Mathematical Social Sciences 76 (2015) 19–30

Contents lists available at ScienceDirect

Mathematical Social Sciences

journal homepage: www.elsevier.com/locate/econbase

# Complexity of manipulation and bribery in judgment aggregation for uniform premise-based quota rules[☆]

Dorothea Baumeister[a,*], Gábor Erdélyi[b], Olivia J. Erdélyi[c], Jörg Rothe[a]

[a] Institut für Informatik, Heinrich-Heine-Universität Düsseldorf, 40225 Düsseldorf, Germany
[b] School of Economic Disciplines, University of Siegen, 57076 Siegen, Germany
[c] Institut für Versicherungsrecht, Heinrich-Heine-Universität Düsseldorf, 40225 Düsseldorf, Germany

## HIGHLIGHTS

- We extend previous work on the complexity of manipulation in judgment aggregation.
- We consider incomplete judgment sets and various notions of preferences on them.
- We introduce bribery in judgment aggregation and study its complexity.

ARTICLE INFO

*Article history:*
Received 28 January 2014
Received in revised form
17 March 2015
Accepted 23 March 2015
Available online 3 April 2015

ABSTRACT

Endriss et al. (2012) initiated the complexity-theoretic study of problems related to judgment aggregation. We extend their results on the manipulation of 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 (2007). This complements previous work on the complexity of manipulation in judgment aggregation that focused on Hamming-distance-respecting preferences only, which we also study here. Furthermore, inspired by work on bribery in voting (Faliszewski and Rothe, in press), we introduce and study the closely related issue of bribery in judgment aggregation.



## 1. Introduction

Judgment Aggregation is the task of aggregating individual judgment sets of possibly interconnected logical propositions (see the surveys by List and Puppe (2009) and List (2012), and the book chapter by Endriss (in press)) 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 and bribery – have been studied intensely, with a particular focus on the complexity of the related problems. In particular, (coalitional) *manipulation* refers to (a group of) strategic voters casting their votes insincerely to reach their desired outcome. In *bribery* (see, e.g., Faliszewski et al., 2009a,b and the book chapter by Faliszewski and Rothe, in press) an external agent seeks to reach her desired outcome by bribing – within a given budget – some voters to alter their votes. Strategic behavior has been studied to a far lesser extent in judgment aggregation than in voting so far.

Mechanisms for collective decision making that are susceptible to strategic behavior, be it from the agents involved as in manipulation or from external authorities or actors as in bribery, are obviously not desirable, as that undermines the trust we have in them. We therefore have a strong interest in accurately assessing how vulnerable a mechanism for collective decision making is to these internal or external influences. Unfortunately, in many concrete settings of social choice, "perfect" mechanisms do not exist. For example, the Gibbard–Satterthwaite theorem says that no

☆ This paper combines and extends the results from preliminary conference versions that appear in the proceedings of the *2nd* and the *3rd International Conference on Algorithmic Decision Theory* (ADT'11 and ADT'13) (Baumeister et al., 2011, 2013) and of the *4th* and the *5th International Workshop on Computational Social Choice* (COMSOC'12 and COMSOC'14) (Baumeister et al., 2012, 2014), and was presented at the *ESSLLI Workshop on Logical Models of Group Decision Making* (ESSLLI-LMGD'13).

* Corresponding author. Tel.: +49 211 8111634.
*E-mail address:* baumeister@cs.uni-duesseldorf.de (D. Baumeister).
*URLs:* http://ccc.cs.uni-duesseldorf.de/~baumeister (D. Baumeister),
http://www.wiwi.uni-siegen.de/dt/team/erdelyi/ (G. Erdélyi),
http://ccc.cs.uni-duesseldorf.de/~rothe (J. Rothe).

http://dx.doi.org/10.1016/j.mathsocsci.2015.03.006
0165-4896/

reasonable voting system can be "strategy-proof" (Gibbard, 1973; Satterthwaite, 1975) (see also the generalization by Duggan and Schwartz, 2000), and Dietrich and List (2007a) 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. While this approach evidently makes sense in voting, one may wonder whether it is also applicable to judgment aggregation scenarios. After all, these scenarios originate from and are motivated by juridical issues in court proceedings (Kornhauser and Sager, 1986) where usually only a few judges collectively make decisions. However, there are also large-scale examples of real-world judgment aggregation scenarios. Suppose that a community is going to make a decision on whether or not to build a new prison and that all citizens are invited to participate in this decision making via an online platform.[1] After discussing the pros and cons, the question boils down to their judgment of the following two propositions: (A) The community's income from municipal taxes is large enough to afford building a prison, and (B) the crime rate in the community has raised so much that this indeed is needed. The prison will be built if both (A) and (B) are affirmed by the citizens. This example is very similar to other examples motivating judgment aggregation (see, e.g., Example 1 in Section 2 on a decision by three judges in a soccer match), but its point is that it can be realistic to have a large number of judges, and it therefore makes sense to apply computational complexity also to judgment aggregation problems. Besides these examples judgment aggregation is also used in artificial intelligence, notably as a tool for collective decision making in systems with interacting autonomous agents, which also may involve a very large number of participating judges. For further applications of judgment aggregation in computer science, see the book chapter by Endriss (in press).

### 1.1. Our contributions

As mentioned above, much work on the complexity of manipulation and bribery problems has been done in voting, but only a few results are known for these problems in judgment aggregation. Most notably, Endriss et al. (2012) 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 the class of uniform premise-based quota rules and to other notions of preference that have been introduced by Dietrich and List (2007a). In particular, we will study incomplete judgment sets and top-respecting and closeness-respecting preferences in addition to Hamming-distance-respecting preferences. We also introduce exact variants of manipulation where the manipulator's goal is to achieve not only a better, but a *best* outcome for a given subset of her desired set. This gives rise to a number of manipulation problems for each judgment aggregation rule, which is why we here focus on only one class of such rules, the uniform premise-based quota rules. Extending this work to other judgment aggregation rules (such as those mentioned in Section 1.2 below), to allow a comparison of these rules in terms of their resistance to manipulation, is left for future work.

A main result of this paper is presented in Theorem 10, which says that for each rational quota and for any fixed number of at least three judges, the uniform premise-based quota rule is hard to manipulate for Hamming-distance-respecting preferences in terms of the parameterized complexity class W[2] (see Section 2.2) when parameterized by the maximum number of changes in the premises needed in the manipulator's desired set. We also provide many complexity results for manipulation with respect to top-respecting and closeness-respecting preferences (see Table 6 in Section 5 for an overview).

In addition, we here initiate the algorithmic and complexity-theoretic study of bribery problems in judgment aggregation. Again, these problems are each closely related to the corresponding bribery problems in voting, yet are specifically tailored to judgment aggregation scenarios. Table 7 in Section 5 gives an overview of our results on the complexity of bribery problems for judgment aggregation with the premise-based procedure.

This paper combines and extends previous work by Baumeister et al. (2011, 2012, 2013, 2014b) that appeared in the proceedings of ADT'11, COMSOC'12, ADT'13, and COMSOC'14. The present version contains some additional results and it provides more discussion and a number of notational improvements.

### 1.2. Related work

Manipulation and bribery are two forms of strategic actions that have been studied extensively in voting (see the references below), yet much less so in judgment aggregation. Endriss et al. (2012) were the first to study manipulation in judgment aggregation from a computational point of view, and we here extend their work as described in Section 1.1. In voting theory, another way of tampering with elections is control, and Baumeister et al. (2013, 2012) have studied certain types of control in judgment aggregation where an external agent seeks to influence the outcome by altering the structure of the judgment aggregation process by adding, deleting, or replacing judges. Dietrich (2014) studied the agenda manipulation problem, where one tries to influence the outcome by carefully choosing the formulas in the agenda. In the case of sequential judgment aggregation procedures, the order of the formulas in the agenda is important and may give the opportunity to rule manipulation, see the work of List (2004) and Dietrich and List (2007b).

For our complexity-theoretic analysis of manipulation and bribery in judgment aggregation, we will focus on the *uniform premise-based quota rules* (see Section 2.1 for a formal definition). In these rules, the agenda is divided into *premises* and *conclusions*, the outcome for each of the premises is determined by a given quota (just one – a *uniform* quota assigned to each of the premises), and the outcome for each conclusion is then derived from the outcome for the premises in a consistent way. That is, under a uniform premise-based quota rule we collectively accept those conclusions that logically follow from the premises we collectively accept according to the given quota. By contrast, in *conclusion-based procedures* (see, e.g., Kornhauser and Sager, 1986; List and Pettit, 2002; Dietrich, 2006), the collective decision is made for the conclusions of the agenda only. Another approach are the *distance-based procedures* (see the work of Miller and Osherson, 2009) where a collective outcome minimizes the distance (according to a certain predetermined metric) to the given individual judgment sets. Lang et al. (2011) define and study judgment aggregation procedures based on minimization, which are inspired by voting

[1] Implementing real-world, large-scale judgment aggregation scenarios such as this one on an online platform is not fictitious. For example, the interdisciplinary graduate college "Online Partizipation" (www.fortschrittskolleg.de) of HHU Düsseldorf and other institutions (including as practice partners registered societies such as Liquid Democracy e.V., limited liability companies such as Zebralog GmbH & Co KG and polidia GmbH as well as the municipal councils of Bonn, Köln, and Münster, among others) investigates such settings to explore the technical, societal, economical, and juridical aspects of a self-organizing society. A central goal in this project, which is funded by the NRW Ministry for Innovation, Science, and Research, is to build an internet platform (www.normsetting.org) that can be used for online discussions and deliberations, and hundreds of participants have been involved in previous pilot projects.

theory and knowledge representation. In a *sequential procedure* (List, 2004), the formulas of the agenda are considered in some predefined order, and then some rule (e.g., the majority rule) is applied to achieve a consistent outcome. In voting theory, the well-known Condorcet principle requires that an election is won by a candidate that beats all other candidates in a pairwise comparison. Similarly, judgment aggregation procedures based on the Condorcet set have been proposed by Nehring et al. (2014). Everaere et al. (2014) propose judgment aggregation procedures that are based on the support (i.e., number of votes) that each member of the agenda receives.

Bartholdi et al. (1989) and Bartholdi and Orlin (1991) were the first to study the complexity of manipulation problems in voting, and Conitzer et al. (2007) extended their model to also study coalitional manipulation in weighted elections. Since then, much work has been done to classify a variety of manipulation problems for many voting rules in terms of their complexity, see, e.g., the recent surveys and book chapters by Faliszewski and Procaccia (2010), Faliszewski et al. (2010), Faliszewski and Rothe (in press), Conitzer and Walsh (in press), Brandt et al. (2012), and Baumeister et al. (2010).

Bribery in the context of voting has been introduced and studied in depth by Faliszewski et al. (2009a), and variants of bribery problems have been investigated, for example, by Elkind et al. (2009) and Faliszewski et al. (2009b, in press) (see, e.g., the survey by Faliszewski and Rothe, in press for an overview and many more references). In particular, Christian et al. (2007) (see also, e.g., Bredereck et al., 2012; Binkele-Raible et al., 2014) studied the related problem of optimal lobbying, which may be seen as a simplified variant of judgment aggregation. We will apply their hardness result for optimal lobbying in the proof of Theorem 17 in Section 4 when we will be concerned with bribery in judgment aggregation.

### 1.3. Organization

This paper is organized as follows. In Section 2 we provide the basic framework of judgment aggregation and define the relevant notions formally, and we provide some background from complexity theory. In Section 3 we study the complexity of manipulation in judgment aggregation for premise-based quota rules, and in Section 4 that of bribery in judgment aggregation for the premise-based procedure. Finally, Section 5 summarizes our results and presents a number of interesting open problems for future research.

## 2. Preliminaries

In this section, we present the formal framework of judgment aggregation and provide some background on complexity theory.

### 2.1. Formal framework of judgment aggregation

We adopt the judgment aggregation framework described by Endriss et al. (2012) (see also their previous conference papers by Endriss et al., 2010a,b). Let $PS$ be the set of all propositional variables and let $\mathcal{L}_{PS}$ 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 $\overline{\alpha}$ denote the complement of $\alpha$, i.e., $\overline{\alpha} = \neg\alpha$ if $\alpha$ is not negated, and $\overline{\alpha} = \beta$ if $\alpha = \neg\beta$. The judges have to judge over all formulas in the *agenda* $\Phi$, which is a finite, nonempty subset of $\mathcal{L}_{PS}$ without doubly negated formulas. The agenda is required to be closed under complementation, i.e., $\overline{\alpha} \in \Phi$ if $\alpha \in \Phi$.

A *judgment set for an agenda* $\Phi$ 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 formulas in the agenda accepted by the collective as the result of a judgment aggregation procedure. A judgment set $J$ is

- *complete* if for all $\alpha \in \Phi$, $\alpha \in J$ or $\overline{\alpha} \in J$;
- *complement-free* if for no $\alpha \in \Phi$, $\alpha \in J$ and $\overline{\alpha} \in J$;
- *consistent* if there is an assignment that makes all formulas in $J$ true;
- *rational* if it is complete and consistent.

If a judgment set is rational, it is obviously complement-free. We denote the set of all rational judgment sets in $\Phi$ by $\mathcal{J}(\Phi)$.

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

The famous doctrinal paradox (Kornhauser and Sager, 1986) 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 aggregation 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 penalty area. The three different individual judgments and the evaluation according to the majority rule are shown in Table 1(a), where a 1 stands for "yes" and a 0 for "no."

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)).

Endriss et al. (2012) 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 and bribery also for the class of uniform premise-based quota rules, which were defined in a more general way by Dietrich and List (2007b).

**Definition 2** (*Uniform Premise-based Quota Rule*). The agenda $\Phi$ is divided into two disjoint subsets $\Phi = \Phi_p \uplus \Phi_c$, where $\Phi_p$ is the set of premises and $\Phi_c$ is the set of conclusions. We assume both $\Phi_p$ and $\Phi_c$ to be closed under complementation. The premises $\Phi_p$ are again divided into two disjoint subsets, $\Phi_p = \Phi_1 \uplus \Phi_2$, such that $\Phi_1$ and $\Phi_2$ each contain exactly one member of each pair $\{\varphi, \overline{\varphi}\} \subseteq \Phi_p$. Assign the quota $q \in \mathbb{Q}, 0 \leq q < 1$, to each literal $\varphi \in \Phi_1$. The quota for each literal $\varphi \in \Phi_2$ is then derived by $q' = 1 - q$. Let $|S|$ denote the cardinality of set $S$ and $\models$ the satisfaction relation. A *uniform premise-based quota rule* is defined

[2] Endriss et al. (2012, 2010b) studied the question of whether one can guarantee for a specific agenda that the outcome is always rational. 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 rational outcome for the majority rule is an intractable problem.

**Table 1**
Example illustrating the doctrinal paradox and how to prevent it by the premise-based procedure.

| (a) Doctrinal paradox with the majority rule | | | | (b) Avoiding it with the premise-based procedure | | | | |
|---|---|---|---|---|---|---|---|---|
| | Penalty area | Foul | Penalty | | Penalty area | Foul | | Penalty |
| Referee 1 | 1 | 1 | 1 | Referee 1 | 1 | 1 | | 1 |
| Referee 2 | 1 | 0 | 0 | Referee 2 | 1 | 0 | | 0 |
| Referee 3 | 0 | 1 | 0 | Referee 3 | 0 | 1 | | 0 |
| *Majority* | 1 | 1 | 0 | *PBP* | 1 | 1 | $\Rightarrow$ | 1 |

to be a function $UPQR_q : \mathcal{J}(\Phi)^n \to 2^{\Phi}$ such that, for $\Phi = \Phi_p \uplus \Phi_c$, each profile $\mathbf{J} = (J_1, \ldots, J_n)$ of individual judgment sets is mapped to the collective judgment set

$$UPQR_q(\mathbf{J}) = \triangle \cup \{\psi \in \Phi_c \mid \triangle \models \psi\}, \quad \text{where}$$
$$\triangle = \{\varphi \in \Phi_1 \mid |\{i \mid \varphi \in J_i\}| > nq\} \cup \{\varphi \in \Phi_2 \mid |\{i \mid \varphi \in J_i\}| \geq nq'\}.$$

To guarantee rational outcomes for this procedure, it is enough to require that $\Phi$ is closed under propositional variables and that $\Phi_p$ consists of all literals. Note that this implies that all premises are independent. To be contained in the collective judgment set, a literal $\varphi \in \Phi_1$ needs to be in $\lfloor nq + 1 \rfloor$ individual judgment sets, and a literal $\varphi \in \Phi_2$ in $\lceil nq' \rceil$ individual judgment sets. Note that $\lfloor nq + 1 \rfloor + \lceil nq' \rceil = n + 1$ ensures that either $\psi \in UPQR_q(\mathbf{J})$ or $\overline{\psi} \in UPQR_q(\mathbf{J})$ for every $\psi \in \Phi$. Note further that the quota $q = 1$ for a variable $\varphi \in \Phi_1$ is not allowed here, as $n + 1$ judges having it in the collective judgment set were then needed for $\varphi \in \Phi_1$ to be in the collective judgment set, which is impossible. However, $q = 0$ is allowed, as in that case $\varphi \in \Phi_1$ needs to be in at least one individual judgment set and $\overline{\varphi} \in \Phi_2$ needs to be in $n$ individual judgment sets, which is possible. For $q = 1/2$ and the case of an odd number of judges, we obtain the premise-based procedure defined by Endriss et al. (2012), and we will denote it by *PBP*.

### 2.2. Background on complexity theory

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 (Papadimitriou, 1995; Rothe, 2005). Downey and Fellows (1999) (see also the textbooks by Flum and Grohe, 2006; Niedermeier, 2006) introduced *parameterized* complexity theory; in their framework it is possible to do a more fine-grained multi-dimensional complexity analysis. In particular, NP-complete problems may be easy (i.e., fixed-parameter tractable) with respect to certain parameters confining the seemingly unavoidable combinatorial explosion. If this parameter is reasonably small, a fixed-parameter tractable problem can be solved efficiently in practice, despite its NP-hardness. Formally, a *parameterized decision problem* is a set $L \subseteq \Sigma^* \times \mathbb{N}$, where $\Sigma^*$ is the set of all words that can be built from the alphabet $\Sigma$. Such a problem is *fixed-parameter tractable* (FPT) if there is a constant $c$ such that for each input $(x, k)$ of size $m = |(x, k)|$ we can determine in time $O(f(k) \cdot m^c)$ whether $(x, k)$ is in $L$, where $f$ is a function depending only on the parameter $k$. The main hierarchy of parameterized complexity classes is:

$$\mathrm{FPT} = \mathrm{W}[0] \subseteq \mathrm{W}[1] \subseteq \mathrm{W}[2] \subseteq \cdots \subseteq \mathrm{W}[\ell] \subseteq \mathrm{XP},$$

where FPT and W[1] are strong parameterized analogues of P and NP, and XP is the class of parameterized decision problems that can be solved in time $O(m^{g(k)})$ for some function $g$. Detailed definitions of the classes in the W-hierarchy can be found in the book by Downey and Fellows (1999).

We say that a parameterized problem $A$ *parameterized reduces* to a parameterized problem $B$ if each instance $(x, k)$ of $A$ can be transformed in time $\mathcal{O}(g(k) \cdot |x|^c)$ (for some function $g$ and some constant $c$) into an instance $(x', k')$ of $B$ such that $(x, k) \in A$ if and only if $(x', k') \in B$, where $k' = g(k)$.

In our results, we will focus on only the class W[2] (see, e.g., Downey and Fellows, 1999; Flum and Grohe, 2006; Niedermeier, 2006 for more background), which refers to problems that are considered to be fixed-parameter intractable.

In order to show that a parameterized problem is W[2]-hard, we will give parameterized reductions from the W[2]-complete problem $k$-Dominating-Set (see Downey and Fellows, 1999). To define this problem, let $G = (V, E)$ be an undirected graph with vertex set $V = \{v_1, \ldots, v_n\}$ and edge set $E$. Define $N(v_i)$ as the closed neighborhood of vertex $v_i$, i.e., the set of all vertices adjacent to $v_i$ and the vertex $v_i$ itself. Then, $V' \subseteq V$ is a dominating set for $G$ if and only if $N(v_i) \cap V' \neq \emptyset$ for each $i$, $1 \leq i \leq n$. The *size of a dominating set* is the number of its vertices. Define the parameterized decision problem:

| $k$-Dominating-Set | |
|---|---|
| **Given:** | A graph $G = (V, E)$ and a positive integer $k \leq \|V\|$. |
| **Parameter:** | $k$. |
| **Question:** | Is there a dominating set of size at most $k$ in $G$? |

When considered as an unparameterized problem (i.e., as a plain NP-complete decision problem), we write Dominating-Set.

## 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 (2006) 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*. In the following section, we will discuss the interesting issue of how to define preferences from judgments, which will be central to how we will formalize strategy-proofness and our manipulation problems later on.

### 3.1. Defining preferences from judgments

Strategy-proofness and manipulation have been studied in a wide variety of fields—such as voting (see, e.g., Gibbard, 1973; Satterthwaite, 1975; Faliszewski and Procaccia, 2010; Conitzer and Walsh, in press), mechanism design (see, e.g., Alon et al., 2013), game theory (see, e.g., Takamiya, 2013), fair division (see, e.g., Svensson, 1999; Klaus and Miyagawa, 2013), etc. In judgment aggregation, manipulability and (the game-theoretic concept of) strategy-proofness were first introduced by Dietrich and List (2007a). 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 the number of judgment sets 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 (2007a), 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. (2012) that focused on Hamming-distance-respecting 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 (2012).

As mentioned above, we apply the notions introduced by Dietrich and List (2007a) to study various types of preferences. We will express weak preferences by weak orders, denoted by $\succsim$,[3] over $\mathcal{J}(\Phi)$. As is common, for all $X, Y \in \mathcal{J}(\Phi)$, define $X \succ Y$ by $X \succsim Y$ and $Y \not\succsim X$, and define $X \sim Y$ by $X \succsim Y$ and $Y \succsim X$. We say $X$ *is weakly preferred to* $Y$ whenever $X \succsim Y$, and we say $X$ *is preferred to* $Y$ whenever $X \succ Y$.

**Definition 3.** Let $\Phi$ be an agenda and $\mathcal{J}(\Phi)$ the set of all rational judgment sets in $\Phi$. Let $U$ be the set of all weak orders over $\mathcal{J}(\Phi)$. Given some (possibly incomplete) judgment set $J$, define

1. the set of *unrestricted $J$-induced* (*weak*[4]) *preferences* as the set $U_J$ of weak orders $\succsim$ in $U$ such that for all $X, Y \in \mathcal{J}(\Phi)$, $X \sim Y$ whenever $X \cap J = Y \cap J$;
2. the set of *top-respecting $J$-induced (weak) preferences* as $TR_J \subseteq U_J$ such that $\succsim \in TR_J$ if and only if for all $X \in \mathcal{J}(\Phi)$ with $X \cap J \neq J$, it holds that $J \succ X$;
3. the set of *closeness-respecting $J$-induced (weak) preferences* as $CR_J \subseteq U_J$ such that $\succsim \in CR_J$ if and only if for all $X, Y \in \mathcal{J}(\Phi)$ with $Y \cap J \subseteq X \cap J$, we have $X \succsim Y$; and
4. the set of *Hamming-distance-respecting $J$-induced (weak) preferences* as $HD_J \subseteq U_J$ such that $\succsim \in HD_J$ if and only if for all $X, Y \in \mathcal{J}(\Phi)$, we have $X \succsim Y$ if and only if $HD(X, J) \leq HD(Y, J)$, where the Hamming distance $HD(S, T)$ between two (possibly incomplete) judgment sets $S$ and $T$ is the number of disagreements on positive formulas that occur in both judgment sets.

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 $J$ to be the most preferred one. This also holds for closeness-respecting preferences, but in addition judgment sets that have additional agreement are preferred. The Hamming-distance-respecting preferences, which focus on the total number of disagreements, are the most restrictive ones. Summing up, we have

$$HD_J \subseteq CR_J \subseteq TR_J \subseteq U_J.$$

**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 |

**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.

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-respecting 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.

### 3.2. Possible and necessary strategy-proofness and possible, necessary, and exact manipulation

Konczak and Lang (2005) introduced the notions of necessary and possible winner in voting:[5] A *necessary winner* is a candidate who wins for *every* extension of a given partial preference profile to a complete profile, and a *possible winner* is a candidate who can be made a winner by *some* complete extension of a given partial preference profile. Inspired by their notions,[6] we now introduce the

[3] A *weak order* is transitive and total (and thus, in particular, reflexive).

[4] We occasionally abuse our notation by omitting to explicitly state "weak", using the term *preference* to refer both to "weak preference" in the sense of $\succsim$ and to "preference" in the sense of $\succ$. The intended meaning will always be clear from the context.

[5] For more background, we refer to the work of, e.g., Xia and Conitzer (2011), Betzler and Dorn (2010), and Baumeister and Rothe (2012) and to the book chapters by Boutilier and Rosenschein (in press) and Baumeister and Rothe (in press).

[6] See also the remotely related notions of "necessary envy-freeness" versus "possible envy-freeness" in fair division that are due to Bouveret et al. (2010) (see also the papers by Brams et al., 2004; Brams and King, 2005 and Baumeister et al., 2014a).

notions of necessary and possible strategy-proofness in judgment aggregation.

Just as Dietrich and List (2007a), we study settings where the desired set is incomplete, to also capture their "reason-oriented" and "outcome-oriented" preferences. However, we will not generally restrict the desired set to the premises or the conclusions; rather, we allow arbitrary incomplete desired 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 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 $J$, $X$, and $Y$ be judgment sets for the same agenda $\Phi$, where $J$ is possibly incomplete. Let $T_J \in \{U_J, TR_J, CR_J\}$ be a type of $J$-induced preferences.

- A judge *necessarily weakly prefers $X$ to $Y$ for type $T_J$* if $X \succsim Y$ for all $\succsim \in T_J$.
- A judge *possibly weakly prefers $X$ to $Y$ for type $T_J$* if there is some $\succsim \in T_J$ with $X \succsim Y$.
- The notions of *possible/necessary preference for type $T_J$* are defined analogously, except with $\succsim$ replaced by $\succ$.
- A judgment aggregation rule $F$ is *necessarily/possibly strategy-proof for induced weak preferences of type $T \in \{U, TR, CR\}$* if for all profiles $(J_1, \ldots, J_n)$, for each $i$, $1 \le i \le n$, and for each $J_i^* \in \mathcal{J}(\Phi)$, judge $i$ necessarily/possibly weakly prefers the outcome $F(J_1, \ldots, J_n)$ to the outcome $F(J_1, \ldots, J_{i-1}, J_i^*, J_{i+1}, \ldots, J_n)$ for type $T_{J_i}$.

The stronger notion of *necessary strategy-proofness* corresponds to the "strategy-proofness" condition defined by Dietrich and List (2007a), whereas the weaker notion of *possible strategy-proofness* is introduced here. Note that since the Hamming distance-respecting weak preferences are a complete relation, we simply say that $F$ is *strategy-proof (for Hamming-distance-respecting weak 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.

In our notation, a result of Dietrich and List (2007a) says that an aggregation rule is necessarily strategy-proof for closeness-respecting preferences if and only if it is independent and monotonic. *Independence* means that the collective decision on each formula 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 for closeness-respecting weak preferences. An aggregation function is *monotonic* if additional support for some formula 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.

Define the related manipulation problems for uniform premise-based quota rules and a given preference type $T$.[7]

[7] We define our manipulation and bribery problems as decision problems as is most common in complexity theory. It would also be interesting to study the search version of these problems that, instead of just deciding whether a manipulative action is possible, actually compute such an action. Note that the decision version of a problem straightforwardly reduces to its search version, and since most of our results are about hardness for decision, the search version is hard as well. Interestingly, Hemaspaandra et al. (2013) studied search versus decision for manipulation in voting and found cases where, under plausible assumptions, it is easy to decide whether manipulation is possible, yet hard to actually compute a successful manipulation.

*$UPQR_q$-$T$*-NECESSARY-MANIPULATION

| | |
|---|---|
| **Given:** | An agenda $\Phi$, a profile $\mathbf{J} = (J_1, \ldots, J_n) \in \mathcal{J}(\Phi)^n$, and the manipulator's desired consistent (possibly incomplete) set $J \subseteq J_n$. |
| **Question:** | Does there exist a judgment set $J^* \in \mathcal{J}(\Phi)$ such that $UPQR_q(J_1, \ldots, J_{n-1}, J^*) \succ UPQR_q(J_1, \ldots, J_n)$ for all $\succsim \in T_J$? |

In $UPQR_q$-$T$-POSSIBLE-MANIPULATION, we for the same input ask whether there exists a judgment set $J^* \in \mathcal{J}(\Phi)$ such that

$$UPQR_q(J_1, \ldots, J_{n-1}, J^*) \succ UPQR_q(J_1, \ldots, J_n)$$

for *some* $\succsim \in T_J$. In the case of Hamming-distance-respecting 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 set. Here, the question is whether there is some judgment set $J^* \in \mathcal{J}(\Phi)$ such that

$$J \subseteq UPQR_q(J_1, \ldots, J_{n-1}, J^*).$$

### 3.3. Results

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

**Theorem 6.** *For each rational quota $q$, $0 \le q < 1$, $UPQR_q$-EXACT-MANIPULATION is* NP*-complete, for every fixed number $n \ge 3$ of 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 $\varphi$ be a given formula in conjunctive normal form, where the clauses are built from the set $A = \{\alpha_1, \ldots, \alpha_m\}$ 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 $\varphi$. Now construct an agenda $\Phi$ that consists of the variables in $A$ and their negations, an additional variable $\beta$ 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 $i$, $1 \le i \le m$, $\neg\beta$, and $\neg(\varphi \vee \beta)$. The third judge is the manipulative one and her individual judgment set, $J_3$, contains $A$, $\beta$, and $(\varphi \vee \beta)$. Her 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, $\beta$ 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 $\varphi$ to true. This implies that there is a satisfying assignment for $\varphi$ 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 $q$, $0 \le q < 1$, the following hold:*

1. *UPQR*$_q$-Exact-Manipulation $\leq^{\mathrm{p}}_{\mathrm{m}}$ *UPQR*$_q$-*T*-Necessary-Manipulation *for each type* $T \in \{TR, CR\}$.
2. *UPQR*$_q$-Exact-Manipulation $\leq^{\mathrm{p}}_{\mathrm{m}}$ *UPQR*$_q$-*T*-Possible-Manipulation *for each type* $T \in \{U, TR, CR\}$.
3. *UPQR*$_q$-Exact-Manipulation $\leq^{\mathrm{p}}_{\mathrm{m}}$ *UPQR*$_q$-HD-Manipulation.

**Proof.** For the exact problem, we have an agenda $\Phi$, some profile $\mathbf{J} = (J_1, \dots, J_n)$, and some desired 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 $J \setminus UPQR_q(\mathbf{J}) \neq \emptyset$ and consider the following problem. Fix some $T \in \{TR, CR, HD\}$, let the agenda $\Phi'$ be the union of $\Phi$, the formula $\varphi = \alpha_1 \wedge \dots \wedge \alpha_m$, and its negation. Let $\mathbf{J}' \in \mathcal{J}(\Phi')^n$ be the consistent extensions of $\mathbf{J}$. In particular, $J'_n = J_n \cup \{\varphi\}$. Let the desired set be $J' = \{\varphi\}$, and we are looking for a modified judgment set $J_n'^*$ such that for all/some $\succsim \in T_{J'}$ we have

$$UPQR_q(J'_1, \dots, J'_{n-1}, J_n'^*) \succ UPQR_q(J'_1, \dots, J'_{n-1}, J'_n).$$

Since $J'$ consists of the single formula $\varphi$, there are only two different collective outcomes when restricted to $J'$. Since $\varphi \in J'$, it obviously holds that $\varphi \succ \neg\varphi$ for $\succsim \in T_{J'}$ for all $T \in \{TR, CR, HD\}$, and hence 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

$$J \subseteq UPQR_q(J_1, \dots, J_{n-1}, J_n^*)$$

if and only if there is some $J_n'^*$ with

$$\varphi \in UPQR_q(J'_1, \dots, J'_{n-1}, J_n'^*),$$

the reduction works in all cases. □

Note that this reduction requires an *incomplete* desired 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 $T \in \{TR, CR\}$, and *UPQR*$_q$-*T*-Possible-Manipulation for $T \in \{U, TR, CR\}$ whenever the desired set of the manipulator is incomplete. Alternatively, the reduction given by Endriss et al. (2012) in fact shows NP-completeness of *PBP*-HD-Manipulation, even if the manipulator's desired set is complete. By contrast, if the manipulator's desired set is complete, the possible manipulation problem turns out to be easy to solve for unrestricted and top-respecting preferences.

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

**Proof.** This result holds, since a *UPQR*$_q$-U-Possible-Manipulation instance is positive exactly if it is possible to achieve an outcome different from the actual one, because it may possibly be preferred. And if the desired set of the manipulator is complete, the conclusions can only change when there is a change in the premises, hence there must be some premise from the desired 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 set is not the actual outcome. □

For closeness-respecting preferences, however, possible manipulation for uniform premise-based quota rules is NP-hard, even if the desired set of the manipulator is complete.

**Proposition 9** (*See* Selker, 2014). *For each rational quota* $q$, $0 \leq q < 1$, *UPQR*$_q$-*CR*-Possible-Manipulation *is* NP-*complete, even if the manipulator's desired set is complete.*

Proposition 9 is due to Selker (2014), and the idea is to modify the reduction from the proof of Theorem 2 in the paper by Endriss et al. (2010a) to work for an arbitrary rational value of $q$, $0 \leq q < 1$. Every outcome that is possibly preferred to the actual outcome with respect to closeness-respecting preferences must have at least one additional agreement with the manipulator's desired set. In this modified reduction, the conclusion is the only possibility to obtain this agreement, and similar arguments as in the proof of Endriss et al. (2010a) work here as well (for more details, see the thesis by Selker, 2014).

We now consider the parameterized complexity of the manipulation problem for Hamming-distance-respecting preferences.

**Theorem 10.** *For each rational quota* $q$, $0 \leq q < 1$, *and for any fixed number* $n \geq 3$ *of judges,* *UPQR*$_q$-HD-Manipulation *is* W[2]-*hard when parameterized by the maximum number of changes in the premises needed in the manipulator's desired set.*

**Proof.** We start by giving the details for $q = 1/2$ and three judges, and later explain how this proof can be extended to capture any other rational quota values $q$, $0 \leq q < 1$, and any fixed number of judges greater than three.

The proof will be by a reduction from the W[2]-complete problem $k$-Dominating-Set, which was defined in Section 2.2. Given an instance of $k$-Dominating-Set, a graph $G = (V, E)$ with vertex set $V = \{v_1, \dots, v_n\}$, we will now describe how to construct a manipulation instance for judgment aggregation. Let the agenda $\Phi$ contain

- the variables $v_1, \dots, v_n$, $y$ and their negations,[8]
- the formula $\varphi_i = v_i^1 \vee \dots \vee v_i^j \vee y$ and its negation, where $\{v_i^1, \dots, v_i^j\}$ is the closed neighborhood of $v_i$, for each $i$, $1 \leq i \leq n$, and
- $n-1$ syntactic variations of each of these formulas and its negation (which in effect means to give each formula $\varphi_i$ a weight of $n$),[9]
- the formula $v_1 \vee \dots \vee v_n$, its negation, and
- $n^2 - k - 2$ syntactic variations of this formula and its negation (this can again be seen as giving this formula a weight of $n^2 - k - 1$).

The set of judges is $N = \{1, 2, 3\}$, with the individual judgment sets $J_1, J_2$, and $J_3$ (where $J_3$ is the judgment set of the manipulative judge), and the collective judgment set as shown in Table 3. Note that the Hamming distance between $J_3$ and the collective judgment set is $1 + n^2$.

We claim that there is an alternative judgment set for $J_3$ that yields a smaller Hamming distance to the collective outcome if and only if there is a dominating set of size at most $k$ for $G$. In addition, the number of premises that differ in the new judgment set and $J_3$ is less than $k$.

($\Leftarrow$) Assume that there is a dominating set $V'$ of $G$ with $|V'| = k$. (If $|V'| < k$, we simply add any $k - |V'|$ vertices to obtain a dominating set of size exactly $k$.) Regarding the premises, the judgment

8 We use the same identifiers $v_1, \dots, v_n$ for the vertices of $G$ and the variables in $\Phi$, specifying the intended meaning only if it is not clear from the context.

9 A syntactic variation of a formula can, for example, be obtained by an additional conjunction with the constant 1.

**Table 3**
Construction for the first part of the proof of Theorem 10.

| Judgment set | $v_1$ | $\cdots$ | $v_n$ | $y$ | | $\varphi_1$ | $\cdots$ | $\varphi_n$ | $v_1 \vee \cdots \vee v_n$ |
|---|---|---|---|---|---|---|---|---|---|
| $J_1$ | 1 | $\cdots$ | 1 | 0 | | 1 | $\cdots$ | 1 | 1 |
| $J_2$ | 0 | $\cdots$ | 0 | 0 | | 0 | $\cdots$ | 0 | 0 |
| $J_3$ | 0 | $\cdots$ | 0 | 1 | | 1 | $\cdots$ | 1 | 0 |
| $UPQR_{1/2}$ | 0 | $\cdots$ | 0 | 0 | $\Rightarrow$ | 0 | $\cdots$ | 0 | 0 |

**Table 4**
Construction for the second part of the proof of Theorem 10.

| Judgment set | $v_1$ | $\cdots$ | $v_n$ | $y$ | | $\varphi_1'$ | $\cdots$ | $\varphi_n'$ | $v_1 \vee \cdots \vee v_n$ |
|---|---|---|---|---|---|---|---|---|---|
| $J_1, \ldots, J_{m-1}$ | 0 | $\cdots$ | 0 | 1 | | 0 | $\cdots$ | 0 | 0 |
| $J_m$ | 0 | $\cdots$ | 0 | 0 | | 1 | $\cdots$ | 1 | 0 |
| $UPQR_{1/2}$ | 0 | $\cdots$ | 0 | 1 | $\Rightarrow$ | 0 | $\cdots$ | 0 | 0 |

set of the manipulator contains the variables $v_i \in V'$ and also the literal $y$. Then the collective outcome also contains the variables $v_i \in V'$, and since $V'$ is a dominating set, each $\varphi_i$, $1 \leq i \leq n$, evaluates to true and the formula $v_1 \vee \cdots \vee v_n$ is also evaluated to true. The Hamming distance to the original judgment set of the manipulator is then $k+1+(n^2-k-1) = n^2$. Hence, the manipulation was successful, and the number of entries changed in the judgment set of the manipulator is exactly $k$.

($\Rightarrow$) Now assume that there is a successful manipulation with judgment set $J'$. The manipulator can change only the premises in the agenda to achieve a better outcome for her. A change for the literal $y$ changes nothing in the collective outcome. Hence, the changes must be within the set $\{v_1, \ldots, v_n\}$. Including $j$ of the $v_i$ to $J'$ has the effect that these $v_i$ are included in the collective judgment set, and that all variations of the formula $v_1 \vee \cdots \vee v_n$ and of those $\varphi_i$ that are evaluated to true are also included in the collective judgment set. If $\ell$ formulas $\varphi_i$ are evaluated to true in the collective judgment set, the Hamming distance to $J_3$ is $j + 1 + (n^2 - n\ell) + (n^2 - k - 1)$. Since the manipulation was successful, the Hamming distance can be at most $n^2$. If $\ell < n$, it must hold that $j \leq k - n$, which is not possible given that $k \leq n$ and $j > 0$. Hence, $\ell = n$ and $j = k$. Then exactly $k$ literals $v_i$ are set to true, and since this satisfies all $\varphi_i$, they must correspond to a dominating set of size $k$, concluding the proof for the quota $q = 1/2$ and three judges.

This proof can be adapted to work for any fixed number $m \geq 3$ of judgment sets $S_1, \ldots, S_m$ and for any rational value of $q$, with $1 \leq mq < m$. The agenda remains the same, but $S_1, \ldots, S_{\lfloor mq \rfloor}$ are each equal to the judgment set $J_1$ and $S_{\lfloor mq \rfloor + 1}, \ldots, S_{m-1}$ are each equal to the judgment set $J_2$. The judgment set $S_m$ of the manipulative judge equals the judgment set $J_3$, and the quota is $q$ for every positive variable and $1-q$ for every negative variable. The number of affirmations every positive formula needs to have in the collective judgment set is then $\lfloor mq \rfloor + 1$. Then the same argumentation as above applies.

For the remaining case, where $0 \leq mq < 1$, the construction must be slightly modified. The formulas $\varphi_1, \ldots, \varphi_n$ are replaced by $\varphi_i' = (v_i^1 \wedge \cdots \wedge v_i^j) \vee \neg y$, where $\{v_i^1, \ldots, v_i^j\} = N(v_i)$ for each $i$, $1 \leq i \leq n$, and the individual judgment sets $J_1, \ldots, J_m$ are shown as in Table 4, where $J_m$ is the judgment set of the manipulative judge. Then by similar arguments as above there is a successful manipulation if and only if the given graph has a dominating set of size at most $k$.

Since the number of premises changed by the manipulator depends only on the size $k$ of the dominating set, W[2]-hardness for $UPQR_q$-Manipulation holds for this parameter. □

Since the reduction in the proof of Theorem 10 is from the NP-complete problem Dominating Set and since $UPQR_q$-HD-Manipulation is in NP, NP-completeness of $UPQR_q$-HD-Manipulation follows immediately from this proof.[10]

**Corollary 11.** *For each rational quota $q$, $0 \leq q < 1$, and for any fixed number $n \geq 3$ of judges, $UPQR_q$-HD-Manipulation is NP-complete.*

As mentioned above, studying the case of a fixed total number of judges is very natural. The parameter we have considered for the manipulation problem in Theorem 10 is the "maximum number of changes in the premises needed in the manipulator's judgment set." Hence, this theorem shows that the problem remains hard even if the number of premises the manipulator can change is bounded by a fixed constant. This is also very natural, since the manipulator may wish to report a judgment set that is as close as possible to her sincere judgment set, because for a completely different judgment set it might be discovered too easily that she was judging strategically.

The W[2]-hardness result stated in Theorem 10 implies that there is little hope to find a polynomial-time algorithm for the general problem, even when the number of participating judges is fixed. In contrast, Proposition 12 below tells us that if the agenda is simple and contains no conclusions, $UPQR_q$ is even strategy-proof, and thus $UPQR_q$-Manipulation trivially is in P.

**Proposition 12.** 1. *If the agenda contains only premises and Hamming-distance-respecting preferences are assumed, then $UPQR_q$, $0 \leq q < 1$, is strategy-proof.*
2. *If the desired set of the manipulator is complete and she tries to exactly reach her desired outcome, then $UPQR_q$, $0 \leq q < 1$, is strategy-proof.*

**Proof.** In both cases the premises are considered independently. Let $n$ be the number of judges. If some $\varphi$ from the premises is contained in the judgment set $J$ of the manipulator, and $\varphi$ 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 $\varphi$ to $\overline{\varphi}$ in her 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 (2007a), since they impose different conditions on nonmanipulability.

**Proposition 13.** *If the desired 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 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 $\alpha$ that is contained in the collective judgment set, but $\overline{\alpha}$ is contained in the desired set. Obviously, it can never be the case that the manipulator switching from $\overline{\alpha}$ to $\alpha$ would cause $\overline{\alpha}$ to be in the collective judgment set. Hence there can be no additional agreement among the premises. Since the desired 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. □

[10] Alternatively, NP-hardness of $UPQR_q$-Manipulation could have been shown by a suitable modification of the proof of Theorem 2 in the paper by Endriss et al. (2010a). However, this modified reduction would not be appropriate to establish W[2]-hardness (as achieved in Theorem 10), since the corresponding parameterized version of SAT is not known to be W[2]-hard.

**Proposition 14.** *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. Bribery in judgment aggregation

We now turn to bribery in judgment aggregation. We will show NP-completeness for ten variants of bribery problems for judgment aggregation with the premise-based procedure and W[2]-hardness for one variant. We will first motivate bribery in judgment aggregation and define these problem variants and will then present our results.

### 4.1. Definitions and motivation

Recall again the example from Table 1(b) illustrating how the doctrinal paradox can be avoided by the premise-based procedure. Suppose there is some party (such as the club of the soccer team committing the alleged foul, or someone who bet some amount of money on the outcome of this soccer match) interested in influencing the referees' decision so that no penalty is given. This party might invest some amount of money (within a given budget) to bribe some of the referees. For example, changing any of the referees' "yes" votes in the "penalty area" or "foul" columns of Table 1(b) into a "no" would be enough to change the conclusive "yes" into a "no" under the premise-based procedure. This is a typical *bribery scenario*.

Manipulation, bribery, and lobbying are usually considered to be undesirable, and most of the recent literature on these topics is devoted to exploring the barriers to prevent such actions in terms of the computational complexity of the corresponding decision problems. Here, we introduce bribery in judgment aggregation and study its computational properties.

Define the following two problems for uniform premise-based quota rules.

| $UPQR_q$-Bribery | |
|---|---|
| **Given:** | An agenda $\Phi$, a profile $\mathbf{J} \in \mathcal{J}(\Phi)^n$, a consistent (possibly incomplete) set $J \subseteq \widehat{J} \in \mathcal{J}(\Phi)$ desired by the briber, and a positive integer $k$. |
| **Question:** | Is it possible to change up to $k$ individual judgment sets in $\mathbf{J}$ such that for the resulting new profile $\mathbf{J}'$ it holds that $HD(UPQR_q(\mathbf{J}'), J) < HD(UPQR_q(\mathbf{J}), J)$? |

| $UPQR_q$-Exact-Bribery | |
|---|---|
| **Given:** | An agenda $\Phi$, a profile $\mathbf{J} \in \mathcal{J}(\Phi)^n$, a consistent (possibly incomplete) set $J \subseteq \widehat{J} \in \mathcal{J}(\Phi)$ desired by the briber, and a positive integer $k$. |
| **Question:** | Is it possible to change up to $k$ individual judgment sets in $\mathbf{J}$ such that for the resulting new profile $\mathbf{J}'$ it holds that $J \subseteq UPQR_q(\mathbf{J}')$? |

Faliszewski et al. (2009b) introduced microbribery for voting systems. We adopt their notion so as to apply to uniform premise-based quote rules in judgment aggregation. In microbribery for judgment aggregation, if the briber's budget is $k$, she is not allowed to change up to $k$ entire judgment sets but instead can change up to $k$ premise entries in the given profile (the conclusions change automatically if necessary).

| $UPQR_q$-Microbribery | |
|---|---|
| **Given:** | An agenda $\Phi$, a profile $\mathbf{J} \in \mathcal{J}(\Phi)^n$, a consistent (possibly incomplete) set $J \subseteq \widehat{J} \in \mathcal{J}(\Phi)$ desired by the briber, and a positive integer $k$. |
| **Question:** | Is it possible to change up to $k$ premise entries in the given individual judgment sets in $\mathbf{J}$ such that for the resulting new profile $\mathbf{J}'$ it holds that $HD(UPQR_q(\mathbf{J}'), J) < HD(UPQR_q(\mathbf{J}), J)$? |

| $UPQR_q$-Exact-Microbribery | |
|---|---|
| **Given:** | An agenda $\Phi$, a profile $\mathbf{J} \in \mathcal{J}(\Phi)^n$, a consistent (possibly incomplete) set $J \subseteq \widehat{J} \in \mathcal{J}(\Phi)$ desired by the briber, and a positive integer $k$. |
| **Question:** | Is it possible to change up to $k$ premise entries in the individual judgment sets in $\mathbf{J}$ such that for the resulting new profile $\mathbf{J}'$ it holds that $J \subseteq UPQR_q(\mathbf{J}')$? |

### 4.2. Results

We will study the complexity of various bribery problems for the premise-based procedure *PBP* (i.e., $UPQR_{1/2}$ with an odd number of judges). We will establish NP-completeness for bribery, microbribery, and exact microbribery, and a W[2]-hardness result for exact bribery with respect to a natural parameter. We start with bribery.

**Theorem 15.** *PBP*-Bribery *is* NP-*complete, even when the total number of judges ($n \geq 3$ odd) or the number of judges that can be bribed is a fixed constant.*

**Proof.** Membership in NP is easy to see. We will show NP-hardness by slightly modifying the construction from the proof of Theorem 10. We start by considering the case where the briber is allowed to bribe exactly one judge. The notation and the agenda from that proof remain unchanged, but the individual judgment sets are slightly different. The first two judges remain unchanged, but the third judge has the same judgment set as the second one, and the desired set $J$ is equal to $J_3$ from the proof of Theorem 10 as shown in Table 5.

Since the quota is $1/2$, two affirmations are needed to be in the collective judgment set. Again the briber cannot benefit from bribing one judge to switch from $\neg y$ to $y$ in her individual judgment set. Hence the change must be in the set of variables $\{v_1, \ldots, v_n\}$ from the second or the third judge. By a similar argument as in the proof of Theorem 10, there is a successful bribery action if and only if there is a dominating set of size at most $k$ for the given graph.

Now we consider the case that the briber is allowed to bribe more than one judge. If the briber is allowed to bribe $k$ judges, we construct an instance with $2k+1$ judges, where one judgment set is equal to $J_1$ and the remaining $2k$ individual judgment sets are equal to $J_2$. It is again not possible for the briber to change the entry for $y$, and the briber must change the entry for any $v_i$ in the judgment sets from $k$ judges to obtain a different collective outcome. This construction works by similar arguments as above. □

Next, we turn to microbribery. Here, the briber can change only up to a fixed number of entries in the individual judgment sets. We again prove NP-completeness when the number of judges or the number of microbribes allowed is a fixed constant.

**Theorem 16.** *PBP*-Microbribery *is* NP-*complete, even when the total number of judges ($n \geq 3$ odd) or the number of microbribes allowed is a fixed constant.*

**Proof.** The proof that *PBP*-Microbribery is NP-hard is similar to the proof of Theorem 15. The agenda $\Phi$ is defined as in the proof of

**Table 5**
Construction for the proof of Theorem 15.

| Judgment set | $v_1$ | $\cdots$ | $v_n$ | $y$ | | $\varphi_1$ | $\cdots$ | $\varphi_n$ | $v_1\vee\cdots\vee v_n$ |
|---|---|---|---|---|---|---|---|---|---|
| $J_1$ | 1 | $\cdots$ | 1 | 0 | | 1 | $\cdots$ | 1 | 1 |
| $J_2, J_3$ | 0 | $\cdots$ | 0 | 0 | | 0 | $\cdots$ | 0 | 0 |
| *PBP* | 0 | $\cdots$ | 0 | 0 | $\Rightarrow$ | 0 | $\cdots$ | 0 | 0 |
| $J$ | 0 | $\cdots$ | 0 | 1 | | 1 | $\cdots$ | 1 | 0 |

Theorem 10. Let $c \in \mathbb{N}$ be a fixed constant. The number of judges is $2c + 1$, where the individual judgment sets of $c$ judges are of type $J_1$ and the remaining $c + 1$ individual judgment sets are of type $J_2$. The briber's desired outcome is the judgment set $J_3$ from the proof of Theorem 10. The number of affirmations needed to be in the collective judgment set is at least $c + 1$, and the number of entries the briber is allowed to change is at most $k$. Since none of the judges have $y$ in their individual judgment sets, the briber cannot change the collective outcome for $y$ to 1. Hence all entries that can be changed are for the variables $v_1, \dots, v_n$. Obviously, setting the value for one $v_i$ in one of the judges of type $J_2$ to 1 causes $v_i$ to be in the collective judgment set and all other changes have no effect on the collective judgment set. By similar arguments as in the proof of Theorem 10, there is a successful microbribery action if and only if the given graph has a dominating set of size at most $k$. Since membership in NP is obvious, the proof is complete. □

The next result immediately follows from the fact that Optimal-Lobbying (restricted to an odd number of voters) is a special case of *PBP*-Exact-Bribery. The formal definition of Optimal-Lobbying is as follows.

| Optimal-Lobbying | |
|---|---|
| **Given:** | An $m \times n$ 0–1 matrix $L$ (whose rows represent the voters, whose columns represent the referenda, and whose 0–1 entries represent No/Yes votes), a positive integer $k \le m$, and a target vector $x \in \{0, 1\}^n$. |
| **Question:** | Is there a choice of $k$ rows in $L$ such that by changing the entries of these rows the resulting matrix has the property that, for each $j$, $1 \le j \le n$, the $j$th column has a strict majority of ones (respectively, zeros) if and only if the $j$th entry of the target vector $x$ of The Lobby is one (respectively, zero)? |

Christian et al. (2007) have shown that Optimal-Lobbying is W[2]-complete when parameterized by the number $k$ of rows The Lobby can change, and we will apply this in our next result on exact bribery for the premise-based judgment aggregation procedure.

**Theorem 17.** *PBP-Exact-Bribery is W[2]-hard when parameterized by the number of judges that can be bribed.*

**Proof.** Observe that an exact bribery instance with only premises in the agenda and with a complete desired set $J$ is exactly the Optimal-Lobbying problem for an odd number of voters. Since this problem is W[2]-complete for the parameter number of rows that can be changed, *PBP*-Exact-Bribery inherits the W[2]-hardness lower bound, where the parameter is the number of judges that can be bribed. □

Note that the unparameterized version of this reduction also establishes that *PBP*-Exact-Bribery is NP-hard; all (unparameterized) problems considered here are easily seen to be in NP.

**Theorem 18.** *PBP-Exact-Microbribery is NP-complete, even when the total number of judges ($n \ge 3$ odd) or the number of microbribes allowed is a fixed constant.*

**Proof.** Consider the construction in the proof of Theorem 16, and change the agenda such that there are only $n^2 - 2$ (instead of $n^2 - k - 2$) syntactic variations of the formula $v_1 \vee \cdots \vee v_n$ (i.e., this can be seen as giving a weight of $n^2 - 1$ to this formula), and that the desired set $J$ is incomplete and contains all conclusions. By similar arguments as above, a successful microbribery of $k$ entries is possible if and only if there is a dominating set for $G$ of size at most $k$. □

As for the manipulation problem, Theorems 15, 16 and 18 are concerned with a fixed number of judges. It turns out that even in this case Bribery, Microbribery, and Exact-Microbribery are NP-complete for *PBP*. Furthermore, we consider the case of a fixed number of judges allowed to bribe for *PBP*-Bribery, the corresponding parameter for its exact variant, and the case where the number of microbribes allowed is a fixed constant for *PBP*-Microbribery and its exact variant. Both parameters concern the budget of the briber. Since the briber aims at spending as little money as possible, it is also natural to consider these cases. But again, NP-completeness was shown even when the budget is a fixed constant and in one case W[2]-hardness for this parameter, so bounding the budget does not help to solve the problem easily. Although the exact microbribery problem is computationally hard in general for the aggregation procedure *PBP*, there are some interesting naturally restricted instances where it is computationally easy.

**Theorem 19.** *If the desired set $J$ is complete or if it is incomplete but contains all of the premises or only premises, then PBP-Exact-Microbribery is in* P.

**Proof.** We give only an informal description of the algorithm that computes a successful microbribery. Our algorithm takes as an input a complete profile **J**, a consistent judgment set $J$, and a positive integer $k$. For each premise in $J$, compute the minimum number of entries that have to be flipped in order to make the collective judgment on that premise equal to the desired set's entry on that premise. Note that this can be done in linear time, since it is a simple counting. Let $d_i$ denote the number of entries needed to flip for premise $i$. If $\sum_i d_i \le k$, output the entries which have to be flipped and halt. Otherwise, output "bribery impossible" and halt.

Clearly, this algorithm works in polynomial time. The output is correct, since if we need less than $k$ flips in the premises, the premises are evaluated exactly as they are in $J$, and the conclusions follow automatically, since we are using a premise-based procedure. □

## 5. Conclusions and open questions

We have studied the complexity of problems related to manipulation and bribery in judgment aggregation for the uniform premise-based quota rules. In particular, for manipulation, we have extended the results of Endriss et al. (2012) 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 (2007a).

Table 6 gives an overview of our results for manipulation problems with uniform premise-based quota rules. In this table, "DS" stands for "desired set" and "NP-c" for "NP-complete."

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, 2005; Xia and Conitzer, 2011). Note that distinguishing between these notions does not apply to the exact problem variants, nor to manipulation problems based on Hamming-distance-respecting preferences.

**Table 6**
Overview of results for manipulation problems with uniform premise-based quota rules.

| | U | TR | CR | HD | Exact |
|---|---|---|---|---|---|
| Possible-Manipulation for incomplete DS | NP-c | NP-c | NP-c | NP-c | NP-c |
| Necessary-Manipulation for incomplete DS | Possibly strategy-proof | NP-c | NP-c | | |
| Possible-Manipulation for complete DS | in P | in P | NP-c (Selker, 2014) | NP-c[a] | Strategy-proof |
| Necessary-Manipulation for complete DS | Possibly strategy-proof | Possibly strategy-proof | Possibly strategy-proof | | |

[a] This result is due to Endriss et al. (2012) for the special case of the premise-based procedure.

**Table 7**
Overview of results for bribery problems with the premise-based procedure.

| Problem Type | General problem | # of judges | # of bribes | # of microbribes |
|---|---|---|---|---|
| Bribery | NP-c | NP-c | NP-c | ✗ |
| Exact-Bribery | NP-c | ? | W[2]-hard | ✗ |
| Microbribery | NP-c | NP-c | ✗ | NP-c |
| Exact-Microbribery | NP-c | NP-c | ✗ | NP-c |

Inspired by previous work on bribery in voting (see, e.g., Faliszewski et al., 2009a; Elkind et al., 2009; Faliszewski et al., 2009b; Faliszewski and Rothe, in press), we have introduced the notion of bribery in judgment aggregation. For the premise-based procedure, we have studied the complexity of bribery and microbribery problems, both in the exact variant (where the briber's goal is to have exactly the desired set occur in the collective judgment set) and in the variant where the briber simply seeks to reach a collective judgment set that is closer to the desired set in terms of their Hamming distance. Our results, stated in Table 7, show that these problems are intractable in general, and even when certain parameters – the number of judges, bribes, or microbribes – are fixed. Note that ✗ in Table 7 indicates that the corresponding problem/parameter pair is not compatible. Only one case remains open: What is the complexity of exact bribery with the premise-based procedure for a fixed number of judges? Also, it would be interesting to study bribery for other (classes of) judgment aggregation procedures, such as conclusion-based procedures (Kornhauser and Sager, 1986; List and Pettit, 2002; Dietrich, 2006), distance-based procedures (Miller and Osherson, 2009), procedures based on minimization (Lang et al., 2011), sequential procedures (List, 2004), procedures based on the Condorcet set (Nehring et al., 2014), or procedures based on the number of votes (Everaere et al., 2014).

## Acknowledgments

We are grateful to the anonymous reviewers of *Mathematical Social Sciences* for their extremely helpful, thorough, and constructive reports, as well as to the anonymous ADT'11, COMSOC'12, ADT'13, ESSLLI-LMGD'13, and COMSOC'14 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, by NRF (Singapore) grant NRF-RF 2009-08, and by an NRW grant for gender-sensitive universities and the project "Online Partizipation", both funded by the NRW Ministry for Innovation, Science, and Research. This work was done in part while the second, third, and fourth author were visiting the University of Rochester and while the fourth author was visiting Stanford University, and we thank the hosts, Lane A. Hemaspaandra and Yoav Shoham, for their warm hospitality.

## References

Alon, N., Feldman, M., Procaccia, A., Tennenholtz, M., 2013. Strategyproof approximation of the minimax on networks. Math. Oper. Res. 35 (3), 513–526.
Bartholdi III, J., Orlin, J., 1991. Single transferable vote resists strategic voting. Soc. Choice Welf. 8 (4), 341–354.
Bartholdi III, J., Tovey, C., Trick, M., 1989. The computational difficulty of manipulating an election. Soc. Choice Welf. 6 (3), 227–241.
Baumeister, D., Bouveret, S., Lang, J., Nguyen, N., Nguyen, T., Rothe, J., 2014a. Scoring rules for the allocation of indivisible goods. In: Proceedings of the 21st European Conference on Artificial Intelligence. IOS Press, pp. 75–80.
Baumeister, D., Erdélyi, G., Erdélyi, O., Rothe, J., Bribery and control in judgment aggregation. In: F. Brandt, P. Faliszewski (Eds.), Proceedings of the 4th International Workshop on Computational Social Choice, AGH University of Science and Technology, Kraków, Poland, Sept. 2012. pp. 37–48.
Baumeister, D., Erdélyi, G., Erdélyi, O., Rothe, J., 2013. Computational aspects of manipulation and control in judgment aggregation. In: Proceedings of the 3rd International Conference on Algorithmic Decision Theory. In: Lecture Notes in Artificial Intelligence, #8176. Springer-Verlag, pp. 71–85.
Baumeister, D., Erdélyi, G., Erdélyi, O., Rothe, J., 2014b. Complexity of manipulative attacks in judgment aggregation for premise-based quota rules. In: A. Procaccia, T. Walsh (Eds.), Proceedings of the 5th International Workshop on Computational Social Choice, Carnegie Mellon University.
Baumeister, D., Erdélyi, G., Hemaspaandra, E., Hemaspaandra, L., Rothe, J., 2010. Computational aspects of approval voting. In: Laslier, J., Sanver, R. (Eds.), Handbook on Approval Voting. Springer, pp. 199–251. chapter 10.
Baumeister, D., Erdélyi, G., Rothe, J., 2011. 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. In: Lecture Notes in Artificial Intelligence, #6992. Springer-Verlag, pp. 1–15.
Baumeister, D., Rothe, J., 2012. Taking the final step to a full dichotomy of the possible winner problem in pure scoring rules. Inf. Process. Lett. 112 (5), 186–190.
Baumeister, D., Rothe, J., 2015. Preference aggregation by voting. In: Rothe, J. (Ed.), Economics and Computation. An Introduction to Algorithmic Game Theory, Computational Social Choice, and Fair Division. Springer-Verlag, chapter 4 (in press).
Betzler, N., Dorn, B., 2010. Towards a dichotomy for the possible winner problem in elections based on scoring rules. J. Comput. System Sci. 76 (8), 812–836.
Binkele-Raible, D., Erdélyi, G., Fernau, H., Goldsmith, J., Mattei, N., Rothe, J., 2014. The complexity of probabilistic lobbying. Discrete Optim. 11 (1), 1–21.
Boutilier, C., Rosenschein, J., 2015. Incomplete information and communication in voting. In: Brandt, F., Conitzer, V., Endriss, U., Lang, J., Procaccia, A. (Eds.), Handbook of Computational Social Choice. Cambridge University Press, chapter 10 (in press).
Bouveret, S., Endriss, U., Lang, J., 2010. Fair division under ordinal preferences: Computing envy-free allocations of indivisible goods. In: Proceedings of the 19th European Conference on Artificial Intelligence. IOS Press, pp. 387–392.
Brams, S., Edelman, P., Fishburn, P., 2004. Fair division of indivisible items. Theory and Decision 5 (2), 147–180.
Brams, S., King, D., 2005. Efficient fair division—help the worst off or avoid envy? Ration. Soc. 17 (4), 387–421.
Brandt, F., Conitzer, V., Endriss, U., 2012. Computational social choice. In: Weiß, G. (Ed.), Multiagent Systems, second edition. MIT Press, pp. 213–283.
Bredereck, R., Chen, J., Hartung, S., Kratsch, S., Niedermeier, R., Suchý, O., 2012. A multivariate complexity analysis of lobbying in multiple referenda. In: Proceedings of the 26th AAAI Conference on Artificial Intelligence. AAAI Press, pp. 1292–1298.
Christian, R., Fellows, M., Rosamond, F., Slinko, A., 2007. On complexity of lobbying in multiple referenda. Rev. Econ. Des. 11 (3), 217–224.
Conitzer, V., Sandholm, T., Lang, J., 2007. When are elections with few candidates hard to manipulate? J. ACM 54 (3), Article 14.
Conitzer, V., Walsh, T., 2015. Barriers to manipulation in voting. In: Brandt, F., Conitzer, V., Endriss, U., Lang, J., Procaccia, A. (Eds.), Handbook of Computational Social Choice. Cambridge University Press, chapter 6 (in press).
Dietrich, F., 2006. Judgment aggregation: (im)possibility theorems. J. Econom. Theory 126 (1), 286–298.
Dietrich, F., 2014. Judgment aggregation and agenda manipulation. Manuscript.
Dietrich, F., List, C., 2007a. Strategy-proof judgment aggregation. Econ. Philos. 23 (3), 269–300.
Dietrich, F., List, C., 2007b. Judgment aggregation by quota rules: Majority voting generalized. J. Theor. Polit. 19 (4), 391–424.
Downey, R., Fellows, M., 1999. Parameterized Complexity. Springer-Verlag.
Duggan, J., Schwartz, T., 2000. Strategic manipulability without resoluteness or shared beliefs: Gibbard–Satterthwaite generalized. Soc. Choice Welf. 17 (1), 85–93.
Elkind, E., Faliszewski, P., Slinko, A., 2009. Swap bribery. In: Proceedings of the 2nd International Symposium on Algorithmic Game Theory. In: Lecture Notes in Computer Science, #5814. Springer-Verlag, pp. 299–310.
Endriss, U., 2015. Judgment aggregation. In: Brandt, F., Conitzer, V., Endriss, U., Lang, J., Procaccia, A. (Eds.), Handbook of Computational Social Choice. Cambridge University Press, chapter 17 (in press).

Endriss, U., Grandi, U., Porello, D., 2010a. Complexity of winner determination and strategic manipulation in judgment aggregation. In: V. Conitzer and J. Rothe (Eds.), Proceedings of the 3rd International Workshop on Computational Social Choice, Universität Düsseldorf, pp. 139–150.

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

Endriss, U., Grandi, U., Porello, D., 2012. Complexity of judgment aggregation. J. Artificial Intelligence Res. 45, 481–514.

Everaere, P., Konieczny, S., Marquis, P., 2014. Counting votes for aggregating judgments. In: Proceedings of the 13th International Joint Conference on Autonomous Agents and Multiagent Systems, IFAAMAS, pp. 1177–1184.

Falik, D., Dokow, E., 2012. Models of manipulation on aggregation of binary evaluations. In: F. Brandt and P. Faliszewski (Eds.), Proceedings of the 4th International Workshop on Computational Social Choice, AGH University of Science and Technology, Kraków, Poland, pp. 191–202.

Faliszewski, P., Hemaspaandra, E., Hemaspaandra, L., 2009a. How hard is bribery in elections? J. Artificial Intelligence Res. 35, 485–532.

Faliszewski, P., Hemaspaandra, E., Hemaspaandra, L., 2010. Using complexity to protect elections. Commun. ACM 53 (11), 74–82.

Faliszewski, P., Hemaspaandra, E., Hemaspaandra, L., Rothe, J., 2009b. Llull and Copeland voting computationally resist bribery and constructive control. J. Artificial Intelligence Res. 35, 275–341.

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

Faliszewski, P., Reisch, Y., Rothe, J., Schend, L., 2015. Complexity of manipulation, bribery, and campaign management in Bucklin and fallback voting. J. Auton. Agents Multi-Agent Syst. in press.

Faliszewski, P., Rothe, J., 2015. Control and bribery. In: Brandt, F., Conitzer, V., Endriss, U., Lang, J., Procaccia, A. (Eds.), Handbook of Computational Social Choice. Cambridge University Press, chapter 7 (in press).

Flum, J., Grohe, M., 2006. Parameterized Complexity Theory. In: EATCS Texts in Theoretical Computer Science, Springer-Verlag.

Gibbard, A., 1973. Manipulation of voting schemes. Econometrica 41 (4), 587–601.

Hemaspaandra, E., Hemaspaandra, L., Menton, C., 2013. Search versus decision for election manipulation problems. In: Proceedings of the 30th Annual Symposium on Theoretical Aspects of Computer Science. In: LIPIcs, vol. 20. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, pp. 377–388.

Klaus, B., Miyagawa, E., 2013. Strategy-proofness, solidarity, and consistency for multiple assignment problems. Internat. J. Game Theory 30 (3), 421–435.

Konczak, K., Lang, J., 2005. Voting procedures with incomplete preferences. In: Proceedings of the Multidisciplinary IJCAI-05 Workshop on Advances in Preference Handling, pp. 124–129.

Kornhauser, L., Sager, L., 1986. Unpacking the court. Yale Law J. 96 (1), 82–117.

Lang, J., Pigozzi, G., Slavkovik, M., Torre, J., 2011. Judgment aggregation rules based on minimization. In: Proceedings of the 13th Conference on Theoretical Aspects of Rationality and Knowledge. ACM Press, pp. 238–246.

List, C., 2004. A model of path-dependence in decisions over multiple propositions. Am. Polit. Sci. Rev. 93 (3), 495–513.

List, C., 2006. The discursive dilemma and public reason. Ethics 116 (2), 362–402.

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

List, C., Pettit, P., 2002. Aggregating sets of judgments: An impossibility result. Econ. Philos. 18 (1), 89–110.

List, C., Puppe, C., 2009. Judgment aggregation. In: Anand, P., Pattanaik, P., Puppe, C. (Eds.), The Handbook of Rational and Social Choice. Oxford University Press, pp. 457–482. chapter 19.

Miller, M., Osherson, D., 2009. Methods for distance-based judgment aggregation. Soc. Choice Welf. 32 (4), 575–601.

Nehring, K., Pivato, M., Puppe, C., 2014. The condorcet set: Majority voting over interconnected propositions. J. Econom. Theory 151, 286–303.

Niedermeier, R., 2006. Invitation to Fixed-Parameter Algorithms. Oxford University Press.

Papadimitriou, C., 1995. Computational Complexity, second edition. Addison-Wesley.

Rothe, J., 2005. Complexity Theory and Cryptology. An Introduction to Cryptocomplexity. In: EATCS Texts in Theoretical Computer Science, Springer-Verlag.

Satterthwaite, M., 1975. Strategy-proofness and Arrow's conditions: Existence and correspondence theorems for voting procedures and social welfare functions. J. Econom. Theory 10 (2), 187–217.

Selker, A., 2014. Manipulative Angriffe auf Judgment-Aggregation-Prozeduren (Master's thesis), Heinrich-Heine-Universität Düsseldorf, Institut für Informatik, Düsseldorf, Germany.

Svensson, L., 1999. Strategy-proof allocation of indivisible goods. Soc. Choice Welf. 16, 557–567.

Takamiya, K., 2013. Coalitional unanimity versus strategy-proofness in coalition formation problems. Internat. J. Game Theory 42 (1), 115–130.

Xia, L., Conitzer, V., 2011. Determining possible and necessary winners given partial orders. J. Artificial Intelligence Res. 41, 25–67.