A Note on Embedding Voting Rules into Judgment Aggregation

Ulle Endriss
Institute for Logic, Language and Computation
University of Amsterdam

Working Paper, 28 October 2015 (this version: 12 December 2015)

Abstract

I propose a model of judgment aggregation in which we can distinguish rationality constraints (to be respected by the individual agents when supplying their judgments) from feasibility constraints (to be respected by the outcome returned by the judgment aggregation rule). This allows for a particularly natural way of embedding a wide variety of different voting rules into judgment aggregation.

1 Introduction

The field of judgment aggregation is concerned with the design and analysis of procedures for combining the judgments of several agents regarding the truth of a number of logically related statements into a collective judgment (List and Puppe, 2009; Grossi and Pigozzi, 2014; Endriss, 2016). Inspired by questions arising in legal theory (Kornhauser and Sager, 1993), it has received significant attention in both philosophy and economics, starting with the seminal contribution of List and Pettit (2002). More recently, due both to its potential for applications in areas such as multiagent systems and crowdsourcing, and due to the interesting algorithmic questions it raises, judgment aggregation is also receiving increasing attention in computer science and artificial intelligence (Endriss, 2016).

While early work on judgment aggregation has focused on applications of the axiomatic method, inspired by closely related work in other areas of social choice theory, to derive a host of impossibility and characterisation results (List and Puppe, 2009), more recently attention has shifted towards designing practically useful aggregation rules (Dietrich and List, 2007a; Miller and Osherson, 2009; Lang et al., 2011; Lang and Slavkovik, 2013; Dietrich, 2014; Nehring et al., 2014; Porello and Endriss, 2014; Endriss and Grandi, 2014; Costantini et al., 2016). A natural question that arises in this context is whether it is possible to translate known voting rules, i.e., rules that aggregate preferences over alternatives to return a winning alternative (Brams and Fishburn, 2002), into judgment aggregation rules. That this should be possible in principle is clear from the standard embedding of Arrovian preference aggregation into judgment aggregation (Dietrich and List, 2007b), under which we ask agents to pass judgments on preference statements such as $x \succcurlyeq y$. For some voting rules, this embedding suggests immediate generalisations to the case of judgment aggregation. Clear examples are the *Kemeny Rule* and the *Slater Rule* (Miller and Osherson, 2009), as well as the *Young Rule* and *Tideman's Ranked Pairs Rule* (Lang et al., 2011). But for others, such as the *Plurality*, the *Borda*, and the *Copeland Rule*, it is much less obvious how to generalise them to judgment aggregation. For example, while there are proposals for a Borda Rule for judgment aggregation (Dietrich, 2014; Duddy et al., 2015), these proposals are significantly more involved and arguably less uncontroversial than the direct embeddings for the other rules mentioned earlier.

A crucial step forward towards opening up new opportunities for embedding voting rules into judgment aggregation has recently been taken by Lang and Slavkovik (2013). While the simple known embeddings all involve an integrity constraint simulating the properties of preference orders to constrain the range of allowed outcomes of a judgment aggregation rule, they noted that by imposing a nonstandard integrity constraint instead, we can reach a wider range of voting rules.

In this paper, I argue that the most natural manner in which to develop this idea to its full potential is to work in a novel judgment aggregation framework that explicitly allows us to distinguish between *rationality constraints* (to be respected by the individual agents when supplying their judgments) and *feasibility constraints* (to be respected by the outcome returned by the judgment aggregation rule). In all existing frameworks of judgment aggregation, on the other hand, there is only a single type of constraint (sometimes explicitly represented and sometimes left implicit), governing what is permissible for both the input and the output.

The notions of rationality and feasibility are of great interest also outside the specific concern of recovering common voting rules in judgment aggregation and allow us to model a wider range of application scenarios than what has hitherto been possible. For example, we may consider it irrational for an individual to support both a law to subsidise local health education and a law providing tax breaks to international fast food chains, while we may very well accept this combination as a compromise when returned by an aggregation rule. At the same time, for the outcomes of an aggregation rule we may be bound by budget considerations, while we may not want to think of individuals as being concerned with such feasibility constraints when pondering their personal judgments.

Existing work on judgment aggregation uses either the term 'consistency' (see, e.g., List and Pettit, 2002) or 'rationality' (see, e.g., Grandi and Endriss, 2013) to define what individual judgments are permissible. Feasibility is not discussed as a separate concept in this literature. Rather, the common assumption is that also for the outcome we should want to impose the same rationality/consistency requirements. This desideratum is what is known as *collective rationality*. On the other hand, the concept of a feasibility constraint on aggregation outcomes is prominent in work on logic-based belief merging (Konieczny and Pino Pérez, 2002), a framework closely related to judgment aggregation. But in belief merging, one usually makes no assumptions regarding individual rationality (other than respecting the laws of classical logic). My proposal to distinguish rationality from feasibility constraints is conceptually similar to work of Porello (2014) who proposes to use different logical calculi to assess the consistency of individual judgments (what I call rationality) and the consistency of collective judgments (what I call feasibility).

The remainder of this paper is organised as follows. Section 2 introduces the framework of judgment aggregation with rationality and feasibility constraints. Section 3 then defines four judgment aggregation rules that all guarantee feasible outcomes and that are *majoritarian* in the sense of respecting the wishes of the majority as much as is possible without violating the feasibility constraint. In Section 4, these aggregation rules are then shown to reduce to a range of different voting rules when combined with specific rationality and feasibility constraints. Finally, Section 5 concludes with a short discus ion of possible future directions for work on the novel framework proposed here.

2 A Framework for Judgment Aggregation

I now introduce a framework for judgment aggregation that allows us to distinguish between rationality and feasibility constraints. It is essentially the same framework as that of *binary aggregation with integrity constraints* (Grandi, 2012; Grandi and Endriss, 2013), except that now there are *two* integrity constraints, one to model rationality and one to model feasibility. As is well known, all frameworks for judgment aggregation can in principle be translated into each other. Thus, the extension proposed here may instead also be applied to *formula-based judgment aggregation* (List and Pettit, 2002), including its generalisations to nonclassical logics (Dietrich, 2007), to the *binary aggregation* framework of Dokow and Holzman (2010), and to the *abstract aggregation* framework of Nehring and Puppe (2010).

2.1 Agendas and Agents

An *agenda* is a finite set Φ. Each agenda item $\varphi \in \Phi$ may be thought of as a question that can be answered with 'yes' or 'no', or as a proposition that can be accepted or rejected. A *judgment* is a function $J : \Phi \to \{0,1\}$, mapping each agenda item to either 0 (to indicate rejection) or 1 (to indicate acceptance). We often write $\{0,1\}^\Phi$ as a shorthand for $\Phi \to \{0,1\}$, the space of all possible judgments. By a slight abuse of notation, for any two judgments $J : \Phi \to \{0,1\}$ and $J' : \Phi \to \{0,1\}$, let their *intersection* $J \cap J' := \{\varphi \in \Phi \mid J(\varphi) = J'(\varphi)\}$ be defined as the set of agenda items on which they agree.

Let $\mathcal{N} = \{1,\ldots,n\}$ be a finite set of *agents*. For ease of exposition, in this paper we will always assume that n is an odd number. A *profile* $\boldsymbol{J} = (J_1,\ldots,J_n) \in (\{0,1\}^\Phi)^n$ is a vector of judgments, one for each agent. That is, in profile $\boldsymbol{J}$ each agent $i \in \mathcal{N}$ provides us with her judgment J_i regarding the items in the agenda.

2.2 Aggregation Rules

An *aggregation rule* is a function F that takes as input a profile of judgments and that returns as output a single collective judgment that is supposed to represent a suitable compromise between the judgments made by the individual agents. In fact, most aggregation rules allow for the possibility of ties between two or more judgments in the output. Thus, formally an aggregation rule is a function $F : (\{0,1\}^\Phi)^n \to \mathcal{P}(\{0,1\}^\Phi)$, mapping every given profile of judgments to a set of judgments. Here, $\mathcal{P}(\cdot)$ denotes the powerset.

The *majority rule* F_{maj} is the aggregation rule that accepts exactly those agenda items that are accepted by a majority of the individual agents. As we assume n to be odd, there always is just a single majority winner. To define the majority rule formally, let me first introduce some further terminology and notation. For any profile $\boldsymbol{J}$ and agenda item φ, let $n_\varphi^{\boldsymbol{J}} := |\{i \in \mathcal{N} \mid J_i(\varphi) = 1\}|$ denote the number of agents who accept φ in $\boldsymbol{J}$. Now, for any given profile $\boldsymbol{J}$, the corresponding *majority judgment* is defined as the judgment $\mathrm{Maj}^{\boldsymbol{J}} : \Phi \to \{0,1\}$ with $\mathrm{Maj}^{\boldsymbol{J}}(\varphi) = 1$ if and only if $n_\varphi^{\boldsymbol{J}} > \frac{n}{2}$. The majority rule then is defined as the aggregation rule $F_{\mathrm{maj}} : \boldsymbol{J} \mapsto \{\mathrm{Maj}^{\boldsymbol{J}}\}$.

2.3 Rationality and Feasibility Constraints

By a slight abuse of notation, we may identify each agenda item $\varphi \in \Phi$ with a propositional variable. Let $\mathcal{L}(\Phi)$ be the *propositional language* over this set of propositional variables. That is, $\mathcal{L}(\Phi)$ is the set of all well-formed formulas of propositional logic that we can construct using the propositional variables in Φ and the familiar connectives $\neg$ (negation), $\wedge$ (conjunction), and $\vee$ (disjunction). We will think of such formulas as constraints on judgments. In particular, we say that a judgment $J : \Phi \to \{0,1\}$ *satisfies* a constraint $\Gamma \in \mathcal{L}(\Phi)$, denoted $J \models \Gamma$, if Γ evaluates to *true* under the assignment of propositional variables to truth values induced by J in the natural manner, for the usual semantics of classical propositional logic (van Dalen, 2013). In other words, this notion of satisfaction is defined recursively as follows:

- $J \models \varphi$ for propositional variables $\varphi \in \Phi$ if and only if $J(\varphi) = 1$.
- $J \models \neg\Gamma$ if and only if $J \not\models \Gamma$ (i.e., if it is not the case that $J \models \Gamma$).
- $J \models \Gamma \wedge \Gamma'$ if and only if both $J \models \Gamma$ and $J \models \Gamma'$.
- $J \models \Gamma \vee \Gamma'$ if and only if $J \models \Gamma$ or $J \models \Gamma'$ (or both).

Thus, for example, if J is given by $(\varphi_1 \mapsto 1, \varphi_2 \mapsto 0, \varphi_3 \mapsto 1)$, then $J \models \varphi_1 \vee \neg\varphi_3$ but $J \not\models \varphi_2$. For a given formula $\Gamma \in \mathcal{L}(\Phi)$, we write $\mathrm{Mod}(\Gamma) := \{J \in \{0,1\}^\Phi \mid J \models \Gamma\}$ for the set of *models* of Γ, i.e., for the set of judgments that satisfy Γ.

We use constraints to limit the range of admissible judgments in two different ways, namely to define the range of judgments we consider rational for an agent to supply, and to define the range of judgments we consider feasible outcomes of aggregation. Let $\Gamma^{\mathrm{rat}} \in \mathcal{L}(\Phi)$ be a *rationality constraint* and let $\Gamma^{\mathrm{feas}} \in \mathcal{L}(\Phi)$ be a *feasibility constraint*. Then we say that the aggregation rule F is *consistent* with respect to Γ^{rat} and Γ^{feas}, if F maps every rational profile to a (set of) feasible outcome(s):

$$J_i \models \Gamma^{\mathrm{rat}} \text{ for all } i \in \mathcal{N} \quad \Rightarrow \quad J \models \Gamma^{\mathrm{feas}} \text{ for all } J \in F(J_1, \ldots, J_n)$$

3 Majoritarian Aggregators

In this section, I define a number of natural judgment aggregation rules that are based on the concept of majoritarianism.

The plain majority rule F_{maj} is not consistent. Counterexamples for consistency occur even in case rationality and feasibility constraints coincide. This is the famous *doctrinal paradox* (Kornhauser and Sager, 1993; List and Pettit, 2002). When the feasibility constraint is defined independently of the rationality constraint, this situation can only get worse. Only aggregation rules that explicitly restrict outcomes to judgments in $\mathrm{Mod}(\Gamma^{\mathrm{feas}})$ have a chance of being consistent. In this paper, I therefore restrict attention to such rules.

While respecting the majority opinion on all agenda items may not be possible due to the requirement of being consistent, intuitively speaking, we should try to respect these majorities as much as possible. But how can we make this notion of "as much as possible" precise? Let us say that judgment J is *majority-dominated* by judgment J' in the context of profile $\boldsymbol{J}$, if we have $J \cap \mathrm{Maj}^{\boldsymbol{J}} \subset J' \cap \mathrm{Maj}^{\boldsymbol{J}}$, i.e., if J' and the majority judgment agree on a strict superset of the set of agenda items that J and the majority judgment agree on. Now say that a consistent aggregation rule F is *majoritarian*, if it never returns an outcome that is majority-dominated by another feasible outcome. Formally, relative to rationality constraint Γ^{rat} and feasibility constraint Γ^{feas}, F is majoritarian, if, for all profiles $\boldsymbol{J} \in \mathrm{Mod}(\Gamma^{\mathrm{rat}})$, every judgment in $F(\boldsymbol{J})$ is majority-undominated in $\mathrm{Mod}(\Gamma^{\mathrm{feas}})$.

Next, I will define four consistent and majoritarian aggregation rules, all of which are known rules in the context of standard judgment aggregation. For all four definitions, I make the implicit assumption that some feasibility constraint Γ^{feas} has been fixed. The first rule simply returns the set of all majority-undominated judgments:

$$\text{max-set}(\boldsymbol{J}) \quad := \quad \{J \in \mathrm{Mod}(\Gamma^{\mathrm{feas}}) \mid J \cap \mathrm{Maj}^{\boldsymbol{J}} \subset J' \cap \mathrm{Maj}^{\boldsymbol{J}} \text{ for no } J' \in \mathrm{Mod}(\Gamma^{\mathrm{feas}})\}$$

The max-set rule has been called the *maximal subagenda rule* by Lang et al. (2011) and its outcome has been called the *Condorcet admissible set* by Nehring et al. (2014). While it returns those judgments that are maximal—with respect to set-inclusion—regarding the majorities they respect, the rule defined next seeks to maximise the *number* of agenda items on which the majority opinion is being respected:

$$\text{max-num}(\boldsymbol{J}) \quad := \quad \underset{J \in \mathrm{Mod}(\Gamma^{\mathrm{feas}})}{\mathrm{argmax}} \; |J \cap \mathrm{Maj}^{\boldsymbol{J}}|$$

The max-num rule has been called the *endpoint rule* by Miller and Osherson (2009) and the *maximum-cardinality subagenda rule* by Lang et al. (2011). If we also take the strengths of the majorities that are being respected into account, then we naturally arrive at a rule that maximises the overall number of opinions that are being respected, summing both over individual agents and over agenda items:

$$\text{max-sum}(\boldsymbol{J}) \quad := \quad \underset{J \in \mathrm{Mod}(\Gamma^{\mathrm{feas}})}{\mathrm{argmax}} \sum_{i \in \mathcal{N}} |J \cap J_i|$$

The max-sum rule is most often referred to as *'the' distance-based rule* (Endriss et al., 2012; Pigozzi, 2006). It has been called the *prototype rule* by Miller and Osherson (2009) and the *maximum-weight subagenda rule* by Lang et al. (2011).

Our fourth rule, which I call *greedy-max*, I will only define informally here. We first order the agenda items φ in terms of their majority strengths $\max\{n_\varphi^{\boldsymbol{J}}, n - n_\varphi^{\boldsymbol{J}}\}$. If two agenda items have the same majority strength, we rank them according to some pre-defined tie-breaking rule. We then attempt to accept or reject agenda items in that order, only going against the majority view if doing otherwise would render the outcome infeasible. This results in a single winning judgment (because n is odd and there can be no ties). Under a refinement of this rule, called 'parallel universe tie-breaking', we declare all those judgments winners that would win under *some* tie-breaking rule. This rule has been called the *ranked agenda rule* by Lang et al. (2011) and the *support-based rule* by Porello and Endriss (2014).

4 Recovery of Common Voting Rules

In this section, we will see that our four majoritarian aggregation rules allow us to recover a number of well-known voting rules when we combine them with certain rationality and feasibility constraints.

Let $\mathcal{X}$ be a finite set of *alternatives*. In voting theory, each agent expresses a preference over these alternatives (Brams and Fishburn, 2002). A voting rule then is a function mapping such a profile of preferences to a winning alternative (or possibly to a set of winning alternatives, in case there is a tie). We will be concerned with two groups of voting rules. The first group of rules requires each agent to express their preferences in the form of a *linear preference order*, i.e., a binary relation on $\mathcal{X}$ that is complete, antisymmetric, and transitive:

- **Borda:** Each alternative x receives as many points from agent i as i ranks other alternatives below x. The alternatives with the most points overall win.
- **Condorcet:** A *Condorcet winner* is an alternative that wins in a majority contest against every other alternative. Analogously, a *Condorcet loser* is an alternative that loses against every other alternative in a majority contest. Note that there can be at most one Condorcet winner and at most one Condorcet loser. The (refined) Condorcet rule elects the Condorcet winner if it exists and otherwise it elects all alternatives, except for the Condorcet loser (in case one exists). Note that the standard definition of the Condorcet rule does not exclude the Condorcet loser in the second case, but arguably this refinement is in fact the most natural implementation of the Condorcet principle in the form of a voting rule.
- **Top Cycle:** The Top Cycle if the smallest nonempty subset X of $\mathcal{X}$ such that every alternative in X wins all majority contests against all alternatives not in X. This thus is a voting rule that often produces ties between many winning alternatives (but when a Condorcet winner exists, then the Top Cycle is a singleton).
- **Copeland:** For each alternative, compute how many majority contests it wins against other alternatives and how many it loses. The *Copeland score* of an alternative is the difference between these two numbers. The Copeland winner is the alternative with the largest Copeland score. As we assume the number of agents to be odd (so majority contests will never end in a tie), we can equivalently elect the alternative(s) with the highest number of won majority contests (and ignore lost contests).
- **Maximin:** Find for each alternative the pairwise majority contest in which it is defeated most decisively. The winners for the Maximin Rule are those alternatives for which this worst defeat is least extreme.
- **Kemeny:** Define the *distance* between two preference orders as the number of pairs of alternatives on which they disagree. The Kemeny Rule first computes the collective preference orders that minimise the sum of the distances to the individual preference orders and then returns the top alternatives of those orders.
- **Slater:** Compute the *majority graph* (with alternatives being the vertices and the majority relation being the edges). Then compute the largest (in terms of cardinality) transitive subgraphs of the majority graph and return as winners their top alternatives.

- **Ranked Pairs:** Under Tideman's Ranked Pairs Rule, we sort the pairs of alternatives in order of their majority strengths, and break ties using some tie-breaking rule. Then lock in pairwise orderings in that order, skipping pairs whenever locking them in would create a cycle. The top alternative in the order produced wins. (A 'parallel universe tie-breaking' variant of this rule is defined in the natural manner.)

Two further voting rules only require each agent to declare which alternative they consider the best alternative:

- **Plurality:** Elect the alternatives listed most frequently as top alternatives by the individual agents.
- **Majority:** If one of the alternatives is listed as the top alternative by more than half of the agents, elect that alternative. Otherwise, declare a tie between all alternatives.

It is a well-known fact that preference aggregation can be embedded into judgment aggregation (Dietrich and List, 2007b; Grandi, 2012). For every pair $x, y \in \mathcal{X}$ we introduce an agenda item $p_{x \succcurlyeq y}$. Intuitively speaking, by accepting $p_{x \succcurlyeq y}$ you declare that you consider x at least as desirable as y. We can now write down an integrity constraint that expresses that judgments should correspond to linear preference orders:

$$\Gamma^{\text{pref}} \quad := \quad \underbrace{\bigwedge_{x,y} (p_{x \succcurlyeq y} \vee p_{y \succcurlyeq x})}_{\text{completeness}} \ \wedge \ \underbrace{\bigwedge_{x \neq y} \neg (p_{x \succcurlyeq y} \wedge p_{y \succcurlyeq x})}_{\text{antisymmetry}} \ \wedge \ \underbrace{\bigwedge_{x,y,z} (p_{x \succcurlyeq y} \wedge p_{y \succcurlyeq z} \to p_{x \succcurlyeq z})}_{\text{transitivity}}$$

If we impose Γ^{pref} as a rationality constraint, then we are essentially asking each agent to supply us with a linear preference order on the alternatives. If we make Γ^{pref} also the feasibility constraint, then our judgment aggregation rule corresponds to a preference aggregation rule, mapping profiles of preferences to a collective preference order (or possibly a set of tied winning preference orders). Any such rule induces a voting rule by stipulating that the winners are the top-ranked alternatives in the collective preference orders returned.

We can also write down a formula that expresses that judgments should identify one of the alternatives as the single top alternative, without imposing any other constraints:

$$\Gamma^{\text{top}} \quad := \quad \bigvee_{x} \underbrace{\left[\bigwedge_{y \neq x} (p_{x \succcurlyeq y} \wedge \neg p_{y \succcurlyeq x}) \right]}_{x \text{ dominates all}}$$

If we impose Γ^{top} as a feasibility constraint, then we are essentially asking the judgment rule to directly determine a winning alternative rather than to let it compute a winning preference order for us, from which we then extract the winning alternative in an additional step. While modelling voting rules (or rather: preference aggregation rules) as judgment aggregation rules mapping profiles that satisfy Γ^{pref} to outcomes that do the same has been the standard approach, at least implicitly, in most existing work on embedding voting rules into judgment aggregation, the central idea of Lang and Slavkovik (2013) has been to instead require outcomes to satisfy (a constraint very similar to) Γ^{top}.

Γ^{rat} / Γ^{feas}	**max-set**	**max-num**	**max-sum**	**greedy-max**
Γ^{pref} / Γ^{pref}	Top Cycle	Slater	Kemeny	Ranked Pairs
Γ^{pref} / $\Gamma^{\text{top}}_{\sim}$	Condorcet	Copeland	Borda	Maximin
$\Gamma^{\text{top}}_{\not\sim}$ / Γ^{pref}	Majority	Majority	Plurality	Plurality
$\Gamma^{\text{top}}_{\not\sim}$ / $\Gamma^{\text{top}}_{\sim}$	Majority	Majority	Plurality	Plurality

Table 1: Common voting rules as special cases of judgment aggregation rules

Note that for Γ^{top} the aggregation rule is free to accept or reject any of the agenda items fixing the relationship between non-winning alternatives. Clearly, for all our rules, it will always be optimal to simply accept all of those other agenda items. Thus, when used as a feasibility constraint, Γ^{top} will have the same effect as the following constraint, which makes it explicit that all agenda items speaking about the relative rankings of losing alternatives should be accepted:

$$\Gamma^{\text{top}}_{\sim} \quad := \quad \bigvee_{x} \Bigg[\underbrace{\bigwedge_{y \neq x} (p_{x \succcurlyeq y} \wedge \neg p_{y \succcurlyeq x})}_{x \text{ dominates all}} \;\wedge \underbrace{\bigwedge_{y \neq x, z \neq x} (p_{y \succcurlyeq z} \wedge p_{z \succcurlyeq y})}_{\text{indifference between rest}} \Bigg]$$

However, sometimes we may want to explicitly rule this out. For example, if we want to use a constraint such as Γ^{top} as rationality constraint, it seems that the most faithful translation from the language of ballots used, for instance, for the Plurality Rule would be to use the following constraint:

$$\Gamma^{\text{top}}_{\not\sim} \quad := \quad \bigvee_{x} \Bigg[\underbrace{\bigwedge_{y \neq x} (p_{x \succcurlyeq y} \wedge \neg p_{y \succcurlyeq x})}_{x \text{ dominates all}} \;\wedge \underbrace{\bigwedge_{x \neq y \neq z \neq x} (\neg p_{y \succcurlyeq z} \wedge \neg p_{z \succcurlyeq y})}_{\text{incomparability between rest}} \Bigg]$$

In what follows, I will use $\Gamma^{\text{top}}_{\not\sim}$ as a rationality constraint and $\Gamma^{\text{top}}_{\sim}$ as a feasibility constraint. Table 1 provides an overview of the results we can obtain this way, showing the voting rule recovered for each of our four judgment aggregation rules when we instantiate them with the relevant combinations of rationality and feasibility constraints. I offer some intuitions below, but full proofs are omitted in this short note. In fact, many of these results have already been known, be it in a slightly different formal framework. My claim here is that the present framework brings out these results in a particularly clear form.

The results in the first row ($\Gamma^{\text{pref}}/\Gamma^{\text{pref}}$) are all stated by Lang and Slavkovik (2013) as well, with the result for Slater arguably and that for Kemeny certainly qualifying as folk theorems. Both results are strongly alluded to by Miller and Osherson (2009). The correspondence for Kemeny is proved explicitly by Endriss et al. (2012). The results for Ranked Pairs is immediate, and due to Lang and Slavkovik (2013); the result for the Top Cycle Rule is less obvious, and due to the same authors. In the second row, corresponding results for Copeland and Maximin are obtained by Lang and Slavkovik (2013) as well, although they work with a different version of $\Gamma^{\text{top}}_{\sim}$, requiring the binary relation over the losing alternatives to be asymmetric. As a consequence of this modelling choice, Lang and Slavkovik (2013) fail to recover the Borda Rule. To see that the result claimed here really is correct, observe that the binary relation corresponding to $\Gamma_{\sim}^{\text{top}}$ is a graph that is almost completely connected; the only missing edges are those leading up to the winning alternative. Thus, if m is the number of alternatives, for every individual ballot putting alternative x at rank k, then max-sum awards $\binom{m}{2} - (k-1)$ points to the judgment corresponding to a graph putting x at the top. This is simply the standard Borda score, with the constant $\binom{m}{2} + 1 - m$ added to it, so we indeed are recovering the Borda Rule. Zwicker (2015), in a different technical context, makes essentially the same observation.[1] The remaining results, on Condorcet, Majority, and Plurality, are relatively straightforward.

5 Future Directions

This study may be taken further in a number of ways. First, the framework of judgment aggregation with rationality and feasibility constraints is more flexible than existing frameworks and can be expected to fit a wider range of application domains, which is an opportunity to be explored.

Second, there are other majoritarian rules that one could define and that may be of interest. One such rule is what one might call the *max-prod rule*, which works just like the max-sum rule, except that we multiply rather than add up the scores coming from the individual judgments. A second such rule is what I propose to call the *lexi-max rule*, which accepts the consistent judgments for which the ordered vector of majority strengths is lexicographically optimal. In other words, this is a refinement of greedy-max, which we may think of as a computationally less demanding approximation of the ideal described by the lexi-max rule. This kind of rule appears not to have been studied before, either in voting or in judgment aggregation, although it promises to having attractive properties (except for being algorithmically very demanding).

Third, in the context of embedding voting rules into judgment aggregation, we may consider other types of rationality and feasibility constraints. Maybe the most obvious choice would be a constraint that forces a dichotomous order over the alternatives, splitting them into 'good' and 'bad' alternatives. This way it may be possible to recover *approval voting* (Brams and Fishburn, 2002). Indeed, as noted elsewhere (Endriss et al., 2009), approval voting and the Borda Rule can be transformed into each other by only changing the language of allowed ballots. When using a constraint imposing dichotomous orders as the feasibility constraint, it may be possible to recover *multiwinner voting rules* (Elkind et al., 2014).

Acknowledgements. I am grateful to Umberto Grandi, Jérôme Lang, and Marija Slavkovik for several enlightening discussions on the topics of this paper.

[1] The results for max-sum are closely related to results on characterising voting rules by means of a consensus class (of profiles in which there is a clear-cut winner) and a notion of distance between profiles (Elkind and Slinko, 2016). Specifically, the result for Borda corresponds to a characterisation of the Borda Rule in terms of the *unanimous winner* consensus class and the *Kendall tau distance* (Farkas and Nitzan, 1979).

References

S. J. Brams and P. C. Fishburn. Voting procedures. In K. J. Arrow, A. K. Sen, and K. Suzumura, editors, *Handbook of Social Choice and Welfare*. North-Holland, 2002.

M. Costantini, C. Groenland, and U. Endriss. Judgment aggregation under issue dependencies. In *Proceedings of the 30th AAAI Conference on Artificial Intelligence (AAAI-2016)*, 2016.

D. van Dalen. *Logic and Structure*. Springer-Verlag, 5th edition, 2013.

F. Dietrich. A generalised model of judgment aggregation. *Social Choice and Welfare*, 28(4): 529–565, 2007.

F. Dietrich. Scoring rules for judgment aggregation. *Social Choice and Welfare*, 42(4):873–911, 2014.

F. Dietrich and C. List. Judgment aggregation by quota rules: Majority voting generalized. *Journal of Theoretical Politics*, 19(4):391–424, 2007a.

F. Dietrich and C. List. Arrow's Theorem in judgment aggregation. *Social Choice and Welfare*, 29(1):19–33, 2007b.

E. Dokow and R. Holzman. Aggregation of binary evaluations. *Journal of Economic Theory*, 145(2):495–511, 2010.

C. Duddy, A. Piggins, and W. S. Zwicker. Aggregation of binary evaluations: A Borda-like approach. *Social Choice and Welfare*, 2015. In press.

E. Elkind and A. Slinko. Rationalizations of voting rules. In F. Brandt, V. Conitzer, U. Endriss, J. Lang, and A. D. Procaccia, editors, *Handbook of Computational Social Choice*. Cambridge University Press, 2016.

E. Elkind, P. Faliszewski, P. Skowron, and A. Slinko. Properties of multiwinner voting rules. In *Proceedings of the 13th International Conference on Autonomous Agents and Multiagent Systems (AAMAS-2014)*, 2014.

U. Endriss. Judgment aggregation. In F. Brandt, V. Conitzer, U. Endriss, J. Lang, and A. D. Procaccia, editors, *Handbook of Computational Social Choice*. Cambridge University Press, 2016.

U. Endriss and U. Grandi. Binary aggregation by selection of the most representative voter. In *Proceedings of the 28th AAAI Conference on Artificial Intelligence (AAAI-2014)*, 2014.

U. Endriss, M. S. Pini, F. Rossi, and K. B. Venable. Preference aggregation over restricted ballot languages: Sincerity and strategy-proofness. In *Proceedings of the 21st International Joint Conference on Artificial Intelligence (IJCAI-2009)*, 2009.

U. Endriss, U. Grandi, and D. Porello. Complexity of judgment aggregation. *Journal of Artificial Intelligence Research*, 45:481–514, 2012.

D. Farkas and S. Nitzan. The Borda rule and Pareto stability: A comment. *Econometrica*, 47(5):1305–1306, 1979.

U. Grandi. *Binary Aggregation with Integrity Constraints*. PhD thesis, ILLC, University of Amsterdam, 2012.

U. Grandi and U. Endriss. Lifting integrity constraints in binary aggregation. *Artificial*

Intelligence, 199–200:45–66, 2013.

D. Grossi and G. Pigozzi. *Judgment Aggregation: A Primer*. Synthesis Lectures on Artificial Intelligence and Machine Learning. Morgan & Claypool Publishers, 2014.

S. Konieczny and R. Pino Pérez. Merging information under constraints: A logical framework. *Journal of Logic and Computation*, 12(5):773–808, 2002.

L. A. Kornhauser and L. G. Sager. The one and the many: Adjudication in collegial courts. *California Law Review*, 81(1):1–59, 1993.

J. Lang and M. Slavkovik. Judgment aggregation rules and voting rules. In *Proceedings of the 3rd International Conference on Algorithmic Decision Theory (ADT-2013)*. Springer-Verlag, 2013.

J. Lang, G. Pigozzi, M. Slavkovik, and L. van der Torre. Judgment aggregation rules based on minimization. In *Proceedings of the 13th Conference on Theoretical Aspects of Rationality and Knowledge (TARK-2011)*, 2011.

C. List and P. Pettit. Aggregating sets of judgments: An impossibility result. *Economics and Philosophy*, 18(1):89–110, 2002.

C. List and C. Puppe. Judgment aggregation: A survey. In P. Anand, P. Pattanaik, and C. Puppe, editors, *Handbook of Rational and Social Choice*. Oxford University Press, 2009.

M. K. Miller and D. Osherson. Methods for distance-based judgment aggregation. *Social Choice and Welfare*, 32(4):575–601, 2009.

K. Nehring and C. Puppe. Abstract Arrowian aggregation. *Journal of Economic Theory*, 145(2):467–494, 2010.

K. Nehring, M. Pivato, and C. Puppe. The Condorcet set: Majority voting over interconnected propositions. *Journal of Economic Theory*, 151:268–303, 2014.

G. Pigozzi. Belief merging and the discursive dilemma: An argument-based account of paradoxes of judgment aggregation. *Synthese*, 152(2):285–298, 2006.

D. Porello. Logics for collective reasoning. In *Proceedings of the European Conference on Social Intelligence (ECSI-2014)*. CEUR, 2014.

D. Porello and U. Endriss. Ontology merging as social choice: Judgment aggregation under the open world assumption. *Journal of Logic and Computation*, 24(6):1229–1249, 2014.

W. S. Zwicker. Aggregating binary relations: Universal scoring rules via inner product. Presentation at Dagstuhl Seminar 15241 on Computational Social Choice: Theory and Applications, Leibniz Zentrum für Informatik, 2015.