Complexity of Sequential Rules in Judgment Aggregation[1]

Dorothea Baumeister, Linus Boes and Robin Weishaupt

Abstract

The task in judgment aggregation is to find a collective judgment set based on the views of individual judges about a given set of propositional formulas. One way of guaranteeing consistent outcomes is the use of sequential rules. In each round, the decision on a single formula is made either because the outcome is entailed by the already obtained judgment set, or, if this is not the case, by some underlying rule, e.g. the majority rule. Such rules are especially useful for cases, where the agenda is not fixed in advance, and formulas are added one by one. This paper investigates the computational complexity of winner determination under a family of sequential rules, and the manipulative influence of the processing order on the final outcome.

1 Introduction

Judgment Aggregation (JA) is the task of aggregating individual judgments over logical formulas into a collective judgment set. The doctrinal paradox by Kornhauser and Sager [16] shows that if the majority rule is used, the outcome may be inconsistent, even if all underlying individual judgment sets are consistent. Since then research related to JA has been undertaken in different disciplines. The book chapter by Endriss [9] provides an overview of recent research on JA in computational social choice, where for example computer science methods are used to analyze problems originating from social choice. The investigation of JA from a computational complexity point of view has been initiated by Endriss et al. [13]. They focused on the winner problem, manipulation, and safety of the agenda problems. Subsequently, e.g. Baumeister et al. [3], Endriss and de Haan [11], and de Haan and Slavkovik [7] studied the complexity of different JA problems.

An important task is to generate consistent collective outcomes, that can, for example, be obtained through the use of sequential rules, see List [20]. A sequential rule works in rounds and uses some underlying JA rule, for example the majority rule as proposed by Dietrich and List [8] (see also Peleg and Zamir [23]). In each round the decision on one specific formula is made by checking whether either the formulas already contained in the collective outcome logically entail an assignment for the formula at hand, or otherwise, the outcome of the underlying rule for this formula will be taken. This is reasonable, since sequential rules occur naturally by incremental decision-making. Since many real-world decisions (e.g., contract agreements) are binding, while reversing may be either favorable but expensive or impracticable, reasoning happens gradually. List [20] discusses similar use cases of such path-dependent rules in detail. We focus on sequential rules that rely on underlying quota rules, where a formula is included in the collective outcome if a certain fraction of the judges approves it. This includes the two extreme cases where a single approval is sufficient or where an approval of all judges is needed or the common case of a majority of 2/3. Such a majority is needed for Senate votes on a presidential Impeachment, for the College of Cardinals in the papal conclave, or in some cases for constitutional amendments. Political referenda are examples of more diverse quotas.

[1] Also puplished in *Proceedings of the 20th Conference on Autonomous Agents and Multiagent Systems*, 2021 [2].

Since JA may also be used in security applications, as mentioned by Jamroga and Slavkovik [15], it is particularly important to have consistent collective judgment sets that are efficiently computable. The complexity of winner determination for different JA rules has been studied by Endriss et al. [13] for the premise-based procedure and the distance-based procedure and by de Haan and Slavkovik [7] for scoring and distance based rules. Along with many other rules, both, Endriss and de Haan [11] and Lang and Slavkovik [19], studied winner determination for the ranked agenda rule[2] and the maxcard subagenda rule[3], which are closely related to some of our results. In this paper we investigate the computational complexity of several problems related to winner determination for sequential JA rules that use a specific quota rule as the underlying rule. Furthermore, we study the problem of manipulative design, i.e., the question whether there is an order in which the formulas should be processed that yields some desired outcome. Additionally, we study majority-preservation for sequential JA rules, see Lang and Slavkovik [19]. The idea for sequential rules is to maintain a maximal agreement with the outcome of the majority rule (or any other underlying rule), when applied sequentially. In this context we identify a correlation between majority-preservation of sequential rules and non-sequential rules—in particular the maximum subagenda rule[4] (see Definition 5) and the maxcard subagenda rule (see Footnote 6). Our results range from membership in P to completeness in the second level of the polynomial hierarchy.

Compared to previous work on the ranked agenda rule (sequential majority rule, where the processing order is based on the majority support), see Endriss and de Haan [11] and Lang and Slavkovik [19], our results generalize and supplement respective complexity results, since lower bounds hold for any quota and even for a constant number of judges, implying para-NP-hardness. Similar results for the ranked agenda rule were recently published independently by Endriss et al. [12]. Additionally, we established matching upper bounds for all sequential rules that rely on a complete and complement-free rule.

2 Preliminaries

The technical framework mainly follows the definitions in Endriss [9]. In JA we talk about a group $[r]$ of $r \in \mathbb{N}$ **judges**, where $[r]$ denotes the set $\{1, \ldots, r\}$. The judges judge over an **agenda** Φ, which consists of boolean formulas in standard propositional logic. In order to avoid double negations let $\sim\varphi$ denote the complement of φ, i.e., $\sim\varphi = \neg\varphi$ if φ is not negated, and $\sim\varphi = \psi$ if $\varphi = \neg\psi$. Thereby, we assume Φ to be finite, nonempty and closed under complement, i.e., for every $\varphi \in \Phi$ it holds that $\sim\varphi \in \Phi$. Furthermore, we assume Φ to be **nontrivial**, i.e., there exist at least two formulas $\{\varphi, \psi\} \subseteq \Phi$, such that $\{\varphi, \psi\}, \{\sim\varphi, \psi\}, \{\varphi, \sim\psi\}$ and $\{\sim\varphi, \sim\psi\}$ are consistent, and we foreclose tautologies and contradictions from Φ. We split the agenda Φ into two disjoint subsets Φ_+ and Φ_-, where for all $\varphi \in \Phi_+$ it holds that $\sim\varphi \in \Phi_-$. Having the agenda introduced, we define an **individual judgment** $J \subseteq \Phi$ as a subset of Φ. We say that J is **complete**, if it holds for all $\varphi \in \Phi$ that $\varphi \in J$ or $\sim\varphi \in J$ is true. We say that J is **complement-free**, if it holds for all $\varphi \in \Phi$ that $|\{\varphi, \sim\varphi\} \cap J| \leq 1$. Lastly, we define J to be **consistent**, if there exists a boolean assignment for the formulas in J, such that all formulas are satisfied at the same time. We denote the set of all complete and consistent judgments over Φ by $\mathcal{J}(\Phi)$. For the set of judges $[r]$ we denote their **profile** of individual judgments over Φ as

[2] Also known in JA as Tideman's ranked pairs (see Endriss and de Haan [11]) and in similar variations as support-based procedure (see Porello and Endriss [24]) or leximax rule (see Lang et al. [18]).

[3] Also known in JA as Slater rule (see Endriss and de Haan [11]), max-num rule (see Endriss [10]) or endpoint rule (for the hamming distance as metric, see Miller et al. [21]).

[4] Also known in JA as maximal Condorcet rule (see Lang et al. [18]), while the outcome is also denoted as Condorcet admissible set (see Nehring et al. [22]).

$P = (P_1, \ldots, P_r) \in \mathcal{J}(\Phi)^r$. We define a (resolute) **judgment aggregation rule** for an agenda Φ and r judges, as a function $R\colon \mathcal{J}(\Phi)^r \to 2^{\Phi}$, mapping a profile $P \in \mathcal{J}(\Phi)^r$ of individual judgments to a subset $R(P)$ of Φ. We say that R is **complete/complement-free/consistent**, if for every profile $P \in \mathcal{J}(\Phi)^r$ it holds that $R(P)$ is complete/complement-free/consistent. Furthermore, we say that R is **anonymous** if it is independent of the order of judges, i.e., $R(P) = R(P_{\pi(1)}, \ldots, P_{\pi(r)})$ for all $P \in \mathcal{J}(\Phi)^r$permutation $\pi\colon [r] \to [r]$. Now, we define a family of JA rules. Within the subsequent definition we define a special case of the quota rules as defined by Dietrich and List [8].

Definition 1 (Quota Rules)**.** *Let $\Phi = \Phi_+ \cup \Phi_-, \Phi_+ \cap \Phi_- = \emptyset$ be an agenda, $P \in \mathcal{J}(\Phi)^r$ a profile of individual judgments and $q \in [0, 1]$. We define a **quota rule with quota** q as a JA rule F_q satisfying*

1. $\forall \varphi \in \Phi_+ \colon \varphi \in F_q(P) \Leftrightarrow |\{i \in [r] \mid \varphi \in P_i\}| \geq \lceil q(r + 1)\rceil$ *and*
2. $\forall \varphi \in \Phi_- \colon \varphi \in F_q(P) \Leftrightarrow |\{i \in [r] \mid \varphi \in P_i\}| \geq \lfloor (1 - q)(r + 1)\rfloor$.

Since $\lceil q(r + 1)\rceil + \lfloor (1 - q)(r + 1)\rfloor = r + 1$ holds for all $0 \leq q \leq 1$, it follows by the results from Dietrich and List [8] that all quota rules as previously defined are complete and complement-free. $\mathcal{F}$ denotes the set of all quota rules.

For an odd number of judges the majority rule equals the quota rule with quota $q = 1/2$. The difference for an even number of judges is that in case of a tie for some formula φ the quota rule executes some tie-breaking mechanism by choosing the corresponding formula from Φ_-, whereas the majority rule neglects completeness and does neither include this formula nor its negation.

We study **sequential judgment aggregation rules** in this paper. The basic idea is to ensure consistency by checking in each round whether the formulas contained in the collective outcome already fix the value for the formula at hand. This is formally denoted by the entailment relation, where $a \models b$ means that the value for b is determined by a. To begin, we define the subsequently studied sequential JA rules in a general way.

Definition 2 (Sequential $\mathcal{K}$-Judgment Aggregation Rule)**.** *Let $\mathcal{K}$ be a complete and complement-free JA rule. Furthermore, let Φ be an agenda, $P \in \mathcal{J}(\Phi)^r$ a profile and $\pi = (\varphi_1, \ldots, \varphi_m)$ an order over Φ_+. In order to obtain the aggregated judgment $S\mathcal{K}(P, \pi)$ of the **sequential $\mathcal{K}$-judgment aggregation rule**, we proceed as follows for $1 \leq i \leq m$:*

1. *If either $(\varphi_1^* \wedge \ldots \wedge \varphi_{i-1}^*) \models \varphi_i$ or $(\varphi_1^* \wedge \ldots \wedge \varphi_{i-1}^*) \models {\sim}\varphi_i$ holds, where $\varphi_j^* \in \{\varphi_j, {\sim}\varphi_j\}$ is the formula added in the j-th iteration to $S\mathcal{K}(P, \pi)$, we add φ_i or ${\sim}\varphi_i$ respectively to $S\mathcal{K}(P, \pi)$,*
2. *otherwise, we add $\{\varphi_i, {\sim}\varphi_i\} \cap \mathcal{K}(P)$ to $S\mathcal{K}(P, \pi)$.*

After m iterations we obtain the final aggregated judgment $S\mathcal{K}(P, \pi)$.

As an example consider an agenda Φ with $\Phi_+ = \{a, b, a \wedge b\}$ and three judges with $J_1 = \{\neg a, b, \neg(a \wedge b)\}$, $J_2 = \{a, \neg b, \neg(a \wedge b)\}$, and $J_3 = \{a, b, a \wedge b\}$. The majority rule returns the inconsistent judgment set $\{a, b, \neg(a \wedge b)\}$. Now, consider the sequential majority rule with order $\pi = (a, a \wedge b, b)$. In the first two steps a and $\neg(a \wedge b)$ are added to the outcome by majority, then the decision for b is entailed by the formulas already considered and $\neg b$ is included.

Observe that by our definition (i) any output $S\mathcal{K}(P, \pi)$ is complete and consistent with respect to the agenda Φ and (ii) if $\mathcal{K}$ is anonymous then $S\mathcal{K}$ is anonymous, too. Combining (i) and (ii) with List's impossibility result [20], we obtain for underlying anonymous rules $\mathcal{K}$ that the resulting judgment of a sequential JA rule $S\mathcal{K}$ depends on the processing order over Φ_+. Therefore, all previously defined (anonymous) sequential JA rules are path-dependent.

Whenever we address a sequential JA rule with respect to some JA rule $\mathcal{K}$, we assume $\mathcal{K}$ to be complement-free and complete. Subsequently, we introduce one more notation to exactly express partially aggregated judgments in order to simplify notation.

Definition 3 (Partially Aggregated Judgment)**.** *Let Φ be an agenda, $P \in \mathcal{J}(\Phi)^r$ a profile for r judges, π an order over Φ_+ and $\psi \in \Phi$. We define the* ***partially aggregated judgment*** *$S\mathcal{K}^{\psi}(P, \pi) \subset S\mathcal{K}(P, \pi)$ as the subset of the final aggregated judgment, for which the order π was processed until, but excluding ψ or $\sim\psi$ respectively.*

Observe that for every $\psi \in \Phi$ either ψ itself or $\sim\psi$ appears in π, ensuring that the previous definition is well-defined. In the following, we will focus on sequential JA rules based on quota rules. For the remaining parts of the paper, we assume that the reader is familiar with the basics of computational complexity such as the classes P, NP, the polynomial hierarchy as well as polynomial-time many-one reductions $\leq_{\mathrm{m}}^{\mathrm{p}}$. SAT denotes the satisfiability problem and $\overline{\mathrm{SAT}}$ its complement. We consider the class $\Delta_2^p = \mathrm{P}^{\mathrm{NP}}$, containing problems, which can be solved by a deterministic turing machine in polynimial time, that may also query an NP-oracle, which returns the answer to an NP-complete problem in one computational step. Similarly, for problems in $\Theta_2^p = \mathrm{P}^{\mathrm{NP[log]}}$ the NP-oracle may only be queried a logarithmic number of times. Finally, we also consider $\Sigma_2^p = \mathrm{NP}^{\mathrm{NP}}$ and $\Pi_2^p = \mathrm{coNP}^{\mathrm{NP}}$. For further reading, we refer to the textbook by Arora and Barak [1].

3 The Winner Problem

The use of JA rules in artificial intelligence technologies raises important computational questions. As the number of judges and/or the number of formulas in the agenda may be high, it is important to design fast algorithms to determine the collective outcome. The computational study of the winner problem for JA was initiated by Endriss et al. [13]. They showed that it is polynomial-time solvable for quota rules and the premise-based procedure, while it is Θ_2^p-complete for the distance-based procedure. Endriss and de Haan [11] showed that the winner problem is Θ_2^p-complete for some JA rules related to known voting rules (e.g., the maxcard subagenda rule), Δ_2^P-complete for the ranked agenda rule with a fixed tie-breaking and Σ_2^p-complete without a fixed tie-breaking. Lang and Slavkovik [19] defined a slightly different problem for winner determination and obtained completeness results in Θ_2^p (e.g., for the maxcard subagenda rule) and Π_2^p (e.g., for the ranked agenda rule without tie-breaking) for majority-preserving rules. We will emphasize relationships to the former results at relevant passages. The formal definition of the winner problem for a sequential JA rule $S\mathcal{K}$ is as follows.

$S\mathcal{K}$-Winner ($S\mathcal{K}$W)	
Instance:	An agenda Φ, a profile $P \in \mathcal{J}(\Phi)^r$, an order π over Φ_+, and a formula $\varphi \in \Phi$.
Question:	Is $\varphi \in S\mathcal{K}(P, \pi)$ true?

We analyze the computational complexity of this problem, starting with its upper bound.

Theorem 1. *$S\mathcal{K}$-Winner is in Δ_2^p if $\mathcal{K}$ is efficiently computable.*

Proof. Let $\mathcal{I} = (\Phi, P, \pi, \varphi)$ be a $S\mathcal{K}$W instance and denote the order by $\pi = (\varphi_1, \ldots, \varphi_m)$. Without loss of generality we may assume $\varphi = \varphi_j$ for one $j \in \{1, \ldots, m\}$, because if $\varphi = \sim\varphi_k$ for some $k \in \{1, \ldots, m\}$, we simply solve the instance $\mathcal{I}' = (\Phi, P, \pi, \sim\varphi)$ and invert its result.

First, we compute $\mathcal{K}(P) = \{\varphi'_1, \ldots, \varphi'_m\}$ in polynomial time. Now, for φ_1 we will use the result of $\mathcal{K}$ based on P to decide whether to add φ_1 or $\sim\varphi_1$ to $S\mathcal{K}(P, \pi)$. Furthermore, denote by $\varphi_1^*, \ldots, \varphi_{i-1}^*$ the elements added to $S\mathcal{K}^{\varphi_i}(P, \pi)$ in the first $i-1$ iterations. Note, that we add any φ_i' approved by $\mathcal{K}$, if and only if we cannot deduce $\sim\varphi_i'$ from the partially aggregated judgment. Consequently, in the i-th iteration, we ask whether $(\varphi_1^* \wedge \ldots \wedge \varphi_{i-1}^*) \models \sim\varphi_i'$ holds, which is equivalent to asking whether there is no satisfying assignment for $(\varphi_1^* \wedge \ldots \wedge \varphi_{i-1}^*) \wedge \varphi_i'$, which can be verified in coNP. Consequently, asking an NP-oracle whether this formula is satisfiable implies that $\sim\varphi_i'$ is not entailed by previously added formulas. In this case, we may add $\varphi_i' \in \mathcal{K}(P)$ directly to $S\mathcal{K}(P, \pi)$, since it is irrelevant for our purpose whether φ_i' is deduced or added by application of $\mathcal{K}$. Therefore, we require one NP-query per iteration, except for $i = 1$. In the worst case, we have $j = m$ and must pose $m - 1$ consecutive NP-queries over m iterations during our computation. Note that $m - 1$ is in $\mathcal{O}(|\mathcal{I}|)$ and thus, we can solve $\mathcal{I}$ in Δ_2^p. Thereby, it follows that $S\mathcal{K}$W$\in \Delta_2^p$ holds. □

In the construction above all queries rely on previous iterations and therefore, cannot be parallelized. Hence, Θ_2^p membership does not follow, which is in line with the general assumption of $\Theta_2^P \subset \Delta_2^p$. Now, having shown an upper bound for the computational complexity of the general winner problem, we like to introduce a lower bound for the computational complexity of the winner problem with respect to quota rules from $\mathcal{F}$. In order to do so, we first introduce the Δ_2^p-complete problem Odd Max Satisfiability, as defined by Krentel [17] (see also Große et al. [14]).

Odd Max Satisfiability (OMS)	
Instance:	A set $X = \{x_1, \ldots, x_n\}$ of boolean variables and a boolean formula $\alpha(x_1, \ldots, x_n)$.
Question:	Is α satisfiable and $x_n = 1$ in α's lexicographically maximum satisfying assignment $x_1 \ldots x_n \in \{0,1\}^n$?

Theorem 2. *Let $F_q \in \mathcal{F}$. Then, SF_q-*Winner *is Δ_2^p-complete.*

Proof. From the previous theorem we know that $SF_q\mathrm{W} \in \Delta_2^p$ holds, since F_q is efficiently computable, complement-free and complete. Therefore, it is sufficient to show OMS $\leq_m^p SF_q$-Winner.

Let $\mathcal{I} = (X, \alpha)$ be an OMS instance with $X = \{x_1, \ldots, x_n\}$. We construct in time polynomial in $|\mathcal{I}|$ a SF_qW instance $\mathcal{I}' = (\Phi, P, \pi, \varphi)$ as follows. Thereby, we separate the construction into two cases depending on the value of F_q's quota q. Due to space constraints, we only present the proof for $q \leq 1/3$, the remaining case can be shown by a similar approach.

Assume $q \leq 1/3$. We introduce β_1, β_2, and γ as new variables and define $\Phi_+ = \{\beta_1, \beta_2, \alpha', \alpha' \wedge x_1, \ldots, \alpha' \wedge x_n\}$ with $\alpha' = (\alpha \wedge \gamma) \vee \neg\beta_1 \vee \neg\beta_2$. Furthermore, we define the order π over Φ_+ as $\pi = (\beta_1, \beta_2, \alpha', \alpha' \wedge x_1, \ldots, \alpha' \wedge x_n)$ and the judges' profile P as follows.

P	β_1	β_2	α'	$\alpha' \wedge x_1$	$\ldots$	$\alpha' \wedge x_n$
P_1	0	1	1	1	$\ldots$	1
P_2	1	0	1	1	$\ldots$	1

We add $\psi \in \Phi_+$ to the aggregated judgment $F_q(P)$ if and only if $|\{i \in [r] \mid \psi \in P_i\}| \geq \lceil q(r+1) \rceil$ holds. For $r = 2$ and $q \leq 1/3$ we have $\lceil q(r+1) \rceil \leq 1$, so that $F_q(P) = \Phi_+$ holds.

We set $\varphi = \alpha' \wedge x_n$. Furthermore, no consistency condition is violated since α' can be satisfied for every individual judgment via β_1, β_2, even when α is unsatisfiable. In order to prevent α' from turning into a tautology when α is a tautology, we added γ.

Now, we prove that $\mathcal{I} \in$ OMS $\Leftrightarrow \mathcal{I}' \in SF_q$-Winner holds. For the direction from left to right assume that $\mathcal{I}$ is a yes-instance. After the first two iterations of the SF_q-rule we have $SF_q^{\alpha'}(P,\pi) = \{\beta_1, \beta_2\}$. By assumption, there exists a satisfying assignment for α and trivially also for $\neg\gamma$. Therefore, in the third round we can neither entail $\neg\alpha' \in SF_q(P,\pi)$ nor $\alpha' \in SF_q(P,\pi)$. Thus, we add α' by applying the F_q-rule. Consequently, after the third iteration we have $SF_q^{\alpha'\wedge x_1}(P,\pi) = \{\beta_1, \beta_2, \alpha'\}$. From this fact it follows that $SF_q^{\alpha'\wedge x_1}(P,\pi) \models \alpha \wedge \gamma \models \alpha, \gamma$ holds, which is in accordance with our assumption that α is satisfiable. Now, we would like to decide whether to add $\alpha' \wedge x_1$ or $\neg(\alpha' \wedge x_1)$ to $SF_q(P,\pi)$. Given the current aggregated judgment and knowing that $\gamma \equiv \textsc{true}$, it holds that $\alpha' \wedge x_1 = [(\alpha\wedge\gamma)\vee\neg\beta_1\vee\neg\beta_2]\wedge x_1 \equiv \alpha\wedge x_1$. Furthermore, knowing from $\alpha\wedge\gamma \equiv \alpha' \in SF_q(P,\pi)$ that α should be true, we distinguish three cases for $\alpha \wedge x_1$: (i) If $x_1 = 1$ is the only option for a satisfying assignment of α, we can deduce $\alpha' \wedge x_1 \in SF_q(P,\pi)$. (ii) If $x_1 = 0$ is the only option for a satisfying assignment of α, we can deduce $\neg(\alpha' \wedge x_1) \in SF_q(P,\pi)$. (iii) If there are satisfying assignments for α with both, $x_1 = 1$ and $x_1 = 0$, we must apply the F_q-rule and obtain $\alpha' \wedge x_1 \in SF_q(P,\pi)$. Note that the last option always favors the bigger satisfying assignment, i.e., preferring $x_1 = 1$ over $x_1 = 0$. We can apply the previous argument for $j \in \{1,\ldots,n\}$ and deduce for all formulas $\alpha' \wedge x_j$ whether to add them or their corresponding negation $\neg(\alpha' \wedge x_j)$ to $SF_q(P,\pi)$. Doing so yields a maximum satisfying assignment for α, represented by $[x_i = 1] \Leftrightarrow [\alpha' \wedge x_i \in SF_q(P,\pi)]$. By assumption, we know that $x_n = 1$ holds for a maximum satisfying assignment of α. Thus, $\alpha' \wedge x_n \in SF_q(P,\pi)$ holds after the last iteration and therefore, $\mathcal{I}' \in SF_q$-Winner is true.

For the direction from right to left assume now that $\mathcal{I}$ is a no-instance. We study two separate cases.

Case 1: α is satisfiable but for its maximum satisfying assignment $x_n = 0$ holds. In the third iteration we add α' to $SF_q(P,\pi)$. As already argued in the first part of the proof, for $1 \leq j \leq n$ we add $\alpha' \wedge x_j$ to $SF_q(P,\pi)$ if and only if $x_j = 1$ holds in α's maximum satisfying assignment. By assumption, we know that $x_n = 0$ is true in α's maximum satisfying assignment. Therefore, we end up with $\alpha' \wedge x_n \notin SF_q(P,\pi)$ and can conclude that $\mathcal{I}' \notin SF_q$-Winner holds.

Case 2: α is not satisfiable. After the first two iterations of the SF_q-rule we have $SF_q^{\alpha'}(P,\pi) = \{\beta_1, \beta_2\}$. By assumption, in the third iteration it holds that

$$\alpha' = (\alpha \wedge \gamma) \vee \neg\beta_1 \vee \neg\beta_2 \equiv (\textsc{false} \wedge \gamma) \vee \neg\beta_1 \vee \neg\beta_2 \equiv \textsc{false}.$$

Consequently, we deduce that $\neg\alpha'$ must hold and thus add $\neg\alpha'$ to $SF_q(P,\pi)$. Obviously, this leads to the fact that we add $\neg(\alpha' \wedge x_j)$ to $SF_q(P,\pi)$ for $1 \leq j \leq n$. Therefore, we have $\alpha' \wedge x_n \notin SF_q(P,\pi)$ and hence, $\mathcal{I}' \notin SF_q$-Winner.

Finally, we have $\mathcal{I} \in$ OMS if and only if $\mathcal{I}' \in SF_q$-Winner and obtain OMS $\leq_{\mathrm{m}}^{\mathrm{p}} SF_q$-Winner. □

Endriss and de Haan [11] showed that the winner problem for the ranked agenda rule (with fixed tie-breaking) is Δ_2^p-hard. However, the corresponding proof requires a linear number of judges. We note that slightly modifying our previous proof by adding a third judge, supporting both, β_1 and β_2, but no other formula, allows us to reuse the same proof (i.e., the given order π) for the ranked agenda rule. This yields an even stricter result for the ranked agenda's winner problem's complexity, namely para-Δ_2^p-hardness[5] with respect to the number of judges.

Corollary 1. *The winner problem for the ranked agenda rule with fixed tie-breaking is para-Δ_2^p-hard when parameterized by the number of judges.*

Our lower bound proofs in Section 4 may be adapted in a similar way (by adding a third judge only approving corresponding β_j) to also handle the ranked agenda rule.

[5] For further reading on related parameterized complexity classes, we refer to the textbook by de Haan [6].

4 Problems of Manipulative Design

While the usage of sequential rules guarantees consistency, at the same time the gradual aggregation approach leads to problems of manipulative design for anonymous underlying rules. Following the impossibility result by List [20], sequential quota rules are path-dependent, i.e., the aggregated judgment is determined by the processing order of formulas and might be altered at will if said order is chosen accordingly. Therefore, a manipulator has a great deal of control over the processing order. We study how hard it is to compute whether at least one (respectively every) order guarantees a partial judgment to be included into the aggregated one. Although List already proposed said approach as *Manipulation by Agenda Setting*, we deviate in studying two variants. In particular, we study the $S\mathcal{K}$-Winner-Design and the $S\mathcal{K}$-Winner-Robustness problem and will show that it is more inefficient for sequential quota rules to solve proposed problems of manipulative design than the corresponding winner problem. The formal definition of the Winner-Design problem is as follows for a given sequential JA rule $S\mathcal{K}$.

$S\mathcal{K}$-Winner-Design ($S\mathcal{K}$D)	
Instance:	An agenda Φ, a profile $P \in \mathcal{J}(\Phi)^r$, and a set of formulas $J \subseteq \Phi$.
Question:	Is there an order $\pi = (\varphi_1, \ldots, \varphi_m)$ over Φ_+ such that $J \subseteq S\mathcal{K}(P, \pi)$?

Analogously we formulate the almost complementary decision problem S$\mathcal{K}$-Winner-Robustness (S$\mathcal{K}$R). The input remains unchanged but the question is whether $J \subseteq S\mathcal{K}(P, \pi)$ holds for every processing order π over Φ_+. In order to determine the computational complexity of S$\mathcal{K}$D and S$\mathcal{K}$R, we require some notation.

Definition 4. *Let $\mathcal{K}$ be a complete and complement-free JA rule, Φ an agenda, and $P \in \mathcal{J}(\Phi)^r$ a profile for r judges. Furthermore, slightly abusing notation, let $\pi = (\varphi_1, \ldots, \varphi_m)$ be an order over $\mathcal{K}(P)$ and denote by $S\mathcal{K}(P, \pi)$ the corresponding aggregated judgment. Let $K_\pi = \mathcal{K}(P) \cap S\mathcal{K}(P, \pi)$ denote the set of formulas in the aggregated judgment also supported by $\mathcal{K}$, and $D_\pi = S\mathcal{K}(P, \pi) \setminus \mathcal{K}(P)$ those not supported by $\mathcal{K}$. For $K_\pi = \{k_1, \ldots, k_p\}$ and $D_\pi = \{d_1, \ldots, d_{m-p}\}$ let $(K_\pi, D_\pi) = (k_1, \ldots, k_p, d_1, \ldots, d_{m-p})$ denote an order, where all formulas in K_π are permuted arbitrarily at the first p places.*

This enables us to formulate the following lemma, whose proof is deferred to Appendix B.

Lemma 1. *Let $\mathcal{K}$ be a complete and complement-free JA rule, Φ an agenda and $P \in \mathcal{J}(\Phi)^r$ a profile for r judges. Then, for every order of the form $\pi' = (K_\pi, D_\pi)$ it holds that $S\mathcal{K}(P, \pi') = S\mathcal{K}(P, \pi)$.*

The intuition is, that we can rearrange every order π in such a way that all formulas supported by $\mathcal{K}$ are at the beginning of π and all remaining formulas follow afterwards. Hence, instead of looking for a specific order it is sufficient to search for a consistent subset $K \subseteq \mathcal{K}(P)$, such that $K \models \bigwedge_{\varphi \in J} \varphi$ holds. Doing so enables us to solve a $S\mathcal{K}$-Winner-Design instance by setting $\pi = (K, J, \ldots)$.

Note that for $q = 1/2$, the problems SF_q-Winner-Design and SF_q-Winner-Robustness are closely related to the winner determination problem for the ranked agenda rule without fixed tie-breaking as studied by Endriss and de Haan [11] and Lang and Slavkovik [19]. Both investigate hardness for similar decision problems, where the succession of the elements for the processing order is fixed as a decreasing sequence over the number of supporting judges per element (i.e., for the order $\pi = (\varphi_1, \ldots, \varphi_m)$ over $F_{1/2}(P)$ it holds that $|\{i \in [r] \mid \varphi_j \in P_i\}| \geq |\{i \in [r] \mid \varphi_{j+1} \in P_i\}|$). We continue to study the complexity for two widely separated cases, namely manipulative design for complete judgment sets (Section 4.1) and for single formulas (Section 4.2). An overview of our results is given in Table 1.

Table 1: Summary of complexity results for different problems for sequential JA rules SF_q.

WINNER	WINNER-DESIGN		WINNER-ROBUSTNESS		SUPPORTED-JUDGMENT
	$J \in \mathcal{J}(\Phi)$	$\varphi \in \Phi$	$J \in \mathcal{J}(\Phi)$	$\varphi \in \Phi$	
Δ_2^p-c.	coNP-c.	Σ_2^p-c.	in P	Π_2^p-c.	NP-c.
Thm. 1, 2	Thm. 3, 4	Thm. 6, 7	Thm. 5	Lem. 2	Thm. 8, 9

4.1 Manipulative Design for Judgment Sets

First, let us investigate the introduced problems of manipulative design for a given judgment which is complete and consistent. Note that we do not consider inconsistent judgments, since those are neither desirable nor a possible output. The ensuing theorem derives an upper bound of coNP for a broad class of sequential JA rules.

Theorem 3. *For every polynomial-time computable JA rule $\mathcal{K}$ that is complete and complement-free, it holds that $S\mathcal{K}D \in$ coNP if the desired subset of formulas equals a complete and consistent judgment $J \in \mathcal{J}(\Phi)$.*

Proof. We precompute $K = J \cap \mathcal{K}(P)$ and $D = J \setminus \mathcal{K}(P)$ in polynomial time. Since $J \in \mathcal{J}(\Phi)$, K and D are consistent. Following Lemma 1 it is sufficient to verify whether each formula in D can be derived from K, since we then may construct an order of the form $\pi' = (K, D)$. Hence, we have to check whether $\left(\bigwedge_{\varphi \in K} \varphi\right) \models \left(\bigwedge_{\psi \in D} \psi\right)$. This is equivalent to checking whether no assignment is satisfying $\left(\bigwedge_{\varphi \in K} \varphi\right) \wedge \neg \left(\bigwedge_{\psi \in D} \psi\right)$ and hence in coNP. □

For the class of quota rules the following theorem establishes the matching lower bound and proves coNP-hardness.

Theorem 4. *For every quota rule $F_q \in \mathcal{F}$ and every given complete and consistent judgment $J \in \mathcal{J}(\Phi)$ it is* coNP*-complete to solve the corresponding SF_qD problem.*

Proof. Recall that we assume every quota rule F_q to be complete and complement-free for every quota q. To show coNP-hardness, we reduce a $\overline{\text{SAT}}$ instance $\mathcal{I} = (\alpha)$ to a SF_qD instance $\mathcal{I}' = (\Phi, P, J)$. We define $\Phi_q = \{(\alpha \wedge \gamma) \vee \neg\beta_1 \vee \neg\beta_2, \beta_1, \beta_2\}$, where γ, β_1, and β_2 are new literals, and choose $\Phi_+ = \Phi_q$ for $q \leq 1/3$ and $\Phi_- = \Phi_q$ otherwise. We consider a profile consisting of two judges with $P_i = \{(\alpha \wedge \gamma) \vee \neg\beta_1 \vee \neg\beta_2, \beta_i, \neg\beta_{3-i}\}$ for $i \in [2]$. Note that by construction it holds that $F_q(P) = \Phi_q$. Lastly, we set $J = P_1$ and show that equivalence holds. For the direction from left to right assume $\mathcal{I}$ is a YES-instance and thus, α is unsatisfiable. Choosing the order $\pi = ((\alpha \wedge \gamma) \vee \neg\beta_1 \vee \neg\beta_2, \beta_1, \beta_2)$ over Φ_q results in $SF_q(P, \pi) = J$. For the direction from right to left assume $\mathcal{I}$ is a NO-instance and thus, α is satisfiable. Then, $F_q(P)$ is already consistent and $SF_q(P, \pi) = F_q(P) \neq J$ holds for every order π. Together with Theorem 3 we obtain coNP-completeness. □

Turning to the robustness problem, we require that the desired judgment set J is contained in the collective outcome for every possible order. This is only possible if each of the formulas is contained in the collective judgment set of the underlying formula. See Appendix B for the respective proof.

Theorem 5. *For every agenda Φ, profile $P \in \mathcal{J}(\Phi)^r$ and complete and consistent judgment $J \in \mathcal{J}(\Phi)$, the corresponding $S\mathcal{K}R$-instance (Φ, P, J) is satisfiable if and only if $\mathcal{K}(P) = J$ for a complete and complement-free rule $\mathcal{K}$.*

Note that for efficiently computable underlying rules and particularly for sequential quota rules SF_q the corresponding problem is decidable in P.

4.2 Manipulative Design for Single Formulas

Before investigating the complexity of S$\mathcal{K}$D and S$\mathcal{K}$R separately, we want to point out that they are tied closely together, when testing whether a single formula is in the aggregated judgment.

Lemma 2. *For every complete and complement-free rule $\mathcal{K}$, every agenda Φ, every profile $P \in \mathcal{J}(\Phi)^r$ and every formula $\varphi \in \Phi$, it holds that $(\Phi, P, \{\varphi\}) \in S\mathcal{K}R \Leftrightarrow (\Phi, P, \{\sim\varphi\}) \in \overline{S\mathcal{K}D}$.*

The above lemma follows from complement-freeness and completeness and has also been shown by Lang and Slavkovik [19]. In the following, we will only show complexity results for S$\mathcal{K}$D, while results for S$\mathcal{K}$R follow directly. We continue to establish upper bounds.

Theorem 6. *For every polynomial-time computable, complete and complement-free JA rule $\mathcal{K}$ and a judgment $J = \{\varphi\} \subset \Phi$ containing a single formula, it holds that $S\mathcal{K}$-*Winner-Design $\in \Sigma_2^p$.

Proof. In order to solve an instance $\mathcal{I} = (\Phi, P, \{\varphi\})$ of the decision problem $S\mathcal{K}$-Winner-Design, we must determine whether there exists an order π such that $\varphi \in S\mathcal{K}(P, \pi)$ holds. Exploiting our previous observations, we know from Lemma 1 that it is sufficient to identify a consistent subset $K \subseteq \mathcal{K}(P)$ with $K \models \varphi$.

Thus, we first calculate $\mathcal{K}(P)$ in polynomial time and can nondeterministically guess a subset $K = \{\varphi_1, \ldots, \varphi_k\} \subseteq \mathcal{K}(P)$. Next, we verify whether K is consistent by asking our NP-oracle whether there exists a satisfying assignment for $\varphi_1 \wedge \ldots \wedge \varphi_k$. In a last step, we must determine whether $K \models \varphi$ holds. Thereby, we have $(\varphi_1 \wedge \ldots \wedge \varphi_k) \models \varphi$. To determine whether this formula is satisfiable can again be solved in coNP. Consequently, we can pose a second NP-query to find out whether K entails φ, resulting in $\varphi \in S\mathcal{K}(P, \pi)$ for $\pi = (K, \varphi, \ldots)$. Overall, we require a polynomial number of nondeterministic computation steps as well as two NP-oracle queries to calculate an answer for $\mathcal{I}$ and thus, $S\mathcal{K}$-Winner-Design $\in \Sigma_2^p$ holds. □

Combining the former theorem with Lemma 2, we derive the following corollary.

Corollary 2. *For every complete and complement-free JA rule $\mathcal{K}$ computable in polynomial time and a judgment $J = \{\varphi\} \subset \Phi$, it holds that $S\mathcal{K}$-*Winner-Robustness $\in \Pi_2^p$.

We continue by showing matching lower bounds for sequential quota rules.

Theorem 7. *For every quota rule $F_q \in \mathcal{F}$ and a judgment $J = \{\varphi\} \subset \Phi$ consisting of a single formula, it holds that the problem SF_q-*Winner-Design *is Σ_2^p-complete.*

The omitted proof is based on a reduction from Succinct Set Cover (SSC), which was proven to be Σ_2^p-complete by Umans [25]. For the reduction we construct a polynomial-time computable formula that is satisfiable by a given truth assignment if and only if at most k of its boolean variables are set to true via the corresponding truth assignment. For a thorough explanation of the formula's construction as well as for the corresponding proof we refer to Appendix A and Appendix B respectively.

Again, we derive a corollary for SF_qR from the previous theorem and Lemma 2.

Corollary 3. *For every quota rule $F_q \in \mathcal{F}$ and a judgment $J = \{\varphi\} \subset \Phi$, it holds that SF_q-*Winner-Robustness *is Π_2^p-complete.*

Endriss and de Haan [11] investigate the complexity of existential winner-determination for the ranked agenda rule without a fixed tie-breaking which is shown to be Σ_2^p-hard. Similarly to Corollary 1, we may improve this result, as our proof of Theorem 7 can easily be adapted (by adding a third judge only approving β_j) to also hold for the ranked agenda rule without fixed tie-breaking.

Corollary 4. *The winner problem for the ranked agenda rule without fixed tie-breaking is para-Σ_2^p-hard when parameterized by the number of judges.*

4.3 Supported Judgment

We conclude this section by formulating a problem, which formally relates to problems of manipulative design, although it is clearly motivated contrarily. In terms of acceptance, it is desirable for an aggregated judgment to be reasonable for the participating judges. Hence, for sequential JA rules it should be preferable to choose an order such that at least k formulas supported by a rule $\mathcal{K}$ are included in the aggregated judgment.

$S\mathcal{K}$-Supported-Judgment ($S\mathcal{K}$SJ)	
Instance:	An agenda Φ with $\|\Phi_+\| = m$, a profile $P \in \mathcal{J}(\Phi)^r$ for r judges and an integer $k \leq m$.
Question:	Is there an order $\pi = (\varphi_1, \ldots, \varphi_m)$ over Φ_+ such that $\|\mathcal{K}(P) \cap S\mathcal{K}(P, \pi)\| \geq k$ holds?

We start by establishing a general upper bound.

Theorem 8. *$S\mathcal{K}$-*Supported-Judgment *is in* NP *for every efficiently computable rule $\mathcal{K}$.*

The omitted proof relies on Lemma 1 and can be found in Appendix B. For the class of sequential quota rules we provide a matching lower bound by adapting the proof of Theorem 4. Again, we refer to Appendix B for the corresponding proof.

Theorem 9. *SF_q-*Supported-Judgment *is* NP*-complete for every quota rule $F_q \in \mathcal{F}$.*

Lastly, we highlight the significance of Lemma 1 for building a connection between our sequential rules and distance based rules. While it is not directly obvious, for $q = 1/2$, SF_qSJ is related to the maxcard subagenda rule[6] as studied by Lang and Slavkovik [19]. In general, S$\mathcal{K}$SJ coincides with asking whether there exists a complete and consistent judgment $J \in \mathcal{J}(\Phi)$, such that $h(\mathcal{K}(P), J) \leq m - k$ (where $h(\mathcal{K}(P), J)$ denotes the hamming distance between $\mathcal{K}(P)$ and J). If there exists such an order π, for the resulting outcome $S\mathcal{K}(P, \pi)$ it clearly holds that $h(\mathcal{K}(P), S\mathcal{K}(P, \pi)) \leq m - k$. Vice versa, if there exists a judgment $J \in \mathcal{J}(\Phi)$ with $h(\mathcal{K}(P), J) \leq m - k$, we construct a valid order π following Lemma 1 by arbitrarily positioning the supported formulas at the beginning. These observations may be an interesting tool for further research on computational complexity for counting problems.

5 Sequential Rules and the Maximum Subagenda Rule

In this section we describe how we can link the sequential JA rules that we've studied to other well-known majority preserving JA rules. Particularly, we highlight the case with the majority rule as underlying rule to our sequential rule. Hereby, we show that the maximum subagenda rule (MSA), as defined by Lang and Slavkovik [19], exactly outputs the set of

[6]Defined as $MSCA(P) = \arg\min_{J \in \mathcal{J}(\Phi)} h(F_{1/2}(P), J)$ for $P \in \mathcal{J}(\Phi)^n$, where $h(F_{1/2}(P), J)$ is the hamming distance between $F_{1/2}(P)$ and J.

aggregated judgments which can also be derived by the sequential majority rule with suitable processing orders applied. This connection enables us to transfer some of our complexity results to related non-sequential rules. In order to make the most out of this connection, we slightly generalize the MSA rule defined in [19] as described afterwards.

Definition 5 (Generalized Maximum Subagenda Rule)**.** *For an agenda Φ and a set $S \subseteq \Phi$ we define $max(S, \subseteq) \subset 2^S$ as the set consisting of inclusion maximal subsets of S with respect to consistency. More formally, for $S' \subseteq S$ it holds that $S' \in max(S, \subseteq)$ if and only if S' is consistent and there exists no consistent set $S'' \subseteq S$ with $S' \subset S''$. For any complete and resolute JA rule $\mathcal{K}$, we define the (irresolute)* ***generalized maximum subagenda rule*** *$MSA_{\mathcal{K}}\colon \mathcal{J}(\Phi)^r \to 2^{\mathcal{J}(\Phi)}$ as follows. Let $P \in \mathcal{J}(\Phi)^r$ be a profile of judgments and $J \in \mathcal{J}(\Phi)$ a judgment, then $J \in MSA_{\mathcal{K}}(P)$ holds if and only if there exists a set $S \in max(\mathcal{K}(P), \subseteq)$ with $S \subseteq J$.*

The MSA rule is irresolute, i.e., it returns a set of judgments as result, and equals the definition presented by Lang and Slavkovik [19] for $\mathcal{K} = F_{1/2}$. Having the MSA rule defined, we make the subsequent observation, establishing a connection between the MSA rule and our earlier studied sequential quota JA rules.

Theorem 10. *Let $P \in \mathcal{J}(\Phi)^r$ be a profile and $J \in \mathcal{J}(\Phi)$ a complete and consistent judgment. Then, $J \in MSA_{\mathcal{K}}(P)$ holds if and only if there exists an order π over Φ_+ with $S\mathcal{K}(P, \pi) = J$.*

Proof. We begin with the direction from left to right. By definition, $MSA_{\mathcal{K}}(P)$ contains every complete and consistent judgment J, such that there doesn't exist a consistent set $K \subseteq \mathcal{K}(P)$ satisfying $J \cap \mathcal{K}(P) \subset K$. Note that this especially holds for $|K| = |J \cap \mathcal{K}(P)| + 1$, i.e., $J \cap \mathcal{K}(P)$ cannot be extended by a single formula from $\mathcal{K}(P)$. Due to consistency of J there is a satisfying truth assignment for $J \cap \mathcal{K}(P)$. Yet, no such truth assignment satisfies any formula in $\mathcal{K}(P) \setminus J$ and must thus satisfy its complement. Hence, it holds that $J \cap \mathcal{K}(P)$ must entail $J \setminus \mathcal{K}(P)$. Now, following a similar argumentation as in Lemma 1, for $\pi = (J \cap \mathcal{K}(P), J \setminus \mathcal{K}(P))$ we obtain $S\mathcal{K}(P, \pi) = J$ and therefore, the right side holds, too.

For the direction from right to left assume that there is an outcome $J = S\mathcal{K}(P, \pi)$ with $J \notin MSA_{\mathcal{K}}(P)$. Note that J is consistent by definition and hence, its intersection with $\mathcal{K}(P)$ is consistent, too. By assumption, $J \cap \mathcal{K}(P)$ cannot be inclusion maximal in $\mathcal{K}(P)$ with respect to consistency as otherwise $J \in MSA_{\mathcal{K}}(P)$ would follow. Therefore, let $K \in max(\mathcal{K}(P), \subseteq)$, such that $J \cap \mathcal{K}(P) \subset K \subseteq \mathcal{K}(P)$ holds. Now, we construct an order π' where $J \cap \mathcal{K}(P)$ is at the beginning of π', immediately followed by $K \setminus J \cap \mathcal{K}(P)$, and all remaining formulas afterwards. With Lemma 1 it holds that $J = S\mathcal{K}(P, \pi')$ is true. Yet, $K \subseteq S\mathcal{K}(P, \pi')$ holds as well since K is a consistent subset of $\mathcal{K}(P)$ processed at the beginning of π'. Hence we conclude that $K \subseteq J$ must hold, which is a contradiction to $J \cap \mathcal{K}(P) \subset K \subseteq \mathcal{K}(P)$. ☐

The previous theorem can be applied to transfer complexity results for our decision problems in Section 4. For complete and resolute JA rules $\mathcal{K}$, asking whether there exists an order π, such that some condition on the output $S\mathcal{K}(P, \pi)$ is satisfied, coincides with asking whether there is a judgment $J \in MSA_{\mathcal{K}}(P)$ satisfying the same condition. In particular, for $F_q = 1/2$ and a single formula φ the problem SF_q-Winner-Design coincides with the existential MSA-Winner problem, while SF_q-Winner-Robustness coincides with the universal variant.[7]

This observation has multiple consequences. First of all, Lang and Slavkovik [19] showed the universal MSA-Winner problem is Π_2^p-complete, which aligns with our result from

[7] Slightly abusing notation, we consider existential ($\exists J \in MSA_{\mathcal{K}}(P) : \{\varphi\} \subseteq J$) and universal variants of MSA$_{\mathcal{K}}$-Winner ($\forall J \in MSA_{\mathcal{K}}(P) : \{\varphi\} \subseteq J$) for irresolute rules.

Corollary 3. However, the referenced result by Lang and Slavkovik requires a linear number of judges while two judges are sufficient for our proof. Consequently, our proof allows a stricter result than the one by Lang and Slavkovik. On the other hand our results also hold if we do not restrict MSA to the majority rule as underlying JA rule. In particular, upper bounds hold for every complete, efficiently computable, resolute rule, while hardness results hold for every of our quota rules.

The following corollaries follow from Theorems 6, 7 and 10, and only refer to existential problems, which imply related Π_2^p results for the universal variants, by additionally following Lemma 2.

Corollary 5. *For any complete, efficiently computable, resolute JA rule $\mathcal{K}$ it holds that* $\mathrm{MSA}_{\mathcal{K}}$-Winner *is in* Σ_2^p.

Corollary 6. *For every quota rule $F_q \in \mathcal{F}$ and even a constant number of judges it holds that* MSA_{F_q}-Winner *is* Σ_2^p*-complete.*

We explicitly highlight that the previous corollary holds for $q = 1/2$, and thereby enhances previous results on MSA.

6 Conclusion

We introduced the complexity theoretic study of problems related to sequential JA rules with a special focus on quota rules as the underlying rule. Our results are summarized in Table 1. We obtained completeness for a number of different complexity classes which show that the problems differ substantially even though they are very related. The study of sequential rules is very important since they model real-world decision making. To ensure consistency with the already decided formulas, it is important to solve the winner problem. On the other hand, we studied the manipulative power a designer of such a rule possesses. The increase in complexity for the case where a single formula is the desired set indicates that the problem is actually harder than winner determination itself. As a task for future research other problems related to sequential JA rules have to be studied. Our study was mostly limited to the class of quota rules as underlying rules and this should obviously be extended to more diverse underlying rules. De Haan [5] follows an approach to identify new ways of representing agendas via specific boolean formulas, such that the complexity of various problems related to JA becomes tractable, when the agenda is represented in a more limited way. Furthermore, he formulated the determination of the complexity of the winner problem for JA rules that have not been considered yet, as future work. In a second step, the author suggests that one can use the tractable languages identified in his paper to study whether the complexity of the problems for the newly investigated JA rules can be decreased. Within our paper we have done the first part and determined the complexity of the winner problem for complete and consistent sequential JA rules. As future work we like to study how the tractable languages as defined by de Haan [5] affect our complexity results and possibly could even enable lower bounds. These results, when enabling tractability, might have enormous impact on the practical usage of the sequential JA rules we studied, since they are used in various scenarios and situations, as described earlier.

Acknowledgements

This work was supported by DFG grants BA 6270/1-1 and RO 1202/21-1 and by a grant from the NRW Ministry for Innovation, Science, and Research.

References

[1] Sanjeev Arora and Boaz Barak. *Computational Complexity: A Modern Approach.* Cambridge University Press, 2009.

[2] Dorothea Baumeister, Linus Boes, and Robin Weishaupt. Complexity of Sequential Rules in Judgment Aggregation. *Proceedings of the 20th Conference on Autonomous Agents and MultiAgent Systems*, pages 187–195, 2021.

[3] Dorothea Baumeister, Gábor Erdélyi, Olivia J. Erdélyi, and Jörg Rothe. Complexity of Manipulation and Bribery in Judgment Aggregation for Uniform Premise-Based Quota Rules. *Mathematical Social Sciences*, 76:19–30, 2015.

[4] Stephen A. Cook. The Complexity of Theorem-Proving Procedures. In *Proceedings of the 3rd ACM Symposium on Theory of Computing*, pages 151–158. ACM Press, 1971.

[5] Ronald de Haan. Hunting for Tractable Languages for Judgment Aggregation. In *Sixteenth International Conference on Principles of Knowledge Representation and Reasoning*, 2018.

[6] Ronald de Haan. *Parameterized Complexity in the Polynomial Hierarchy - Extending Parameterized Complexity Theory to Higher Levels of the Hierarchy*, volume 11880 of *Lecture Notes in Computer Science.* Springer, 2019.

[7] Ronald de Haan and Marija Slavkovik. Complexity Results for Aggregating Judgments using Scoring or Distance-Based Procedures. In *Proceedings of the 16th Conference on Autonomous Agents and MultiAgent Systems*, pages 952–961. International Foundation for Autonomous Agents and Multiagent Systems, 2017.

[8] Franz Dietrich and Christian List. Judgment Aggregation by Quota Rules: Majority Voting Generalized. *Journal of Theoretical Politics*, 19(4):391–424, 2007.

[9] Ulle Endriss. Judgment Aggregation. In Felix Brandt, Vincent Conitzer, Ulle Endriss, Jérôme Lang, and Ariel D. Procaccia, editors, *Handbook of Computational Social Choice*, chapter 17, pages 399–426. Cambridge University Press, 2016.

[10] Ulle Endriss. Judgment Aggregation with Rationality and Feasibility Constraints. In *Proceedings of the 17th International Conference on Autonomous Agents and MultiAgent Systems*, pages 946–954. International Foundation for Autonomous Agents and Multiagent Systems, 2018.

[11] Ulle Endriss and Ronald de Haan. Complexity of the Winner Determination Problem in Judgment Aggregation: Kemeny, Slater, Tideman, Young. In *Proceedings of the 2015 International Conference on Autonomous Agents and Multiagent Systems*, pages 117–125. International Foundation for Autonomous Agents and Multiagent Systems, 2015.

[12] Ulle Endriss, Ronald de Haan, Jérôme Lang, and Marija Slavkovik. The Complexity Landscape of Outcome Determination in Judgment Aggregation. *Journal of Artificial Intelligence Research*, 69:687–731, 2020.

[13] Ulle Endriss, Umberto Grandi, and Daniele Porello. Complexity of Judgment Aggregation. *Journal of Artificial Intelligence Research*, 45:481–514, 2012.

[14] André Große, Jörg Rothe, and Gerd Wechsung. On Computing the Smallest Four-Coloring of Planar Graphs and Non-Self-Reducible Sets in P. *Information processing letters*, 99(6):215–221, 2006.

[15] Wojciech Jamroga and Marija Slavkovik. Some Complexity Results for Distance-Based Judgment Aggregation. In *AI 2013: Advances in Artificial Intelligence - 26th Australasian Joint Conference, Dunedin, New Zealand, December 1-6, 2013. Proceedings*, pages 313–325, 2013.

[16] Lewis A. Kornhauser and Lawrence G. Sager. Unpacking the Court. *Yale Law Journal*, 96(1):82–117, 1986.

[17] Mark W. Krentel. The Complexity of Optimization Problems. In *Proceedings of the eighteenth annual ACM symposium on Theory of computing*, pages 69–76, 1986.

[18] Jérôme Lang, Gabriella Pigozzi, Marija Slavkovik, Leendert van der Torre, and Srdjan Vesic. A Partial Taxonomy of Judgment Aggregation Rules and their Properties. *Social Choice and Welfare*, 48(2):327–356, 2017.

[19] Jérôme Lang and Marija Slavkovik. How Hard is it to Compute Majority-Preserving Judgment Aggregation Rules? In *Proceedings of the Twenty-first European Conference on Artificial Intelligence*, pages 501–506. IOS Press, 2014.

[20] Christian List. A Model of Path-Dependence in Decisions over Multiple Propositions. *American Political Science Review*, 98(3):495–513, 2004.

[21] Michael K. Miller and Daniel Osherson. Methods for Distance-Based Judgment Aggregation. *Social Choice and Welfare*, 32(4):575, 2009.

[22] Klaus Nehring, Marcus Pivato, and Clemens Puppe. The Condorcet set: Majority Voting over Interconnected Propositions. *Journal of Economic Theory*, 151:268–303, 2014.

[23] Bezalel Peleg and Shmuel Zamir. Judgments Aggregation by a Sequential Majority Procedure. *Mathematical Social Sciences*, 95:37–46, 2018.

[24] Daniele Porello and Ulle Endriss. Ontology Merging as Social Choice: Judgment Aggregation under the Open World Assumption. *Journal of Logic and Computation*, 24(6):1229–1249, 2014.

[25] Christopher Umans. Hardness of Approximating Σ_2^p Minimization Problems. In *40th Annual Symposium on Foundations of Computer Science (Cat. No. 99CB37039)*, pages 465–474. IEEE, 1999.

Dorothea Baumeister
Institut für Informatik
Heinrich-Heine-Universität
Düsseldorf, Germany
Email: `d.baumeister@uni-duesseldorf.de`

Linus Boes
Institut für Informatik
Heinrich-Heine-Universität
Düsseldorf, Germany
Email: `linus.boes@uni-duesseldorf.de`

Robin Weishaupt
Institut für Informatik
Heinrich-Heine-Universität
Düsseldorf, Germany
Email: `robin.weishaupt@uni-duesseldorf.de`

A Counting Technique

Within this section, we introduce a polynomial-time computable technique used to construct a boolean formula ψ_k^B. The formula is able to count the number of satisfied boolean variables for a given boolean assignment T of a set of boolean variables B in the sense that a truth assignment evaluates the formula to true if and only if at most $k \in \mathbb{N}$ of the variables in B for T are TRUE.

In some sense our technique generalizes the already known technique used by Cook in his famous theorem to prove that SAT is NP-complete, cf. [4]. Cook's technique describes an approach how to formulate a boolean formula for a set of boolean variables which is true if and only if exactly one of the boolean variables is true.

Lemma 3. *Let $B = \{x_1, \ldots, x_n\}$ be a set of boolean variables and $k \leq n$. We can construct a formula ψ_k^B from a set of boolean variables B' with $|B'| = nk$ in time polynomial in n, such that ψ_k^B evaluates to* TRUE *if and only if at most k of the n boolean variables in B are set to* TRUE.

Proof. In a first step, we create k copies $\{x_i^1, \ldots, x_i^k\}$ for every boolean variable x_i in B. Then, we define a boolean formula X_i for every $1 \leq i \leq n$ as follows

$$X_i = \left[\bigvee_{j \in [k]} \left(x_i^j \wedge \bigwedge_{\ell \in [k] \setminus \{j\}} \neg x_i^\ell \right)\right] \vee \left[\bigwedge_{j \in [k]} \neg x_i^j\right].$$

Consequently, X_i is satisfied if and only if at most one of the k copies of x_i is satisfied. Note that every X_i can be constructed in time in $\mathcal{O}(n^2)$ since $|X_i| = k(k+1) \leq n(n+1)$ holds.

In a second step, we construct k boolean formulas Y_j for $1 \leq j \leq k$ as follows

$$Y_j = \left[\bigvee_{i \in [n]} \left(x_i^j \wedge \bigwedge_{\ell \in [n] \setminus \{i\}} \neg x_\ell^j \right)\right] \vee \left[\bigwedge_{i \in [n]} \neg x_i^j\right].$$

Thereby, Y_j is satisfied if and only if at most one of the n variables in the j-th set of copies $\{x_1^j, \ldots, x_n^j\}$ is satisfied. Note that we can also construct Y_j in time in $\mathcal{O}(n^2)$ since $|Y_j| = n(n+1)$ holds.

In a third step, we define two more boolean formulas, namely $Y = \bigwedge_{j=1}^k Y_j$ and $X = \bigwedge_{i=1}^n X_i$. Consequently, Y is satisfied if and only if for every j, $1 \leq j \leq k$, at most one variable in the set $\{x_1^j, \ldots, x_n^j\}$ is satisfied. Analogously, X is satisfied if and only if at most one of the copies for every x_i, $1 \leq i \leq n$, is satisfied. Finally, setting $\psi_k^B = Y \wedge X$ obviously completes the construction.

It remains to show the correctness of the construction. To do so, first we explain how to derive a boolean assignment T' for $B' = \{x_1^1, \ldots, x_1^k, \ldots, x_n^1, \ldots, x_n^k\}$ out of a boolean assignment T for $B = \{x_1, \ldots, x_n\}$. Therefore, denote by $\rho(B, T) = \{x \in B \mid T(x) = \text{TRUE}\}$ the set of variables set to TRUE by T. We construct T' as follows. Write $\rho(B, T) = \{x_{i_1}, \ldots, x_{i_m}\}$ for $m \leq n$. For $1 \leq j \leq m$, we set $x_{i_j}^{(j \bmod k)+1}$ to TRUE and all other variables in B' to FALSE.

We prove the directions separately.

"$\Rightarrow$" Assume that $|\rho(B, T)| \leq k$ is true. Let $\rho(B, T) = \{x_{i_1}, \ldots, x_{i_m}\}$ with $m \leq k$. We define the boolean assignment T' for B' as follows. For every $1 \leq j \leq m$, we set $x_{i_j}^j =$ TRUE and all other variables of B' to FALSE. Thereby, we can replace "$(j \bmod k) + 1$" by "j" , since $m \leq k$ holds. Using this assignment, at most one copy per variable is true, because $i_r \neq i_s$ holds for all $1 \leq r < s \leq m$. Consequently, every X_i, $1 \leq i \leq n$

is satisfied. Furthermore, we know that for every set of copies $\{x_1^j, \dots x_n^j\}$ at most one variable is true, as $m \leq k$ holds. Thus, every Y_j, $1 \leq j \leq k$, is satisfied, too. Finally, we have that Y and X are satisfied and thus, $\psi_k^B(T'(B')) = \text{TRUE}$ holds.

"$\Leftarrow$" Assume that $\psi_k^B(T'(B')) = \text{TRUE}$ holds. By construction of ψ_k^B it follows that X and Y are satisfied. Thus, every X_i, $1 \leq i \leq n$, and every Y_j, $1 \leq j \leq k$, is TRUE. Since every X_i is satisfied, it holds for every variable $x_i \in B$ that at most one of the k copies of x_i is TRUE in T'. Because every Y_j is TRUE, it holds that for every set of copies $\{x_1^j, \dots, x_n^j\}$ at most one variable is set to TRUE. Consequently, at most k of all variables in B' can be true and there can't be two copies of the same variable true. Therefore, using the assignment $x_i \Leftrightarrow x_i^1 \vee \dots \vee x_i^k$ for $1 \leq i \leq n$, we have that at most k of the variables in B are true for T. □

We use this technique as follows. Let $B = \{x_1, \dots, x_n\}$ be a set of boolean variables, $k \in \mathbb{N}$ and $\alpha(B)$ some boolean formula over B. At some point, we must know whether a given assignment T satisfies $\alpha(B)$, while no more than k of the boolean variables in B should be set to TRUE. In order to decide this fact efficiently, we first globally replace each variable $x_i \in B$ that appears in α by $\bigvee_{j\in[k]} x_i^j$ and denote the result as α_k. Then, we construct a new boolean formula $\alpha'_k = \alpha_k \wedge \psi_k^B$ and check whether α'_k is true for the corresponding assignment T'. If this is the case, we know that $\alpha(T(B))$ is true, while no more than k of the n variables in B are true for T. In order to keep our notation as simple as possible, we write $\alpha' = \alpha \wedge \psi_k^B$.

B Omitted Proofs

Case $q > 1/3$ for proof of Theorem 2. Assume $q > 1/3$. Define $\Phi_+ = \{\neg\beta_1, \neg\beta_2, \neg\alpha', \neg(\alpha' \wedge x_1), \dots, \neg(\alpha' \wedge x_n)\}$ with α' as set in the case for $q \leq 1/3$. Furthermore, we set the order π over Φ_+ as $\pi = (\neg\beta_1, \neg\beta_2, \neg\alpha', \neg(\alpha' \wedge x_1), \dots, \neg(\alpha' \wedge x_n))$, and the profile P remains unchanged. For $r = 2$ and $q > 1/3$ we have $\lceil q(r+1) \rceil \geq 2$, such that $F_q(P) = \Phi \backslash \Phi_+$ holds.

Note that the aggregated judgment $F_q(P)$ is the same as the one in the case for $q \leq 1/3$. Consequently, we can handle both cases for the remaining part of the proof as a single one and refer to the original proof of Theorem 3.2 for the left part. □

Proof of Lemma 1. First, note that by construction any subset of $S\mathcal{K}(P, \pi)$ must be consistent and thus, K_π and D_π are consistent, too. Now, let $\pi' = (K_\pi, D_\pi)$ be any order as described in the lemma. Since the formulas in K_π do not contradict each other, no negation $\sim\varphi$ of any formula $\varphi \in K_\pi$ can be derived by previously added formulas from K_π. Consequently, we obtain $K_\pi \subseteq S\mathcal{K}(P, \pi')$. For our purpose it is incidental whether formulas of K_π are deduced from previously added formulas or are added by the support of $\mathcal{K}(P)$.

Now, every formula in D_π was deduced by some previously added formulas, which in return were either added because of $\mathcal{K}(P)$ or transitively deduced again by formerly added formulas. Thereby, the formulas added directly because of $\mathcal{K}(P)$, namely $K'_\pi \subseteq K_\pi$, somehow represent a starting point for all further deductions. Since this starting point is also included in the partially aggregated judgment of $S\mathcal{K}(P, \pi')$ after the first $|K_\pi|$ iterations, it must follow that $D_\pi \subseteq S\mathcal{K}(P, \pi')$ holds as well and hence, we obtain $S\mathcal{K}(P, \pi') = S\mathcal{K}(P, \pi)$. □

Proof of Theorem 5. If $\mathcal{K}(P) = J$ holds, then $\mathcal{K}(P)$ is consistent by definition and thus, every formula is in the aggregated judgment for every order. Otherwise, there exists a formula $\varphi \in J \setminus \mathcal{K}(P)$. Processing φ in the first iteration for some order π results in $\sim\varphi \in SF_q(P, \pi)$, such that $J \neq S\mathcal{K}(P, \pi)$ follows. □

Proof of Theorem 7. In order to identify lower bounds for sequential quota rules, let us first define the decision problem SUCCINCT SET COVER (SSC), which was proven to be Σ_2^p-complete by Umans [25]. The instance consists of a collection of 3-DNF formulas $S = \{\varphi_1, \ldots, \varphi_n\}$ over m variables and $k \in \mathbb{N}$. The question is whether there is a subset $N' \subseteq [n]$ with $|N'| \leq k$ and $\bigvee_{i \in N'} \varphi_i \equiv \text{TRUE}$?

Due to Theorem 6 it is enough to show Σ_2^p-hardness. We reduce SUCCINCT SET COVER to SF_q-WINNER-DESIGN. Let $\mathcal{I} = (\{\varphi_1, \ldots, \varphi_n\}, k)$ be a SSC instance. To construct $\mathcal{I}' = (\Phi, P, \{\varphi\})$, we first introduce some auxiliary variables. Let $B = \{x_1, \ldots, x_n\}$ be a set of boolean literals, ψ_k^B defined as described in Appendix A and $\varphi_i' = (\varphi_i \wedge x_i)$ for $1 \leq i \leq n$. For our construction we set $\varphi = \psi_k^B \wedge \left[\left(\bigvee_{i \in [n]} \varphi_i'\right) \vee \gamma\right] \wedge \beta_1 \wedge \beta_2$ and $\Phi_q = B \cup \{\beta_1, \beta_2\} \cup \{\psi_k^B \vee \neg\beta_1 \vee \neg\beta_2, \sim\!\varphi\}$ with new literals β_j and γ. Note that by including γ, the agenda cannot contain any contradictions or tautologies. More precisely, both ψ_k^B, φ and their negations are satisfiable, even if every φ_i is a contradiction. The judges' profile consists of two judgments $P_i = \Phi_q \setminus \{\beta_i\} \cup \{\neg\beta_i\}$ for $i \in \{1, 2\}$ and the individual judgments' consistency is not violated, since $\sim\!\varphi$ is always satisfiable by any $\neg\beta_j$. Finally, we set $\Phi_+ = \Phi_q$ for $q \leq 1/3$ and $\Phi_- = \Phi_q$ otherwise. By construction it holds that $F_q(P) = \Phi_q$ and, slightly abusing notation, we consider any order π over Φ_q instead of Φ_+. Clearly, this construction can be done in polynomial time. Subsequently, we prove $\mathcal{I} \in \text{SSC} \Leftrightarrow \mathcal{I}' \in \text{SF}_q\text{D}$.

($\Rightarrow$) Assume $\mathcal{I}$ is a YES-instance. Consequently, there exists a set $N' = \{i_1, \ldots, i_m\} \subseteq [n]$ with $m \leq k$ such that $\bigvee_{i \in N'} \varphi_i \equiv \text{TRUE}$. As order we choose $\pi = (\beta_1, \beta_2, \psi_k^B \vee \neg\beta_1 \vee \neg\beta_2, x_{i_1}, \ldots, x_{i_m}, \sim\!\varphi, \ldots)$, where the order of the elements after $\sim\!\varphi$ is irrelevant. Applying the SF_q-rule, we may add each formula in the first $m+3$ iterations by using the quota rule F_q, since ψ_k^B and $m \leq k$ variables from B are satisfiable simultaneously, even if both β_j are set to TRUE. Now, we show that $\varphi = \psi_k^B \wedge \left[\left(\bigvee_{i \in [n]} \varphi_i'\right) \vee \gamma\right] \wedge \beta_1 \wedge \beta_2$ may be deduced from the initial assumption by showing that each formula in $\{\psi_k^B, \bigvee_{i \in [n]} \varphi_i', \beta_1, \beta_2\}$ can be deduced separately. First, note that each β_j trivially entails itself and $\beta_1 \wedge \beta_2 \wedge \left(\psi_k^B \vee \neg\beta_1 \vee \neg\beta_2\right) \models \psi_k^B$ holds. For the remaining formula it holds that

$$\begin{aligned}
&\bigwedge_{i \in N'} x_i \Rightarrow \bigvee_{i \in [n]} \varphi_i' \\
\Leftrightarrow\ &\bigvee_{i \in N'} \neg x_i \vee \bigvee_{i \in [n]} (\varphi_i \wedge x_i) \\
\Leftrightarrow\ &\bigvee_{i \in N'} ((\neg x_i \wedge \varphi_i) \vee (\neg x_i \wedge \sim\!\varphi_i)) \vee \bigvee_{i \in [n]} (\varphi_i \wedge x_i) \\
\Leftrightarrow\ &\bigvee_{i \in N'} \varphi_i \vee \bigvee_{i \in N'} (\neg x_i \wedge \sim\!\varphi_i) \vee \bigvee_{i \in [n] \setminus N'} (\varphi_i \wedge x_i), \qquad (1)
\end{aligned}$$

where the left disjunction in (1) already is a tautology by assumption. Consequently, it holds that $SF_q^{\sim\varphi}(P, \pi) \Rightarrow \varphi$. Hence, we conclude $\varphi \in SF_q(P, \pi)$, resulting in $\mathcal{I}' \in SF_q\text{D}$.

($\Leftarrow$) Assume $\mathcal{I}$ is a NO-instance. Consequently, there does not exist any $N' \subseteq [n]$ with $|N'| \leq k$, such that $\bigvee_{i \in N'} \varphi_i \equiv \text{TRUE}$ holds. By contradiction, we assume $\mathcal{I}'$ to still be a YES-instance. Then, there exists an order π over Φ_q such that

$$\varphi = \psi_k^B \wedge \left[\left(\bigvee_{i \in [n]} \varphi_i'\right) \vee \gamma\right] \wedge \beta_1 \wedge \beta_2 \in SF_q(P, \pi)$$

holds. We deduce that $\beta_1, \beta_2 \in SF_q(P, \pi)$ and $\psi_k^B \vee \neg\beta_1 \vee \neg\beta_2 \in SF_q(P, \pi)$ hold as well due to consistency. Hence, at most k of the variables x_i, $1 \leq i \leq n$, are satisfied. Let us denote the satisfied variables by $M = \{x_{i_1}, \dots, x_{i_{k'}}\}$ and the unsatisfied variables by $B \backslash M = \{x_{i_{k'+1}}, \dots, x_{i_n}\}$. Furthermore, we can imply the following out of $\varphi \in SF_q(P, \pi)$:

$$\begin{aligned}
\text{TRUE} &\equiv \left[\bigvee_{i\in[n]} \varphi_i'\right] \vee \gamma \equiv \left[\bigvee_{i\in[n]} (\varphi_i \wedge x_i)\right] \vee \gamma \\
&\equiv \left[\bigvee_{i\in M} (\varphi_i \wedge x_i)\right] \vee \left[\bigvee_{i\in[n]\backslash M} (\varphi_i \wedge x_i)\right] \vee \gamma \\
&\equiv \left[\bigvee_{i\in M} (\varphi_i \wedge \text{TRUE})\right] \vee \left[\bigvee_{i\in[n]\backslash M} (\varphi_i \wedge \text{FALSE})\right] \vee \gamma \\
&\equiv \left[\bigvee_{i\in M} \varphi_i\right] \vee \text{FALSE} \vee \gamma \equiv \left[\bigvee_{i\in M} \varphi_i\right] \vee \gamma.
\end{aligned}$$

Yet, we know that φ must have been entailed by previously added formulas because $\varphi \notin F_q(P)$. Hence, we conclude that for the given order π it holds that $SF_q^{\sim\varphi}(P, \pi) \models \left(\bigvee_{i\in M} \varphi_i\right) \vee \gamma$, although neither γ nor any φ_i shares any literals with formulas from $SF_q^{\sim\varphi}(P, \pi)$. Overall, $\left(\bigvee_{i\in M} \varphi_i\right) \vee \gamma$ can only be entailed if the disjunction contains a tautology. Since γ is a literal, this implies that $\bigvee_{i\in M} \varphi_i \equiv \text{TRUE}$ with $|M| \leq k$ would be a solution to $\mathcal{I}$, which is a contradiction to our assumption. Therefore, such an order π cannot exist and $\mathcal{I}'$ must be a NO-instance, too. $\square$

Proof of Theorem 8. Again, we refer to Lemma 1. Note that it is sufficient to identify k formulas in $\mathcal{K}(P)$ that may be satisfied simultaneously since we can put these at the front of the desired order π. This can be done by nondeterministically guessing an assignment over all literals in every formula of the agenda and afterwards verifying in P whether there are at least k formulas in $\mathcal{K}(P)$ satisfied by the given assignment. $\square$

Proof of Theorem 9. In order to show NP-hardness, we modify the proof of Theorem 4. This time we reduce a SAT instance $\mathcal{I} = (\alpha)$ to a SF_qSJ instance $\mathcal{I}' = (\Phi, P, 3)$, such that $F_q(P) = \{(\alpha \wedge \gamma) \vee \neg\beta_1 \vee \neg\beta_2, \beta_1, \beta_2\}$ holds. By construction $\mathcal{I}'$ is a YES-instance if and only if $F_q(P)$ is consistent, which in return is consistent if and only if α is satisfiable. $\square$