individuals who accept p should contribute a positive score $s_+(p) > 0$, while those who reject p should contribute a negative score $s_-(p) < 0$. The absolute sizes of $s_+(p)$ and $s_-(p)$ should be calibrated such that the sum-total score of p becomes positive (helping the scoring rule to accept p) exactly when the quota rule accepts p , i.e., when at least m_p individuals accept p . Specifically, we set:

$$s_J(p) = \begin{cases} s_+(p) = n + 1 - m_p & \text{for all } J \in \mathcal{J} \text{ containing } p \\ s_-(p) = -m_p & \text{for all } J \in \mathcal{J} \text{ not containing } p. \end{cases} \quad (12)$$

Intuitively, the higher the acceptance threshold m_p is, the smaller the positive contribution $s_+(p)$ is and the larger the negative contribution $s_-(p)$ is (in absolute value); hence, the

²³They are the only independent rules which are anonymous, monotonic and unanimity-preserving.

more individuals accepting p are needed for p 's sum-total score to get positive, and the harder it becomes for the scoring rule to accept p . This scoring does the intended job:

Proposition 5 *For every threshold family $(m_p)_{p \in X}$, the scoring rule w.r.t. scoring (12) matches the quota rule $F_{(m_p)_{p \in X}}$ at all profiles where the quota rule generates rational outputs (and still generates rational outputs at all other profiles).*

As an example, consider our doctrinal paradox agenda $X = \{p, q, r\}^\pm$ with $n = 3$ individuals, and suppose the quota rule departs only slightly from propositionwise majority voting: all propositions t in $X \setminus \{\neg r\}$ keep a majority threshold of $m_t = 2$, but $\neg r$ receives a unanimity threshold $m_{\neg r} = 3$. This quota rule manages to never generate logically inconsistent collective judgment sets,²⁴ but does so at the expense of allowing collective incompleteness. Indeed, for our example profile, the quota rule returns the collective judgment set pq , which is silent on the choice between r and $\neg r$. As illustrated in Table 5, the scoring rule w.r.t. (12) restores collective rationality by leading to the premise-based

	Score of...									
	p	$\neg p$	q	$\neg q$	r	$\neg r$	pqr	$p\neg q\neg r$	$\neg pq\neg r$	$\neg p\neg q\neg r$
Indiv. 1 (pqr)	2	-2	2	-2	2	-3	6	-3	-3	-7
Indiv. 2 ($p\neg q\neg r$)	2	-2	-2	2	-2	1	-2	5	-3	1
Indiv. 3 ($\neg pq\neg r$)	-2	2	2	-2	-2	1	-2	-3	5	1
Group	2	-2	2	-2	-2	-1	2*	-1	-1	-5

Table 5: Scoring (12) for the doctrinal paradox agenda and profile

outcome pqr . To read the table, note that scoring (12) is given by $s_+(t) = 2$ and $s_-(t) = -2$ for all t in $X \setminus \{\neg r\}$, $s_+(\neg r) = 1$ and $s_-(\neg r) = -3$.

How does our scoring rule ‘repair’ those special quota rules which use a uniform threshold $m \equiv m_p$ ($p \in X$), such as majority rule?

Remark 3 *For a uniform threshold $m \equiv m_p$, the scoring rule w.r.t. scoring (12) is the Kemeny rule, or equivalently, the simple scoring rule.*

This remark follows from Proposition 1 and the fact that, for a uniform threshold $m \equiv m_p$, scoring (12) is equivalent to simple scoring by footnote 14.

Finally, I note that the scoring rule w.r.t. (12) is not the only scoring rule which can ‘repair’ the quota rule $F_{(m_p)_{p \in X}}$ – though it might be the most plausible one, as long as we do not wish to introduce additional parameters. If, however, we are prepared to introduce additional parameters, scoring (12) can be generalized: for each $p \in X$ let $\alpha_p > 0$ be a coefficient measuring how important it is that the scoring rule is faithful to the quota rule’s collective judgment on p ; and let scoring be defined by

$$s_J(p) = \begin{cases} s_+(p) = \alpha_p(n + 1 - m_p) & \text{if } p \in J \\ s_-(p) = -\alpha_p m_p & \text{if } p \notin J. \end{cases} \quad (13)$$

The earlier scoring (12) is obviously a special case in which all α_p are 1. Proposition 5 still holds for this generalized kind of propositionwise scoring. The scoring rule will

²⁴This follows from Nehring and Puppe’s (2010) *intersection property*, generalized to possibly incomplete collective judgment sets (Dietrich and List 2007b).

tend to match the quota rule on propositions p with high importance coefficient α_p , while modifying (‘repairing’) the quota rule at propositions p with low α_p .

3.7 Premise- and conclusion-based aggregation

I have just mentioned the possibility of a differential treatment of propositions when ‘repairing’ a quota rule. This possibility is particularly salient in the popular context of premise- or conclusion-based aggregation.²⁵ One may indeed view the classical premise- and conclusion-based rules as two (rival) ways of repairing the simplest of all quota rules – majority rule – by privileging certain propositions over others, namely premise propositions or conclusion propositions, respectively.

Let me put this precisely. Consider majority voting, i.e., the quota rule with a uniform majority threshold $m \equiv m_p$ (the smallest integer above $n/2$). To restore collective rationality, we again endow each proposition $p \in X$ with a ‘coefficient of importance’, but now let this coefficient be determined by whether p has a ‘premise’ or ‘conclusion’ status. Formally, suppose the agenda is partitioned into two negation-closed sets, the set P of ‘premise propositions’ and the set $X \setminus P$ of ‘conclusion propositions’. In the case of our doctrinal paradox agenda $X = \{p, q, r\}^\pm$, we have $P = \{p, q\}^\pm$. Each premise proposition $p \in P$ has the importance coefficient $\alpha_p \equiv \alpha_{\text{premise}}$, and each conclusion proposition $p \in X \setminus P$ has the importance coefficient $\alpha_p \equiv \alpha_{\text{conclusion}}$, for fixed parameters $\alpha_{\text{premise}}, \alpha_{\text{conclusion}} \geq 0$. In this scenario, the scoring (13) becomes equivalent (by footnote 14) to the scoring given by

$$s_J(p) = \begin{cases} \alpha_{\text{premise}} & \text{for accepted premise propositions } p \in J \cap P \\ \alpha_{\text{conclusion}} & \text{for accepted conclusion propositions } p \in J \setminus P \\ 0 & \text{for rejected propositions } p \notin J. \end{cases} \quad (14)$$

By calibrating the two importance coefficients, we can influence the relative weights of premises and conclusions. If we give *far more* importance to premise propositions ($\alpha_{\text{premise}} \gg \alpha_{\text{conclusion}}$) or to conclusion propositions ($\alpha_{\text{conclusion}} \gg \alpha_{\text{premise}}$), the scoring rule reduces to the premise- or conclusion-based rule, respectively. To substantiate this claim, one needs to define both rules. For simplicity, I restrict attention to our doctrinal paradox agenda $X = \{p, q, r\}^\pm$ with $P = \{p, q\}^\pm$ (though more general X and P could be considered²⁶). In this case, assuming for simplicity that the group size n is odd,

- the *premise-based rule* is the aggregation rule which for each profile in \mathcal{J}^n delivers the (unique) judgment set in \mathcal{J} containing each premise proposition accepted by a majority;
- the *conclusion-based rule* is the aggregation rule which for each profile in \mathcal{J}^n delivers the judgment set (or sets) in \mathcal{J} containing the conclusion proposition accepted by a majority.²⁷

These two rules have the following characterizations as scoring rules:

²⁵See for instance List (2004), Dietrich and Mongin (2010) and Nehring and Puppe (2010b).

²⁶Our analysis generalizes easily to any X and P such that (i) the premise propositions in P are logically independent, and (ii) complete judgments across the premise propositions in P uniquely determine the judgments on the conclusion propositions in $X \setminus P$.

²⁷In the literature, the conclusion-based procedure is usually taken to be *silent* on the premises, i.e., to return an incomplete judgment set not in \mathcal{J} . I have replaced this silence by a tie between all compatible judgments on the premise propositions.

Remark 4 For our doctrinal paradox agenda $X = \{p, q, r\}^\pm$ with set of premise propositions $P = \{p, q\}^\pm$, and for an odd group size, the scoring rule w.r.t. scoring (14) is

- the premise-based rule if and only if $\alpha_{\text{premise}} > (n - 2)\alpha_{\text{conclusion}}$,
- the conclusion-based rule if and only if $\alpha_{\text{conclusion}} > \alpha_{\text{premise}} = 0$.

This result lets premise- and conclusion-based aggregation appear in a rather extreme light: each rule is based on somewhat unequal importance coefficients α_{premise} and $\alpha_{\text{conclusion}}$, deeming one type of proposition to be overwhelmingly more important than the other. It might therefore be interesting to consider more equilibrated values of the importance coefficients, so as to achieve a compromise between democracy at the premise level and democracy at the conclusion level.

4 Set scoring rules: assigning scores to entire judgment sets

An interesting generalization of scoring rules is obtained by assigning scores directly to entire judgment sets rather than single propositions. A *set scoring function* – or simply *set scoring* – is a function σ which to every pair of rational judgment sets C and J assigns a real number $\sigma_J(C)$, the *score* of C given J , which measures how well C performs (‘scores’) from the perspective of holding the judgment set J . Formally, $\sigma : \mathcal{J} \times \mathcal{J} \rightarrow \mathbb{R}$. The most elementary example, to be called *naive* set scoring, is given by

$$\sigma_J(C) = \begin{cases} 1 & \text{if } C = J \\ 0 & \text{if } C \neq J. \end{cases} \quad (15)$$

Any set scoring σ gives rise to an aggregation rule F_σ , the *set scoring rule* (or *generalized scoring rule*) w.r.t. σ , which for each profile $(J_1, \dots, J_n) \in \mathcal{J}^n$ selects the collective judgment set(s) C in \mathcal{J} having maximal sum-total score across individuals:

$$F_\sigma(J_1, \dots, J_n) = \operatorname{argmax}_{C \in \mathcal{J}} \sum_{i \in N} \sigma_{J_i}(C).$$

An aggregation rule is a *set scoring rule* simpliciter if it is the set scoring rule w.r.t. some set scoring σ . Set scoring rules generalize ordinary scoring rules, since to any ordinary scoring s corresponds a set scoring σ , given by

$$\sigma_J(C) \equiv \sum_{p \in C} s_J(p),$$

and the ordinary scoring rule w.r.t. s coincides with the set scoring rule w.r.t. σ .

4.1 Naive set scoring and plurality voting

The *plurality rule* is the aggregation rule F which for every profile $(J_1, \dots, J_n) \in \mathcal{J}^n$ declares the most often submitted judgment set(s) as the collective judgment set(s):

$$\begin{aligned} F(J_1, \dots, J_n) &= \text{most frequently submitted judgment set(s)} \\ &= \operatorname{argmax}_{C \in \mathcal{J}} |\{i : J_i = C\}|. \end{aligned}$$

This rule is of course normatively questionable;²⁸ but it deserves our attention, if only because of its simplicity and the recognized importance of plurality voting in social choice theory more broadly. The plurality rule can be construed as a set scoring rule:

Remark 5 *The naive set scoring rule is the plurality rule.*

4.2 Distance-based set scoring

Set scoring rules generalize distance-based aggregation. Given an arbitrary distance function d over \mathcal{J} (not necessarily the Kemeny-distance), all that is needed is to consider what I call *distance-based* set scoring, defined by

$$\sigma_J(C) = -d(C, J). \tag{16}$$

So, C scores high if it is close to the judgment set held, J . This renders sum-score-maximization equivalent to sum-distance-minimization:

Remark 6 *For every given distance function over \mathcal{J} , the distance-based set scoring rule is the distance-based rule.*

So, all distance-based rules can be modelled as set scoring rules (but not vice versa²⁹). As an example, consider the so-called *discrete* distance,³⁰ defined by

$$d(J, K) = \begin{cases} 0 & \text{if } J = K \\ 1 & \text{if } J \neq K. \end{cases}$$

Here, distance-based set scoring (16) is equivalent to naive set scoring (15), since the two differ only by a constant (of one). So, joining Remarks 5 and 6, we may view the plurality rule either as the naive set scoring rule or as the discrete-distance-based rule.

4.3 Averaging rules

Given an ordinary scoring s , we have so far aimed for collective judgments with a high total score. But this is not the only plausible aim or approach. We now turn to an altogether different approach. Rather than using s to assign scores only from each individual's perspective, we now care about how propositions score under the collective judgment set. Instead of wanting the collective judgments to achieve the highest total score from individuals, we now want them to resemble the 'average individual judgments' in the sense that the collective judgments should lead (approximately) to the same scores of propositions as the individual judgments do on average. In short, any proposition p 's collective score should be (approximately) p 's average individual score. This approach has its own, rather different intuitive appeal. But is it really totally different? As will turn out, aggregation

²⁸It ignores the internal structure of judgment sets, hence 'throws away' much information.

²⁹In trying to re-model an arbitrary set scoring rule F_σ as a distance-based rule, one might be tempted to define the 'distance' between J and J' as $d_\sigma(J, J') := \sigma_J(J) - \sigma_J(J')$. If d_σ turns out to define a proper distance function (see fn. 15), then we obtain a distance-based rule F_{d_σ} , which coincides with the set scoring rule F_σ . But for many plausible set scorings σ , d_σ has little in common with a distance function, violating up to all three axioms, notably symmetry and the triangle inequality.

³⁰This metric derives its name from the fact that it induces the discrete topology on whatever set it is defined on (such as \mathbb{R} instead of \mathcal{J}).

rules which follow this approach – I call them ‘averaging rules’ as opposed to ‘scoring rules’ – can be viewed as a particular kind of set scoring rules. This result is essentially a special case of a powerful precursor result by Zwicker (2008), as Marcus Pivato kindly pointed out to me. A comparison with Zwicker (2008) is given at the end of this subsection.

Given an ordinary scoring s , we can represent judgment sets in \mathcal{J} as vectors in \mathbb{R}^X , by identifying each judgment set J in \mathcal{J} with its *score vector*, i.e., the vector in \mathbb{R}^X whose p^{th} component is the score of p , $s_J(p)$.³¹ The score vector corresponding to $J \in \mathcal{J}$ is denoted $J^s \equiv (s_J(p))_{p \in X} \in \mathbb{R}^X$. Having represented judgment sets as vectors of numbers, we can apply standard algebraic and geometric operations, such as adding judgment sets, taking their average, or measuring their distance – where, of course, sums or averages of (score vectors of) judgment sets in \mathcal{J} may be ‘infeasible’, i.e., not correspond to any judgment set in \mathcal{J} .

The *averaging rule* w.r.t. scoring s is defined as the aggregation rule F which for every profile $(J_1, \dots, J_n) \in \mathcal{J}^n$ chooses the collective judgment set(s) whose score vector comes closest to the group’s average score vector $\frac{1}{n} \sum_{i \in N} J_i^s$ in the sense of Euclidean distance in \mathbb{R}^X :

$$F(J_1, \dots, J_n) = \text{j.s. closest to the average individual j.s. in score vector terms} \\ = \operatorname{argmin}_{C \in \mathcal{J}} \left\| C^s - \frac{1}{n} \sum_{i \in N} J_i^s \right\|.$$

Viewed geometrically as an operation in \mathbb{R}^X , the collective score vector is the projection of the average score vector $\frac{1}{n} \sum_i J_i^s$ on the set $\mathcal{J}^s \equiv \{J^s : J \in \mathcal{J}\} \subseteq \mathbb{R}^X$ of feasible score vectors.³²

As an illustration, consider once again reversal scoring for our doctrinal paradox example. Table 6 reports the score vector of each judgment set (including the one not submitted

	p	$\neg p$	q	$\neg q$	r	$\neg r$	distance to group’s average
pqr (indiv. 1)	2	0	2	0	2	0	$\sqrt{58}/3 \approx 2.54$
$p\neg q\neg r$ (indiv. 2)	1	0	0	2	0	2	$\sqrt{37}/3 \approx 2.03$
$\neg pq\neg r$ (indiv. 3)	0	2	1	0	0	2	$\sqrt{37}/3 \approx 2.03$
$\neg p\neg q\neg r$ (no indiv.)	0	1	0	1	0	3	$7/3 \approx 2.33$
group’s average	1	$\frac{2}{3}$	1	$\frac{2}{3}$	$\frac{2}{3}$	$\frac{4}{3}$	0

Table 6: The averaging rule (w.r.t. reversal scoring) for the doctrinal paradox agenda and profile

by any individual), and its distance to the group’s average score vector. By minimizing this distance, the rule delivers a tie between the two conclusion-based outcomes $p\neg q\neg r$ and $\neg pq\neg r$. The premise-based outcome pqr looks worse than ever: it is even farther from the average than the never-submitted outcome $\neg p\neg q\neg r$.

Note that we now have two rival ways of aggregating based on a scoring s : the scoring rule and the averaging rule. Can any connection be established? In fact, the averaging

³¹This identification is one-to-one as long as the scoring has the (very plausible) property that $s_J(p) > s_J(\neg p)$ whenever $p \in J$.

³²Formally, $F(J_1, \dots, J_n)^s = \operatorname{PROJ}_{\mathcal{J}^s}(\frac{1}{n} \sum_i J_i^s)$, where the projection of $x \in \mathbb{R}^X$ on $Y \subseteq \mathbb{R}^X$ is defined as $\operatorname{PROJ}_Y(x) := \operatorname{argmin}_{y \in Y} \|y - x\|$.

rule can be construed as a *set* scoring rule, namely in virtue of the set scoring given by

$$\sigma_J(C) = -\|C^s - J^s\|^2. \quad (17)$$

Here, C is taken to score high if it is close to J in terms of the *squared* Euclidean distance of score vectors.

Proposition 6 *For any scoring s , the averaging rule w.r.t. s is the set scoring rule w.r.t. set scoring (17).*

As an application, let s be simple scoring (2). Here, the set scoring (17) is expressible as an increasing affine transformation of the set scoring corresponding to simple scoring, i.e., of the set scoring σ' given by³³

$$\sigma'_J(C) = \sum_{p \in C} s_J(p) = |C \cap J|.$$

So, the set scoring rule F_σ coincides with the simple scoring rule F_s , and hence with the Kemeny rule $F_{d_{\text{Kemeny}}}$ by Proposition 1. Thus, as a corollary of Propositions 1 and 6, the Kemeny rule can be characterized not just as a scoring rule but also as an averaging rule, both times using the same scoring:

Corollary 2 *The Kemeny rule is both the scoring rule and the averaging rule w.r.t. simple scoring.*

Our averaging rules are special cases of Zwicker’s (2008) *mean proximity rules*. Like an averaging rule, a mean proximity rule represents each possible vote as well as each possible output as a vector of characteristics (real numbers), and determines the winning output(s) by minimizing the distance – in that vector representation – to the average vote. In our case, this vector is $|X|$ -dimensional and contains the support (score) given to each proposition in X . In Zwicker’s more general case, the vector could have any dimensionality and could contain any sorts of characteristics. Moreover, the votes need not be judgment sets, but could for instance be political candidates, characterized, say, by a three-dimensional vector containing age, number of pre-election speeches, and political orientation on a left-right spectrum. One may therefore view our averaging rules as a concrete way to implement Zwicker’s general approach inside judgment aggregation theory: we offer a concrete method of choosing the vector representation of votes, which is left open in Zwicker’s approach.³⁴ Alternative vector representations are of course also imaginable.

4.4 Probability-based set scoring

I close the analysis by taking a brief (skippable) excursion into an important, but different approach to judgment aggregation: the *epistemic* or *truth-tracking* approach. In this approach, each proposition $p \in X$ is taken to have an objective, but unknown truth value (‘true’ or ‘false’), and the goal of aggregation is to track the truth, i.e., to generate true

³³Since $\sigma_J(C) = -\left(\sqrt{|C \Delta J|}\right)^2 = -|C \Delta J| = -2|C \setminus J| = -2(|C| - |C \cap J|) = -|X| + 2|C \cap J|$.

³⁴Zwicker’s Theorem 4.2.1 (more precisely, its proof) essentially implies our Proposition 6 after some translation work.

collective judgments.³⁵ The truth-tracking perspective has a long history elsewhere in social choice theory (e.g., Condorcet 1785, Grofman *et al.* 1983, Austen-Smith and Banks 1996, Dietrich 2006b, Pivato 2011a); but within judgment aggregation theory specifically, rather little work has been done on the epistemic side (e.g., Bovens and Rabinowicz 2006, List 2005, Bozbay *et al.* 2011).

The epistemic approach warrants the use of particular set scoring rules. To show this, I import standard statistical estimation techniques (such as maximum-likelihood estimation), following the path taken by other authors in the context of preference aggregation (e.g., Young 1995) and other aggregation problems (e.g., Dietrich 2006b, Pivato 2011a). My goal is to give no more than a brief introduction to what could be done. The results given below are essentially variants of existing results; see in particular Pivato (2011a).³⁶

For each combination $(J_1, \dots, J_n, T) \in \mathcal{J}^n \times \mathcal{J}$ of $n+1$ judgment sets, let $\Pr(J_1, \dots, J_n, T) > 0$ measure the probability that people submit the profile (J_1, \dots, J_n) and the set of true propositions is T , where of course $\sum_{(J_1, \dots, J_n, T) \in \mathcal{J}^n \times \mathcal{J}} \Pr(J_1, \dots, J_n, T) = 1$. From this joint probability function we can, as usual, derive various marginal and conditional probabilities, such as the probability that the truth is $T \in \mathcal{J}$, $\Pr(T) = \sum_{(J_1, \dots, J_n) \in \mathcal{J}^n} \Pr(J_1, \dots, J_n, T)$, the probability that the profile is (J_1, \dots, J_n) , $\Pr(J_1, \dots, J_n) = \sum_{T \in \mathcal{J}} \Pr(J_1, \dots, J_n, T)$, the conditional probability $\Pr(T|J_1, \dots, J_n) = \frac{\Pr(J_1, \dots, J_n, T)}{\Pr(J_1, \dots, J_n)}$ (called the *posterior* probability of T given the ‘data’ J_1, \dots, J_n), and the conditional probability $\Pr(J_1, \dots, J_n|T) = \frac{\Pr(J_1, \dots, J_n, T)}{\Pr(T)}$ (called the *likelihood* of the ‘data’ J_1, \dots, J_n given T).

The *maximum-likelihood rule* is the aggregation rule $F : \mathcal{J}^n \rightrightarrows \mathcal{J}$ which for each profile $(J_1, \dots, J_n) \in \mathcal{J}^n$ defines the collective judgments such that their truth would make the observed profile (‘data’) maximally likely:

$$F(J_1, \dots, J_n) = \operatorname{argmax}_{T \in \mathcal{J}} \Pr(J_1, \dots, J_n|T).$$

The *maximum-posterior rule* is the aggregation rule $F : \mathcal{J}^n \rightrightarrows \mathcal{J}$ which for each profile $(J_1, \dots, J_n) \in \mathcal{J}^n$ defines the collective judgments such that they have maximal posterior probability of truth conditional on the observed profile (‘data’):

$$F(J_1, \dots, J_n) = \operatorname{argmax}_{T \in \mathcal{J}} \Pr(T|J_1, \dots, J_n).$$

Both of these rules correspond to well-established statistical estimation procedures.

Let us now make two standard, but restrictive assumptions on probabilities. We assume that voters are ‘independent’ and ‘equally competent’ (in analogy to the assumptions of Condorcet’s classical jury theorem³⁷). Formally, for every $T \in \mathcal{J}$,

(IND) the individual judgment sets are independent conditional on T being the true judgment set, i.e., $\Pr(J_1, \dots, J_n|T) = \Pr(J_1|T) \cdots \Pr(J_n|T)$ for all $J_1, \dots, J_n \in \mathcal{J}$ (‘*independence*’)

(COM) for each $J \in \mathcal{J}$, each individual has the same probability, denoted $\Pr(J|T)$, of

³⁵The epistemic perspective is usually contrasted with the *procedural* perspective, which takes the goal of aggregation to be to generate collective judgments which reflect the individuals’ judgments in a procedurally fair way. To illustrate the contrast between the two perspectives, suppose that all individuals hold the same judgment set J . Then J is clearly the right collective judgment set from the perspective of procedural fairness. But from an epistemic perspective, all depends on whether people’s unanimous endorsement of J is sufficient evidence for J being true.

³⁶Proposition 7 follows from proofs in Pivato (2011a), and is also related to Dietrich (2006b).

³⁷The classical Condorcet jury theorem is essentially concerned with a simple judgment aggregation problem with a binary agenda $X = \{p, \neg p\}$.

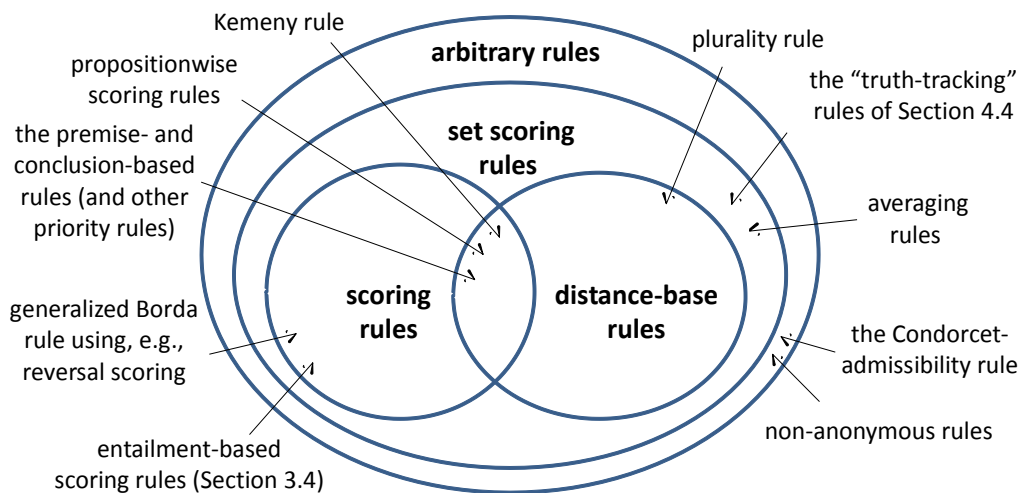


Figure 1: A map of judgment aggregation possibilities

submitting the judgment set J conditional on T being the true judgment set (*‘equal competence’*).

Condition (COM) in particular implies that individuals have the same (conditional) probability of holding the *true* judgment set; but nothing is assumed about the size of this probability of ‘getting it right’. The just-defined aggregation rules turn out to be set scoring rules in virtue of defining the score of $T \in \mathcal{J}$ given $J \in \mathcal{J}$ by, respectively,

$$\sigma_J(T) = \log \Pr(J|T) \tag{18}$$

$$\sigma_J(T) = \log \Pr(J|T) + \frac{1}{n} \log \Pr(T). \tag{19}$$

Proposition 7 *If voters are independent (IND) and equally competent (COM), then*

- *the maximum-likelihood rule is the set scoring rule w.r.t. set scoring (18),*
- *the maximum-posterior rule is the set scoring rule w.r.t. set scoring (19).*

5 Concluding remarks

I hope to have convinced the reader that scoring rules, and more generally set scoring rules, form interesting positive solutions to the judgment aggregation problem. They for instance allow us to generalize Borda aggregation to judgment aggregation (the simplest method being to use reversal scoring). Figure 1 summarizes where we stand, by depicting different classes of rules (scoring rules, set scoring rules, and distance-based rules) and positioning several concrete rules (such as the Kemeny rule). While the positions of most rules in Figure 1 have been established above or follow easily, a few positions are of the order of conjectures. This is so for the placement of our Borda generalization *outside* the class of distance-based rules.³⁸

³⁸For technical correctness, I also note two details about how to read Figure 1. First, for trivial agendas, such as a single-issue agenda $X = \{p, \neg p\}$, several rules of course become equivalent, and distinctions

Though several old and new aggregation rules are scoring rules (or at least set scoring rules), there are important counterexamples. One counterexample is the mentioned rule introduced by Nehring *et al.* (2011) (the so-called Condorcet-admissibility rule, which generates rational judgment set(s) that ‘approximate’ the majority judgment set). Other counterexamples are non-anonymous rules (such as rules prioritizing experts), and rules that return boundedly rational collective judgments (such as rules returning incomplete but still consistent and deductively closed judgments). The last two kinds of counterexamples suggest two generalizations of the notion of a scoring rule. Firstly, scoring might be allowed to depend on the individual; this leads to ‘non-anonymous scoring rules’. Secondly, the search for a collective judgment set with maximal total score might be done within a larger set than the set \mathcal{J} of fully rational judgment sets (such as the set of consistent but possibly incomplete judgment sets); this leads to ‘boundedly rational scoring rules’. The same generalizations could of course be made for set scoring rules.

Many challenges lie ahead of us. For instance, it would be interesting to characterize scoring rules and set scoring rules by some salient axiomatic properties.³⁹ And, in the epistemic context of searching for objectively ‘true’ or ‘correct’ judgments, one might give a maximum-likelihood rationalization of (set) scoring rules, by proving that they are precisely the aggregation rules which generate maximum-likelihood estimators of the truth relative to *some* probability model from a suitable class. Clearly, the paper raises more questions than it answers.

6 References

- Austen-Smith, D., Banks, J. (1996) Information Aggregation, Rationality, and the Condorcet Jury Theorem, *American Political Science Review* 90: 34–45
- Bovens, L., Rabinowicz, W. (2006) Democratic answers to complex questions: an epistemic perspective, *Synthese* 150(1): 131-153
- Bozbay, I., Dietrich, F., Peters, H. (2011) Judgment aggregation in search for the truth, working paper, London School of Economics
- Condorcet, Marquis de (1785) *Essai sur l'application de l'analyse á la probabilité des décisions rendues á la pluralité des voix*
- Conitzer, V., Rognlie, M., Xia, L. (2009) Preference functions that score rankings and maximum likelihood estimation, *IJCAI'09 Proceedings of the 21st international joint conference on Artificial intelligence*, p. 109-115, Morgan Kaufmann Publishers Inc. San Francisco
- Dietrich, F. (2006a) Judgment aggregation: (im)possibility theorems, *Journal of Economic Theory* 126(1): 286-298
- Dietrich, F. (2006b) General representation of epistemically optimal procedures, *Social Choice and Welfare* 26(2): 263-283
- Dietrich, F. (2007) A generalised model of judgment aggregation, *Social Choice and Wel-*

drawn in Figure 1 disappear. More precisely, by positioning a rule outside a class of rules (e.g., by positioning plurality rule outside the class of scoring rules), I am of course not implying that for *all* agendas the rule does not belong to the class, but that for some (in fact, *most*) agendas this is so. Second, in placing propositionwise scoring rules among the distance-based rules, I made a very plausible restriction: $s_+(p) > s_-(p)$ for each $p \in X$.

³⁹To this end, one might re-cast these rules in a setup with anonymous variable-population profiles, in which it turns out that these rules satisfy analogues of the classical axioms of ‘consistency’ (also called ‘reinforcement’) and ‘continuity’ (see Smith 1973, Young 1975 and Myerson 1995).

- fare 28(4): 529-565
- Dietrich, F. (2010) The possibility of judgment aggregation on agendas with subjunctive implications, *Journal of Economic Theory* 145(2): 603-638
- Dietrich, F., List, C. (2007a) Arrow's theorem in judgment aggregation, *Social Choice and Welfare* 29(1): 19-33
- Dietrich, F., List, C. (2007b) Judgment aggregation by quota rules: majority voting generalized, *Journal of Theoretical Politics* 19(4): 391-424
- Dietrich, F., List, C. (2010) Majority voting on restricted domains, *Journal of Economic Theory* 145(2): 512-543
- Dietrich, F., Mongin, P. (2010) The premise-based approach to judgment aggregation, *Journal of Economic Theory* 145(2): 562-582
- Dokow, E., Holzman, R. (2010) Aggregation of binary evaluations, *Journal of Economic Theory* 145: 495-511
- Duddy, C., Piggins, A. (2012) A measure of distance between judgment sets, *Social Choice And Welfare* 39: 855-867
- Eckert, D., Klamler, C. (2009) A geometric approach to paradoxes of majority voting: from Anscombe's paradox to the discursive dilemma with Saari and Nurmi, *Homo Oeconomicus* 26: 471-488
- Grofman, B., Owen, G., Feld, S. L. (1983) Thirteen Theorems in Search of the Truth, *Theory and Decision* 15: 261-278
- Hartmann, S., Pigozzi, G., Sprenger, J. (2010) Reliable methods of judgment aggregation, *Journal for Logic and Computation* 20: 603-617
- Konieczny, S., Pino-Perez, R. (2002) Merging information under constraints: a logical framework, *Journal of Logic and Computation* 12: 773-808
- Kornhauser, L. A., Sager, L. G. (1986) Unpacking the Court. *Yale Law Journal* 96(1): 82-117
- Lang, J., Pigozzi, G., Slavkovik, M., van der Torre, L. (2011) "Judgment aggregation rules based on minimization". In *Proceedings of the 13th Conference on the Theoretical Aspects of Rationality and Knowledge (TARK XIII)*, ACM, pp. 238-246
- List, C. (2004) A Model of Path Dependence in Decisions over Multiple Propositions, *American Political Science Review* 98(3): 495-513
- List, C. (2005) The probability of inconsistencies in complex collective decisions, *Social Choice and Welfare* 24(1): 3-32
- List, C., Pettit, P. (2002) Aggregating sets of judgments: an impossibility result, *Economics and Philosophy* 18(1): 89-110
- List, C., Polak, B. eds. (2010) *Symposium on Judgment Aggregation*, *Journal of Economic Theory* 145(2)
- Miller, M. K., Osherson, D. (2008) Methods for distance-based judgment aggregation, *Social Choice and Welfare* 32(4): 575-601
- Myerson, R. B. (1995) Axiomatic derivation of scoring rules without the ordering assumption, *Social Choice and Welfare* 12 (1): 59-74
- Nehring, K., Pivato, M., Puppe, C. (2011) Condorcet admissibility: indeterminacy and path-dependence under majority voting on interconnected decisions, working paper, University of California at Davis
- Nehring, K., Puppe, C. (2010a) Abstract Arrowian Aggregation, *Journal of Economic Theory* 145: 467-494
- Nehring, K., Puppe, C. (2010b) Justifiable Group Choice, *Journal of Economic Theory* 145: 583-602

- Pettit, P. (2001) Deliberative Democracy and the Discursive Dilemma. *Philosophical Issues* 11: 268-299
- Pigozzi, G. (2006) Belief merging and the discursive dilemma: an argument-based account to paradoxes of judgment aggregation. *Synthese* 152(2): 285-298
- Pivato, M. (2011a) Voting rules as statistical estimators, Working Paper, Trent University, Canada
- Pivato, M. (2011b) Variable-population voting rules, Working Paper, Trent University, Canada
- Smith, J. H. (1973) Aggregation of preferences with variable electorate, *Econometrica* 41: 1027-41
- Wilson R (1975) On the Theory of Aggregation, *Journal of Economic Theory* 10: 89-99
- Young, H. P. (1975) Social choice scoring functions, *SIAM J. Appl. Math.* 28: 824-38
- Young, H. P. (1995) Optimal voting rules, *Journal of Economic Perspectives* 9 (1): 51-64
- Zwicker, W. (1991) The voters' paradox, spin, and the Borda count, *Mathematical Social Sciences* 22, 187-227
- Zwicker, W. (2008) Consistency without neutrality in voting rules: When is a vote an average? *Math. Comput. Modelling* 48: 1357-73
- Zwicker, W. (2011) Towards a "Borda count" for judgment aggregation, extended abstract, presented at the conference 'Judgment Aggregation and Voting', Karlsruhe Institute of Technology

7 Appendix: proofs

Throughout, the complement of a set $J \subseteq X$ is denoted by $\bar{J} := X \setminus J$.

Proof of Proposition 1. The Kemeny-distance between $J, C \in \mathcal{J}$ can be written as

$$d_{\text{Kemeny}}(J, C) = \frac{1}{2} |J \Delta C| = \frac{1}{2} (|X| - (|J \cap C| + |\bar{J} \cap \bar{C}|)).$$

Now, since J and C each contains exactly one member of each pair $\{p, \neg p\} \subseteq X$, we have $p \in J \cap C \Leftrightarrow \neg p \in \bar{J} \cap \bar{C}$, and so, $|J \cap C| = |\bar{J} \cap \bar{C}|$. Hence, $d_{\text{Kemeny}}(J, C) = \frac{1}{2} |X| - |J \cap C|$. So, for each profile $(J_1, \dots, J_n) \in \mathcal{J}^n$, minimizing $\sum_{i \in N} d_{\text{Kemeny}}(J_i, C)$ is equivalent to maximizing $\sum_{i \in N} |J_i \cap C|$. Hence, rewriting each $|J_i \cap C|$ as $\sum_{p \in C} s_{J_i}(p)$ where s is simple scoring (2), it follows that $F_{d_{\text{Kemeny}}}(J_1, \dots, J_n) = F_s(J_1, \dots, J_n)$. ■

Before proving Proposition 2, I start with a lemma.

Lemma 1 *Consider the preference agenda (for any finite set of alternatives A), any classical scoring SCO , and the scoring s given by (4). For all distinct $x, y \in A$ and all $J \in \mathcal{J}$,*

$$SCO_J(x) - SCO_J(y) = s_J(xPy) - s_J(yPx). \quad (20)$$

Proof. This follows easily from (4). ■

Two elements of a set of alternatives A are called *neighbours* w.r.t. a strict linear order \succ over A if they differ and no alternative in A is ranked strictly between them. In the case of the preference agenda (for a set of alternatives A), the strict linear order over A corresponding to any $J \in \mathcal{J}$ is denoted \succ_J .

Proof of Proposition 2. Consider the preference agenda X for a set of alternatives A of finite size k , and let SCO be any classical scoring. I show that $F_{SCO} = F_s$ for each scoring s satisfying (20), and hence for the scoring (4) (since it satisfies (20) by Lemma 1) and the scoring (3) (since a half times it satisfies (20)).

Consider any scoring s satisfying (20). Fix a profile $(J_1, \dots, J_n) \in \mathcal{J}^n$; I show $F_s(J_1, \dots, J_n) = F_{SCO}(J_1, \dots, J_n)$. The proof is in three claims.

Claim 1. For all $a, b \in A$ and $C, C' \in \mathcal{J}$, if $C \setminus C' = \{aPb\}$, then

$$\sum_{i \in N} SCO_{J_i}(a) - \sum_{i \in N} SCO_{J_i}(b) = \sum_{i \in N, p \in C} s_{J_i}(p) - \sum_{i \in N, p \in C'} s_{J_i}(p).$$

Consider $a, b \in A$ and $C, C' \in \mathcal{J}$ such that $C \setminus C' = \{aPb\}$. For each individual $i \in N$, we by (20) have

$$SCO_{J_i}(a) - SCO_{J_i}(b) = s_{J_i}(aPb) - s_{J_i}(bPa),$$

which, noting that $C' = (C \setminus \{aPb\}) \cup \{bPa\}$, implies that

$$SCO_{J_i}(a) - SCO_{J_i}(b) = \sum_{p \in C} s_{J_i}(p) - \sum_{p \in C'} s_{J_i}(p).$$

Summing over all individuals, the claim follows, q.e.d.

Claim 2. $F_s(J_1, \dots, J_n) \subseteq F_{SCO}(J_1, \dots, J_n)$.

Consider any $C \in F_s(J_1, \dots, J_n)$. We have to show that $C \in F_{SCO}(J_1, \dots, J_n)$, i.e., that for all distinct $x, y \in A$,

$$\sum_{i \in N} SCO_{J_i}(x) > \sum_{i \in N} SCO_{J_i}(y) \Rightarrow xPy \in C,$$

or equivalently,

$$yPx \in C \Rightarrow \sum_{i \in N} SCO_{J_i}(y) \geq \sum_{i \in N} SCO_{J_i}(x).$$

Said in yet another way, we have to show that

$$\sum_{i \in N} SCO_{J_i}(x_k) \geq \sum_{i \in N} SCO_{J_i}(x_{k-1}) \geq \dots \geq \sum_{i \in N} SCO_{J_i}(x_1),$$

where I have labelled the alternatives x_1, x_2, \dots, x_k such that $x_k \succ_C x_{k-1} \succ_C \dots \succ_C x_1$. Consider any $t \in \{1, \dots, k-1\}$, and write a for x_{t+1} and b for x_t . Let C' be the judgment set arising from C by replacing aPb with its negation bPa . Now $C' \in \mathcal{J}$; this is because a and b are neighbours w.r.t. \succ_C , which guarantees that C' corresponds to a strict linear order (namely to the same one as for C except that b now ranks above a). Since $C \in F_s(J_1, \dots, J_n)$, C has maximal sum-total score within \mathcal{J} ; in particular,

$$\sum_{i \in N, p \in C} s_{J_i}(p) \geq \sum_{i \in N, p \in C'} s_{J_i}(p),$$

which by Claim 1 implies the desired inequality,

$$\sum_{i \in N} SCO_{J_i}(a) \geq \sum_{i \in N} SCO_{J_i}(b), \text{ q.e.d.}$$

Claim 3. $F_{SCO}(J_1, \dots, J_n) \subseteq F_s(J_1, \dots, J_n)$.

Consider any $C \in F_{SCO}(J_1, \dots, J_n)$. To show that $C \in F_s(J_1, \dots, J_n)$, we consider an arbitrary $C' \in \mathcal{J} \setminus \{C\}$ and have to show that C has an at least as high sum-total score as C' :

$$\sum_{i \in N, p \in C} s_{J_i}(p) \geq \sum_{i \in N, p \in C'} s_{J_i}(p). \quad (21)$$

To prove this, we first transform C gradually into C' in $m \equiv |C' \setminus C|$ steps, where each step consists in a single judgment reversal, i.e., in the replacement of a single proposition xPy ($\in C \setminus C'$) by its negation yPx ($\in C' \setminus C$). This defines a sequence of judgment sets C_0, \dots, C_m , where $C_0 = C$ and $C_m = C'$, and where for each step $t \in \{1, \dots, m\}$ there is a proposition x_tPy_t such that $C_t = (C_{t-1} \setminus \{x_tPy_t\}) \cup \{y_tPx_t\}$. Note that $\{x_tPy_t : t = 1, \dots, m\} = C \setminus C'$. By a standard relation-theoretic argument, we may assume that in each step t the judgment reversal consists in switching the relative order of two *neighbouring* alternatives; i.e., x_t, y_t are neighbours w.r.t. the old and new relations $\succ_{C_{t-1}}$ and \succ_{C_t} . This guarantees that each step t generates a set C_t such that \succ_{C_t} is still a strict linear order, i.e., such that $C_t \in \mathcal{J}$.

Now for each step t , by Claim 1 we have

$$\sum_{i \in N} SCO_{J_i}(x_t) - \sum_{i \in N} SCO_{J_i}(y_t) = \sum_{i \in N, p \in C_{t-1}} s_{J_i}(p) - \sum_{i \in N, p \in C_t} s_{J_i}(p),$$

and also, since $y_tPx_t \notin C$ and $C \in F_{SCO}(J_1, \dots, J_n)$, we have

$$\sum_{i \in N} SCO_{J_i}(y_t) \leq \sum_{i \in N} SCO_{J_i}(x_t);$$

using Claim 1, it follows that

$$\sum_{i \in N, p \in C_{t-1}} s_{J_i}(p) - \sum_{i \in N, p \in C_t} s_{J_i}(p) \geq 0.$$

Summing this inequality over all steps $t \in \{1, \dots, m\}$, we obtain

$$\sum_{i \in N, p \in C_0} s_{J_i}(p) - \sum_{i \in N, p \in C_m} s_{J_i}(p) \geq 0,$$

which is equivalent to the desired inequality (21) since $C_0 = C$ and $C_m = C'$. ■

Proof of Remark 2. Let s' be defined from reversal scoring s in the specified way.

Claim 1. s' and s are equivalent.

Consider any profile $(J_1, \dots, J_n) \in \mathcal{J}^n$. I show for all $C, D \in \mathcal{J}$ that

$$\sum_{i \in N, p \in C} s_{J_i}(p) \geq \sum_{i \in N, p \in D} s_{J_i}(p) \Leftrightarrow \sum_{i \in N, p \in C} s'_{J_i}(p) \geq \sum_{i \in N, p \in D} s'_{J_i}(p).$$

Consider any $C, D \in \mathcal{J}$. I prove that $\Delta \geq 0 \Leftrightarrow \Delta' \geq 0$, where

$$\begin{aligned} \Delta &\equiv \sum_{i \in N, p \in C} s_{J_i}(p) - \sum_{i \in N, p \in D} s_{J_i}(p) \geq 0, \\ \Delta' &\equiv \sum_{i \in N, p \in C} s'_{J_i}(p) - \sum_{i \in N, p \in D} s'_{J_i}(p) \geq 0. \end{aligned}$$

We have

$$\Delta = \sum_{i \in N} \left\{ \sum_{p \in C} s_{J_i}(p) - \sum_{p \in D} s_{J_i}(p) \right\} = \sum_{i \in N} \left\{ \sum_{p \in C \setminus D} s_{J_i}(p) - \sum_{p \in D \setminus C} s_{J_i}(p) \right\}.$$

So, noting that $D \setminus C = \{\neg p : p \in C \setminus D\}$, we obtain

$$\Delta = \sum_{i \in N} \sum_{p \in C \setminus D} (s_{J_i}(p) - s_{J_i}(\neg p)).$$

By an analogous reasoning,

$$\Delta' = \sum_{i \in N} \sum_{p \in C \setminus D} (s'_{J_i}(p) - s'_{J_i}(\neg p)).$$

Hence, using the definition of s' ,

$$\begin{aligned} \Delta' &= \sum_{i \in N} \sum_{p \in C \setminus D} ([s_{J_i}(p) - s_{J_i}(\neg p)] - [s_{J_i}(\neg p) - s_{J_i}(p)]) \\ &= 2 \sum_{i \in N} \sum_{p \in C \setminus D} (s_{J_i}(p) - s_{J_i}(\neg p)) \\ &= 2\Delta. \end{aligned}$$

So, $\Delta \geq 0 \Leftrightarrow \Delta' \geq 0$, q.e.d.

Claim 2. If X is the preference agenda, SCO is classical Borda scoring, $J \in \mathcal{J}$, and $xPy \in X$, then $s'_J(xPy) = SCO_J(x) - SCO_J(y)$.

Let X , SCO , J and xPy be as specified. If $xPy \in J$, then

$$\begin{aligned} s'(xPy) &= s(xPy) \text{ by definition of } s' \\ &= SCO_J(x) - SCO_J(y) \text{ by Remark 1, as } xPy \in J. \end{aligned}$$

If $xPy \notin J$, i.e., $yPx \in J$, then

$$\begin{aligned} s'(xPy) &= -s(yPx) \text{ by definition of } s' \\ &= -(SCO_J(y) - SCO_J(x)) \text{ by Remark 1, as } yPx \in J \\ &= SCO_J(x) - SCO_J(y). \blacksquare \end{aligned}$$

Proof of Proposition 4. Let X be the preference agenda for some set of alternatives A of size $k < \infty$. Let s^{rev} , s^{dis} and s^{irr} be reversal, disjoint-entailment, and irreducible-entailment scoring, respectively. Consider any $J \in \mathcal{J}$, denote the corresponding strict linear order by \succ , let x_1, \dots, x_k be the alternatives in the order given by $x_k \succ x_{k-1} \succ \dots \succ x_1$, and consider any $p \in X$, say $p = x_i P x_{i'} \in X$.

Claim 1. $s_J^{\text{rev}}(p) = s_J^{\text{dis}}(p)$.

By the argument given in footnote 20, $s_J^{\text{rev}}(p) \geq s_J^{\text{dis}}(p)$. I now show that $s_J^{\text{dis}}(p) \geq s_J^{\text{rev}}(p)$. This inequality is trivial if $p \notin J$, since then $s_J^{\text{rev}}(p) = 0$ (as $\neg p \in J$). Now suppose $p \in J$. By Remark 1, $s_J^{\text{rev}}(p) = i - i'$. So we need to show that $s_J^{\text{dis}}(p) \geq i - i'$. Consider the $i - i'$ judgment subsets $S_1, \dots, S_{i-i'} \subseteq J$ defined as follows: for each $j \in \{1, \dots, i - i'\}$,

$$S_j \equiv \{x_i P x_{i-j}, x_{i-j} P x_{i'}\} \subseteq J,$$

where $S_{i-i'}$ is interpreted as the set $\{x_i P x_{i'}\}$ (rather than the set $\{x_i P x_{i'}, x_{i'} P x_{i'}\}$, which is not well-defined since $x_{i'} P x_{i'}$ is not a proposition in X). Since these judgment subsets are pairwise disjoint and each of them entails $p (= x_i P x_{i'})$, we have $s_J^{\text{dis}}(p) \geq i - i'$, q.e.d.

Claim 2. $s_J^{\text{rev}}(p) = s_J^{\text{irr}}(p)$.

If $p \notin J$, then $s_J^{\text{rev}}(p) = s_J^{\text{irr}}(p)$ since $s_J^{\text{rev}}(p) = 0$ (as $\neg p \in J$) and $s_J^{\text{irr}}(p) = 0$ (as J does not entail p). Now suppose $p \in J$. Then, as already mentioned, $s_J^{\text{rev}}(p) = i - i'$ by Remark 1. So we need to show that $s_J^{\text{irr}}(p) = i - i'$. As one may show, each of the just-defined sets $S_1, \dots, S_{i-i'}$ irreducibly entails $p (= x_i P x_{i'})$. So it remains to show that no other judgment subset irreducibly entails p . Suppose $S \subseteq J$ irreducibly entails p . I have to show that $S \in \{S_1, \dots, S_{i-i'}\}$. As is easily checked, the set $S \cup \{\neg p\}$ ($= S \cup \{x_{i'} P x_i\}$) is minimal inconsistent. Hence, this set is *cyclic*, i.e., of the form $S \cup \{\neg p\} = \{y_1 P y_2, y_2 P y_3, \dots, y_{m-1} P y_m, y_m P y_1\}$ for some $m \geq 2$ and some distinct alternatives $y_1, \dots, y_m \in A$ (see Dietrich and List 2010). Without loss of generality, assume $y_1 = x_i$ and $y_m = x_{i'}$, so that $y_m P y_1 = x_{i'} P x_i$ and

$$S = \{y_1 P y_2, y_2 P y_3, \dots, y_{m-1} P y_m\}.$$

If $m = 2$, then $S = \{y_1 P y_2\} = \{x_i P x_{i'}\}$, which equals $S_{i-i'}$, and we are done. If $m = 3$, then $S = \{y_1 P y_2, y_2 P y_3\} = \{x_i P y_2, y_2 P x_{i'}\}$. Since S is by assumption included in J , it follows that J ranks y_2 between x_i and $x_{i'}$. So there is a $j \in \{1, \dots, i - i' - 1\}$ such that $y_2 = x_{i-j}$. Hence, S is the set $\{x_i P x_{i-j}, x_{i-j} P x_{i'}\} = S_j$, and we are done again. Finally, m cannot exceed 3, since otherwise the set $S (= \{x_i P y_2, y_2 P y_3, \dots, y_{m-1} P x_{i'}\})$ would entail $p (= x_i P x_{i'})$ *non-irreducibly*, since the set arising from S by replacing $x_i P y_2$ and $y_2 P y_3$ with their implication $x_i P y_3$ still entails p . ■

Proof of Proposition 5. Consider any threshold family $(m_p)_{p \in X}$ ($\in \{1, \dots, n\}^X$), and define scoring s by (12). Consider a profile $(J_1, \dots, J_n) \in \mathcal{J}^n$ for which $C^* \equiv F_{(m_p)_{p \in X}}(J_1, \dots, J_n)$ belongs to \mathcal{J} . We have to show that $F_s(J_1, \dots, J_n) = C^*$. For each proposition $p \in X$, writing the number of individuals accepting p as $n_p \equiv |\{i : p \in J_i\}|$, the sum-total score of p is given by

$$\begin{aligned} \sum_{i \in N} s_{J_i}(p) &= \sum_{i \in N: p \in J_i} (n + 1 - m_p) + \sum_{i \in N: p \notin J_i} (-m_p) \\ &= n_p(n + 1 - m_p) + (n - n_p)(-m_p) \\ &= n n_p + n_p - n m_p. \\ &= n(n_p - m_p) + n_p; \end{aligned}$$

and so,

$$\sum_{i \in N} s_{J_i}(p) \begin{cases} > 0 & \text{if } n_p \geq m_p, \text{ i.e., if } p \in C^* \\ < 0 & \text{if } n_p < m_p, \text{ i.e., if } p \notin C^*. \end{cases} \quad (22)$$

Now we have $\{C^*\} = \text{argmax}_{C \in \mathcal{J}} \sum_{p \in C, i \in N} s_{J_i}(p)$, because for each $C \in \mathcal{J} \setminus \{C^*\}$,

$$\sum_{p \in C^*, i \in N} s_{J_i}(p) - \sum_{p \in C, i \in N} s_{J_i}(p) = \underbrace{\sum_{p \in C^* \setminus C} \sum_{i \in N} s_{J_i}(p)}_{> 0 \text{ by (22)}} - \underbrace{\sum_{p \in C \setminus C^*} \sum_{i \in N} s_{J_i}(p)}_{< 0 \text{ by (22)}} > 0.$$

So, $F_s(J_1, \dots, J_n) = \{C^*\} \equiv C^*$. ■

Proof of Remark 4. Consider this X and P , let n be odd, and let s be scoring (14). I write α_{pr} for α_{premise} and α_{co} for $\alpha_{\text{conclusion}}$. Whenever I consider a profile $(J_1, \dots, J_n) \in \mathcal{J}^n$, I write $N_t := \{i : t \in J_i\}$ for all $t \in X$, and I write \mathcal{MAJ} , \mathcal{PRE} , \mathcal{CON} and \mathcal{SCO} for the outcome of the majority rule, premise-based rule, conclusion-based rule, and the scoring rule w.r.t. (14), respectively. Note that for all $(J_1, \dots, J_n) \in \mathcal{J}^n$ the sum-total score of a $C = \{p', q', r'\} \in \mathcal{J}$ (where $p' \in \{p, \neg p\}$, $q' \in \{q, \neg q\}$ and $r' \in \{r, \neg r\}$) is given by

$$\sum_{i \in N, t \in C} s_{J_i}(t) = (|N_{p'}| + |N_{q'}|)\alpha_{\text{pr}} + |N_r|\alpha_{\text{co}}. \quad (23)$$

Claim 1. $[\mathcal{PRE} = \mathcal{SCO}$ for all profiles in \mathcal{J}^n] if and only if $\alpha_{\text{pr}} > (n-2)\alpha_{\text{co}}$.

First, assume $\mathcal{PRE} = \mathcal{SCO}$ for all profile in \mathcal{J}^n . As one may check, there is a profile such that $|N_p| = |N_q| = \frac{n+1}{2}$ and $|N_r| = 1$. For this profile, $\mathcal{PRE} = \{p, q, r\}$. So, $\mathcal{SCO} = \{p, q, r\}$. Hence, the sum-total score of $\{p, q, r\}$ exceeds that of $\{\neg p, q, \neg r\}$. By (23), these two sum-total scores can be written, respectively, as

$$\begin{aligned} \sum_{i \in N, t \in \{p, q, r\}} s_{J_i}(t) &= \frac{n+1}{2}\alpha_{\text{pr}} + \frac{n+1}{2}\alpha_{\text{pr}} + \alpha_{\text{co}} = (n+1)\alpha_{\text{pr}} + \alpha_{\text{co}} \\ \sum_{i \in N, t \in \{\neg p, q, \neg r\}} s_{J_i}(t) &= \frac{n-1}{2}\alpha_{\text{pr}} + \frac{n+1}{2}\alpha_{\text{pr}} + (n-1)\alpha_{\text{co}} = n\alpha_{\text{pr}} + (n-1)\alpha_{\text{co}}. \end{aligned}$$

Hence,

$$(n+1)\alpha_{\text{pr}} + \alpha_{\text{co}} > n\alpha_{\text{pr}} + (n-1)\alpha_{\text{co}},$$

or equivalently, $\alpha_{\text{pr}} > (n-2)\alpha_{\text{co}}$.

Conversely, assume $\alpha_{\text{pr}} > (n-2)\alpha_{\text{co}}$. Consider any profile. We have to show that $\mathcal{PRE} = \mathcal{SCO}$.

Case 1: $\mathcal{MAJ} \in \mathcal{J}$. Check that it follows that $\mathcal{PRE} = \mathcal{MAJ}$, and also that $\mathcal{SCO} = \mathcal{MAJ}$. So, $\mathcal{PRE} = \mathcal{SCO}$.

Case 2: $\mathcal{MAJ} \notin \mathcal{J}$. Check that it follows that $\mathcal{MAJ} = \{p, q, \neg r\}$. Hence $\mathcal{PRE} = \{p, q, r\}$. We thus have to show that $\mathcal{SCO} = \{p, q, r\}$, i.e., that

$$\begin{aligned} \Delta_1 &\equiv \sum_{i \in N, t \in \{p, q, r\}} s_{J_i}(t) - \sum_{i \in N, t \in \{\neg p, q, \neg r\}} s_{J_i}(t) > 0, \\ \Delta_2 &\equiv \sum_{i \in N, t \in \{p, q, r\}} s_{J_i}(t) - \sum_{i \in N, t \in \{p, \neg q, \neg r\}} s_{J_i}(t) > 0, \\ \Delta_3 &\equiv \sum_{i \in N, t \in \{p, q, r\}} s_{J_i}(t) - \sum_{i \in N, t \in \{\neg p, \neg q, \neg r\}} s_{J_i}(t) > 0. \end{aligned}$$

By (23),

$$\Delta_1 = (|N_p| - |N_{\neg p}|)\alpha_{\text{pr}} + (|N_r| - |N_{\neg r}|)\alpha_{\text{co}} = (2|N_p| - n)\alpha_{\text{pr}} + (2|N_r| - n)\alpha_{\text{co}}. \quad (24)$$

In this, as $p \in \mathcal{MAJ}$ we have $|N_p| \geq (n+1)/2$; and further, as $p, q \in \mathcal{MAJ}$ the sets N_p and N_q each contain a majority, so that $N_p \cap N_q \neq \emptyset$, which (since $N_p \cap N_q \subseteq N_r$) implies $|N_r| \geq 1$. Using these lower bounds for $|N_p|$ and $|N_r|$, we obtain

$$\Delta_1 \geq ((n+1) - n)\alpha_{\text{pr}} + (2 - n)\alpha_{\text{co}} = \alpha_{\text{pr}} + (2 - n)\alpha_{\text{co}} > 0.$$

The proof that $\Delta_2 > 0$ is analogous. Finally, by (23),

$$\Delta_3 = (|N_p| - |N_{\neg p}|)\alpha_{\text{pr}} + (|N_q| - |N_{\neg q}|)\alpha_{\text{pr}} + (|N_r| - |N_{\neg r}|)\alpha_{\text{co}}.$$

Since $|N_q| > |N_{\neg q}|$ (as $q \in \mathcal{MAJ}$), it follows using (24) that $\Delta_3 > \Delta_1$, and hence, that $\Delta_3 > 0$, q.e.d.

Claim 2. $[\mathcal{CON} = \mathcal{SCO}$ for all profiles in J^n] if and only if $\alpha_{co} > \alpha_{pr} = 0$.

Unlike in the proof of the Claim, there may be ties, and so we treat \mathcal{CON} and \mathcal{SCO} as subsets of J , not elements. First, if $\alpha_{co} > \alpha_{pr} = 0$, then it is easy to show that $\mathcal{CON} = \mathcal{SCO}$ for each profile. Conversely, suppose it is not the case that $\alpha_{co} > \alpha_{pr} = 0$. Then either $\alpha_{co} = \alpha_{pr} = 0$ or $\alpha_{pr} > 0$. In the first case, clearly $\mathcal{CON} \neq \mathcal{SCO}$ for some profiles, since \mathcal{SCO} is always J . In the second case, again $\mathcal{CON} \neq \mathcal{SCO}$ for some profiles: for instance, if each individual submits $\neg pq \neg r$ then $\mathcal{SCO} = \{\neg pq \neg r\}$ while $\mathcal{CON} = \{\neg pq \neg r, p \neg q \neg r, \neg p \neg q \neg r\}$. ■

Proof of Proposition 6. It will sometimes be convenient to write a vector $D = (D_1, \dots, D_n) \in \mathbb{R}^n$ as $\langle D_i \rangle$. The mean and variance of this vector D are denoted and defined by, respectively,

$$\bar{D} \equiv \frac{1}{n} \sum_{i \in N} D_i \text{ and } Var(D) \equiv \frac{1}{n} \sum_{i \in N} (D_i - \bar{D})^2.$$

In this notation, the average square deviation of a constant $c \in \mathbb{R}$ from the components in D is $\overline{\langle (c - D_i)^2 \rangle}$ and satisfies

$$\overline{\langle (c - D_i)^2 \rangle} = (c - \bar{D})^2 + Var(D), \quad (25)$$

by the following argument borrowed from statistics:

$$\begin{aligned} \overline{\langle (c - D_i)^2 \rangle} &= \overline{\langle (c - \bar{D} + \bar{D} - D_i)^2 \rangle} \\ &= \overline{\langle (c - \bar{D})^2 + 2(c - \bar{D})(\bar{D} - D_i) + (\bar{D} - D_i)^2 \rangle} \\ &= (c - \bar{D})^2 + 2(c - \bar{D})\overline{\langle \bar{D} - D_i \rangle} + \overline{\langle (\bar{D} - D_i)^2 \rangle} \\ &= (c - \bar{D})^2 + 0 + Var(D). \end{aligned}$$

Now consider any scoring s and let the set scoring σ be defined by (17). Consider any profile $(J_1, \dots, J_n) \in J^n$ and any $C \in \mathcal{J}$. Under σ , the sum-total score of C can be written as

$$\begin{aligned} \sum_{i \in N} \sigma_{J_i}(C) &= - \sum_{i \in N} \|C^s - J_i^s\|^2 \\ &= - \sum_{i \in N} \sum_{p \in X} (C_p^s - J_{ip}^s)^2 \\ &= -n \sum_{p \in X} \frac{1}{n} \sum_{i \in N} (C_p^s - J_{ip}^s)^2. \end{aligned}$$

Here, the inner expression can be re-expressed as

$$\frac{1}{n} \sum_{i \in N} (C_p^s - J_{ip}^s)^2 = \overline{\langle (C_p^s - J_{ip}^s)^2 \rangle} = (C_p^s - \overline{\langle J_{ip}^s \rangle})^2 + Var(\langle J_{ip}^s \rangle),$$

where the last equality applies (25) with $c = C_p^s$ and $D = \langle J_{ip}^s \rangle$. It follows that

$$\begin{aligned} \sum_{i \in N} \sigma_{J_i}(C) &= -n \sum_{p \in X} \left\{ (C_p^s - \overline{\langle J_{ip}^s \rangle})^2 + \text{Var}(\langle J_{ip}^s \rangle) \right\} \\ &= -n \sum_{p \in X} (C_p^s - \overline{\langle J_{ip}^s \rangle})^2 + d \text{ (for some } d \text{ independent of } C) \\ &= -n \left\| C - \overline{\langle J_i^s \rangle} \right\|^2 + d. \end{aligned}$$

Maximizing this expression w.r.t. $C \in \mathcal{J}$ is equivalent to minimizing its strictly decreasing transformation $\left\| C - \overline{\langle J_i^s \rangle} \right\|$ w.r.t. $C \in \mathcal{J}$. So, the set scoring rule w.r.t. σ delivers the same collective judgment set(s) C as the averaging rule w.r.t. s . ■

Proof of Proposition 7. Assume (IND) and (COM) and consider a profile $(J_1, \dots, J_n) \in \mathcal{J}^n$.

Firstly, using (IND), the likelihood of the profile given $C \in \mathcal{J}$ can be written as

$$\Pr(J_1, \dots, J_n | T) = \prod_{i \in N} \Pr(J_i | T).$$

Maximizing this expression (w.r.t. $T \in \mathcal{J}$) is equivalent to maximizing its logarithm,

$$\sum_{i \in N} \log \Pr(J_i | T),$$

which is precisely the sum-total score of T under set scoring (18).

Secondly, writing π for the profile's probability $\Pr(J_1, \dots, J_n)$, the posterior probability of $T \in \mathcal{J}$ given the profile can be written as

$$\Pr(T | J_1, \dots, J_n) = \frac{1}{\pi} \Pr(T) \Pr(J_1, \dots, J_n | T) = \frac{1}{\pi} \Pr(T) \prod_{i \in N} \Pr(J_i | T).$$

Maximizing this expression (w.r.t. $T \in \mathcal{J}$) is equivalent to maximizing its logarithm, and hence, to maximizing

$$\log \Pr(T) + \sum_{i \in N} \log \Pr(J_i | T) = \sum_{i \in N} (\log \Pr(J_i | T) + \frac{1}{n} \log \Pr(T)),$$

which is the sum-total score of T under set scoring (19). ■