

Proportional Belief Merging

Adrian Haret, Martin Lackner, Andreas Pfandler, and Johannes P. Wallner

{haret, lackner, pfandler, wallner}@dbai.tuwien.ac.at

TU Wien

Vienna, Austria

Abstract

In this paper we introduce proportionality to belief merging. Belief merging is a framework for aggregating information presented in the form of propositional formulas, and it generalizes many aggregation models in social choice. In our analysis, two incompatible notions of proportionality emerge: one similar to standard notions of proportionality in social choice, the other more in tune with the logic-based merging setting. Since established merging operators meet neither of these proportionality requirements, we design new proportional belief merging operators. We analyze the proposed operators against established rationality postulates, finding that current approaches to proportionality from the field of social choice are, at their core, incompatible with standard rationality postulates in belief merging. We provide characterization results that explain the underlying conflict, and provide a complexity analysis of our novel operators.

1 Introduction

Proportionality is one of the central fairness notions studied in social choice theory (Black 1958; Dummett 1984; Monroe 1995), arising whenever a collective decision should reflect the amount of support in favor of a set of issues. Thus, notions of proportionality are key when it is desirable that preferences of larger groups have more influence on the outcome, while preferences of smaller groups are not neglected.

The idea of proportional representation shows up in many application scenarios: it is a key ingredient of parliamentary elections (Balinski and Young 1982) and, more generally, of multiwinner voting, i.e., the task of electing a committee of multiple candidates (Faliszewski et al. 2017). Recent work has set out to extend the notion of proportionality from mathematically simple formalisms (mainly the apportionment setting) to more general settings, with significant progress in areas such as approval-based multiwinner voting (Aziz et al. 2017; Sánchez-Fernández et al. 2017), ordinal multiwinner voting (Elkind et al. 2017a; 2017b), proportional rankings (Skowron et al. 2017), and multi-attribute committees (Lang and Skowron 2018).

In this paper we introduce proportionality to the very general framework of *belief merging* (Konieczny, Lang, and

Marquis 2004; Konieczny and Pino Pérez 2002; 2011), which allows agents to combine their individual positions on a set of issues in order to obtain a collective solution, with the added option of imposing constraints on admissible outcomes. Though the agents' individual positions are called *beliefs*, the belief merging framework is versatile enough that it can accommodate a broad range of attitudes (e.g., beliefs, preferences, judgments, goals or items of knowledge), as long as these, together with the constraint and the outcome, can be expressed as formulas in a logical language. The key challenges of such a process are that agents may hold mutually conflicting beliefs, and that beliefs may reflect complex interdependencies between issues. The theory of belief merging then offers (i) a range of methods, called *belief merging operators*, for aggregating beliefs and (ii) postulates used to assess the rationality of the operators.

The most prominent belief merging operators studied so far tend to fall into two main categories: operators following the majority opinion, and which can be said to embody a *utilitarian* stance; and operators that place particular emphasis on the worst-off agents, and which can be said to be based on an *egalitarian* viewpoint. Our aim is to find a compromise between these two opposing positions, which, in a belief merging scenario, translates as the following desideratum: if a large enough proportion of the agents share common beliefs, then these beliefs should be reflected at the collective level, to a degree matching their proportion. Despite its intuitive appeal, such a proportionality requirement has yet to find its way in the study of belief merging operators.

In defining proportional belief merging operators we rely on the Proportional Approval Voting (PAV) rule, studied in multiwinner voting scenarios and known to satisfy particularly strong proportionality requirements (Aziz et al. 2017). Based on PAV, we introduce three belief merging operators: the PAV operator, the bounded PAV operator and the harmonic Hamming operator. All these operators fall into the class of *satisfaction-based operators*, introduced by us (Section 4) as an alternative to the standard way of representing merging operators, which is distance-based.

We look at the proposed belief merging operators from three perspectives. Firstly, in Section 5, the operators are placed against the standard belief merging IC-postulates. We

show that any belief merging operator directly extending PAV cannot be compatible with all IC-postulates; in particular, such an operator will not satisfy postulate IC₂, which stipulates that any admissible agreement among agents shall be part of the merged result. We also provide a characterization of operators that fail IC₂ based on properties of the ranking a satisfaction-based operator induces. This provides an alternative view of why the PAV approach is inconsistent with IC₂. However, we show that what we call *the bounded PAV operator* can be characterized as the only merging operator (of a certain natural class) that extends PAV and satisfies all other postulates. While the harmonic Hamming operator is defined via the harmonic sum used by PAV, it does not generalize PAV. Thus, the aforementioned impossibility does not hold; indeed, the harmonic Hamming operator satisfies all standard IC postulates IC_{0–8}.

Secondly, in Section 6, we introduce two basic proportionality postulates for the belief merging domain. The first one, *classical proportionality*, is the kind of proportionality requirement typically studied in social choice settings, in particular in the apportionment setting (Balinski and Young 1982). This notion is based on the assumption that agents derive utility from positive occurrences, i.e., from approved candidates being selected in the collective choice. The second notion, *binary proportionality*, is closer to the logical nature of belief merging. Here, no difference is made between positive and negative agreement: the agents' utility derives from the (Hamming) distance between their preferences and the collective choice. We show that these two notions are mutually exclusive and contradict each other. Furthermore, we show by example that established belief merging operators satisfy neither of these two postulates. In contrast, the aforementioned PAV and bounded PAV operators satisfy classical proportionality and the harmonic Hamming operator satisfies binary proportionality.

Thirdly, in Section 7, we study the complexity of our proposed merging operators. Our results are that our novel operators fall into similar complexity classes as established merging operators, which shows that the introduction of proportionality comes at a moderate computational cost.

As mentioned before, belief merging can be seen as a general framework. In Section 3 we make this argument precise for approval-based committee elections, and in Section 8 we show that our work has implications for other settings: in particular, it yields new proportional goal-based voting rules (Novaro et al. 2018) and approval-based multiwinner rules with a variable number of winners (Kilgour 2016; Faliszewski, Slinko, and Talmon 2017), and gives insights for proportional judgment aggregation (List and Puppe 2009; Everaere, Konieczny, and Marquis 2015; 2017).

2 Belief Merging

We assume a set \mathcal{A} of m propositional atoms, with \mathcal{L} the set of propositional formulas generated from \mathcal{A} using the usual connectives. An *interpretation* w is a truth-value assignment to atoms in \mathcal{A} , and we denote by \mathcal{U} the set of all interpretations over the set \mathcal{A} . We typically write interpretations as words, standing in for the set of atoms assigned to true. If v and w are interpretations, the *symmetric difference* $v\Delta w$

between v and w is defined as $v\Delta w = (v \setminus w) \cup (w \setminus v)$. The *Hamming* and *drastic* distances d_H and d_D , respectively, are defined as $d_H(v, w) = |v\Delta w|$ and $d_D(v, w) = 0$, if $v = w$, and 1 otherwise. If $\varphi \in \mathcal{L}$ is a propositional formula and w is an interpretation, w is a *model* of φ if w satisfies φ , with $[\varphi]$ being the set of models of φ . If $\varphi_1, \varphi_2 \in \mathcal{L}$, we say that $\varphi_1 \models \varphi_2$ if $[\varphi_1] \subseteq [\varphi_2]$, and that $\varphi_1 \equiv \varphi_2$ if $[\varphi_1] = [\varphi_2]$. A formula φ is *consistent*, or *satisfiable*, if $[\varphi] \neq \emptyset$.

A *propositional profile* $P = (\varphi_1, \dots, \varphi_n)$ is a finite tuple of consistent propositional formulas, with each formula φ_i assumed to correspond to an agent i . If P_1 and P_2 are profiles, $P_1 + P_2$ is the profile obtained by appending P_2 to P_1 . If there is no danger of ambiguity, we write $P + \varphi_i$ instead of $P + (\varphi_i)$. A *merging operator* Δ is a function mapping a profile P of consistent formulas and a propositional formula μ , called *the constraint*, to a propositional formula $\Delta_\mu(P)$. Operators Δ^1 and Δ^2 are *equivalent* if $\Delta_\mu^1(P) \equiv \Delta_\mu^2(P)$, for any P and μ . The following postulates are typically taken to provide a core set of rationality constraints any merging operator Δ is expected to satisfy (Konieczny and Pino Pérez 2002; 2011):

- (IC₀) $\Delta_\mu(P) \models \mu$.
- (IC₁) If μ is consistent, then $\Delta_\mu(P)$ is consistent.
- (IC₂) If $\bigwedge P \wedge \mu$ is consistent, then $\Delta_\mu(P) \equiv \bigwedge P \wedge \mu$.
- (IC₃) If $P_1 \equiv P_2$ and $\mu_1 \equiv \mu_2$, then $\Delta_{\mu_1}(P_1) \equiv \Delta_{\mu_2}(P_2)$.
- (IC₄) If $\varphi_1 \models \mu$ and $\varphi_2 \models \mu$, then $\Delta_\mu(\varphi_1, \varphi_2) \wedge \varphi_1$ is consistent if and only if $\Delta_\mu(\varphi_1, \varphi_2) \wedge \varphi_2$ is consistent.
- (IC₅) $\Delta_\mu(P_1) \wedge \Delta_\mu(P_2) \models \Delta_\mu(P_1 + P_2)$.
- (IC₆) If $\Delta_\mu(P_1) \wedge \Delta_\mu(P_2)$ is consistent, then $\Delta_\mu(P_1 + P_2) \models \Delta_\mu(P_1) \wedge \Delta_\mu(P_2)$.
- (IC₇) $\Delta_{\mu_1}(P) \wedge \mu_2 \models \Delta_{\mu_1 \wedge \mu_2}(P)$.
- (IC₈) If $\Delta_{\mu_1}(P) \wedge \mu_2$ is consistent, then $\Delta_{\mu_1 \wedge \mu_2}(P) \models \Delta_{\mu_1}(P) \wedge \mu_2$.

These postulates are best understood as axiomatizing a decision procedure based on the aggregation of information coming from different sources (the formulas in P), under a constraint μ that must be satisfied by the result (postulate IC₀). The result should be consistent (postulate IC₁), independent of the syntax of the formulas involved (postulate IC₃), include outcomes that are unanimously accepted across subprofiles (postulates IC_{5–6}) and coherent when varying the constraint (postulates IC_{7–8}). Additionally, postulate IC₂ requires that if there is any agreement between the formulas in P and μ , then the merged result is nothing more than the agreed upon outcomes; and postulate IC₄ stipulates that merging two formulas φ_1 and φ_2 should be fair, in the sense that if the result contains outcomes consistent with one of the formulas, it should contain results consistent with the other as well. We will see that the latter two postulates are problematic for proportionality-driven merging operators.

Standard ways of constructing merging operators that satisfy postulates IC_{0–8} are based on the idea of minimizing overall distance to the profile $P = (\varphi_1, \dots, \varphi_n)$, and rely on two ingredients (Konieczny and Pino Pérez 2002; 2011). The first ingredient is a notion of *pseudo-distance*

d_H	$x_1x_2x_3x_4$	$x_1x_2x_3x_4$	$x_1x_2x_3x_4$	$y_1y_2y_3y_4$	Σ	gmax
$x_1x_2x_3x_4$	0	0	0	8	8	(8, 0, 0, 0)
$x_1x_2x_3y_1$	2	2	2	6	12	(6, 2, 2, 2)
$x_1x_2y_1y_2$	4	4	4	4	16	(4, 4, 4, 4)
$x_1y_1y_2y_3$	6	6	6	2	20	(6, 6, 6, 2)
$y_1y_2y_3y_4$	8	8	8	0	24	(8, 8, 8, 0)

Table 1: Hamming distances for $\Delta^{H,\Sigma}$ and $\Delta^{H,gmax}$.

$d: \mathcal{U} \times \mathcal{U} \rightarrow \mathbb{R}_{\geq 0}$ between interpretations, typically either Hamming distance d_H or drastic distance d_D . The distance $d(\varphi, w)$ from a formula φ to an interpretation w is then defined as $d(\varphi, w) = \min_{v \in [\varphi]} d(v, w)$. The collective distance w.r.t. profile P is obtained using the second ingredient, an *aggregation function* $f: \mathbb{R}_{\geq 0}^n \rightarrow \mathbb{R}_{\geq 0}$ that, for any integer n , maps a vector of n real numbers to a real number, and is defined as $d^f(P, w) = f(d(\varphi_1, w), \dots, d(\varphi_n, w))$. Typical aggregation functions are the sum Σ and gmax. By $f = \text{gmax}$ vectors are ordered in descending order. For this aggregation function, the resulting ordered vectors are compared according to a lexicographic order. The *distance-based merging operator* $\Delta^{d,f}$ is defined, for any profile P and formula μ , as a formula $\Delta_{\mu}^{d,f}(P)$ such that $[\Delta_{\mu}^{d,f}(P)] = \text{argmin}_{w \in [\mu]} d^f(P, w)$, i.e., as a formula whose models are the models of μ at minimal collective distance to P . When $d = d_D$, the operators $\Delta^{D,\Sigma}$ and $\Delta^{D,gmax}$ are equivalent and we will refer to them as Δ^D . Thus, we recall three main distance-based operators ($\Delta^{H,\Sigma}$, $\Delta^{H,gmax}$ and Δ^D), all of which are known to satisfy postulates IC_{0-8} (Konieczny, Lang, and Marquis 2002; Konieczny and Pino Pérez 2011).

Example 1. For the set of atoms $\mathcal{A} = X \cup Y$, where $X = \{x_1, \dots, x_4\}$ and $Y = \{y_1, \dots, y_4\}$, take a profile $P = (\varphi_1, \varphi_2, \varphi_3, \varphi_4)$ with $\varphi_i = (x_1 \wedge x_2 \wedge x_3 \wedge x_4) \wedge (\neg y_1 \wedge \neg y_2 \wedge \neg y_3 \wedge \neg y_4)$, for $i \in \{1, 2, 3\}$, $\varphi_4 = (\neg x_1 \wedge \neg x_2 \wedge \neg x_3 \wedge \neg x_4) \wedge (y_1 \wedge y_2 \wedge y_3 \wedge y_4)$. We obtain that $[\varphi_i] = \{x_1x_2x_3x_4\}$, for $i \in \{1, 2, 3\}$ and $[\varphi_4] = \{y_1y_2y_3y_4\}$. Additionally, take a constraint μ such that $[\mu] = \{x_1x_2x_3x_4, x_1x_2x_3y_1, x_1x_2y_1y_2, x_1y_1y_2y_3, y_1y_2y_3y_4\}$. Table 1 displays Hamming distances between models of μ and formulas in P as well as the aggregated distances, for the Σ and gmax aggregation functions. We have $d_H^{\Sigma}(P, x_1x_2x_3x_4) < d_H^{\Sigma}(P, x_1x_2x_3y_1)$ and $d_H^{\text{gmax}}(P, x_1x_2y_1y_2) < d_H^{\text{gmax}}(P, x_1x_2x_3y_1)$, since the overall distance (4, 4, 4, 4) lexicographically dominates (6, 2, 2, 2). Optimal outcomes are written in bold, i.e., $[\Delta_{\mu}^{H,\Sigma}(P)] = \{x_1x_2x_3x_4\}$ and $[\Delta_{\mu}^{H,gmax}(P)] = \{x_1x_2y_1y_2\}$. We also obtain that $[\Delta_{\mu}^D(P)] = \{x_1x_2x_3x_4\}$.

Example 1 illustrates a general feature of the standard merging operators: $\Delta^{H,\Sigma}$ sees optimal outcomes in utilitarian terms and thereby favors the majority opinion, while $\Delta^{H,gmax}$ attempts to improve the standing of the worse off agent, thereby favoring an egalitarian outcome. While such approaches may produce, on occasion, proportional outcomes, they are in no way guaranteed to do so in general.

3 Approval-Based Committee Elections as Instances of Belief Merging

Notions of proportionality have been systematically studied in the social choice literature, notably in the case of

PAV	$x_1x_2x_3x_4$	$x_1x_2x_3x_4$	$x_1x_2x_3x_4$	$y_1y_2y_3y_4$	Σ
$x_1x_2x_3x_4$	$h(4)$	$h(4)$	$h(4)$	$h(0)$	6.25
$x_1x_2x_3y_1$	$h(3)$	$h(3)$	$h(3)$	$h(1)$	6.5
$x_1x_2y_1y_2$	$h(2)$	$h(2)$	$h(2)$	$h(2)$	6
$x_1y_1y_2y_3$	$h(1)$	$h(1)$	$h(1)$	$h(3)$	4.83
$y_1y_2y_3y_4$	$h(0)$	$h(0)$	$h(0)$	$h(4)$	2.08

Table 2: PAV scores for a selection of committees of size 4.

Approval-Based Committee (ABC) elections (Faliszewski et al. 2017). An ABC election consists of a set of candidates C , a desired size of the committee k , and a preference profile $A = (A_1, \dots, A_n)$. The preference profile A contains approval ballots, i.e., $A_i \subseteq C$ is the set of candidates agent i approves of. An ABC voting rule outputs one or more size- k subsets of C , the chosen committee(s). The ABC voting rule of interest to us is called *Proportional Approval Voting* (PAV) (Thiele 1895). It is based on the *harmonic function* $h: \mathbb{N} \rightarrow \mathbb{R}$, defined as $h(\ell) = \sum_{i=1}^{\ell} \frac{1}{i}$ with the added convention that $h(0) = 0$. Given a committee w of size k , the PAV-score of w w.r.t. A is $\text{PAV}(A, w) = \sum_{i=1}^n h(|A_i \cap w|)$. The PAV rule applied to the preference profile A , for a desired size k of the committee, is defined as $\text{PAV}_k(A) = \text{argmax}_{w \subseteq C, |w|=k} \text{PAV}(A, w)$, i.e., it outputs committees of size k that maximize the PAV score w.r.t. A .

Example 2. Take a set $C = X \cup Y$ of candidates, where X and Y are as in Example 1, and a preference profile $A = (A_1, A_2, A_3, A_4)$ with $A_i = [\varphi_i]$, where φ_i are as in Example 1. Suppose $k = 4$, i.e., the task is to choose committees of size 4. Intuitively, a proportional outcome would consist of three candidates from X and one from Y , to reflect the fact that supporters of X outnumber supporters of Y in A by a ratio of 3 : 1. Indeed, this is exactly the type of outcome PAV will select. A committee maximizing the overall PAV score w.r.t. A is $x_1x_2x_3y_1$, as depicted in Table 2.

In Example 2 we have identified models of propositional formulas with sets of approved candidates in an ABC election. Indeed, we may pursue this analogy further and show that any ABC election can be rephrased as a belief merging instance. Given an instance of an ABC election, we associate to C the set of propositional atoms $\mathcal{A}_C = C$. To agent i 's approval ballot $A_i \subseteq C$ we associate the propositional formula: $\varphi_{A_i} = \bigwedge_{x \in A_i} x \wedge \bigwedge_{x \in C \setminus A_i} \neg x$, the sole model of which is A_i . To the preference profile A we associate the propositional profile $P_A = (\varphi_{A_1}, \dots, \varphi_{A_n})$. To obtain solutions that adhere to the cardinality constraint k , we choose μ_k to be a formula whose models are all subsets of \mathcal{A}_C of size k . By postulates IC_0 and IC_1 , $[\Delta_{\mu_k}(P_A)]$ consists of a non-empty set of models of size k , which can be seen as the winning committees of the ABC election.

In general, any ABC election for size- k committees can be seen as a belief merging instance where the profile consists of formulas with exactly one model and the constraint μ has models of fixed size k . A merging operator Δ extends PAV if for all preference profiles A , it holds that $\text{PAV}_k(A) = [\Delta_{\mu_k}(P_A)]$, i.e., the output of the PAV voting rule is the set of interpretations, or sets of atoms, returned by $\Delta_{\mu_k}(P_A)$.

$$\begin{array}{l}
\text{approval-based} \left\{ \begin{array}{l} s_{\text{AV}}(v, w) = |v \cap w| \\ s_{\text{PAV}}(v, w) = h(|v \cap w|) \\ s_{\text{bPAV}}(v, w) = 2h(|v \cap w|) - h(|w|) \end{array} \right. \\
\text{binary sat.-based} \left\{ \begin{array}{l} s_{\text{hH}}(v, w) = h(m - d_{\text{H}}(v, w)) \\ s_{\text{hD}}(v, w) = h(m - d_{\text{D}}(v, w)) \end{array} \right.
\end{array}$$

Figure 1: Proposed satisfaction measures.

In the following we will introduce merging operators that extend PAV and another one that is inspired by it.

4 Satisfaction-based Merging Operators

The framework of ABC elections presented in Section 3 can be used as a springboard for designing proportional belief merging operators. By conceiving ways in which an agent derives utility from a possible outcome, it becomes possible to reason about the social welfare of merging, i.e., the utility of the agents’ society as a whole.

The key notion in doing so is a *satisfaction measure* $s: \mathcal{U} \times \mathcal{U} \rightarrow \mathbb{R}$, quantifying the amount of satisfaction $s(v, w)$ of interpretation v with interpretation w . The satisfaction $s(\varphi, w)$ of a formula φ with w is then defined as $s(\varphi, w) = \max_{v \in [\varphi]} s(v, w)$. Finally, the *collective satisfaction* $s(P, w)$ of a profile P with w is defined as $s(P, w) = \sum_{\varphi \in P} s(\varphi, w)$. The *satisfaction-based merging operator* Δ^s outputs a formula $\Delta_\mu^s(P)$ such that $[\Delta_\mu^s(P)] = \text{argmax}_{w \in [\mu]} s(P, w)$, i.e., a formula whose models are exactly the models of μ that maximize satisfaction of P .

Note that we can convert a distance-based merging operator $\Delta^{d, \Sigma}$ (see Section 2) into an equivalent satisfaction-based operator by inverting the pseudo distance d , i.e., by defining a satisfaction measure s as $s(v, w) = m - d(v, w)$, for any interpretations v and w (remember that m is the number of atoms in \mathcal{A}). The resulting satisfaction-based operator is s.t. $\Delta_\mu^s(P) \equiv \Delta_\mu^{d, \Sigma}(P)$, for any profile P and μ . Note that since d is a pseudo distance and thus symmetric (i.e., $d(v, w) = d(w, v)$), for any interpretations v and w , the satisfaction measure s defined on the basis of it is also symmetric. This being said, we do not require satisfaction measures to be symmetric in general. Consequently, satisfaction-based operators as defined here form a more general class than distance-based operators $\Delta^{d, \Sigma}$, where d is a pseudo-distance. It is worth mentioning that the satisfaction-based approach we propose here is not a mere stylistic variant of the distance-based view; it also encourages a different viewpoint on merging, where the goal is to find an outcome making agents happy, subject to fairness norms. Scenarios where this viewpoint is warranted occur most of all in settings requiring social considerations.

The concrete satisfaction measures we propose are defined, for any interpretations v and w , in Figure 1, and generate two groups of operators. The *approval-based* operators, consisting of the AV operator Δ^{AV} , the PAV operator Δ^{PAV} and the *bounded* PAV operator Δ^{bPAV} , mimic the behavior of an ABC voting rule (see Section 3) in that they compute satisfaction of v with w based on how many atoms v and w have in common, similarly to how satisfaction of

an approval ballot A_i with a potential committee w is based on how many approved candidates in A_i find themselves in w . Note that, while an ABC voting rule is defined only for committees of fixed size, the merging operators we propose select among interpretations of any size. Nonetheless, it is straightforward to see that if the allowed outcomes (here, models of the constraint μ) are restricted to a given size, then the operators Δ^{PAV} and Δ^{bPAV} are equivalent and extend, in the sense described in Section 3, the PAV voting rule.

The operator Δ^{AV} is put forward as a benchmark approval-based operator, based on a satisfaction measure that simply counts the atoms v and w have in common: in particular, Δ^{AV} does not incorporate any proportionality ideas. Consequently the Δ^{AV} operator does not extend PAV, and, as shown in Section 6, does not meet any of the proportionality requirements we propose. The Δ^{PAV} operator refines Δ^{AV} by using the harmonic function h , which is known to behave well w.r.t. proportionality requirements (Aziz et al. 2017). Intuitively, the harmonic function reflects the “diminishing returns” of added satisfaction: the difference between $h(x)$ and $h(x + 1)$ gets smaller as x increases. Thus, the operator Δ^{PAV} is a prime candidate for a proportional satisfaction-based merging operator. Nonetheless, Δ^{PAV} has several shortcomings, which serve as motivation for the remaining operators.

One drawback of Δ^{PAV} is that it favors larger interpretations if available (Example 3), i.e., it tries to increase agents’ satisfaction by setting as many atoms to true as possible. Such an inflationary strategy may be undesirable in practice and, in a belief merging setting, interferes with postulate IC₄.

Example 3. For $\mathcal{A} = \{x_1, x_2\}$, $P = (\varphi_1, \varphi_2)$, with $[\varphi_1] = \{x_1\}$ and $[\varphi_2] = \{x_1x_2\}$, and μ such that $[\mu] = \{x_1, x_1x_2\}$, we obtain that $[\Delta_\mu^{\text{PAV}}(P)] = \{x_1x_2\}$, contradicting IC₄. The same result is obtained for Δ^{AV} , but $[\Delta_\mu^{\text{bPAV}}(P)] = \{x_1, x_1x_2\}$, which is in accordance with IC₄.

To curb the inflationary tendencies of Δ^{PAV} , operator Δ^{bPAV} introduces a penalty on interpretations depending on their size, in the process ensuring satisfaction of postulate IC₄ as well. Indeed, as Section 5 shows, Δ^{bPAV} is the *only* operator from a fairly broad class that manages to balance proportionality and fairness, as formalized by postulate IC₄. Note, however, that s_{bPAV} is not symmetric.

Example 4. It holds that $s_{\text{bPAV}}(x, xy) < s_{\text{bPAV}}(xy, x)$. Intuitively, this means it is worse to get y if it is not wanted than to not get it if it is wanted.

A related problem with Δ^{PAV} stems from the fact that $s_{\text{PAV}}(v, w)$ is obtained by counting only atoms v and w have in common. Hence, Δ^{PAV} has a bias towards positive literals, which turns out to interfere with postulate IC₂.

The *binary satisfaction-based* operators, consisting of the *harmonic drastic operator* Δ^{hD} and the *harmonic Hamming operator* Δ^{hH} , are introduced in an attempt to deal with the effect of unwanted atoms while, at the same time, providing proportional outcomes. The satisfaction measures they are based on penalize interpretations w if they include atoms for which an explicit preference is not stated. This is done by inverting familiar notions of distance, which pay attention to atoms appearing in one of the interpretation but not

in the other, and leads to an equal treatment of positive and negative literals. The harmonic function h is applied to this satisfaction notion, with the idea of ensuring proportionality. The operators that emerge are worth investigating: neither of them extends PAV (as hinted at in Example 5), but from this point onward their properties diverge. Though Δ^{hH} does not extend PAV, it still ends up having interesting proportionality properties, formalized in Section 6.

Example 5. *It holds that $s_{\text{PAV}}(x_1, x_1) = s_{\text{PAV}}(x_1, x_1x_2)$. The assumption behind s_{PAV} is that an agent who wants x_1 is equally satisfied with x_1x_2 as it is with x_1 , i.e., is not bothered by the presence of x_2 . This leads to non-satisfaction of IC_2 : for $\mathcal{A} = \{x_1, x_2\}$, $P = (\varphi)$, $[\varphi] = \{x_1\}$ we obtain that $[\Delta_{\top}^{\text{PAV}}(P)] = \{x_1, x_1x_2\}$, contrary to IC_2 . On the other hand, $s_{\text{hH}}(x_1, x_1) = h(2)$, while $s_{\text{hH}}(x_1, x_1x_2) = h(1)$ and $[\Delta_{\mu}^{\text{hH}}(P)] = \{x_1\}$, in accordance with IC_2 . Thus, according to s_{hH} , x_1x_2 provides less satisfaction to x_1 than x_1 alone, i.e., the agent is bothered by the presence of x_2 , for which an explicit preference was not stated.*

The operator Δ^{hD} turns out to be so coarse in its assessment of satisfaction as to become, as Proposition 1 shows, indistinguishable from existing merging operators defined using drastic distance d_{D} . As a result, the Δ^{hD} operator is not responsive to proportionality requirements.

Proposition 1. *The satisfaction-based operator Δ^{hD} is equivalent to the distance-based operator Δ^{D} .*

What emerges is a landscape with three merging operators relevant to the issue of proportionality, i.e., Δ^{PAV} , Δ^{bPAV} and Δ^{hH} . Out of these, Δ^{bPAV} and Δ^{hH} address, each in its own way, problems arising with the Δ^{PAV} operator: Δ^{bPAV} penalizes interpretations according to their size, while Δ^{hH} uses an approach reminiscent from logic, where positive and negative literals are treated equally. As we will see in Sections 5 and 6, the proposed solutions involve various trade-offs between proportionality and the IC postulates.

5 IC Postulates: Possibility and Impossibility

In this section we analyze the merging operators introduced in Section 4 in light of the standard merging postulates IC_{0-8} . The first result shows that any satisfaction-based operator satisfies a core set of IC-postulates.

Proposition 2. *If s is a satisfaction measure, then the merging operator Δ^s satisfies postulates $\text{IC}_{0-1,3,5-8}$.*

Proposition 2 applies to both the approval-based and the harmonic distance-based operators. What remains, then, is an understanding of how the new satisfaction measures interact with postulates IC_2 and IC_4 , and we settle the issue by characterizing the types of satisfaction measures compliant with these postulates. If $v \neq w$, the following properties prove to be relevant:

$$\begin{aligned} (\text{S}_1) \quad & s(v, v) > s(v, w); & (\text{S}_2) \quad & s(v, v) > s(w, v); \\ (\text{S}_3) \quad & s(v, v) = s(w, w); & (\text{S}_4) \quad & s(v, w) = s(w, v). \end{aligned}$$

They formalize the intuition that satisfaction is symmetric (S_4), maximal when one gets *exactly* what one wants, and trailing off as the outcome diverges from one's most desired outcome (S_{1-3}). Theorem 1 shows that properties S_{1-3} capture satisfaction measures compliant with postulate IC_2 .

Theorem 1. *A satisfaction-based merging operator Δ^s satisfies postulate IC_2 iff s satisfies properties S_{1-3} .*

Since the satisfaction measures s_{AV} , s_{PAV} or s_{bPAV} satisfy none of the properties S_{1-3} , Theorem 1 implies that the approval-based operators Δ^{AV} , Δ^{PAV} or Δ^{bPAV} do not satisfy postulate IC_2 . On the other hand, the satisfaction measures s_{hD} and s_{hH} do satisfy properties S_{1-3} , showing that the corresponding operators satisfy postulate IC_2 .

As mentioned in Section 4, we do not require satisfaction measures to be symmetric and, indeed, s_{bPAV} is not symmetric (though the other satisfaction measures are). The following result shows that, in the presence of postulate IC_2 , symmetry is connected to postulate IC_4 .

Theorem 2. *If a satisfaction-based merging operator Δ^s satisfies postulate IC_2 , then Δ^s satisfies postulate IC_4 if and only if s also satisfies property S_4 (i.e., is symmetric).*

Since the satisfaction measures s_{hD} and s_{hH} are symmetric and, as implied by Theorem 1, satisfy properties S_{1-3} , we get by Theorem 2 that they also satisfy postulate IC_4 . Together with Proposition 2, this yields the full picture for the binary satisfaction-based operators Δ^{hH} and Δ^{hD} .

Corollary 1. *The operators Δ^{hH} and Δ^{hD} satisfy postulates IC_{0-8} .*

For the approval-based operators, satisfaction of postulates IC_2 and IC_4 is clarified by another perspective on satisfaction measures. A satisfaction measure s is a *counting index* if there exists a function $\sigma: \mathbb{N} \times \mathbb{N} \rightarrow \mathbb{R}$, called *the witness of s* , such that $\sigma(0, 0) = 0$ and $s(v, w) = \sigma(|v \cap w|, |w|)$, for any interpretations v and w . Theorem 3 shows that counting indices do not fit with postulate IC_2 .

Theorem 3. *If s is a counting index, the satisfaction-based merging operator Δ^s does not satisfy postulate IC_2 .*

It is straightforward to see that the approval-based satisfaction measures introduced in Section 4 are counting indices. Thus, by Theorem 3, none of the operators they generate satisfies postulate IC_2 . For postulate IC_4 , however, the situation is different. Example 3 shows that the Δ^{AV} and Δ^{PAV} operators do not satisfy postulate IC_4 , though Δ^{bPAV} manages to evade the counter-example. In fact, it turns out that not only does the operator Δ^{bPAV} satisfy postulate IC_4 , but a much stronger result can be shown: it is the *only* operator based on a counting index that does so.

Theorem 4. *If Δ^s is a satisfaction-based merging operator such that s is a counting index with σ as witness, extends PAV and satisfies postulate IC_4 , then $\sigma(x, y) = 2h(x) - h(y)$, for any $x, y \in \mathbb{R}$.*

It deserves emphasis that Δ^{bPAV} manages to satisfy postulate IC_4 even though s_{bPAV} is not a symmetric satisfaction measure: since Δ^{bPAV} does not satisfy IC_2 , Theorem 2 does not apply. Indeed, none of the approval-based operators manages to satisfy both IC_2 and IC_4 . This suggests that there is a trade-off between the kind of proportionality these operators stand for and the satisfaction of IC_2 and IC_4 .

It is relevant that approval-based operators can consider interpretations of various sizes: reflection on Examples 3

and 5 shows that they are at the root of the problematic situations. Interestingly, it turns out that fixing the size of the models of the constraint μ yields merging operators that behave well w.r.t. the IC postulates.

Theorem 5. *If all models of the constraint μ have some fixed size k , then the approval-based merging operators Δ^{AV} , Δ^{PAV} and Δ^{bPAV} satisfy all postulates IC_{0-8} .*

6 Two Types of Proportionality

Here we formalize two notions of proportionality, arising out of two different ways of conceptualizing satisfaction with respect to a possible outcome. For the sake of clarity, we define these notions for rather restricted profiles.

A formula φ is *complete* if it has exactly one model, and a profile P is complete if all its formulas are complete. We write $P = (v_1, \dots, v_n)$ to denote the complete profile with $[\varphi_i] = \{v_i\}$, for all $i \in \{1, \dots, n\}$. A complete profile $P = (v_1, \dots, v_n)$ is *simple* if $v_1 \cup \dots \cup v_n = \mathcal{A}$, and either $v_i = v_j$ or $v_i \cap v_j = \emptyset$, for every $i, j \in \{1, \dots, n\}$.¹ A complete profile $P = (v_1, \dots, v_n)$ is ℓ -*simple* if it is simple and $|\{v_1, \dots, v_n\}| = \ell$, i.e., P contains ℓ distinct sets. If v_1, \dots, v_ℓ constitutes a partition of \mathcal{A} , and p_1, \dots, p_ℓ are positive integers, we write $(v_1^{p_1}, \dots, v_\ell^{p_\ell})$ to denote the ℓ -simple profile: $(\underbrace{v_1, \dots, v_1}_{p_1 \text{ times}}, \underbrace{v_2, \dots, v_2}_{p_2 \text{ times}}, \dots, \underbrace{v_\ell, \dots, v_\ell}_{p_\ell \text{ times}})$. If

$P = (v_1^{p_1}, \dots, v_\ell^{p_\ell})$ is an ℓ -simple profile with $\sum_{i=1}^{\ell} p_i = n$, we say that P is k -*integral* if $\frac{k \cdot p_i}{n}$ is an integer, for every $i \in \{1, \dots, \ell\}$. Intuitively, for a model w of μ of size k , the fraction $\frac{k \cdot p_i}{n}$ denotes the intended satisfaction if proportionality is taken into account: out of the k atoms selected, the share of group i should be the relative size of the group.

We propose two proportionality postulates, formulated for simple profiles $P = (v_1^{p_1}, \dots, v_\ell^{p_\ell})$. As in Section 3, constraint μ_k has as its models all interpretations of size k .

(IC_{cp}) For any $k \in \{1, \dots, m\}$ and $w \in [\Delta_{\mu_k}(P)]$, it holds that if P is k -integral and $|v_j| \geq \frac{k \cdot p_j}{n}$ for each j , $1 \leq j \leq \ell$, then $|v_i \cap w| = \frac{k \cdot p_i}{n}$, for all $i \in \{1, \dots, \ell\}$.

(IC_{bp}) If $P = (v_1^{p_1}, v_2^{p_2})$ is simple and there is a $w \in [\mu]$ such that $m - d_{\text{H}}(v_i, w) = \frac{m \cdot p_i}{n}$ for $i \in \{1, 2\}$, then this equality holds for all $w' \in [\Delta_{\mu}(P)]$.

We refer to IC_{cp} and IC_{bp} as postulates of *weak classical proportionality* and *weak binary proportionality*, respectively, as they refer to different sources of satisfaction. Postulate IC_{cp} talks about *classical satisfaction*, in which agent i 's satisfaction with an interpretation w is given by $|v_i \cap w|$, just like the satisfaction with a committee in an ABC election is measured by the number of approved committee members. This is the kind of satisfaction notion typically used in a social choice context. Postulate IC_{bp} talks about *binary satisfaction*, in which agent i 's satisfaction with w is given by $m - d_{\text{H}}(v_i, w)$, i.e., by the closeness between v_i and w . This type of satisfaction, alluded to already in Section 4, follows

¹In the context of ABC voting, such profiles are referred to as party-list profiles (Lackner and Skowron 2018b).

from a logical viewpoint where positive and negative variable assignments are treated equally. This approach is better suited to deal with interpretations of varying sizes than the classical one, and thus postulate IC_{bp} allows such interpretations to be selected.

Intuitively, both postulates stipulate ‘shares’ groups of agents shall receive (under a classical or binary viewpoint) that meet proportionality based on the relative size of the groups. For IC_{cp} we restrict to μ_k , with k atoms to be distributed proportionally by each solution w (like for ABC elections). Postulate IC_{bp} states that in the presence of at least one admissible $w \in [\mu]$ that meets the proportionality requirements, *all* solutions shall meet said requirements (otherwise μ permits no proportional solution). Note that if $P = (v_1^{p_1}, v_2^{p_2})$ satisfies the conditions of IC_{bp}, then P is m -integral, and the binary satisfaction of v_1 and v_2 adds up to m , i.e., $m - d_{\text{H}}(v_1, w) + m - d_{\text{H}}(v_2, w) = m$. Postulate IC_{bp} demands that this total satisfaction m is split proportionally.

Example 6. *For $\mathcal{A} = X \cup Y$, with $X = \{x_1, \dots, x_6\}$ and $Y = \{y_1, y_2\}$, take the simple profile $P = (v_1^3, v_2^1)$, with $v_1 = x_1 \dots x_6$ and $v_2 = y_1 y_2$, and a constraint μ_4 , with models of size 4. According to IC_{cp}, an optimal outcome contains three variables from X and one from Y , e.g., the interpretation $w = x_1 x_2 x_3 y_1$. Such an outcome is in the spirit of classical proportionality.*

According to postulate IC_{bp}, an optimal outcome w would be such that $d_{\text{H}}(v_1, w) = 2$ and $d_{\text{H}}(v_2, w) = 6$, e.g., $w' = x_1 x_2 x_3 x_4$, $w'' = x_1 x_2 x_3 x_4 x_5 y_1$ or $w''' = x_1 x_2 x_3 x_4 x_5 x_6 y_1 y_2$. Note, IC_{bp} allows interpretations of varying sizes to be selected. If the size is restricted to 4 (i.e., the constraint is μ_4), then the outcome narrows down to interpretations such as w' , consisting of four atoms from X .

Example 6 shows that classical and binary proportionality may require different interpretations to be selected on the same input. Thus, even though our notions of proportionality apply only to simple profiles, they set up a clear boundary for distinguishing among the different merging operators.

Theorem 6. *The merging operators Δ^{PAV} and Δ^{bPAV} satisfy postulate IC_{cp}, Δ^{hH} satisfies postulate IC_{bp}, while $\Delta^{\text{H}, \Sigma}$, $\Delta^{\text{H}, \text{gmax}}$, Δ^{hD} and Δ^{AV} satisfy neither IC_{cp} nor IC_{bp}.*

The proposed merging operators Δ^{PAV} and Δ^{bPAV} are representative of the notion of classical proportionality, while Δ^{hH} is representative for binary proportionality. Theorem 7 shows that these notions are thoroughly incompatible.

Theorem 7. *There is no merging operator that satisfies IC₁ and both IC_{cp} and IC_{bp}.*

The literature on belief merging suggests other properties concerned, in some way or another, with fairness of merging operators (Konieczny and Pino Pérez 2002; Everaere, Konieczny, and Marquis 2010; 2014). In the interest of brevity, we mention only that these properties are largely orthogonal to the proportionality requirements we study here. The one exception is the majority axiom (Konieczny and Pino Pérez 2002), which our operators do satisfy.

7 Computational Complexity

To investigate the complexity of our novel merging operators, we look at the standard decision problem studied in this context (Koniieczny, Lang, and Marquis 2004): given an operator Δ , a profile P , an integrity constraint μ , and a Boolean formula ψ , determine whether $\Delta_\mu(P) \models \psi$ holds. That is, the decision problem asks whether formula ψ follows from the merged result. The hardness results we use and recall here hold when $\psi = a$ is an atom. The two main complexity classes appearing here are Δ_2^P and Θ_2^P , denoting the classes of decision problems solvable via a deterministic polynomial time algorithm with access to an NP oracle, with the latter class having the additional restriction that at most logarithmically many oracle calls may be made. Many standard merging operators are complete for one of these two classes (Koniieczny, Lang, and Marquis 2004).

We show that our novel operators fit into this picture; we obtain Θ_2^P hardness and Δ_2^P membership for all new operators, except for Δ^{AV} , which we show to be Θ_2^P -complete. That is, our introduction of proportionality leads to neither milder nor significantly more complex operators. Hardness for Θ_2^P can be shown by adapting an existing reduction, originally from belief revision (Eiter and Gottlob 1992, Theorem 6.9). Finally, membership diverges for Δ^{hH} , Δ^{PAV} , Δ^{bPAV} and Δ^{AV} , since the first three operators induce an exponential set of possible satisfaction scores for interpretations—in contrast to Δ^{AV} that only induces a polynomial set.

Theorem 8. *Deciding whether a formula follows from the result of merging operator Δ^s is Θ_2^P -complete for $s = \text{AV}$, and both Θ_2^P -hard and in Δ_2^P , for $s \in \{\text{PAV}, \text{bPAV}, \text{hH}\}$.*

We conjecture that merging under Δ^{hH} , Δ^{PAV} and Δ^{bPAV} is Δ_2^P -complete.

8 Applications Beyond Belief Merging

In this section we briefly discuss how our results can be transferred to other, related formalisms.

Variable Approval-Based Committee Elections In contrast to ABC elections as introduced in Section 3, it is sometimes desirable to have flexibility with respect to the size of the committee by not fixing its size in advance (Kilgour 2016; Faliszewski, Slinko, and Talmon 2017). We refer to ABC voting rules without a size constraint as *variable ABC voting rules*. Note that a merging operator defines a variable ABC rule by setting $\mu = \top$. It is easy to see that the AV and PAV operators are not sensible in this context, as $w = \mathcal{A}$ (i.e., setting all atoms to true) is always an optimal model. However, the Δ^{bPAV} and Δ^{hH} operators present themselves as novel additions to this framework, being *proportional* variable ABC rules.

Goal-Based Voting Goal-based voting (Novaro et al. 2018) is a formalism similar to belief merging but with a focus on resolute rules (i.e., rules return only one model) and with different postulates. All proposed operators in this paper can be viewed as goal-based voting rules (subject to tie-breaking), and our proportionality postulates can be adapted for this setting as well. To the best of our knowledge, our

Table 3: Summary of results. New results in gray, for all others see (Koniieczny, Lang, and Marquis 2004)). Per Theorem 5, for results marked with * the \times becomes \checkmark when models of the constraint μ are assumed to have fixed size.

	IC _{0,1,3,5-8}	IC ₂	IC ₄	IC _{cp}	IC _{bp}	Complexity
$\Delta^{\text{H},\Sigma}$	\checkmark	\checkmark	\checkmark	\times	\times	$\Theta_2^P\text{-c}$
$\Delta^{\text{H},\text{gmax}}$	\checkmark	\checkmark	\checkmark	\times	\times	$\Delta_2^P\text{-c}$
$\Delta^{\text{hD}} \equiv \Delta^{\text{D}}$	\checkmark	\checkmark	\checkmark	\times	\times	$\Theta_2^P\text{-c}$
Δ^{hH}	\checkmark	\checkmark	\checkmark	\times	\checkmark	in $\Delta_2^P, \Theta_2^P\text{-h}$
Δ^{AV}	\checkmark	\times^*	\times^*	\times	\times	$\Theta_2^P\text{-c}$
Δ^{PAV}	\checkmark	\times^*	\times^*	\checkmark	\times	in $\Delta_2^P, \Theta_2^P\text{-h}$
Δ^{bPAV}	\checkmark	\times^*	\checkmark	\checkmark	\times	in $\Delta_2^P, \Theta_2^P\text{-h}$

proposed merging operators yield the first proportional goal-based voting rules. It would be particularly interesting to see whether Theorem 4 can be replicated by axioms from the goal-based voting setting (instead of postulate IC₄).

Judgment Aggregation Judgment aggregation (JA) is another formalism for aggregating beliefs, distinct from belief merging but overlapping in certain respects (Everaere, Koniieczny, and Marquis 2015; 2017). Even though they differ in important aspects, the main ideas in our paper can be transferred to JA. While propositional variables are the basic building blocks for belief merging, it might be more suitable to take the agenda (a set of propositional formulas) as the basis for defining proportionality in JA. This allows for the definition of proportional JA operators. Further work is required to analyze the resulting JA operators.

9 Discussion

In this paper we have initiated the study of proportional belief merging operators. We have presented three proportional operators: the PAV operator and the bounded PAV operator, both satisfying IC_{cp}, and the harmonic Hamming operator satisfying IC_{bp}. We summarize our results in Table 3.

Apart from the questions posed in Section 8, the current work suggests several directions for future research. While the two proportionality postulates we proposed apply only to certain instances, even weak proportionality postulates have proven sufficient for axiomatic characterizations (Lackner and Skowron 2018b) and in our paper these two postulates are sufficient to distinguish proportional from non-proportional operators. On the other hand, stronger postulates are desirable to determine to which degree proportionality guarantees can be given. This has recently been investigated in the context of approval-based committee elections (Aziz et al. 2017; 2018; Sánchez-Fernández et al. 2017); this line of work can serve as a basis for a similar analysis for belief merging operators.

Finally, manipulation and strategic voting, common concerns in social choice theory, have received some attention in the belief merging framework as well (Everaere, Koniieczny, and Marquis 2007; Haret and Wallner 2019). It can be expected that proportional belief merging operators are prone to strategic voting, as in ABC voting even weak forms of

proportionality and strategy-proofness have been shown to be incompatible (Peters 2018). Still, it has been found that the percentage of manipulable instances depends strongly on the choice of voting rules (Lackner and Skowron 2018a), indicating that a detailed analysis of vulnerabilities is an interesting avenue for future work.

Acknowledgements

This work was supported by the Austrian Science Fund (FWF): P30168-N31, P30930-N35, P31890-N31, and W1255-N23.

References

- Aziz, H.; Brill, M.; Conitzer, V.; Elkind, E.; Freeman, R.; and Walsh, T. 2017. Justified representation in approval-based committee voting. *Soc. Choice Welf.* 48(2):461–485.
- Aziz, H.; Elkind, E.; Huang, S.; Lackner, M.; Fernández, L. S.; and Skowron, P. 2018. On the complexity of extended and proportional justified representation. In *Proc. AAAI*, 902–909. AAAI Press.
- Balinski, M., and Young, H. P. 1982. *Fair Representation: Meeting the Ideal of One Man, One Vote*. Yale University Press.
- Black, D. 1958. *The Theory of Committees and Elections*. Cambridge University Press.
- Dummett, M. 1984. *Voting Procedures*. Oxford University Press.
- Eiter, T., and Gottlob, G. 1992. On the complexity of propositional knowledge base revision, updates, and counterfactuals. *Artif. Intell.* 57(2-3):227–270.
- Elkind, E.; Faliszewski, P.; Laslier, J.-F.; Skowron, P.; Slinko, A.; and Talmon, N. 2017a. What Do Multi-winner Voting Rules Do? An Experiment Over the Two-Dimensional Euclidean Domain. In *Proc. AAAI*, 494–501. AAAI Press.
- Elkind, E.; Faliszewski, P.; Skowron, P.; and Slinko, A. 2017b. Properties of multiwinner voting rules. *Soc. Choice Welf.* 48(3):599–632.
- Everaere, P.; Konieczny, S.; and Marquis, P. 2007. The strategy-proofness landscape of merging. *J. Artif. Intell. Res.* 28:49–105.
- Everaere, P.; Konieczny, S.; and Marquis, P. 2010. Disjunctive merging: Quota and gmin merging operators. *Artif. Intell.* 174(12-13):824–849.
- Everaere, P.; Konieczny, S.; and Marquis, P. 2014. On egalitarian belief merging. In *Proc. KR*, 121–130. AAAI Press.
- Everaere, P.; Konieczny, S.; and Marquis, P. 2015. Belief merging versus judgment aggregation. In *Proc. AAMAS*, 999–1007. ACM.
- Everaere, P.; Konieczny, S.; and Marquis, P. 2017. An introduction to belief merging and its links with judgment aggregation. In Endriss, U., ed., *Trends in Computational Social Choice*. AI Access. chapter 7, 123–143.
- Faliszewski, P.; Skowron, P.; Slinko, A.; and Talmon, N. 2017. Multiwinner voting: A new challenge for social choice theory. In Endriss, U., ed., *Trends in Computational Social Choice*. AI Access. chapter 2, 27–47.
- Faliszewski, P.; Slinko, A.; and Talmon, N. 2017. The complexity of multiwinner voting rules with variable number of winners. *CoRR* abs/1711.06641.
- Haret, A., and Wallner, J. P. 2019. Manipulating skeptical and credulous consequences when merging beliefs. In *Proc. JELIA*, volume 11468 of *LNCS*, 133–150. Springer.
- Kilgour, M. D. 2016. Approval elections with a variable number of winners. *Theory Decis.* 81(2):199–211.
- Konieczny, S., and Pino Pérez, R. 2002. Merging information under constraints: A logical framework. *J. Log. Comput.* 12(5):773–808.
- Konieczny, S., and Pino Pérez, R. 2011. Logic based merging. *J. Philosophical Logic* 40(2):239–270.
- Konieczny, S.; Lang, J.; and Marquis, P. 2002. Distance based merging: A general framework and some complexity results. In *Proc. KR 2002*, 97–108. Morgan Kaufmann.
- Konieczny, S.; Lang, J.; and Marquis, P. 2004. DA^2 merging operators. *Artif. Intell.* 157(1-2):49–79.
- Lackner, M., and Skowron, P. 2018a. Approval-based multi-winner rules and strategic voting. In *Proc. IJCAI*, 340–346. ijcai.org.
- Lackner, M., and Skowron, P. 2018b. Consistent approval-based multi-winner rules. In *Proc. EC*, 47–48. ACM. For details, see arXiv long version arXiv:1704.02453.
- Lang, J., and Skowron, P. 2018. Multi-attribute proportional representation. *Artif. Intell.* 263:74–106.
- List, C., and Puppe, C. 2009. Judgment aggregation: A survey. In Anand, P.; Pattanaik, P.; and Puppe, C., eds., *The Handbook of Rational and Social Choice*. Oxford University Press. chapter 19.
- Monroe, B. 1995. Fully proportional representation. *Am. Polit. Sci. Rev.* 89(4):925–940.
- Novaro, A.; Grandi, U.; Longin, D.; and Lorini, E. 2018. Goal-based collective decisions: Axiomatics and computational complexity. In *Proc. IJCAI*, 468–474. ijcai.org.
- Peters, D. 2018. Proportionality and strategyproofness in multiwinner elections. In *Proc. AAMAS*, 1549–1557. IFAA-MAS / ACM.
- Sánchez-Fernández, L.; Elkind, E.; Lackner, M.; Fernández, N.; Fisteus, J. A.; Val, P. B.; and Skowron, P. 2017. Proportional justified representation. In *Proc AAAI*, 670–676. AAAI Press.
- Skowron, P.; Lackner, M.; Brill, M.; Peters, D.; and Elkind, E. 2017. Proportional rankings. In *Proc IJCAI*, 409–415. ijcai.org.
- Thiele, T. N. 1895. Om flerfoldtsvalg. In *Oversigt over det Kongelige Danske Videnskabernes Selskabs Forhandlinger*. 415–441.