Complexity of Manipulative Attacks in Judgment Aggregation for Premise-Based Quota Rules[1]

Dorothea Baumeister, Gábor Erdélyi, Olivia J. Erdélyi, and Jörg Rothe

Abstract

Endriss et al. [26] initiated the complexity-theoretic study of problems related to judgment aggregation. We extend their results for manipulating two specific judgment aggregation procedures to a whole class of such procedures, namely to uniform premise-based quota rules. In addition, we consider incomplete judgment sets and the notions of top-respecting and closeness-respecting preferences introduced by Dietrich and List [21]. This complements previous work on the complexity of manipulation in judgment aggregation that focused on Hamming-distance-induced preferences only, which we also study here. We also introduce the notion of control by bundling judges and study it in terms of its computational complexity.

1 Introduction

Judgment Aggregation is the task of aggregating individual judgment sets of possibly interconnected logical propositions (see the surveys by List and Puppe [40] and List [39], and the bookchapter by Endriss [23]) and can therefore be seen as an important framework for collective decision making. Decision-making processes are often susceptible to various types of interference, be it internal or external. In social choice theory and in computational social choice, ways of influencing the outcome of elections—such as manipulation, bribery, and control—have been studied intensely, with a particular focus on the complexity of the related problems (see, e.g., the early work of Bartholdi et al. [4, 3, 5] and the recent surveys and bookchapters by Faliszewski et al. [31, 29, 32], Brandt et al. [15], and Baumeister et al. [9]). In particular, (coalitional) *manipulation* (see, e.g., [4, 3, 18], the survey by Faliszewski and Procaccia [31], and the bookchapter by Conitzer and Walsh [19]) refers to (a group of) strategic voters casting their votes insincerely to reach their desired outcome; in *bribery* (see, e.g., [28, 30] and the bookchapter by Faliszewski and Rothe [32]) an external agent seeks to reach her desired outcome by bribing (without exceeding a given budget) some voters to alter their votes; and in *control* (see, e.g., [5, 34] and the bookchapters by Faliszewski and Rothe [32] and Baumeister et al. [9]) an external agent (usually called the "Chair") seeks to change the structure of an election (e.g., by adding/deleting/partitioning either candidates or voters) in order to reach her desired outcome. In judgment aggregation, strategic behavior has been studied to a far lesser extent than in voting so far.

Decision-making mechanisms or systems that are susceptible to strategic behavior, be it from the agents involved as in manipulation or from external authorities or actors as in bribery and control, are obviously not desirable, as that undermines the trust we have in these systems. We therefore have a strong interest in accurately assessing how vulnerable a system for decision-making processes is to these internal or external influences. Unfortunately, in many concrete settings of social choice, "perfect" systems are impossible to exist. For example, the Gibbard–Satterthwaite theorem says that no reasonable voting system can be "strategy-proof" [33, 43] (see also the generalization by Duggan and Schwartz [22]), many natural voting systems are not "immune" to most or even all of the standard types of control [5, 34, 32], and Dietrich and List [21] give an analogue of the Gibbard–Satterthwaite theorem in judgment aggregation.

To avoid this obstacle, a common approach in computational social choice is to apply methods from theoretical computer science to show that undesirable strategic behavior is blocked, or at least hindered, by the corresponding task being a computationally intractable problem. Again, much

[1]This paper is based on and extends the work presented at ADT'13 [8].

work has been done in this regard for manipulation, bribery, and control problems in voting, but only a few results are known for these problems in judgment aggregation. Most notably, Endriss et al. [26] recently initiated the algorithmic and complexity-theoretic study of the winner determination problem and the manipulation problem in judgment aggregation, and we here extend their work for manipulation to other judgment aggregation procedures and to other notions of preference that have been introduced by Dietrich and List [21]. Baumeister et al. [10, 6, 7] investigated bribery and control in judgment aggregation. Extending their work on control, we here introduce and study control by bundling judges.

This paper is organized as follows. In Section 2, we provide the basic framework of judgment aggregation and define the relevant notions formally. We study the complexity of manipulation in judgment aggregation in Section 3 and that of control by bundling judges in Section 4. Finally, Section 5 summarizes our results and presents open problems for future research.

2 Preliminaries

We adopt the judgment aggregation framework described by Endriss et al. [26] (see also their previous conference papers [25, 24]). Let PS be the set of all propositional variables and let $\mathscr{L}_{PS}$ be the set of propositional formulas built from PS, where the following connections can be used in their usual meaning: disjunction ($\vee$), conjunction ($\wedge$), implication ($\rightarrow$), equivalence ($\leftrightarrow$), and the boolean constants 1 and 0. To avoid double negations, let $\sim\alpha$ denote the complement of α, i.e., $\sim\alpha = \neg\alpha$ if α is not negated, and $\sim\alpha = \beta$ if $\alpha = \neg\beta$. The judges have to judge over all formulas in the *agenda* Φ, which is a finite, nonempty subset of $\mathscr{L}_{PS}$ without doubly negated formulas. The agenda is required to be closed under complementation, i.e., $\sim\alpha \in \Phi$ if $\alpha \in \Phi$.

A *judgment set for an agenda* Φ is a subset $J \subseteq \Phi$. It is said to be an *individual judgment set* if it is the set of propositions in the agenda accepted by an individual judge. A *collective judgment set* is the set of propositions in the agenda accepted by all judges as the result of a judgment aggregation procedure. A judgment set J is (1) *complete* if for all $\alpha \in \Phi$, $\alpha \in J$ or $\sim\alpha \in J$; (2) *complement-free* if for no $\alpha \in \Phi$, α and $\sim\alpha$ are in J; and (3) *consistent* if there is an assignment that makes all formulas in J true. If a judgment set is complete and consistent, it is obviously complement-free. We denote the set of all complete and consistent subsets of Φ by $\mathscr{J}(\Phi)$.

A judgment aggregation procedure is a function $F : \mathscr{J}(\Phi)^n \rightarrow 2^\Phi$ that maps a profile of n individual complete and consistent judgment sets to one collective judgment set. We will call a procedure *complete* (*complement-free*, *consistent*) if the collective judgment set is always *complete* (*complement-free*, *consistent*).

The famous doctrinal paradox [37] in judgment aggregation says that if the majority rule is used, the collective judgment set can be inconsistent even if all individual judgment sets are consistent. One way of circumventing the doctrinal paradox is to impose restrictions on the agenda.[2] For example, the premise-based judgment procedure preserves consistency (and thus avoids the doctrinal paradox) by first applying the majority rule individually to the premises, and then logically deriving the result for the conclusions from the result of the premises.

Example 1 *Consider, for example, a controversial penalty situation in a soccer match with three referees having different views of the situation. According to the rules, a team must get a penalty if they have been fouled in the penalty area. The first referee says that there was a foul in the penalty area; the second referee says that what he observed in the penalty area in fact was a dive, not a foul, so there is no penalty; and the third one denies a penalty as well, since he has seen a foul outside the*

[2]Endriss et al. [26, 24] studied the question of whether one can guarantee for a specific agenda that the outcome is always complete and consistent. They established necessary and sufficient conditions on the agenda to satisfy these criteria, and they studied the complexity of deciding whether a given agenda satisfies these conditions. They also showed that deciding whether an agenda guarantees a complete and consistent outcome for the majority rule is an intractable problem.

penalty area. The three different individual judgments and the evaluation according to the majority rule are shown in Table 1(a).

Table 1: Example of the doctrinal paradox and how to prevent it by the premise-based procedure

(a) Doctrinal paradox with the majority rule

	penalty area	foul	penalty
Referee 1	yes	yes	yes
Referee 2	yes	no	no
Referee 3	no	yes	no
Majority	yes	yes	no

(b) Avoiding it with the premise-based procedure

	penalty area	foul		penalty
Referee 1	yes	yes		yes
Referee 2	yes	no		no
Referee 3	no	yes		no
PBP	yes	yes	$\Rightarrow$	yes

Applying the majority rule here leads to the inconsistent outcome that there was a foul in the penalty area, but there is no penalty. By contrast, this can be avoided by using the premise-based procedure (see Table 1(b)), where penalty area and foul are the premises, and the conclusion is penalty area and foul which equals the penalty decision.

Endriss et al. [26] introduced and studied the winner and the manipulation problem for two specific judgment aggregation procedures that always guarantee consistent outcomes: the premise-based procedure and the distance-based procedure. We will study the complexity of manipulation also for the more general class of premise-based quota rules as defined by Dietrich and List [20].

Definition 2 (Premise-based Quota Rule) *The agenda Φ is divided into two disjoint subsets $\Phi = \Phi_p \uplus \Phi_c$, where Φ_p is the set of premises and Φ_c is the set of conclusions. We assume both Φ_p and Φ_c to be closed under complementation. The premises Φ_p are again divided into two disjoint subsets, $\Phi_p = \Phi_1 \uplus \Phi_2$, such that either $\varphi \in \Phi_1$ and $\sim\varphi \in \Phi_2$, or $\sim\varphi \in \Phi_1$ and $\varphi \in \Phi_2$. Assign a quota $q_\varphi \in \mathbb{Q}$, $0 \leq q_\varphi < 1$, to each literal $\varphi \in \Phi_1$. The quota for each literal $\varphi \in \Phi_2$ is then derived by $q'_\varphi = 1 - q_\varphi$.*

Let $\|S\|$ denote the cardinality of set S and $\models$ the satisfaction relation. A premise-based quota rule is defined to be a function $PQR : \mathscr{J}(\Phi)^n \to 2^\Phi$ such that, for $\Phi = \Phi_p \uplus \Phi_c$, each profile $\mathbf{J} = (J_1, \dots, J_n)$ of individual judgment sets is mapped to the collective judgment set

$$\begin{aligned} PQR(\mathbf{J}) &= \triangle_q \cup \{\varphi \in \Phi_c \mid \triangle_q \models \varphi\}, \text{ where} \\ \triangle_q &= \{\varphi \in \Phi_1 \mid \|\{i \mid \varphi \in J_i\}\| > nq_\varphi\} \cup \{\varphi \in \Phi_2 \mid \|\{i \mid \varphi \in J_i\}\| > \lceil nq'_\varphi - 1 \rceil\}. \end{aligned}$$

To guarantee complete and consistent outcomes for this procedure, it is enough to require that Φ is closed under propositional variables and that Φ_p consists of all literals. The number of affirmations needed to be in the collective judgment set is $\lfloor nq_\varphi + 1 \rfloor$ for literals $\varphi \in \Phi_1$ and $\lceil nq'_\varphi \rceil$ for literals $\varphi \in \Phi_2$. Note that $\lfloor nq_\varphi + 1 \rfloor + \lceil nq'_\varphi \rceil = n + 1$ ensures that either $\varphi \in PQR(\mathbf{J})$ or $\sim\varphi \in PQR(\mathbf{J})$ for every $\varphi \in \Phi$. Note further that the quota $q_\varphi = 1$ for a literal $\varphi \in \Phi_1$ is not allowed here, as $n+1$ affirmations were then needed for $\varphi \in \Phi_1$ to be in the collective judgment set, which is impossible. However, $q_\varphi = 0$ is allowed, as in that case $\varphi \in \Phi_1$ needs at least one affirmation and $\sim\varphi \in \Phi_2$ needs n affirmations, which is possible. In the special case of *uniform premise-based quota rules*, there is one quota q for every literal in Φ_1, and the quota $q' = 1 - q$ for every literal in Φ_2. We will focus on such rules and denote them by $UPQR_q$. For $q = 1/2$ and the case of an odd number of judges, we obtain the premise-based procedure defined by Endriss et al. [26], and we will denote it by *PBP*.

We assume that the reader is familiar with the basic concepts of complexity theory, with complexity classes such as P and NP, and the notions of hardness and completeness with respect to the polynomial-time many-one reducibility (denoted by $\leq_{\mathrm{m}}^{\mathrm{p}}$); see, e.g., the textbooks [41, 42].

3 Manipulation in Judgment Aggregation

Recall the example from Table 1(b) illustrating how the doctrinal paradox can be avoided by the premise-based procedure. From a similar example List [38] concludes that in a premise-based procedure the judges might have an incentive to report insincere judgments. Suppose that in the example from Table 1(b) all soccer referees are absolutely sure that they are right, so they all want the aggregated outcome to be identical to their own conclusions. In this case, referee 3 knows that insincerely changing her judgment on whether there was a foul from "yes" to "no" would aggregate with the other individual judgments on this issue to a "no" by majority and thus would deny the penalty in conclusion. For the same reason, referee 2 might have an incentive to give an insincere judgment of the "penalty area" question. This is a typical *manipulation scenario*.

Strategy-proofness and manipulation have been studied in a wide variety of fields—such as voting (see, e.g., [33, 43, 31, 19]), mechanism design (see, e.g., [2]), game theory (see, e.g., [45]), fair division (see, e.g., [44, 35]), etc. In judgment aggregation, manipulability and (the game-theoretic concept of) strategy-proofness were first introduced by Dietrich and List [21]. We focus on their notion of strategy-proofness, since their (non)manipulability condition is not always appropriate in our setting. They define nonmanipulability on a given subset of the agenda by considering every proposition in this subset independently, whereas we will consider the subset as a whole.

The incentive of a manipulative attack is always to achieve a "better" result by agents (voters, players, etc.) providing untruthful information. In judgment aggregation, this untruthful information is the manipulator's individual judgment set and the result is the collective outcome of a judgment aggregation procedure. However, it is not at all obvious what a "*better*" result is. To compare two collective judgment sets, a preference over all possible judgment sets would be needed, but such preferences are rarely elicited, and they may be exponentially large in the number of formulas in the agenda. One way to avoid this obstacle, is to derive an order from a given individual judgment set. Based on the notions introduced by Dietrich and List [21], we in particular consider incomplete judgment sets and the notions of top-respecting and closeness-respecting preferences. Since most judgment aggregation rules are not strategy-proof, we study the computational complexity of the corresponding decision problems. This complements and continues previous work on the complexity of manipulation in judgment aggregation, which has been initiated by Endriss et al. [26] that focused on Hamming-distance-induced preferences, which we also study here. For a very general framework of manipulation in (both preference and judgment) aggregation, see the work of Falik and Dokow [27].

As mentioned above, we apply the notions introduced by Dietrich and List [21] to study various types of preferences. If for two judgment sets $X,Y \in \mathscr{J}(\Phi)$, X is preferred to Y for a given type of preference T and some individual judgment set J, we write $X \succ_T^J Y$.

Definition 3 *Given some individual judgment set J, we define preferences to be (strictly)*

- unrestricted (U) *if there is no restriction on* $\succ_{\mathrm{U}}^J$*;*
- top-respecting (TR) *if* $J \succ_{\mathrm{TR}}^J X$ *for all* $X \in \mathscr{J}(\Phi) \setminus \{J\}$*;*
- closeness-respecting (CR) *if for all* $X,Y \in \mathscr{J}(\Phi)$*, we have* $X \succ_{\mathrm{CR}}^J Y$ *if* $Y \cap J \subset X \cap J$*;*
- Hamming-distance-induced (HD) *if for all* $X,Y \in \mathscr{J}(\Phi)$*,* $X \succ_{\mathrm{HD}}^J Y$ *if and only if* $\mathrm{HD}(X,J) < \mathrm{HD}(Y,J)$*, where the Hamming distance* $\mathrm{HD}(X,Y)$ *between two (possibly incomplete) judgment sets X and Y is the number of disagreements on propositions that occur in both judgment sets.*

By allowing equalities, the Hamming-distance-induced preference is the only complete relation among the above. Intuitively, unrestricted preferences capture the setting where we know nothing about the individual preferences. The slightly more restricted case of top-respecting preferences at least requires the given judgment set to be the most preferred one. This also holds for closeness-respecting preferences, but in addition judgment sets that have additional agreement are preferred. In contrast, the Hamming-distance-induced preferences focus only on the total number of disagreements. Hence, for $X, Y \in \mathscr{J}(\Phi)$, if $X \succ^J_{\mathrm{TR}} Y$ then $X \succ^J_{\mathrm{CR}} Y$, and if $X \succ^J_{\mathrm{CR}} Y$ then $X \succ^J_{\mathrm{HD}} Y$.

Example 4 *For variables a, b, c, and d, let the agenda contain the formulas*

$$a, \quad b, \quad c, \quad d, \quad a \vee b, \quad b \vee c, \quad a \vee c, \quad b \vee d,$$

and their negations. The individual judgment sets of three judges are shown in Table 2. A 0 *indicates that the negation of the formula is in the judgment set, and a* 1 *indicates that the formula itself is contained in the judgment set.*

Table 2: Applying the premise-based judgment aggregation procedure

	a	b	c	d		$a \vee b$	$b \vee c$	$a \vee c$	$b \vee d$
Judge 1	1	1	0	0		1	1	1	1
Judge 2	0	0	0	0		0	0	0	0
Judge 3	1	0	1	1		1	1	1	1
PBP	1	0	0	0	$\Rightarrow$	1	0	1	0

The result according to the premise-based procedure is also given in the table. Now assume that the third judge is trying to manipulate and reports the untruthful individual judgment set $\{a, b, c, d\}$ *and the corresponding conclusions. Then the collective outcome equals the individual judgment set of the first judge.*

- *If the manipulator has unrestricted preferences, we do not know whether she prefers this new outcome or not.*
- *If she has closeness-respecting preferences, we again do not know whether she prefers the new outcome, since the agreement on* $\neg b$ *is no longer given. However, if she is interested only in the conclusions, then she does prefer the new outcome, since the agreement on* $a \vee b$ *and* $a \vee c$ *is preserved and there are the two additional agreements on* $b \vee c$ *and* $b \vee d$*.*
- *The same holds for top-respecting preferences: If the manipulator is interested in the whole collective judgment set, we do not know which outcome is better for her, but restricted to the conclusions the new outcome equals her initial individual judgment set and thus is preferred to all other outcomes.*
- *If the manipulator has Hamming-distance-induced preferences, we know that the new outcome is preferred to the old one, since before the manipulation the Hamming distance was* 4*, but now it is only* 3*.*

Konczak and Lang [36] (see also the work of Xia and Conitzer [46]) introduced the notions of necessary and possible winner in voting. A possible winner is a candidate who can be made a winner by *some* extension of a given partial preference profile to a complete profile, and a necessary winner is a candidate who wins for *every* complete extension of a given partial preference profile. Inspired by their notions,[3] we now introduce the notions of necessary and possible strategy-proofness in judgment aggregation.

Just as Dietrich and List [21], we study settings where the desired judgment set is incomplete, to also capture their "reason-oriented" and "outcome-oriented" preferences. However, we will not

[3] See also the remotely related notions of "possible envy-freeness" vs "necessary envy-freeness" in fair division that are due to Bouveret et al. [12] (see also the papers by Brams et al. [13, 14]).

generally restrict the desired judgment set to the premises or the conclusions; rather, we allow arbitrary incomplete desired judgment sets (which still must have a consistent extension to the whole agenda). In this case, we restrict the preferences to the formulas that occur in the desired judgment set. Since we want to compare two preferences with each other, but most of the induced preferences will be incomplete, we distinguish the cases where the relation between them is known or unknown.

Definition 5 *Let* $T \in \{\mathrm{U}, \mathrm{TR}, \mathrm{CR}\}$ *be a type of induced preferences and J, X, and Y individual judgment sets. (1) A judge* necessarily prefers X to Y for type T and individual judgment set J *if* $X >_T^J Y$ *for all complete extensions of* $\succ_T^J$. *(2) A judge* possibly prefers X to Y for type T and individual judgment set J *if* $X >_T^J Y$ *for some complete extension of* $\succ_T^J$.

A judgment aggregation rule F is necessarily/possibly strategy-proof with respect to induced preferences of type $T \in \{\mathrm{U}, \mathrm{TR}, \mathrm{CR}\}$ *if for all profiles* $(J_1, \ldots, J_n)$ *and each i,* $1 \leq i \leq n$, *judge i necessarily/possibly prefers the outcome* $F(J_1, \ldots, J_n)$ *to the outcome* $F(J_1, \ldots, J_{i-1}, J_i^*, J_{i+1}, \ldots, J_n)$ *(with respect to preferences of type T and the individual judgment set* J_i*) for any* $J_i^* \in \mathscr{J}(\Phi)$ *with* $F(J_1, \ldots, J_n) \neq F(J_1, \ldots, J_{i-1}, J_i^*, J_{i+1}, \ldots, J_n)$.

Definition 5 applies to complete desired judgment sets J_i only. More generally, the definition can easily be extended to incomplete desired judgment sets $J \subseteq J_i$ as well.

The stronger notion of *necessary strategy-proofness* corresponds to the "strategy-proofness" condition defined by Dietrich and List [21], whereas the weaker notion of *possible strategy-proofness* is introduced here. Note that since the Hamming-distance-induced preferences are a complete relation, we simply say that F is *strategy-proof (with respect to Hamming-distance-induced preferences)* if for each individual judge the actual outcome is at least as good as all outcomes obtained by reporting a different individual judgment set.

The result of Dietrich and List [21] says that an aggregation rule that satisfies the "universal domain" condition is necessarily strategy-proof with respect to nonstrict closeness-respecting preferences if and only if it is independent and monotonic. *Universal domain* is satisfied if the domain of the aggregation function is the set of all possible profiles from $\mathscr{J}(\Phi)^n$, which obviously is true for $UPQR_q$. *Independence* means that the collective decision on each proposition only relies on the individual judgments of this proposition. Since $UPQR_q$ derives the outcome for the conclusions from the outcome of the premises, it is not independent and hence not necessarily strategy-proof with respect to nonstrict closeness-respecting preferences. An aggregation function is *monotonic* if additional support for some proposition that is currently accepted may never result in a nonacceptance for this formula, provided everything else remains unchanged. In the case where the agenda contains solely premises, $UPQR_q$ is independent and monotonic, and hence necessarily strategy-proof also for the case of strict closeness-respecting preferences.

Define the related manipulation problems for uniform premise-based quota rules and a given preference type T.

$UPQR_q$-T-NECESSARY-MANIPULATION	
Given:	An agenda Φ, a profile $\mathbf{J} = (J_1, \ldots, J_n) \in \mathscr{J}(\Phi)^n$, and the manipulator's desired consistent (possibly incomplete) judgment set $J \subseteq J_n$.
Question:	Does there exist a judgment set $J^* \in \mathscr{J}(\Phi)$ such that $UPQR_q(J_1, \ldots, J_{n-1}, J^*)\|_J >_T^J UPQR_q(J_1, \ldots, J_n)\|_J$ for all extensions $>_T^J$ that are consistent with $\succ_T^J$?

Here, $UPQR_q(J_1, \ldots, J_n)|_J$ denotes the restriction of $UPQR_q(J_1, \ldots, J_n)$ to the formulas that occur, negated or not, in the manipulator's desired judgment set J.

In $UPQR_q$-T-POSSIBLE-MANIPULATION, we for the same input ask whether there exists a judgment set $J^* \in \mathscr{J}(\Phi)$ such that $UPQR_q(J_1, \ldots, J_{n-1}, J^*)|_J >_T^J UPQR_q(J_1, \ldots, J_n)|_J$ for *some* extension $>_T^J$ that is consistent with $\succ_T^J$. In the case of Hamming-distance-induced preferences we will simply say $UPQR_q$-HD-MANIPULATION, since the relation between two given judgment sets is always known.

Furthermore, we introduce and study the exact variant, $UPQR_q$-EXACT-MANIPULATION, where the manipulator seeks to achieve not only a better, but a *best* outcome for a given subset of her desired judgment set. Here, the question is whether there is some judgment set $J^* \in \mathscr{J}(\Phi)$ such that $J \subseteq UPQR_q(J_1, \dots, J_{n-1}, J^*)$.

We start by showing that exact manipulation is hard to achieve for uniform premise-based quota rules.

Theorem 6 *For each rational quota q, $0 \leq q < 1$, $UPQR_q$-EXACT-MANIPULATION is* NP-*complete, even for only three judges.*

Proof. We will only present the proof for $q = 1/2$; the remaining cases can be shown by slightly adapting this proof.

The proof for $q = 1/2$ is by a reduction from the NP-complete satisfiability problem. Let φ be a given formula in conjunctive normal form, where the clauses are built from the set $A = \{\alpha_1, \dots, \alpha_m\}$ of variables. The question is whether there is a satisfying assignment for this formula. Without loss of generality, we may assume that neither setting all variables to true, nor setting all variables to false is a satisfying assignment for φ. Now construct an agenda Φ that consists of the variables in A and their negations, an additional variable β and its negation, and the formula $\varphi \vee \beta$ and its negation. The profile $\mathbf{J} = (J_1, J_2, J_3)$ consists of three individual judgment sets. The first one, J_1, contains A, $\neg\beta$, and $\neg(\varphi \vee \beta)$, and the second one, J_2, contains $\neg\alpha_i$ for each i, $1 \leq i \leq m$, $\neg\beta$, and $\neg(\varphi \vee \beta)$. The third judge is the manipulative one and his individual judgment set, J_3, contains A, β, and $(\varphi \vee \beta)$. His desired outcome consists of the conclusion $\varphi \vee \beta$ only. It holds that

$$UPQR_{1/2}(\mathbf{J}) = A \cup \{\neg\beta\} \cup \{\neg(\varphi \vee \beta)\}.$$

Note also that the third judge is decisive for every formula in A, and that independently of the individual judgment set of the manipulator, β is never contained in the collective judgment set. Hence, the only way to obtain the conclusion $\varphi \vee \beta$ in the collective outcome is to evaluate the formula φ to true. This implies that there is a satisfying assignment for φ if and only if the individual judgment set of the third judge can be modified such that $\varphi \vee \beta$ is contained in the collective outcome. ❑

Next, we provide generic relations between the various manipulation problems we have defined.

Theorem 7 *For each uniform premise-based quota rule with rational quota q, $0 \leq q < 1$,*

1. *$UPQR_q$-*EXACT-MANIPULATION $\leq_m^p$ *$UPQR_q$-T-*NECESSARY-MANIPULATION *for each type $T \in \{\mathrm{TR}, \mathrm{CR}\}$,*
2. *$UPQR_q$-*EXACT-MANIPULATION $\leq_m^p$ *$UPQR_q$-T-*POSSIBLE-MANIPULATION *for each type $T \in \{\mathrm{U}, \mathrm{TR}, \mathrm{CR}\}$, and*
3. *$UPQR_q$-*EXACT-MANIPULATION $\leq_m^p$ *$UPQR_q$-*HD-MANIPULATION.

Proof. For the exact problem, we have an agenda Φ, some profile $\mathbf{J} = (J_1, \dots, J_n)$, and some desired judgment set $J = \{\alpha_1, \dots, \alpha_m\} \subseteq J_n$, and we are looking for a modified judgment set J_n^* such that $J \subseteq UPQR_q(J_1, \dots, J_{n-1}, J_n^*)$. In the trivial case that $J \subseteq UPQR_q(\mathbf{J})$, $J_n^* = J_n$ obviously fulfills the requirement, so we can construct an arbitrary yes-instance for the corresponding manipulation problem. We will prove all three assertions via the same reduction, but using different arguments.

Assume that $J \smallsetminus UPQR_q(\mathbf{J}) \neq \emptyset$ and consider the following problem. Fix some $T \in \{\mathrm{TR}, \mathrm{CR}, \mathrm{HD}\}$, let the agenda Φ' be the union of Φ, the formula $\varphi = \alpha_1 \wedge \dots \wedge \alpha_m$, and its negation. Let $\mathbf{J}' \in \mathscr{J}(\Phi')^n$ be the consistent extensions of $\mathbf{J}$. In particular, $J_n' = J_n \cup \{\varphi\}$. Let the desired judgment set be $J' = \{\varphi\}$, and we are looking for a modified judgment set $J_n'^*$ such that for all/some extensions $>_T^{J'}$ of $\succ_T^{J'}$, we have $UPQR_q(J_1', \dots, J_{n-1}', J_n'^*)|_{J'} >_T^{J'} UPQR_q(J_1', \dots, J_{n-1}', J_n')|_{J'}$. Since J' consists of the single formula φ, there are only two different collective outcomes when restricted to

J'. Since $\varphi \in J'$, it obviously holds that $\varphi \succ_T^{J'} \neg\varphi$ for all $T \in \{\text{TR}, \text{CR}, \text{HD}\}$, and since in this case $\succ_T^{J'}$ is complete, there is no difference between the notions of necessary and possible preference. In the case of unrestricted preferences and the possible manipulation problem, we ask whether there is some different outcome, since they all may be possibly preferred. Since there is some J_n^* with $J \subseteq UPQR_q(J_1, \ldots, J_{n-1}, J_n^*)$ if and only if there is some $J_n'^*$ with $\varphi \in UPQR_q(J_1', \ldots, J_{n-1}', J_n'^*)$, the reduction works in all cases. ❑

Note that this reduction requires a *partial* desired judgment set of the manipulator for $UPQR_q$-T-NECESSARY-MANIPULATION, $UPQR_q$-T-POSSIBLE-MANIPULATION, and $UPQR_q$-HD-MANIPULATION. Together with Theorem 6 (and the obvious NP upper bounds of these problems), this implies NP-completeness of $UPQR_q$-HD-MANIPULATION, $UPQR_q$-T-NECESSARY-MANIPULATION for $T \in \{\text{TR}, \text{CR}\}$, and $UPQR_q$-T-POSSIBLE-MANIPULATION for $T \in \{\text{U}, \text{TR}, \text{CR}\}$ whenever the desired judgment set of the manipulator is incomplete. Alternatively, the reduction given by Endriss et al. [26] in fact shows NP-completeness for *PBP*-HD-MANIPULATION even if the desired judgment set of the manipulator is complete. By contrast, if the manipulator's desired judgment set is complete, the possible manipulation problem turns out to be easy to solve for unrestricted and top-respecting preferences.

Proposition 8 *For $T \in \{\text{U}, \text{TR}\}$ and for each rational quota q, $0 \leq q < 1$, $UPQR_q$-T-*POSSIBLE-MANIPULATION *can be solved in polynomial time if the desired judgment set of the manipulator is complete.*

Proof. This result holds, since a $UPQR_q$-U-POSSIBLE-MANIPULATION instance is positive exactly if there is some premise from the desired judgment set for which the manipulator is decisive, i.e., the collective outcome depends on the decision of the manipulator. For a $UPQR_q$-TR-POSSIBLE-MANIPULATION instance to be positive, it must additionally be required that the desired judgment set is not the actual outcome. ❑

Proposition 9 *If the desired judgment set of the manipulator is complete and he tries to exactly reach his desired outcome, then $UPQR_q$, $0 \leq q < 1$, is strategy-proof.*

Proof. Note that the premises are considered independently. Let n be the number of judges. If some φ from the premises is contained in the judgment set J of the manipulator, and φ does not have $\lfloor n \cdot q + 1 \rfloor$ (respectively, $\lceil n(1-q) \rceil$) affirmations without considering J, it cannot reach the required number of affirmations if the manipulator switches from φ to $\sim\varphi$ in his judgment set. ❑

Finally, we state a result on possible strategy-proofness for the premise-based procedure. Note that this does not contradict the results of Dietrich and List [21], since they impose different conditions on nonmanipulability and nonstrict preferences.

Proposition 10 *If the desired judgment set of the manipulator is complete and top-respecting or closeness-respecting preferences are assumed, then $UPQR_q$, $0 \leq q < 1$, is possibly strategy-proof.*

Proof. In case of possible strategy-proofness, there may be no alternative outcome resulting from an untruthful individual judgment set of the manipulator that is necessarily preferred to the actual outcome. If closeness-respecting preferences are assumed, a judgment set that is necessarily preferred to the actual collective outcome must preserve all agreements between the desired judgment set and the actual outcome. If top-respecting preferences are assumed, a judgment set that is necessarily preferred to the actual collective outcome must equal the manipulators individual true judgment set.

Now consider a premise α that is contained in the collective judgment set, but $\sim\alpha$ is contained in the desired judgment set. Obviously, it can never be the case that the manipulator switching from

$\sim\alpha$ to α would cause $\sim\alpha$ to be in the collective judgment set. Hence there can be no additional agreement among the premises. Since the desired judgment set is complete and the outcome for the conclusions depends solely on the outcome of the premises, $UPQR_q$ is possibly strategy-proof in both cases. ❑

Proposition 11 *Assuming unrestricted preferences,* $UPQR_q$*,* $0 \leq q < 1$*, is possibly strategy-proof.*

Proof. In case of unrestricted preferences, we know nothing about the preference of the manipulator. Hence, the actual outcome is always possibly preferred to all outcomes that result from a different individual judgment set of the manipulator. ❑

4 Control by Bundling Judges

Previous work on control in judgment aggregation (see [6, 7]) considered the problems of control by adding, deleting, or replacing judges. While adding and deleting judges is inspired by the corresponding control problems in voting, explicit examples for such control actions in judgment aggregation have been given, and the third type, control by replacing judges, was motivated by real-world examples from international arbitration.

We here introduce another type of control motivated by real-world scenarios, *control by bundling judges*, which is remotely akin to control by partitioning voters in voting. A prominent natural example for control by bundling judges can be found in European legislation. Certain European legislative acts, such as Directives, give considerable freedom to Member States regarding the concrete implementation of these acts. Yet, in some cases uniform implementation is crucial, so the basic act confers implementing powers on the European Commission or the Council of the European Union to adopt the required implementing acts.[4] The exercise of implementing powers through the Commission and Council is controlled by the member states through so-called comitology committees in accordance with previously specified rules.[5] The committees are set up by the basic act in question.[6] Some of these committees are concerned with such a broad range of issues that they are divided into subcommittees, each of which is dealing with different issues. When preparing implementing acts covering several issues, each subcommittee votes on the issues assigned to it, and the implementing act is shaped according to the decisions of the different subcommittees.[7]

The formal definition for the Hamming-distance-induced version control by bundling judges is as follows. In the problem definition below, we will use the notation $\Delta = \bigcup_{1 \leq i \leq k} UPQR_q(\mathbf{J}|_{\Phi_p^i, N_i})$, where $UPQR_q(\mathbf{J}|_{\Phi_p^i, N_i})$ is the collective judgment set obtained by restricting the premises Φ_p of the agenda to its part Φ_p^i in a partition (see the problem definition below) and the set of judges to $N_i \subseteq N$.

The formal definition is as follows.

$UPQR_q$-CONTROL BY BUNDLING JUDGES	
Given:	An agenda Φ, where the premises are partitioned into k subsets $\Phi_p^1, \dots, \Phi_p^k$, a complete profile $\mathbf{J} \in \mathscr{J}(\Phi)^n$, and a consistent judgment set $J \subseteq \widehat{J} \in \mathscr{J}(\Phi)$ (not necessarily complete).
Question:	Is there a partition $\{N_1, \dots, N_k\}$ of the n judges such that $$\mathrm{HD}(J, \Delta \cup \{\varphi \in \Phi_c \mid \Delta \models \varphi\}) < \mathrm{HD}(J, UPQR_q(\mathbf{J}))?$$

[4] Article 291 of the Treaty on the Functioning of the European Union.

[5] Regulation (EU) No 182/2011 of the European Parliament and of the Council of 16 February 2011 laying down the rules and general principles concerning mechanisms for control by Member States of the Commission's exercise of implementing powers (Implementing Acts Regulation).

[6] Recital 6 of the Preamble of Implementing Acts Regulation.

[7] One example is the Customs Code Committee, see Articles 1 (1) and 5 (7) (8) of the Rules of procedure for the Customs Code Committee.

In *UPQR*$_q$-EXACT CONTROL BY BUNDLING JUDGES we ask, for the same input, whether there is a partition $\{N_1,\dots,N_k\}$ of the n judges such that

$$J \subseteq \Delta \cup \{\varphi \in \Phi_c \mid \Delta \models \varphi\}).$$

Example 12 *Consider the same variables a,b,c, and d and the same individual judgment sets as in Example 4. Let the quota be $q = 1/2$ for every positive literal in the agenda. Assume that the set of premises is partitioned into $\Phi_1^p = \{a,b\}$ and $\Phi_2^p = \{c,d\}$, and that the desired judgment set J contains*

$$a \vee b, \quad b \vee c, \quad \neg(a \vee c), \quad \text{and} \quad b \vee d.$$

Note that this is a consistent judgment set, since it can be reached by accepting b and the negation of all other variables. The Hamming distance between the current collective outcome and J is 3. But if we partition the set of judges into two groups, where the first judge forms the first group and the last two judges are in the second group, the outcome is as shown in Table 3, where the individual judgments for a single variable not belonging to the group who decides over this variable are marked with $\not 1$ or $\not 0$. Recall that the negative literal is contained in the collective judgment set in case of a tie by convention.

Table 3: Example for CONTROL BY BUNDLING JUDGES

	a	b	c	d		$a\vee b$	$b\vee c$	$a\vee c$	$b\vee d$
Judge 1	1	1	$\not 0$	$\not 0$		1	1	1	1
Judge 2	$\not 0$	$\not 0$	0	0		0	0	0	0
Judge 3	$\not 1$	$\not 0$	1	1		1	1	1	1
UPQR$_{1/2}$	1	1	0	0	$\Rightarrow$	1	1	1	1

*After bundling the judges, the Hamming distance between the collective outcome and J has decreased to 1. Hence, this is a positive instance of UPQR$_{1/2}$-*CONTROL BY BUNDLING JUDGES. *However, since it is not possible to bundle the judges into two groups to obtain exactly J as a subset of the collective outcome, it is a negative instance of UPQR$_{1/2}$-*EXACT CONTROL BY BUNDLING JUDGES.

Remotely related bundling problems in judgment aggregation have recently been studied by Alon et al. [1]. However, their setting is different from ours. They consider judgment aggregation over independent variables, and only the variables are bundled in their bundling attacks. It is assumed that then all judges decide over all bundles by deciding uniformly for all variables contained in the same bundle. Furthermore, the goal in their model is to always accept all positive variables, that is, a complete desired judgment set. This setting in fact covers a restriction of judgment aggregation known as optimal lobbying (see the papers by Christian et al. [17], Binkele-Raible et al. [11], and Bredereck et al. [16]).

To study the computational complexity of bundling judges, we adopt the terminology introduced by Bartholdi, Tovey, and Trick [5] for control problems in voting and adapt it to judgment aggregation. Let $\mathscr{C}$ be a given control type. *UPQR*$_q$ is said to be *immune* to control by $\mathscr{C}$ if it is never possible for the chair to successfully control the judgment aggregation procedure via $\mathscr{C}$-control. *UPQR*$_q$ is said to be *susceptible* to control by $\mathscr{C}$ if it is not immune. *UPQR*$_q$ is said to be *resistant* to control by $\mathscr{C}$ if it is susceptible and the corresponding decision problem is NP-hard. *UPQR*$_q$ is said to be *vulnerable* to control by $\mathscr{C}$ if it is susceptible and the corresponding decision problem is in P.

UPQR$_q$-CONTROL BY BUNDLING JUDGES is somewhat similar to the problem *UPQR*$_q$-CONTROL BY DELETING JUDGES defined in [6, 7]. We will exploit this in the following proof. Note that it does not make sense to consider uniform *constant* premise-based quota rules for control by bundling judges: If we have a constant number of judges and then partition the group of judges, bundling them to smaller groups, it wouldn't be reasonable to have the original constant number of judges carry over to the smaller groups.

Theorem 13 *$UPQR_{1/2}$ is resistant to exact control by bundling judges and to control by bundling judges.*

Proof. The proof will be by a reduction from the related problem $UPQR_{1/2}$-EXACT CONTROL BY DELETING JUDGES. Given an agenda $\Phi = \Phi_p \cup \Phi_c$, a complete profile $\mathbf{J} \in \mathscr{J}(\Phi)^n$, and a positive integer k as a bound on the number of judges that may be deleted. The quota 1/2 holds for every positive literal in the agenda. We assume that the desired judgment set is J. Now, we construct an instance of $UPQR_{1/2}$-EXACT CONTROL BY BUNDLING JUDGES, resistance for $UPQR_{1/2}$-CONTROL BY BUNDLING JUDGES then follows easily. Without loss of generality, we assume that $n \geq k+2$. The agenda is $\Phi' = \Phi \cup \{\alpha, \neg\alpha\}$, and the premises are divided into two subsets. The first one consists of Φ_p, and the second one is $\{\alpha, \neg\alpha\}$. The quota 1/2 again holds for every positive literal in the agenda. The profile $\mathbf{S} \in \mathscr{J}(\Phi')^{n+k+1}$ contains the individual judgment sets from $\mathbf{J}$, each extended by $\neg\alpha$. Furthermore, there are $k+1$ new individual judgment sets that each contain $\varphi \in \Phi_p$ if and only if $\sim\varphi \in J$, they each contain α, and the conclusions are evaluated accordingly. These $k+1$ new judges will be denoted by N'. The desired judgment set is $J' = J \cup \{\alpha\}$. We show that it is possible to obtain the desired judgment set J by deleting at most k judges from $\mathbf{J}$ if and only if the judges from $\mathbf{S}$ can be bundled into two groups such that the desired outcome is J'.

From left to right, assume that there is a subset $\mathbf{T}' \subseteq \mathbf{J}$, $\|\mathbf{T}'\| \leq k$, such that $UPQR_{1/2}(\mathbf{J} \setminus \mathbf{T}') = J$. Then the judges can be bundled as follows. The $k+1$ new judges and the judges corresponding to $\mathbf{T}'$ decide over α. Then obviously α is contained in the collective outcome, hence the constructed instance is a positive one for $UPQR_{1/2}$-EXACT CONTROL BY BUNDLING JUDGES.

From right to left, assume that the judges can be bundled into N_1 and N_2 such that the collective outcome is J'. Hence, it holds that $UPQR_{1/2}(\mathbf{S}|_{\Phi,N_1}) = J$. We will show that $\|N_2 \setminus N'\| \leq k$ and $UPQR_{1/2}(\mathbf{S}|_{\Phi,N_1 \setminus N'}) = J$. Since α is contained in the collective judgment set and since there are only $k+1$ judges having α in their individual judgment set, at most k of the initial judges can be in N_2. Due to the premise-based procedure, it is enough to show that $UPQR_{1/2}(\mathbf{S}|_{\Phi_p,N_1}) = UPQR_{1/2}(\mathbf{S}|_{\Phi_p,N_1 \setminus N'})$. This holds trivially, since for all judges from N' it holds that $\varphi \in \Phi_p$ is contained in the individual judgment set if and only if $\sim\varphi \in J$. ❑

5 Conclusions and Open Questions

We have studied the complexity of problems related to manipulation and a new control type in judgment aggregation. In particular, for manipulation, we have extended the results of Endriss et al. [26] from two specific judgment aggregation procedures to the class of uniform premise-based quota rules. Moreover, our results also apply to incomplete judgment sets and the notions of top-respecting and closeness-respecting preferences that are due to Dietrich and List [21]. Table 4 gives an overview of our results for manipulation problems with uniform premise-based quota rules. In this table, "DJS" stands for "desired judgment set" and "NP-c" for "NP-complete."

More specifically, we have introduced and studied the notions of necessary and possible strategy-proofness in judgment aggregation, which are inspired by the notions of necessary and possible winner in voting (see the work of Konczak and Lang [36] and Xia and Conitzer [46]). Note that distinguishing between these notions does not apply to the exact problem variants, nor to manipulation problems based on Hamming-distance-induced preferences. Only one of the cases considered in Table 4 remains open: the complexity of possible manipulation for complete desired judgment sets with respect to closeness-respecting preferences.

Table 4: Overview of results for manipulation problems with uniform premise-based quota rules

	U	TR	CR	HD	EXACT
POSSIBLE-MANIPULATION for incomplete DJS	NP-c	NP-c	NP-c	NP-c	NP-c
NECESSARY-MANIPULATION for incomplete DJS	possibly strategy-proof	NP-c	NP-c		
POSSIBLE-MANIPULATION for complete DJS	in P	in P	?	NP-c[8]	strategy-proof
NECESSARY-MANIPULATION for complete DJS	possibly strategy-proof	possibly strategy-proof	possibly strategy-proof		

Finally, we have introduced and studied a new control scenario in judgment aggregation, control by bundling judges, in addition to the previously studied scenarios of control by adding, deleting, or replacing judges [6, 7].

Acknowledgments: We thank the anonymous ADT'13 reviewers for their helpful comments and literature pointers. This work was supported in part by DFG grants RO-1202/15-1, ER 738/1-1, and ER 738/2-1 and an NRW grant for gender-sensitive universities.

References

[1] N. Alon, D. Falik, R. Meir, and M. Tennenholtz. Bundling attacks in judgment aggregation. In *Proceedings of the 27th AAAI Conference on Artificial Intelligence*, pages 39–45. AAAI Press, 2013.

[2] N. Alon, M. Feldman, A. Procaccia, and M. Tennenholtz. Strategyproof approximation of the minimax on networks. *Mathematics of Operations Research*, 35(3):513–526, 2013.

[3] J. Bartholdi III and J. Orlin. Single transferable vote resists strategic voting. *Social Choice and Welfare*, 8(4):341–354, 1991.

[4] J. Bartholdi III, C. Tovey, and M. Trick. The computational difficulty of manipulating an election. *Social Choice and Welfare*, 6(3):227–241, 1989.

[5] J. Bartholdi III, C. Tovey, and M. Trick. How hard is it to control an election? *Mathematical and Computer Modelling*, 16(8/9):27–40, 1992.

[6] D. Baumeister, G. Erdélyi, O. Erdélyi, and J. Rothe. Bribery and control in judgment aggregation. In F. Brandt and P. Faliszewski, editors, *Proceedings of the 4th International Workshop on Computational Social Choice*, pages 37–48. AGH University of Science and Technology, Kraków, Poland, September 2012.

[7] D. Baumeister, G. Erdélyi, O. Erdélyi, and J. Rothe. Control in judgment aggregation. In *Proceedings of the 6th European Starting AI Researcher Symposium*, pages 23–34. IOS Press, August 2012.

[8] D. Baumeister, G. Erdélyi, O. Erdélyi, and J. Rothe. Computational aspects of manipulation and control in judgment aggregation. In *Proceedings of the 3rd International Conference on Algorithmic Decision Theory*, pages 71–85. Springer-Verlag *Lecture Notes in Artificial Intelligence #8176*, November 2013.

[8] This result is due to Endriss et al. [26] for the special case of the premise-based procedure.

[9] D. Baumeister, G. Erdélyi, E. Hemaspaandra, L. Hemaspaandra, and J. Rothe. Computational aspects of approval voting. In J. Laslier and R. Sanver, editors, *Handbook on Approval Voting*, chapter 10, pages 199–251. Springer, 2010.

[10] D. Baumeister, G. Erdélyi, and J. Rothe. How hard is it to bribe the judges? A study of the complexity of bribery in judgment aggregation. In *Proceedings of the 2nd International Conference on Algorithmic Decision Theory*, pages 1–15. Springer-Verlag *Lecture Notes in Artificial Intelligence #6992*, October 2011.

[11] D. Binkele-Raible, G. Erdélyi, H. Fernau, J. Goldsmith, N. Mattei, and J. Rothe. The complexity of probabilistic lobbying. *Discrete Optimization*, 11(1):1–21, 2014.

[12] S. Bouveret, U. Endriss, and J. Lang. Fair division under ordinal preferences: Computing envy-free allocations of indivisible goods. In *Proceedings of the 19th European Conference on Artificial Intelligence*, pages 387–392. IOS Press, August 2010.

[13] S. Brams, P. Edelman, and P. Fishburn. Fair division of indivisible items. *Theory and Decision*, 5(2):147–180, 2004.

[14] S. Brams and D. King. Efiicient fair division—help the worst off or avoid envy? *Rationality and Society*, 17(4):387–421, 2005.

[15] F. Brandt, V. Conitzer, and U. Endriss. Computational social choice. In G. Weiß, editor, *Multiagent Systems*, pages 213–283. MIT Press, second edition, 2012.

[16] R. Bredereck, J. Chen, S. Hartung, S. Kratsch, R. Niedermeier, and O. Suchý. A multivariate complexity analysis of lobbying in multiple referenda. In *Proceedings of the 26th AAAI Conference on Artificial Intelligence*, pages 1292–1298. AAAI Press, 2012.

[17] R. Christian, M. Fellows, F. Rosamond, and A. Slinko. On complexity of lobbying in multiple referenda. *Review of Economic Design*, 11(3):217–224, 2007.

[18] V. Conitzer, T. Sandholm, and J. Lang. When are elections with few candidates hard to manipulate? *Journal of the ACM*, 54(3):Article 14, 2007.

[19] V. Conitzer and T. Walsh. Barriers to manipulation. In F. Brandt, V. Conitzer, U. Endriss, J. Lang, and A. Procaccia, editors, *Handbook of Computational Social Choice*, chapter 6. Cambridge University Press, to appear.

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

[21] F. Dietrich and C. List. Strategy-proof judgment aggregation. *Economics and Philosophy*, 23(3):269–300, 2007.

[22] J. Duggan and T. Schwartz. Strategic manipulability without resoluteness or shared beliefs: Gibbard–Satterthwaite generalized. *Social Choice and Welfare*, 17(1):85–93, 2000.

[23] U. Endriss. Judgment aggregation. In F. Brandt, V. Conitzer, U. Endriss, J. Lang, and A. Procaccia, editors, *Handbook of Computational Social Choice*, chapter 17. Cambridge University Press, to appear.

[24] U. Endriss, U. Grandi, and D. Porello. Complexity of judgment aggregation: Safety of the agenda. In *Proceedings of the 9th International Joint Conference on Autonomous Agents and Multiagent Systems*, pages 359–366. IFAAMAS, May 2010.

[25] U. Endriss, U. Grandi, and D. Porello. Complexity of winner determination and strategic manipulation in judgment aggregation. In V. Conitzer and J. Rothe, editors, *Proceedings of the 3rd International Workshop on Computational Social Choice*, pages 139–150. Universität Düsseldorf, September 2010.

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

[27] D. Falik and E. Dokow. Models of manipulation on aggregation of binary evaluations. In F. Brandt and P. Faliszewski, editors, *Proceedings of the 4th International Workshop on Computational Social Choice*, pages 191–202. AGH University of Science and Technology, Kraków, Poland, September 2012.

[28] P. Faliszewski, E. Hemaspaandra, and L. Hemaspaandra. How hard is bribery in elections? *Journal of Artificial Intelligence Research*, 35:485–532, 2009.

[29] P. Faliszewski, E. Hemaspaandra, and L. Hemaspaandra. Using complexity to protect elections. *Communications of the ACM*, 53(11):74–82, 2010.

[30] P. Faliszewski, E. Hemaspaandra, L. Hemaspaandra, and J. Rothe. Llull and Copeland voting computationally resist bribery and constructive control. *Journal of Artificial Intelligence Research*, 35:275–341, 2009.

[31] P. Faliszewski and A. Procaccia. AI's war on manipulation: Are we winning? *AI Magazine*, 31(4):53–64, 2010.

[32] P. Faliszewski and J. Rothe. Control and bribery. In F. Brandt, V. Conitzer, U. Endriss, J. Lang, and A. Procaccia, editors, *Handbook of Computational Social Choice*, chapter 7. Cambridge University Press, to appear.

[33] A. Gibbard. Manipulation of voting schemes. *Econometrica*, 41(4):587–601, 1973.

[34] E. Hemaspaandra, L. Hemaspaandra, and J. Rothe. Anyone but him: The complexity of precluding an alternative. *Artificial Intelligence*, 171(5–6):255–285, 2007.

[35] B. Klaus and E. Miyagawa. Strategy-proofness, solidarity, and consistency for multiple assignment problems. *International Journal of Game Theory*, 30(3):421–435, 2013.

[36] K. Konczak and J. Lang. Voting procedures with incomplete preferences. In *Proceedings of the Multidisciplinary IJCAI-05 Workshop on Advances in Preference Handling*, pages 124–129, July/August 2005.

[37] L. Kornhauser and L. Sager. Unpacking the court. *Yale Law Journal*, 96(1):82–117, 1986.

[38] C. List. The discursive dilemma and public reason. *Ethics*, 116(2):362–402, 2006.

[39] C. List. The theory of judgment aggregation: An introductory review. *Synthese*, 187(1):179–207, 2012.

[40] C. List and C. Puppe. Judgment aggregation. In P. Anand, P. Pattanaik, and C. Puppe, editors, *The Handbook of Rational and Social Choice*, chapter 19, pages 457–482. Oxford University Press, 2009.

[41] C. Papadimitriou. *Computational Complexity*. Addison-Wesley, second edition, 1995.

[42] J. Rothe. *Complexity Theory and Cryptology. An Introduction to Cryptocomplexity*. EATCS Texts in Theoretical Computer Science. Springer-Verlag, 2005.

[43] M. Satterthwaite. Strategy-proofness and Arrow's conditions: Existence and correspondence theorems for voting procedures and social welfare functions. *Journal of Economic Theory*, 10(2):187–217, 1975.

[44] L. Svensson. Strategy-proof allocation of indivisible goods. *Social Choice and Welfare*, 16:557–567, 1999.

[45] K. Takamiya. Coalitional unanimity versus strategy-proofness in coalition formation problems. *International Journal of Game Theory*, 42(1):115–130, 2013.

[46] L. Xia and V. Conitzer. Determining possible and necessary winners given partial orders. *Journal of Artificial Intelligence Research*, 41:25–67, 2011.

Dorothea Baumeister and Jörg Rothe
Institut für Informatik
Heinrich-Heine-Universität Düsseldorf
40225 Düsseldorf, Germany
Email: `{baumeister,rothe}@cs.uni-duesseldorf.de`

Gábor Erdélyi
School of Economic Disciplines
University of Siegen
57076 Siegen, Germany
Email: `erdelyi@wiwi.uni-siegen.de`

Olivia J. Erdélyi
Institut für Versicherungsrecht
Heinrich-Heine-Universität Düsseldorf
40225 Düsseldorf, Germany
Email: `olivia.erdelyi@uni-duesseldorf.de`