Technical Report AC-TR-17-005

March 2017

Complexity Results for Manipulation, Bribery and Control of the Kemeny Judgment Aggregation Procedure

Ronald de Haan

This is the author's copy of a paper that appeared in the proceedings of AAMAS 2017.

www.ac.tuwien.ac.at/tr

Complexity Results for Manipulation, Bribery and Control of the Kemeny Judgment Aggregation Procedure

Ronald de Haan
Technische Universität Wien
Vienna, Austria
dehaan@ac.tuwien.ac.at

ABSTRACT

An important criterium for social choice methods is their resistance against various types of strategic behavior. Seminal results in the social choice literature indicate that absolute resistance is in many cases impossible. For this reason, it has often been argued that computational intractability could be used as an obstruction for strategic behavior for different procedures.

In this paper, we study the computational complexity of strategic behavior for the Kemeny procedure in the setting of judgment aggregation. In particular, we investigate problems related to (1) strategic manipulation, (2) bribery, and (3) control (by adding or deleting issues). We show that these problems are complete for the second level of the Polynomial Hierarchy. Our results hold for two different judgment aggregation frameworks and for different notions of preference over judgment sets. The hardness results that we establish hold up even under various restrictions, such as unidimensional alignment of the profile.

CCS Concepts

•**Theory of computation** → **Problems, reductions and completeness;** •**Computing methodologies** → *Artificial intelligence;*

Keywords

Judgment aggregation, computational complexity, computational social choice, strategic behavior, manipulation, bribery, control

1. INTRODUCTION

An important topic in the research field of computational social choice is the (im)possibility of strategic behavior in collective decision making. This is epitomized by the eminence of results such as the Gibbard-Satterthwaite Theorem [20, 30], that identifies various conditions under which strategic voting (or manipulation) is—in principle—unavoidable. Manipulation in voting is a typical example of strategic behavior, and involves individuals reporting insincere preferences with the aim of obtaining a group outcome that is preferable for them.

Appears in: *Proc. of the 16th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2017), S. Das, E. Durfee, K. Larson, M. Winikoff (eds.), May 8–12, 2017, São Paulo, Brazil.*

Since strategic behavior in collective decision making is generally considered to be (socially) undesirable, a lot of research effort has been invested in diagnosing what social choice procedures are resistant to strategic behavior, and under what conditions. An important research direction along these lines investigates how computational complexity can be used to establish that various social choice procedures are (in many cases) practically immune to strategic behavior [2, 10]. For example, in many cases, it is in principle possible to manipulate voting rules (by reporting insincere preferences), but determining what insincere preference leads to a better outcome is computationally so demanding that it prevents manipulative behavior from being a useful policy.

In this paper, we investigate the computational complexity of several types of strategic behavior in the setting of *judgment aggregation*. Judgment aggregation studies collective decision making on a set of issues that are logically related [16]—in this setting, individuals report their judgments on the issues, and an aggregation procedure is applied to combine these individual opinions into a single collective opinion. One judgment aggregation that is often considered in the literature, and on which we focus in this paper, is the *(generalized) Kemeny procedure*[1] (see, e.g., [16]).

The types of strategic behavior that we study in this paper are *manipulation*, *bribery* and *control*. Manipulation involves an individual reporting an insincere judgment with the aim to get a group outcome that is better for this individual. Bribery involves an external party achieving some desired collective outcome by persuading a number of individuals to report insincere judgments. Control involves an external party attaining some desired collective outcome by changing the set of issues in the judgment aggregation scenario—some issues are removed or added.

Contributions.

We show that the problem of deciding whether a desired outcome can be achieved by means of these three types of strategic behavior for the Kemeny judgment aggregation procedure is Σ_2^p-complete (Theorems 2, 5 and 8). These results hold for two different judgment aggregation frameworks, and for several notions of preference relations over judgments. Moreover, we show that this result extends to the setting of group manipulation, where a coalition of individuals reports insincere judgments to obtain a better collective outcome. These results are a good starting point for more detailed com-

[1]This procedure is also known as the Prototype-Hamming procedure [26], the distance-based procedure [18], the median rule [27], MWA [23], and the simple scoring rule [11].

plexity investigations, e.g., using parameterized complexity or approximation methods.

Our Σ_2^p-completeness result for manipulation of the Kemeny judgment aggregation procedure, in particular, answers an open question from the literature [29, p. 382].

Related Work.

For the Kemeny judgment aggregation procedure, the complexity of computing collective judgments has been investigated in the literature [18, 19, 23], also from a parameterized complexity point of view [22]. Moreover, the computational complexity of several problems related to strategic behavior in judgment aggregation—manipulation, bribery and control—has been considered for a class of judgment aggregation procedures known as uniform premise-based quota rules [3, 4, 5, 6, 7, 8, 18]. Our results are complementary to previous work from the literature. All the contributions that are mentioned above either studied different problems for the same judgment aggregation procedure, or studied the same problems for different judgment aggregation procedures.

Roadmap.

We begin in Section 2 with revisiting relevant notions from judgment aggregation and computational complexity. Then, in Section 3, we discuss the different kinds of strategic behavior that we consider, and we specify the problems that we investigate. In Section 4, we present the computational complexity results that we establish. Finally, in Section 5, we conclude and discuss directions for future research.

2. PRELIMINARIES

We begin by reviewing several relevant concepts from judgment aggregation and computational complexity theory.

2.1 Judgment Aggregation

We discuss the two formal judgment aggregation frameworks that we use in this paper: *formula-based judgment aggregation* (as used in, e.g., [13, 18, 23]) and *constraint-based judgment aggregation* (as used in, e.g., [17, 21]). The former we present in detail, whereas for the latter, we only discuss the main features of the framework—the reason for this is that the proofs that we present in detail in this paper are for the formula-based judgment aggregation framework. Moreover, we describe the concept of unidimensional alignment of profiles in judgment aggregation.

2.1.1 Formula-Based Judgment Aggregation

We begin with the framework of formula-based judgment aggregation. An *agenda* is a finite, non-empty set Φ of formulas that does not contain any doubly-negated formulas and that is closed under complementation. Moreover, if $\Phi = \{\varphi_1, \dots, \varphi_n, \neg\varphi_1, \dots, \neg\varphi_n\}$ is an agenda, then we let $[\Phi] = \{\varphi_1, \dots, \varphi_n\}$ denote the *pre-agenda* associated to the agenda Φ. For every propositional formula φ, we let $\sim\varphi$ denote the *complement* of φ. We denote the bitsize of the agenda Φ by $\text{size}(\Phi) = \sum_{\varphi \in \Phi} |\varphi|$. A *judgment set* J for an agenda Φ is a subset $J \subseteq \Phi$. We call a judgment set J *complete* if $\varphi \in J$ or $\sim\varphi \in J$ for all $\varphi \in \Phi$; and we call it *consistent* if there exists an assignment that makes all formulas in J true. Intuitively, the consistent and complete judgment sets are the opinions that individuals and the group can have.

We associate with each agenda Φ an integrity constraint Γ, that can be used to further restrict the set of feasible opinions. Such an *integrity constraint* consists of a single propositional formula. We say that a judgment set J is Γ-*consistent* if there exists a truth assignment that simultaneously makes all formulas in J and Γ true. Let $\mathcal{J}(\Phi, \Gamma)$ denote the set of all complete and Γ-consistent subsets of Φ. We say that finite sequences $\boldsymbol{J} \in \mathcal{J}(\Phi, \Gamma)^+$ of complete and Γ-consistent judgment sets are *profiles*, and where convenient we equate a profile $\boldsymbol{J} = (J_1, \dots, J_p)$ with the (multi)set $\{J_1, \dots, J_p\}$. Moreover, for $i \in [p] = \{1, \dots, p\}$, we let $\boldsymbol{J}_{-i}$ denote the profile $(J_1, \dots, J_{i-1}, J_{i+1}, \dots, J_p)$.

A *judgment aggregation procedure* (or *rule*) for the agenda Φ and the integrity constraint Γ is a function F that takes as input a profile $\boldsymbol{J} \in \mathcal{J}(\Phi, \Gamma)^+$, and that produces a non-empty set of non-empty judgment sets. We call a judgment aggregation procedure F *resolute* if for any profile $\boldsymbol{J}$ it returns a singleton, i.e., $|F(\boldsymbol{J})| = 1$; otherwise, we call F *irresolute*. We call a judgment aggregation procedure F *complete* and Γ-*consistent*, if J is complete and Γ-consistent, respectively, for every $\boldsymbol{J} \in \mathcal{J}(\Phi, \Gamma)^+$ and every $J \in F(\boldsymbol{J})$.

The *Kemeny aggregation procedure* is based on the Hamming distance $d(J, J') = |\{\, \varphi \in [\Phi] : \varphi \in (J \backslash J') \cup (J' \backslash J) \,\}|$ between two complete judgment sets J, J'. Intuitively, the Hamming distance $d(J, J')$ counts the number of issues on which two judgment sets disagree. Let J be a single Γ-consistent and complete judgment set, and let $(J_1, \dots, J_p) = \boldsymbol{J} \in \mathcal{J}(\Phi, \Gamma)^+$ be a profile. We define the distance between J and $\boldsymbol{J}$ to be $d(J, \boldsymbol{J}) = \sum_{i \in [p]} d(J, J_i)$. Then, we let the outcome $\text{Kemeny}_{\Phi,\Gamma}(\boldsymbol{J})$ of the Kemeny rule be the set of those $J^* \in \mathcal{J}(\Phi, \Gamma)$ for which there is no $J \in \mathcal{J}(\Phi, \Gamma)$ such that $d(J, \boldsymbol{J}) < d(J^*, \boldsymbol{J})$. (If Φ and Γ are clear from the context, we often write $\text{Kemeny}(\boldsymbol{J})$ to denote $\text{Kemeny}_{\Phi,\Gamma}(\boldsymbol{J})$.) Intuitively, the Kemeny rule selects those complete and Γ-consistent judgment sets that minimize the cumulative Hamming distance to the judgment sets in the profile. The Kemeny rule is irresolute, complete and Γ-consistent.

2.1.2 Constraint-Based Judgment Aggregation

We continue with a brief description of the framework of constraint-based judgment aggregation, and focus on the difference with the formula-based judgment aggregation framework. Instead of using an agenda Φ to model the issues, a finite set $\mathcal{I} = \{x_1, \dots, x_n\}$ of propositional variables is used. Accordingly, truth assignments $\alpha : \mathcal{I} \to \mathbb{B}$ to these variables are used instead of complete judgment sets to represent opinions—we use $\mathbb{B}$ to denote the set $\{0, 1\}$ of truth values. The logical relation between the issues is modelled using an integrity constraint Γ, which is a propositional formula containing only variables $\mathcal{I}$. This means that no additional variables (beyond those in $\mathcal{I}$) can be used to specify the logical relation between issues—in contrast, in the formula-based framework, one is free to introduce additional variables.

The notions of Γ-consistent opinions, profiles, and judgment aggregation procedures are then defined entirely similarly to the case of formula-based judgment aggregation. In particular, the Kemeny judgment aggregation procedure is defined similarly as for the case of formula-based judgment aggregation. For more details, we refer to papers in the literature that feature the constraint-based judgment aggregation framework (e.g., [17, 21]).

2.1.3 Unidimensional Alignment

A property of judgment aggregation scenarios (more specifically, of profiles) that is of use for computing the outcome of various judgment aggregation procedures more efficiently is that of *unidimensional alignment* [24]. We describe this property for the setting of formula-based judgment aggregation. For the setting of constraint-based judgment aggregation, it is defined entirely similarly.

A profile $\boldsymbol{J} = (J_1, \dots, J_p)$ over an agenda Φ is *unidimensionally aligned* if there exists a bijection $\pi : \{J_1, \dots, J_p\} \rightarrow [p]$ such that for each $\varphi \in [\Phi]$ it holds that there are no $J'_1, J'_2, J'_3 \in \boldsymbol{J}$ such that both $\pi(J'_1) < \pi(J'_2) < \pi(J'_3)$ and either (1) $\varphi \in J'_1$, $\neg\varphi \in J'_2$ and $\varphi \in J'_3$ or (2) $\neg\varphi \in J'_1$, $\varphi \in J'_2$ and $\neg\varphi \in J'_3$. In other words, when ordering the judgment sets in $\boldsymbol{J}$ according to the permutation π, all judgment sets containing φ appear either to the left or to the right of all judgment sets containing $\neg\varphi$.

For any profile $\boldsymbol{J}$ (containing an odd number of judgment sets) that is unidimensionally aligned, the majority outcome is consistent. Namely, the majority outcome is identical to the median judgment set according to the bijection π that witnesses the unidimension alignment of $\boldsymbol{J}$. Therefore, in such cases, computing the outcome of the Kemeny judgment aggregation procedure, for instance, is easy, because it coincides with the majority outcome in those cases where the majority outcome is consistent.

2.2 Computational Complexity

Next, we review some notions from computational complexity. We assume the reader to be familiar with basic concepts. For more details, we refer to textbooks (see, e.g., [1]).

We briefly review the classes of the Polynomial Hierarchy (PH) [25, 28, 31, 32]. In order to do so, we consider quantified Boolean formulas. A *quantified Boolean formula (in prenex form)* is a formula of the form $Q_1 x_1 Q_2 x_2 \dots Q_n x_n . \psi$, where all x_i are propositional variables, each Q_i is either an existential or a universal quantifier, and ψ is a (quantifier-free) propositional formula over the variables $x_1, \dots, x_n$. Truth for such formulas is defined in the usual way.

To consider the complexity classes of the PH, we restrict the number of quantifier alternations occurring in quantified Boolean formulas, i.e., the number of times where $Q_i \neq Q_{i+1}$. We consider the complexity classes Σ^{p}_k, for each $k \geq 1$. Let $k \geq 1$ be an arbitrary, fixed constant. The complexity class Σ^{p}_k consists of all decision problems for which there exists a polynomial-time reduction to the problem QSAT_k, that is defined as follows. Instances of the problem QSAT_k are quantified Boolean formulas of the form $\exists x_1 \dots \exists x_{\ell_1} \forall x_{\ell_1+1} \dots \forall x_{\ell_2} \dots Q_k x_{\ell_{k-1}+1} \dots Q_k x_{\ell_k} . \psi$, where $Q_k = \exists$ if k is odd and $Q_k = \forall$ if k is even, where $1 \leq \ell_1 \leq \dots \leq \ell_k$, and where ψ is quantifier-free. The problem is to decide if the quantified Boolean formula is true.

Alternatively, one can characterize the class Σ^{p}_2 using nondeterministic polynomial-time algorithms with access to an oracle for an NP-complete problem. Let O be a decision problem. A Turing machine $\mathbb{M}$ with access to an O *oracle* is a Turing machine with a dedicated *oracle tape* and dedicated states q_{query}, q_{yes} and q_{no}. Whenever $\mathbb{M}$ is in the state q_{query}, it does not proceed according to the transition relation, but instead it transitions into the state q_{yes} if the oracle tape contains a string x that is a yes-instance for the problem O, i.e., if $x \in O$, and it transitions into the state q_{no} if $x \notin O$. Intuitively, the oracle solves arbitrary instances of O in a single time step. The class Σ^{p}_2 consists of all decision problems that can be solved in polynomial time by a nondeterministic Turing machine that has access to an O-oracle, for some $O \in \mathrm{NP}$.

3. MANIPULATION, BRIBERY AND CONTROL IN JUDGMENT AGGREGATION

In this section, we describe the different strategic behavior scenarios that we investigate—manipulation, bribery and control. Moreover, we formally define the decision problems that we use to model the different kinds of strategic behavior. (We consider decision problems because they are technically more convenient to analyze than search problems. The hardness results that we obtain for the decision problems imply that no efficient algorithm exists to solve search variants of the decision problems that we consider.)

In order to precisely state what we mean with the different kinds of strategic behavior, we need to specify a notion of preference over judgments. We consider two different notions of preference: *subset-based preferences* and *preferences based on (weighted) Hamming distances*—we focus on the former in our presentation of the problems and the results.

3.1 Preferences over Judgment Sets

The different types of strategic behavior that we will consider all involve the incentive to obtain a "better" outcome. Therefore, in order to study strategic behavior in judgment aggregation, it is essential to define a notion of preference over opinions—i.e., when is one opinion preferred over another.

In the worst case, the number of possible opinions that play a role is exponential in the number of issues—e.g., for m issues there could be up to 2^m possible opinions. As a result, it is unreasonable to expect agents to explicitly specify a preference relation over all (feasible) opinions. Instead it makes more sense to use a compact specification language to represent a preference relation. Various preference relations over opinions have been studied [6, 12, 14].

In this paper, we consider two types of compactly specified preferences over opinions (or judgment sets). We describe these below, after which we briefly discuss other notions of preferences over opinions that have been considered in the judgment aggregation literature. (In this section, we consider the case of formula-based judgment aggregation; the case of constraint-based judgment aggregation is entirely similar.)

3.1.1 Subset-based Preferences

The first compact method of specifying preferences that we consider is that of *subset-based preferences*. For this preference relation over judgments, an agent with sincere judgment $J \in \mathcal{J}(\Phi, \Gamma)$ specifies a subset $L \subseteq J$ of important issues. Then for judgment sets $J_1, J_2 \in \mathcal{J}(\Phi, \Gamma)$, judgment set J_1 is preferred over judgment set J_2 if $L \subseteq J_1$ and $L \not\subseteq J_2$. In other words, every judgment set that includes L is preferred over every judgment set that does not include L. (Judgment sets that both include L are equally preferable; similarly for judgment sets that both do not include L.)

3.1.2 (Weighted) Hamming Distance Preferences

Another type of preferences that we consider is the class of preferences based on a *weighted Hamming distance*. An agent can specify their preference relation over complete and Γ-consistent judgment sets $J \in \mathcal{J}(\Phi, \Gamma)$ by providing a weight function $w : [\Phi] \to \mathbb{N}$ that produces a weight $w(\varphi)$ for each formula $\varphi \in [\Phi]$. Intuitively, for each $\varphi \in [\Phi]$, the weight $w(\varphi)$ indicates how important it is for the agent that the outcome agrees with their truthful opinion on the issue φ. (Alternatively, one could consider weight functions that produce rational or real weights.) Then, for two complete judgment sets J_1 and J_2 the weighted Hamming distance $d(J_1, J_2, w)$ is defined by letting $d(J_1, J_2, w) = \sum\{ w(\varphi) : \varphi \in (J_1 \setminus J_2) \cup (J_2 \setminus J_1) \}$. That is, for each formula $\varphi \in [\Phi]$ that J and J' disagree on, the weighted Hamming distance is increased by $w(\varphi)$.

Using this notion of weighted Hamming distance, we can define a preference relation for an agent. Suppose that the agent's truthful opinion is given by a complete and Γ-consistent judgment set J. Moreover, suppose that the agent's view on the relative importance of the separate issues is given by a weight function $w : [\Phi] \to \mathbb{N}$. Then the preference relation $\leq_{w,J}$ for this agent is defined as follows. For any two complete and Γ-consistent judgment sets J_1, J_2, it holds that $J_1 \geq_{w,J} J_2$ if and only if $d(J, J_1, w) \leq d(J, J_2, w)$. Correspondingly, a judgment set J_1 is (strictly) preferred over another judgment set J_2 if and only if the weighted Hamming distance from J_1 to J is (strictly) smaller than the weighted Hamming distance from J_2 to J.

A particular case of the weighted Hamming distance is the *unweighted Hamming distance*. That is, the case where $w(\varphi) = 1$ for all $\varphi \in [\Phi]$. Whenever the weight function w is the constant function that always returns 1, we drop the "w" from the notation—that is, the unweighted Hamming distance between two judgment sets J_1 and J_2 is denoted by $d(J_1, J_2)$.

Subset-based preferences for some $L \subseteq J$ with $|L| = 1$ can also be seen as a particular case of weighted Hamming distances: namely, take $w(\varphi_L) = 1$ for the unique $\varphi_L \in [\Phi]$ with $\{\varphi_L, {\sim}\varphi_L\} \cap L \neq \emptyset$, and $w(\varphi) = 0$ for all other $\varphi \in [\Phi]$.

3.1.3 Other preference relations

In the literature, there have been various proposals for notions of preference over opinions. For example, the phenomenon of manipulation in judgment aggregation has been studied in the settings (1) where one judgment set is preferred over another if it agrees with a fixed optimal judgment set on at least one issue where the other judgment set disagrees [14], and (2) where one judgment set is preferred over a second judgment set if it agrees with a fixed optimal judgment set on at least one issue where the second judgment set disagrees, and for all issues it holds that if the second judgment set agrees with the optimal judgment set then the first judgment set also agrees with the optimal [14]. Other preference relations that have been investigated are top-respecting preferences and closeness-respecting preferences. The class of top-respecting preferences contains all preferences that prefer a single most preferred judgment set over all other judgment sets (and the preference between the other judgment sets is arbitrary) [6, 12]. The class of closeness-respecting preferences contains preferences that additionally satisfy the condition of closeness: if one judgment set agrees with the most preferred judgment on a superset of issues compared to another judgment set, then the one judgment is preferred over the other [6, 12].

3.2 Manipulation

The first form of strategic behavior in judgment aggregation that we consider is *manipulation*. This involves an individual aiming to influence the outcome of the aggregation procedure in their favor by reporting an insincere judgment. We model this using the following decision problem.

> MANIPULATION(Kemeny)
> *Instance:* An agenda Φ with an integrity constraint Γ, a profile $(J_1, \dots, J_p) = \boldsymbol{J} \in \mathcal{J}(\Phi, \Gamma)^+$, and a subset $L \subseteq J_1$.
> *Question:* Is there a complete and consistent judgment set $J' \in \mathcal{J}(\Phi, \Gamma)$ such that for all $J^*_{\text{new}} \in \text{Kemeny}(\boldsymbol{J}_{-1}, J')$ it holds that $L \subseteq J^*_{\text{new}}$?

3.3 Bribery

Another form of strategic behavior in judgment aggregation is *bribery*. In this setting, an external agent wishes to influence the outcome of a judgment aggregation scenario by bribing a number of individuals.

The briber has a set $L \subseteq \Phi$ of desired conclusions that they want to attain in the collective opinion. Additionally, the briber has a budget that suffices to bribe at most k individuals. For all bribed individuals, the briber can specify an arbitrary (complete and Γ-consistent) judgment set. The question is to determine whether the briber can pick up to k individuals and specify judgment sets for these individuals so that the outcome of the judgment aggregation procedure is better (with respect to L) than without bribing. We model this using the following decision problem.

> BRIBERY(Kemeny)
> *Instance:* An agenda Φ with an integrity constraint Γ, a profile $\boldsymbol{J} \in \mathcal{J}(\Phi, \Gamma)^+$, a set $L \subseteq \Phi$, and an integer $k \in \mathbb{N}$.
> *Question:* Is it possible to change up to k individual judgment sets in $\boldsymbol{J}$, resulting in a new profile $\boldsymbol{J}'$, so that for all $J^*_{\text{new}} \in \text{Kemeny}(\boldsymbol{J}')$ it holds that $L \subseteq J^*_{\text{new}}$?

3.4 Control by Adding or Removing Issues

A third form of strategic behavior in judgment aggregation is *control*. In this setting, an external agent wishes to influence the outcome of by influencing the conditions of a judgment aggregation scenario. Here, we consider control by (1) adding or (2) deleting issues.

We begin with the scenario of (1) control by adding issues. A number of individuals each have an opinion for an agenda Φ in the presence of an integrity constraint Γ. That is, we are considering a profile $\boldsymbol{J} \in \mathcal{J}(\Phi, \Gamma)^+$. However, they are performing judgment aggregation only on a selection of issues, specified by an agenda $\Phi' \subseteq \Phi$. (For any $\Psi \subseteq \Phi$, we let the profile $\boldsymbol{J}|_\Psi$ consist of the judgment sets $J|_\Psi$ for each $J \in \boldsymbol{J}$, where $J|_\Psi = J \cap \Psi$—that is, $\boldsymbol{J}|_\Psi = \{ J \cap \Psi : J \in \boldsymbol{J} \}$. Intuitively, $\boldsymbol{J}|_\Psi$ is the restriction of $\boldsymbol{J}$ to the formulas in Ψ.) The external agent wishes to ensure that the outcome of the judgment aggregation procedure includes a set $L \subseteq \Phi$ of desired conclusions, and they want to do so by enlarging the set of issues that the individuals perform judgment aggregation on. Formally, the external agent wants to select an agenda Φ'' with $\Phi' \subseteq \Phi'' \subseteq \Phi$ such that $L \subseteq J^*$ for all $J^* \in \text{Kemeny}(\boldsymbol{J}|_{\Phi''})$. (Obviously, if the external agent wishes to succeed, they need to choose some Φ'' with $L \subseteq \Phi''$.) We model this using the following decision problem.

CONTROL-BY-ADDING-ISSUES(Kemeny)
Instance: An agenda Φ with an integrity constraint Γ, an agenda $\Phi' \subseteq \Phi$, a set $L \subseteq \Phi$ and a profile $\boldsymbol{J} \in \mathcal{J}(\Phi, \Gamma)^{+}$.
Question: Is there an agenda Φ'' with $\Phi' \subseteq \Phi'' \subseteq \Phi$ such that for all $J^{*} \in \text{Kemeny}(\boldsymbol{J}|_{\Phi''})$ it holds that $L \subseteq J^{*}$?

We continue with the scenario of (2) control by deleting issues. In this scenario, the external agent wishes to ensure that the outcome of the judgment aggregation procedure includes a set $L \subseteq \Phi$ of desired conclusions, and they want to do so by restricting the set of issues that the individuals perform judgment aggregation on. Formally, the external agent wants to select an agenda $\Phi'' \subseteq \Phi$ such that $L \subseteq J^{*}$ for each $J^{*} \in \text{Kemeny}(\boldsymbol{J}|_{\Phi''})$. (Again, if the external agent wishes to succeed, they need to choose some Φ'' with $L \subseteq \Phi''$.) We model this using the following decision problem.

CONTROL-BY-REMOVING-ISSUES(Kemeny)
Instance: An agenda Φ with an integrity constraint Γ, a set $L \subseteq \Phi$ and a profile $\boldsymbol{J} \in \mathcal{J}(\Phi, \Gamma)^{+}$.
Question: Is there an agenda $\Phi'' \subseteq \Phi$ such that for all $J^{*} \in \text{Kemeny}(\boldsymbol{J}|_{\Phi''})$ it holds that $L \subseteq J^{*}$?

3.5 Problem Variants

In the discussion of the different types of strategic behavior, in Sections 3.2–3.4, we focused on the formula-based judgment aggregation framework and subset-based preferences. In this section, we briefly describe the decision problems that we consider for the other settings that we consider.

For the setting of manipulation in the presence of weighted and unweighted Hamming distance preferences, the decision problems MANIPULATION(Kemeny; W-HAM) and MANIPULATION(Kemeny; HAM) that we consider are similar to the problem MANIPULATION(Kemeny). However, the input for these problems does not contain a set $L \subseteq J_1$. For the problem MANIPULATION(Kemeny; W-HAM), the input contains a weight function $w : [\Phi] \rightarrow \mathbb{N}$ instead. The question is whether each set $J^{*}_{\text{new}} \in \text{Kemeny}(\boldsymbol{J}_{-1}, J')$ is preferred over each set $J^{*}_{\text{old}} \in \text{Kemeny}(\boldsymbol{J})$—according to the (weighted) Hamming distance preferences.

For the setting of bribery, we consider the problems BRIBERY(Kemeny; W-HAM), BRIBERY(Kemeny; HAM) that are defined analogously. In addition, for these problems, the input contains a reference judgment set $J_0 \in \mathcal{J}(\Phi, \Gamma)$ that underlies the (weighted) Hamming distance preferences for the external agent. For the setting of control, we consider similar decision problems—e.g., CONTROL-BY-ADDING-ISSUES(Kemeny; W-HAM).

For all of these decision problems, we also consider variants for the constraint-based judgment aggregation framework, where the input contains the corresponding elements in the constraint-based framework. For these variants we do not introduce additional notation.

4. COMPLEXITY RESULTS

We now turn to presenting the complexity results that we establish in this paper for the problems of manipulation, bribery and control. We present detailed proofs for the respective problems for the formula-based judgment aggregation framework and for subset-based preferences over judgment sets. We briefly indicate how these results can be extended to work also for other settings—i.e., for the constraint-based judgment aggregation framework and for preferences based on (weighted) Hamming distances. In addition, we indicate how these results can be extended to the setting where the judgment aggregation scenario is manipulated by a group of individuals, rather than a single individual. (The propositions for which we do not give detailed proofs we indicate with a $\star$.)

4.1 Manipulation

We begin with showing Σ^{p}_{2}-completeness for the problem of manipulation. The Σ^{p}_{2}-hardness result that we obtain works even for unidimensionally aligned profiles—that is, for cases where the input contains a unidimensionally aligned profile; the profile that arises from the act of manipulation does not need to be unidimensionally aligned.

LEMMA 1. MANIPULATION(Kemeny) *is in* Σ^{p}_{2}.

PROOF. We describe a nondeterministic polynomial-time algorithm with access to an NP oracle that solves the problem. Let $(\Phi, \Gamma, \boldsymbol{J}, L)$ specify an instance of MANIPULATION(Kemeny). The algorithm proceeds in several steps.

Firstly, (1) the algorithm guesses a complete judgment set J'_1 together with a truth assignment $\alpha : \text{Var}(\Phi, \Gamma) \rightarrow \mathbb{B}$, and it checks whether α satisfies both J'_1 and Γ. This can be done in nondeterministic polynomial time.

Then, (2) the algorithm determines the minimum Hamming distance d_{win} from $(J'_1, \boldsymbol{J}_{-1})$ to any judgment set $J^{*} \in \mathcal{J}(\Phi, \Gamma)$. That is, d_{win} is the Hamming distance from the judgment sets in $(J'_1, \boldsymbol{J}_{-1})$ to the judgment sets $J^{*} \in \text{Kemeny}(J'_1, \boldsymbol{J}_{-1})$. This can be done in (deterministic) polynomial time using $O(\log n + \log p)$ queries to an NP oracle.

Finally, (3) the algorithm determines by using a single query to an NP oracle whether there exists some complete and Γ-consistent judgment set $J^{*}_{\text{new}} \in \mathcal{J}(\Phi, \Gamma)$ such that $d(J^{*}_{\text{new}}, (\boldsymbol{J}_{-1}, J'_1)) = d_{\text{win}}$ and $L \not\subseteq J^{*}_{\text{new}}$. If this is the case, the algorithm rejects; otherwise, the algorithm accepts.

It is straightforward to verify that the algorithm runs in nondeterministic polynomial time. Moreover, the algorithm accepts the input (for some sequence of nondeterministic choices) if and only if there is some complete and consistent judgment set J'_1 such that for all $J^{*} \in \text{Kemeny}(\boldsymbol{J}_{-1}, J'_1)$ it holds that $L \subseteq J^{*}$. $\square$

THEOREM 2. MANIPULATION(Kemeny) *is* Σ^{p}_{2}*-complete. Moreover, hardness holds even when the input is restricted to unidimensionally aligned profiles (for* 3 *individuals) and where* $|L| = 1$.

PROOF. Membership in Σ^{p}_{2} is shown in Lemma 1. We show Σ^{p}_{2}-hardness by reducing from QSAT$_2$. Let $\varphi = \exists X. \forall Y. \psi$ be an instance of QSAT$_2$. We construct an instance of MANIPULATION(Kemeny) as follows. We introduce auxiliary variables x' for each $x \in X$, that is, we introduce the set $X' = \{ x' : x \in X \}$ of variables. We define Φ by letting $[\Phi] = X \cup X' \cup \{w, z\}$. Moreover, we define the integrity constraint Γ as follows:

$$\Gamma = \Big(\neg w \wedge \bigwedge_{x \in X} (x \wedge x')\Big) \vee \Big(\neg w \wedge \bigwedge_{x \in X} (\neg x \wedge \neg x')\Big) \vee \Big((\psi \rightarrow (w \wedge z)) \wedge (\bigwedge_{x \in X} (x \oplus x')\Big).$$

As a result, each complete judgment set J is Γ-consistent if and only if it satisfies one of the following three conditions: (1) J contains all variables $x \in X$ and all variables $x' \in X'$, and J does not contain w; (2) J contains no variables $x \in X$ and no variables $x' \in X'$, and J does not contain w; or (3) for each $x \in X$, J contains exactly one of x and x', and J satisfies the formula $(\psi \rightarrow (w \wedge z))$. (The variable w plays no role in the current reduction; we only include this variable in the construction so that we can use the same construction for the proof of Theorem 5 below.)

Next, we define the profile $\boldsymbol{J} = (J_1, J_2, J_3)$ as shown in Figure 1. The profile $\boldsymbol{J}$ is unidimensionally aligned—this is witnessed by the bijection $\pi : \{J_1, J_2, J_3\} \rightarrow [3]$ where $\pi(J_i) = i$. Finally, we let $L = \{z\}$.

$\boldsymbol{J}$	J_1	J_2	J_3
x	0	0	1
x'	0	0	1
w	0	0	0
z	1	0	0

Figure 1: The profile $\boldsymbol{J}$ in the proof of Theorem 2. Here x ranges over X, and x' ranges over X'.

We show that there is some complete and Γ-consistent judgment set J'_1 such that for all $J^* \in \text{Kemeny}(J'_1, \boldsymbol{J}_{-1})$ it holds that $L \subseteq J^*$ if and only if $\varphi \in \text{QSAT}_2$.

($\Rightarrow$) Suppose that there is a successful manipulation J'_1 such that for all $J^* \in \text{Kemeny}(J'_1, \boldsymbol{J}_{-1})$ it holds that $L \subseteq J^*$. Since both J_2 and J_3 contain $\neg z$, we know that each $J^* \in \text{Kemeny}(J'_1, \boldsymbol{J}_{-1})$ must satisfy condition (3). Otherwise, there would also be some $J^* \in \text{Kemeny}(J'_1, \boldsymbol{J}_{-1})$ with $L \not\subseteq J^*$. From this, we know that J'_1 must satisfy condition (3), because if J'_1 would satisfy condition (1) or (2), the majority outcome would be consistent, and would also satisfy condition (1) or (2). Thus, in this situation, it would not be possible that each $J^* \in \text{Kemeny}(J'_1, \boldsymbol{J}_{-1})$ satisfies condition (3).

Now construct the truth assignment $\alpha : X \rightarrow \{0, 1\}$ by letting $\alpha(x) = 1$ if and only if $x \in J'_1$. We show that $\forall Y.\psi[\alpha]$ is true. In order to do so, consider the majority outcome $\text{Majority}(J'_1, \boldsymbol{J}_{-1})$. We know that $\text{Majority}(J'_1, \boldsymbol{J}_{-1})$ agrees with J'_1 on all variables $w \in X \cup X'$. Moreover, we know that $\text{Majority}(J'_1, \boldsymbol{J}_{-1})$ contains $\neg z$. If $\text{Majority}(J'_1, \boldsymbol{J}_{-1})$ were Γ-consistent, it would be selected by the Kemeny rule, and thus it would not be the case that for each $J^* \in \text{Kemeny}(J'_1, \boldsymbol{J}_{-1})$ it holds that $L \subseteq J^*$. Thus, $\text{Majority}(J'_1, \boldsymbol{J}_{-1})$ is inconsistent with Γ. By construction of Γ, this can only be the case if $\text{Majority}(J'_1, \boldsymbol{J}_{-1}) \models \psi$. Then also $J'_1 \models \psi$, and thus $\forall Y.\psi[\alpha]$ is true.

($\Leftarrow$) Conversely, suppose that $\varphi \in \text{QSAT}_2$. That is, there is some truth assignment $\alpha : X \rightarrow \{0, 1\}$ such that $\forall Y.\psi[\alpha]$ is true. We construct the judgment set J'_1 as follows. For each $x \in X$ we let $x \in J'_1$ if and only if $\alpha(x) = 1$ (and we let $\neg x \in J'_1$ otherwise). For each $x' \in X'$ we let $x' \in J'_1$ if and only if $\alpha(x) = 0$ (and we let $\neg x' \in J'_1$ otherwise). Finally, we let $w \in J'_1$ and $z \in J'_1$.

We show that for each $J^* \in \text{Kemeny}(J'_1, \boldsymbol{J}_{-1})$ it holds that $L \subseteq J^*$, that is, that $z \in J^*$. We know that each $J^* \in \text{Kemeny}(J'_1, \boldsymbol{J}_{-1})$ satisfies one of the conditions (1), (2) or (3). It is straightforward to verify that condition (3) can be satisfied by differing on as few issues as possible with the profile $(J'_1, \boldsymbol{J}_{-1})$. Thus, each $J^* \in \text{Kemeny}(J'_1, \boldsymbol{J}_{-1})$ agrees with J'_1 on the issues $x \in X$ and $x' \in X'$. Then, since each J^* satisfies condition (3), we know that $J^* \models (\psi \rightarrow z)$. By construction of J'_1, we know that $J'_1 \models \psi$. Therefore, for each J^* it holds that $J^* \models \psi$ (because J^* and J'_1 agree on X), and thus that $z \in J^*$. □

4.1.1 *(Weighted) Hamming Distance Preferences*

The above Σ^p_2-completeness result can be extended to the case of (weighted) Hamming distance preferences.

PROPOSITION* 3. MANIPULATION(Kemeny; W-HAM) *and* MANIPULATION(Kemeny; HAM) *are* Σ^p_2*-complete.*

PROOF IDEA. Membership can be shown by straightforwardly extending the algorithm in Lemma 1 to work also for preferences based on weighted Hamming distances. Σ^p_2-hardness for MANIPULATION(Kemeny; W-HAM) follows directly from the proof of Theorem 2, since preferences based on a subset L of size 1 can be seen as a special case of weighted Hamming distance preferences. Σ^p_2-hardness can be shown for MANIPULATION(Kemeny; HAM) by a (tedious and lengthy) reduction from QSAT_2 that is based on the same principles as the hardness proof that we gave for Theorem 2. □

4.1.2 *Constraint-Based Judgment Aggregation*

The Σ^p_2-completeness results of Theorem 2 and Proposition 3 can also straightforwardly be extended to the setting of constraint-based judgment aggregation. The nondeterministic algorithms with access to an NP oracle, used to show membership in Σ^p_2, can be applied also in the constraint-based judgment aggregation framework. Moreover, one can modify the Σ^p_2-hardness proof that we gave for Theorem 2 in such a way that all variables occuring in the agenda Φ and in the integrity constraint Γ occur as separate formulas in the agenda. In this case, we can transform the agenda, the integrity constraint and the profile in polynomial time to the constraint-based framework [17]. Thus, this allows us to show Σ^p_2-hardness also for the constraint-based framework.

4.1.3 *Group Manipulation*

Another question that has been investigated in the judgment aggregation literature is in what cases a judgment aggregation scenario can be manipulated by a group of individuals, rather than by a single individual [9]. In such group manipulation situations, a group of individuals coordinates to express insincere judgments with the aim of obtaining an outcome that is preferred by each of the individuals over the outcome when all individuals report their sincere judgments. The Σ^p_2-completeness results that we described above all carry over straightforwardly to the setting of group manipulation—the membership results can be modified easily, and the hardness results carry over since an individual forms a group of size 1.

When considering group manipulation, the question arises whether no individual in the manipulating coalition can obtain a further improvement by (unilaterally) deviating from the manipulation strategy. The Σ^p_2-hardness results that we obtained also extend to the problem of stable group manipulation [9], where such unilateral deviations obstruct successful manipulation strategies.

4.2 Bribery

We continue with showing Σ^p_2-completeness for the problem of bribery. The Σ^p_2-hardness result that we obtain works even for unidimensionally aligned profiles—again, this does not mean that the manipulated profile needs to be unidimensionally aligned, only the profile given in the input.

LEMMA 4. BRIBERY(Kemeny) *is in* Σ_2^{p}.

PROOF. We describe a nondeterministic polynomial-time algorithm with access to an NP oracle that solves the problem. The algorithm that we describe is similar to the algorithm used in the proof of Lemma 1. Let $(\Phi, \Gamma, \boldsymbol{J}, L, k)$ specify an instance of BRIBERY(Kemeny), where $\boldsymbol{J} = (J_1, \dots, J_p)$. The algorithm proceeds in several steps.

Firstly, (1) the algorithm guesses k indices $\ell_1, \dots, \ell_k \in [p]$ and k complete judgment sets $J'_1, \dots, J'_k$ together with truth assignments $\alpha_1, \dots, \alpha_k : \mathrm{Var}(\Phi, \Gamma) \to \mathbb{B}$, and it checks whether α_j satisfies both J'_j and Γ, for each $j \in [k]$. This can be done in nondeterministic polynomial time. Let $\boldsymbol{J}'$ denote the profile obtained from $\boldsymbol{J}$ by replacing each J_{ℓ_j} by J'_j.

Then, (2) the algorithm determines the minimum Hamming distance d_{win} from $\boldsymbol{J}'$ to any complete and consistent judgment set $J^* \in \mathcal{J}(\Phi, \Gamma)$. That is, d_{win} is the Hamming distance from the judgments in $\boldsymbol{J}'$ to the judgment sets $J^* \in \text{Kemeny}(\boldsymbol{J}')$. This can be done in (deterministic) polynomial time using $O(\log n)$ queries to an NP oracle.

Finally, (3) the algorithm determines by using a single query to an NP oracle whether there exists some complete and Γ-consistent judgment set $J^*_{\text{new}} \in \mathcal{J}(\Phi, \Gamma)$ such that $d(J^*_{\text{new}}, \boldsymbol{J}') = d_{\text{win}}$ and $L \not\subseteq J^*_{\text{new}}$. If this is the case, the algorithm rejects; otherwise, the algorithm accepts.

It is straightforward to verify that the algorithm is correct and that it runs in nondeterministic polynomial time. $\square$

THEOREM 5. BRIBERY(Kemeny) *is* Σ_2^{p}*-complete. Moreover, hardness holds even when the input is restricted to unidimensionally aligned profiles and where* $|L| = 1$.

PROOF. We show Σ_2^{p}-hardness by reducing from QSAT_2. Let $\varphi = \exists X.\forall Y.\psi$ be an instance of QSAT_2. We construct an instance of BRIBERY(Kemeny) as follows. Our construction is very similar to the construction used in the proof of Theorem 2. In fact, we let Φ, Γ and $\boldsymbol{J} = (J_1, J_2, J_3)$ be exactly as defined in the proof of Theorem 2. In particular, this means that $\boldsymbol{J}$ is unidimensionally aligned. Moreover, we let $L = \{w\}$. Finally, we let $k = 1$—i.e., we give the bribing party the option of changing only a single judgment set.

We show that there is some complete and Γ-consistent judgment set J' and a profile $\boldsymbol{J}'$ obtained from $\boldsymbol{J}$ by replacing a single judgment set by J' such that for all $J^* \in \text{Kemeny}(\boldsymbol{J}')$ it holds that $L \subseteq J^*$ if and only if $\varphi \in \text{QSAT}_2$.

($\Rightarrow$) Suppose that there is some complete and Γ-consistent judgment set J' and a profile $\boldsymbol{J}'$ obtained from $\boldsymbol{J}$ by replacing a single judgment set by J' such that for all $J^* \in \text{Kemeny}(\boldsymbol{J}')$ it holds that $L \subseteq J^*$, i.e., $w \in J^*$. By construction of Γ, we know that if $w \in J^*$ then J^* satisfies condition (3). Therefore, the only way that the bribing party can enforce $w \in J^*$ is to change J_1 or J_2; otherwise, each $J^* \in \text{Kemeny}(\boldsymbol{J}')$ would satisfy condition (2). Suppose that J_1 is changed by the bribing party (the case for J_2 is entirely similar).

The only way to ensure that the sets $J^* \in \text{Kemeny}(\boldsymbol{J}')$ satisfy condition (3) is to change J_1 to J'_1 in such a way that for each $x \in X$ it contains exactly one of x and x'. Moreover, to ensure that for each $J^* \in \text{Kemeny}(\boldsymbol{J}')$ it holds that $w \in J^*$, by construction of Γ, the judgment set J'_1 needs to be chosen in such a way that $J'_1 \models \psi$. Now consider the truth assignment $\alpha : X \to \mathbb{B}$ defined by letting $\alpha(x) = 1$ if and only if $x \in J'_1$ for each $x \in X$. Since $J'_1 \models \psi$, we know that $\forall Y.\psi[\alpha]$ is true. Thus, $\varphi \in \text{QSAT}_2$.

($\Leftarrow$) Conversely, suppose that $\varphi \in \text{QSAT}_2$. That is, there is some truth assignment $\alpha : X \to \{0, 1\}$ such that $\forall Y.\psi[\alpha]$ is true. We construct a judgment set J'_1 as follows, and we show that for the profile $\boldsymbol{J}'$ obtained from $\boldsymbol{J}$ by replacing J_1 by J'_1 it holds that for each $J^* \in \text{Kemeny}(\boldsymbol{J}')$ we have that $w \in J^*$. For each $x \in X$, we let $x \in J'_1$ if and only if $\alpha(x) = 1$ (and we let $\neg x \in X$ otherwise). For each $x' \in X'$, we let $x' \in J'_1$ if and only if $\alpha(x) = 0$ (and we let $\neg x' \in X$ otherwise). Moreover, we let $\neg w \in J'_1$ and $\neg z \in J'_1$. Let $\boldsymbol{J}'$ denote the profile obtained from $\boldsymbol{J}$ by replacing J_1 by J'_1, i.e., $\boldsymbol{J}' = (J'_1, J_2, J_3)$. It is straightforward to verify that each $J^* \in \text{Kemeny}(\boldsymbol{J}')$ satisfies condition (3). Moreover, for each $J^* \in \text{Kemeny}(\boldsymbol{J}')$ it holds that J^* agrees with J'_1 on each $x \in X$ and each $x' \in X'$, and thus $J^* \models \psi$. Then, by construction of Γ we know that for each $J^* \in \text{Kemeny}(\boldsymbol{J}')$ it holds that $w \in J^*$. $\square$

4.2.1 *(Weighted) Hamming Distance Preferences & Constraint-Based Judgment Aggregation*

Similarly to the case for manipulation, the above Σ_2^{p}-completeness result can be extended to the case of (weighted) Hamming distance preferences.

PROPOSITION* 6. BRIBERY(Kemeny; W-HAM) *and* BRIBERY(Kemeny; HAM) *are* Σ_2^{p}*-complete.*

PROOF IDEA. Membership can be shown by straightforwardly extending the algorithm in Lemma 4. Σ_2^{p}-hardness for BRIBERY(Kemeny; W-HAM) follows directly from the proof of Theorem 5. Σ_2^{p}-hardness can be shown for BRIBERY-(Kemeny; HAM) by a (tedious and lengthy) reduction from QSAT_2 that is based on the same principles as the hardness proof that we gave for Theorem 5. $\square$

The Σ_2^{p}-completeness results of Theorem 5 and Proposition 6 can also straightforwardly be extended to the setting of constraint-based judgment aggregation. Similarly to the case of manipulation, the algorithm we gave to show membership can be applied also in the constraint-based framework, and the hardness proofs can be modified to work also for the constraint-based judgment aggregation framework.

4.3 Control

Finally, we show Σ_2^{p}-completeness for the two control problems that we consider.

LEMMA 7. CONTROL-BY-ADDING-ISSUES(Kemeny) *and* CONTROL-BY-REMOVING-ISSUES(Kemeny) *are in* Σ_2^{p}.

PROOF. We describe a nondeterministic polynomial-time algorithm with access to an NP oracle that solves the problem. Let $(\Phi, \Gamma, \Phi', L, \boldsymbol{J})$ specify an instance of CONTROL-BY-ADDING-ISSUES(Kemeny). The algorithm guesses an agenda Φ'' such that $\Phi' \subseteq \Phi'' \subseteq \Phi$. Then, the algorithm computes the minimum Hamming distance d_{win} from the profile $\boldsymbol{J}|_{\Phi''}$ to any judgment set that is Γ-consistent and complete for Φ''. This can be done in polynomial time using an NP oracle. Finally, the algorithm uses one more query to the NP oracle to decide if there exists a judgment set J^* that is Γ-consistent and complete for Φ'' that has Hamming distance d_{win} to the profile $\boldsymbol{J}|_{\Phi''}$ and that satisfies that $L \not\subseteq J^*$. The algorithm accepts if and only if no such judgment set J^* exists. It is straightforward to verify

Technical Report AC-TR-17-005

that the algorithm runs in nondeterministic polynomial time. Moreover, the algorithm accepts the input (for some sequence of nondeterministic choices) if and only if there exists an agenda $\Phi' \subseteq \Phi'' \subseteq \Phi$ such that for all $J^* \in \text{Kemeny}(\boldsymbol{J}|_{\Phi''})$ it holds that $L \subseteq J^*$.

The algorithm described above can straightforwardly be modified to work also for the case of CONTROL-BY-REMOVING-ISSUES(Kemeny). $\square$

THEOREM 8. CONTROL-BY-ADDING-ISSUES(Kemeny) *and* CONTROL-BY-REMOVING-ISSUES(Kemeny) *are* Σ_2^p*-complete.*

PROOF. Membership in Σ_2^p is shown in Lemma 7. We show Σ_2^p-hardness by reducing from QSAT_2. Let $\varphi = \exists X.\forall Y.\psi$ be an instance of QSAT_2. Without loss of generality, we may assume that there is a truth assignment $\alpha : X \to \mathbb{B}$ such that $\forall Y.\psi[\alpha]$ is not true, i.e., that $\exists X.\exists Y.\neg\psi$ is true. We construct an instance of CONTROL-BY-ADDING-ISSUES-(Kemeny) as follows. We define Φ' by letting $[\Phi'] = \{ w_{i,j} : i \in [2], j \in [3] \} \cup \{z\}$, and we define Φ by letting $[\Phi] = [\Phi'] \cup \{ x, x' : x \in X \}$. Then, we define the integrity constraint Γ as follows:

$$\Gamma = \left(\bigwedge_{x \in X} (x \wedge x') \wedge \bigwedge_{i\in[2]} \bigvee_{j\in[3]} w_{i,j} \right) \vee \left(\bigwedge_{x\in X} (x \oplus x') \wedge (\psi \to z) \right).$$

As a consequence, a complete judgment set J is Γ-consistent if and only if J is compatible with one of the following two conditions: (1) the variables x and x' are true for each $x \in X$, and at least one $w_{i,j}$ is true for each $i \in [2]$; or (2) for each $x \in X$, exactly one of x and x' are true, and the formula $(\psi \to z)$ is true. Then, we let $L = \{z\}$, and we define the profile $\boldsymbol{J} = (J_1, J_2, J_3)$ as shown in Figure 2.

We show that there is an agenda Φ'' with $\Phi' \subseteq \Phi'' \subseteq \Phi$ such that for all $J^* \in \text{Kemeny}(\boldsymbol{J}|_{\Phi''})$ it holds that $L \subseteq J^*$ if and only if $\varphi \in \text{QSAT}_2$.

($\Rightarrow$) Suppose that there is an agenda Φ'' with $\Phi' \subseteq \Phi'' \subseteq \Phi$ such that for all $J^* \in \text{Kemeny}(\boldsymbol{J}|_{\Phi''})$ it holds that $L \subseteq J^*$. Then these $J^* \in \text{Kemeny}(\boldsymbol{J}|_{\Phi''})$ do not satisfy condition (1); if this were the case, we could get a judgment set that is closer to the profile by negating z in J^*, which contradicts our assumption that $L \subseteq J^*$ for all $J^* \in \text{Kemeny}(\boldsymbol{J}|_{\Phi''})$. Thus, each $J^* \in \text{Kemeny}(\boldsymbol{J}|_{\Phi''})$ is compatible with condition (2).

We show that Φ'' must contain either x or x' for at least one $x \in X$. To derive a contradiction, suppose that Φ'' contains neither x nor x' for any $x \in X$. Then condition (2) can be satisfied by the judgment sets $J^* \in \text{Kemeny}(\boldsymbol{J}|_{\Phi''})$ without including z (since $\exists X.\exists Y.\neg\psi$ is true), which contradicts our assumption that each J^* contains z.

$\boldsymbol{J}$	J_1	J_2	J_3
x	1	1	1
x'	1	1	1
$w_{i,1}$	1	0	0
$w_{i,2}$	0	1	0
$w_{i,3}$	0	0	1
z	0	0	1

Figure 2: The profile $\boldsymbol{J}$ in the proof of Theorem 8. Here x ranges over X, x' ranges over $\{ x' : x \in X \}$, and i ranges over $[2]$.

Next, we show that Φ'' cannot contain both x and x' for any $x \in X$. To derive a contradiction, suppose that $x, x' \in \Phi''$ for some $x \in X$. Then the Hamming distance from $\boldsymbol{J}|_{\Phi''}$ to any complete and Γ-consistent judgment set compatible with condition (1) is smaller than to any complete and Γ-consistent judgment sets that is compatible with condition (2)—even if Φ'' contains $w_{i,j}$ for each $i \in [2]$ and each $j \in [3]$. This would mean that $z \notin J^*$ for each $J^* \in \text{Kemeny}(\boldsymbol{J}|_{\Phi''})$, which contradicts our assumption.

Now define the (partial) truth assignment $\alpha : X \to \mathbb{B}$ by letting $\alpha(x) = 1$ if $x \in \Phi''$, letting $\alpha(x) = 0$ if $x' \in \Phi''$, and letting $\alpha(x)$ be undefined otherwise. We claim that there for every complete truth assignment $\alpha' : X \to \mathbb{B}$ that extends α it holds that $\forall Y.\psi[\alpha']$ is true. To derive a contradiction, suppose that this is not the case. Then each $J^* \in \text{Kemeny}(\boldsymbol{J}|_{\Phi''})$ can satisfy Γ without including z, which contradicts our assumption. Therefore, $\varphi \in \text{QSAT}_2$.

($\Leftarrow$) Conversely, suppose that $\varphi \in \text{QSAT}_2$. That is, there is some truth assignment $\alpha : X \to \mathbb{B}$ such that $\forall Y.\psi[\alpha]$ is true. We construct an agenda Φ'' with $\Phi' \subseteq \Phi'' \subseteq \Phi$ as follows. We let $[\Phi''] = [\Phi'] \cup \{ x : x \in X, \alpha(x) = 1 \} \cup \{ x' : x \in X, \alpha(x) = 0 \}$. It is straightforward to verify that for each $J^* \in \text{Kemeny}(\boldsymbol{J}|_{\Phi''})$ it holds that $z \in J^*$, because each such J^* is compatible (only) with condition (1) and agrees with α, and thus $J^* \models \psi$.

The above reduction can also be used as a reduction from QSAT_2 to CONTROL-BY-REMOVING-ISSUES(Kemeny). $\square$

4.3.1 *Constraint-Based Judgment Aggregation*

Similarly to the cases for manipulation and bribery, the Σ_2^p-completeness result of Theorem 8 can be extended to the constraint-based judgment aggregation framework. Just as in the other cases, the algorithm we gave to show membership can be applied also in the constraint-based framework, and the hardness proofs can be modified to work also for the constraint-based judgment aggregation framework.

5. CONCLUSIONS

In this paper, we investigated the computational complexity of several problems related to several types of strategic behavior in judgment aggregation—namely, manipulation, bribery and control—for the Kemeny judgment aggregation procedure. We showed that deciding whether a successful strategic behavior policy exists is Σ_2^p-complete. These results hold for all types of strategic behavior that we consider, for two formal judgment aggregation frameworks that are commonly considered in the literature, and for several types of preference relations over judgment sets.

These intractability results can be interpreted as a computational barrier against (undesirable) strategic behavior for the Kemeny procedure in judgment aggregation. However, as such worst-case complexity results can be overly negative, it is an important topic for future research to further investigate the computational complexity of these problems for restricted fragments and using more sensitive methods, such as parameterized complexity. Another direction for future research is to analyze the complexity of the different types of strategic behavior for other judgment aggregation procedures.

Acknowledgments.

De Haan is supported by the Austrian Science Fund (FWF), project P26200.

REFERENCES

[1] S. Arora and B. Barak. *Computational Complexity – A Modern Approach*. Cambridge University Press, 2009.

[2] J. Bartholdi, C. A. Tovey, and M. A. Trick. Voting schemes for which it can be difficult to tell who won the election. *Social Choice and Welfare*, 6(2):157–165, 1989.

[3] D. Baumeister, G. Erdélyi, O. J. Erdélyi, and J. Rothe. Bribery and control in judgment aggregation. In *Proceedings of the 4th International Workshop on Computational Social Choice (COMSOC 2012)*, 2012.

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

[5] D. Baumeister, G. Erdélyi, O. J. Erdélyi, and J. Rothe. Computational aspects of manipulation and control in judgment aggregation. In *Proceedings of the 3rd International Conference on Algorithmic Decision Theory (ADT 2013)*, pages 71–85. Springer Verlag, 2013.

[6] D. Baumeister, G. Erdélyi, O. J. Erdélyi, and J. Rothe. Complexity of manipulation and bribery in judgment aggregation for uniform premise-based quota rules. *Mathematical Social Sciences*, 76:19–30, 2015.

[7] 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 Second International Conference on Algorithmic Decision Theory (ADT 2011), Piscataway, NJ, USA, October 26–28, 2011*, volume 6992 of *Lecture Notes in Computer Science*, pages 1–15. Springer Verlag, 2011.

[8] D. Baumeister, G. Erdélyi, and J. Rothe. Judgment aggregation. In J. Rothe, editor, *Economics and Computation*. Springer Verlag, 2016.

[9] S. Botan, A. Novaro, and U. Endriss. Group manipulation in judgment aggregation. In *Proceedings of the 15th International Conference on Autonomous Agents and Multiagent Systems (AAMAS-2016)*. IFAAMAS, May 2016.

[10] V. Conitzer and T. Walsh. Barriers to manipulation in voting. In F. Brandt, V. Conitzer, U. Endriss, J. Lang, and A. Procaccia, editors, *Handbook of Computational Social Choice*. Cambridge University Press, Cambridge, 2015.

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

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

[13] F. Dietrich and C. List. Judgment aggregation without full rationality. *Social Choice and Welfare*, 31(1):15–39, 2008.

[14] E. Dokow and D. Falik. Models of manipulation on aggregation of binary evaluations. In *Proceedings of the 4th International Workshop on Computational Social Choice (COMSOC-2012)*. Kraków, Poland, September 2012. Full version available as [15].

[15] E. Dokow and D. Falik. Models of manipulation on aggregation of binary evaluations. *CoRR*, abs/1201.6388, 2012.

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

[17] U. Endriss, U. Grandi, R. de Haan, and J. Lang. Succinctness of languages for judgment aggregation. In *Proceedings of the Fifteenth International Conference on the Principles of Knowledge Representation and Reasoning (KR 2016), Cape Town, South Africa, April 25-29, 2016*. AAAI Press, 2016.

[18] U. Endriss, U. Grandi, and D. Porello. Complexity of judgment aggregation. *J. Artif. Intell. Res.*, 45:481–514, 2012.

[19] U. Endriss and R. de Haan. Complexity of the winner determination problem in judgment aggregation: Kemeny, Slater, Tideman, Young. In *Proceedings of AAMAS 2015, the 14th International Conference on Autonomous Agents and Multiagent Systems*. IFAAMAS/ACM, 2015.

[20] A. Gibbard. Manipulation of voting schemes: a general result. *Econometrica: journal of the Econometric Society*, pages 587–601, 1973.

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

[22] R. de Haan. Parameterized complexity results for the Kemeny rule in judgment aggregation. In *Proceedings of the 22nd European Conference on Artificial Intelligence (ECAI 2016)*, pages 1502–1510, 2016.

[23] J. Lang and M. Slavkovik. How hard is it to compute majority-preserving judgment aggregation rules? In *21st European Conference on Artificial Intelligence (ECAI 2014)*. IOS Press, 2014.

[24] C. List. A possibility theorem on aggregation over multiple interconnected propositions. *Mathematical Social Sciences*, 45(1):1–13, 2003.

[25] A. R. Meyer and L. J. Stockmeyer. The equivalence problem for regular expressions with squaring requires exponential space. In *SWAT*, pages 125–129. IEEE Computer Soc., 1972.

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

[27] K. Nehring, M. Pivato, and C. Puppe. Condorcet admissibility: indeterminacy and path-dependence under majority voting on interconnected decisions. MPRA, 2011. http://mpra.ub.uni-muenchen.de/32434/.

[28] C. H. Papadimitriou. *Computational Complexity*. Addison-Wesley, 1994.

[29] J. Rothe. *Economics and Computation. An Introduction to Algorithmic Game Theory, Computational Social Choice, and Fair Division*. Springer, 2016.

[30] M. A. 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.

[31] L. J. Stockmeyer. The polynomial-time hierarchy. *Theoretical Computer Science*, 3(1):1–22, 1976.

[32] C. Wrathall. Complete sets and the polynomial-time hierarchy. *Theoretical Computer Science*, 3(1):23–33, 1976.